ROCHESTON®

THE ROCHESTON CYBERSECURITY FRAMEWORK

The Superset Standard for Global Cyber-Resilience

THE ROCHESTON CYBERSECURITY FRAMEWORK

This book is intended for informational and educational purposes only. The views expressed herein are the opinion of the author and should not be taken as professional advice. The author of this book and publisher are not responsible for any loss or damage resulting from the use of this book.

TABLE OF CONTENTS

WHAT IS RCF?

THE ROCHESTON CYBERSECURITY FRAMEWORK (RCF)

For more than thirty years, cybersecurity has been governed by fragmentation. Organizations have been forced to navigate dozens of overlapping frameworks, regional regulations, and industry mandates—each written in different language, structured around different assumptions, and assessed on different timelines. The result has not been stronger security, but duplicated effort, audit fatigue, and a widening gap between what is documented and what actually exists in production environments.

Most legacy frameworks were created for a world that no longer exists. They assume static infrastructure, clearly defined network perimeters, infrequent system change, and periodic audits as a meaningful measure of security. In modern environments—where cloud workloads scale by the minute, identities move continuously, software is deployed daily, and artificial intelligence systems operate autonomously—those assumptions no longer hold.

The Rocheston Cybersecurity Framework (RCF) was created to address this reality.

RCF is not a checklist, a maturity model, or a certification syllabus. It is a cyber-resilience operating framework designed to function in environments that change continuously, adversaries that adapt rapidly, and regulatory expectations that grow more complex every year. RCF treats cybersecurity as a living system—one that must be enforced, validated, and proven in real time rather than asserted once per audit cycle.

At its foundation, RCF is engineered as a superset framework. Instead of managing NIST, ISO/IEC 27001, SOC 2, PCI DSS, HIPAA, and regional regulations as separate programs, RCF harmonizes their overlapping requirements into a single, higher baseline control architecture. When an organization implements RCF correctly, it inherently satisfies the intent and technical safeguards of those standards simultaneously. Compliance becomes an outcome of how the organization operates every day, not a special project repeated for each regulator.

This principle is captured in a single directive:
Implement once. Comply everywhere.

RCF goes beyond unification. It addresses entire categories of risk that legacy standards either partially address or ignore entirely. These include autonomous AI agents, post-quantum cryptographic survivability, cognitive and psychological attack surfaces, orbital and space-dependent infrastructure, sustainable cybersecurity operations, and framework self-evolution. These are not speculative concerns. They are already shaping real-world incidents, regulatory pressure, and national security posture.

The framework is structured around 25 domains, organized into five strategic tiers. Each domain is written as an executable control system, not a theoretical guideline. Every domain defines a mission, architectural blueprint, operating model, and phased implementation roadmap. The roadmap is not aspirational; it produces concrete artifacts and verifiable evidence at each stage. This design ensures that RCF can be implemented, operated, and assessed consistently across organizations of different sizes, industries, and geographies.

RCF rejects the notion that security can be proven through policy statements or self-attestation. Auditors, regulators, courts, and boards do not certify intent—they certify evidence. For that reason, RCF is built around proof-grade security. Controls are validated continuously against live environments. Drift is detected as it occurs. Evidence is collected automatically and can be anchored to immutable records, creating a defensible history of security posture that cannot be silently altered or reconstructed after the fact.

Operationally, RCF is designed to run as a system. Intelligent automation validates control states, correlates signals across domains, and enables autonomous containment and recovery when defined thresholds are met. Human decision-making remains central, but no longer sits in the critical path for every defensive action. Security shifts from reactive response to engineered survivability.

This book is the authoritative specification of the Rocheston Cybersecurity Framework.

It documents the philosophy, structure, and execution of RCF as a unified standard. It explains why the framework exists, how it is architected, how each domain functions, how evidence is produced and preserved, and how the framework evolves over time. It is written to serve multiple audiences simultaneously: executives who require clarity and accountability, engineers who require precision and repeatability, auditors and regulators who require traceable proof, and organizations that must operate in a permanently contested digital environment.

RCF is not designed to replace existing standards through argument. It replaces them through execution. By implementing a single, coherent control architecture, organizations eliminate duplication, reduce risk, and gain resilience that endures beyond any single audit or regulation.

This is cybersecurity designed for continuity, not compliance theater.
This is security that can be proven, not claimed.
This is a framework built to survive change.

This is the Rocheston Cybersecurity Framework.

DOMAIN 1 —
GOVERNANCE & POLICY

MISSION

Governance & Policy is the domain that makes every other RCF domain real. It exists to ensure security is not optional, not dependent on individual heroics, and not something that only appears during audit season. When governance is weak, controls drift, decisions get overridden, exceptions pile up, and the organization ends up with compliance theater instead of resilience. This domain turns security from "advice" into authority, and from authority into enforceable outcomes.

What this domain prevents

This domain prevents the root causes behind most security failures. It prevents unclear ownership where nobody is accountable when a control breaks. It prevents policy sprawl where documents exist but controls are not enforced. It prevents silent risk acceptance where teams quietly "take the risk" without leadership approval. It prevents crisis-driven decision making during incidents because authority was never defined in advance. It prevents audit failures caused by inconsistent enforcement across departments, and it prevents long-term control decay when leadership changes, organizations restructure, or priorities shift.

What "done" looks like

Governance & Policy is done when the organization can prove that security decisions are deliberate, authorized, time-bound, and enforceable. You can immediately identify who owns each control, who approved any exceptions, why the exception exists, and when it expires. Leadership can see real control health rather than self-attestation, and auditors can trace decisions back to accountable executives. Most importantly, security decisions do not require emergency executive intervention because decision rights and escalation paths already exist and are practiced.

Scope boundaries

This domain includes the governance structure, decision authority, policy lifecycle, control ownership assignment, risk acceptance and exception governance, and executive or board-level reporting. It does not include the technical implementation of security controls, risk quantification modeling, identity enforcement mechanics, or incident response execution.

Governance & Policy defines direction and authority so technical execution remains consistent under pressure.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Governance & Policy sits above your technical security stack as the command layer that turns evidence into decisions and decisions into enforceable direction. The model is the same whether you are on-prem, cloud, or hybrid: policies and controls must be centrally managed, approvals must be routed through accountable identity-backed workflows, evidence must be continuously collected, and leadership must see measurable control states rather than narratives.

Required systems, data sources, and integrations

At minimum, you need a central governance repository that holds policies and maps them to controls, a workflow engine for approvals, an identity system that binds approvals to real accountable roles, an evidence platform such as Rocheston Noodles, and a leadership reporting view that shows posture and drift. This domain consumes live signals from the rest of RCF, including control state evidence, exception requests, risk acceptance records, audit findings, remediation status, and organizational ownership mappings. It integrates with AINA for continuous validation of whether real environments match governance expectations, and it uses Noodles as the system of record for evidence visibility and governance dashboards.

Data flows

In a properly built system, controls produce evidence continuously and that evidence flows into Noodles. AINA evaluates the evidence against governance expectations and highlights drift, noncompliance, or degradation as it occurs. Exception and risk decisions are then routed through approval workflows so leadership decisions are captured, traceable, and time-limited. The resulting decisions flow back into technical enforcement and operational remediation, ensuring governance is not separated from reality.

Minimum viable setup vs enterprise setup

A minimum viable setup establishes authority and repeatability: a governance charter, a named accountable security owner, a core policy baseline mapped to controls, manual exception approvals, and a quarterly governance review cadence. An enterprise setup makes governance scalable and defensible under pressure: policies mapped directly to controls, automated exception lifecycle with expiry and renewal, continuous validation through AINA, executive dashboards with trend analysis, and cross-jurisdiction alignment so regional compliance becomes a byproduct of daily operations rather than a separate program.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically the CISO or equivalent. Supporting roles include an executive sponsor such as the CIO, CTO, or COO, legal and compliance leadership for regulatory alignment, domain control owners across the RCF model, and RCCE governance engineers who operationalize governance into a living system that survives change.

Cadence

Governance must run on disciplined routines. Daily operations focus on critical policy violations, expiring exceptions, and governance items that directly increase active risk. Weekly routines focus on reviewing new exception requests and addressing high-risk deviations. Monthly routines focus on posture review, exception aging, and whether policy is producing real outcomes or just paperwork. Quarterly routines focus on executive governance review, reassessment of accepted risks, and approval of policy updates. Annual routines focus on resetting the governance baseline and delivering board-level assurance reporting that is backed by traceable evidence rather than statements of intent.

Required meetings and approvals

Governance should be lean but real. You run a monthly governance council to drive decisions and remove blockers. You require formal approvals for risk acceptance and exceptions because undocumented acceptance is the fastest path to unbounded exposure. You conduct quarterly or annual board or audit committee review depending on the organization's regulatory and operational risk profile, ensuring that material risk decisions are visible at the appropriate leadership level.

Escalation paths

Escalation must be defined before crisis conditions exist. Control violations escalate to the domain owner and control owner for immediate action. Unresolved issues escalate to the CISO for authority-based intervention. Business-impacting risk escalates to the executive sponsor for cross-functional decisions. Material or systemic risk escalates to the board or audit committee so leadership cannot claim surprise after the fact.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish governance authority and accountability so the system is real from day one. You appoint the accountable owner and executive sponsor, define decision rights and enforcement authority in a governance charter, inventory existing policies and controls, assign owners to controls, and set up identity-backed approval workflows. The required outputs are a signed governance charter, a policy and control inventory, a control ownership matrix, and documented workflow definitions. Evidence at the end of this phase is simple but decisive: governance is authorized, accountability is named, and decision authority exists in writing.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize governance and connect it to evidence. You map policies to RCF controls so policy becomes measurable, you implement exception and risk acceptance processes so risk is explicit and time-bound, you connect Rocheston Noodles to centralize evidence visibility, you enable AINA validation signals so governance reflects reality, and you build executive dashboards that show control health and drift. Outputs include policy-to-control mappings, exception records, dashboards, and validation reports. Evidence at the end of this phase includes active approvals and traceable decisions, supported by continuous validation outputs rather than periodic attestations.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make governance survivable under pressure and defensible over time. You enforce exception expiry so risk does not become permanent by accident, you test escalation paths so leadership response is practiced, you run governance failure scenarios so authority is exercised before emergencies, you validate audit readiness, and you anchor governance artifacts into Rosecoin Vault so evidence and decisions cannot be silently altered. Outputs include escalation test results, a governance maturity assessment, and immutable governance evidence records. Evidence at the end of this phase includes time-bound risk acceptance, verified enforcement actions, and tamper-proof proof of governance integrity.

End state

When this domain is complete, Governance & Policy is no longer a document set. It becomes a living control system that continuously directs security, enforces accountability, and proves resilience every day, regardless of personnel changes, organizational pressure, or audit cycles.

DOMAIN 2 — RISK QUANTIFICATION & VALUE

DOMAIN MISSION AND OUTCOMES

Mission

Risk Quantification & Value exists to make cybersecurity risk understandable, comparable, and actionable at the business level. Its purpose is to translate technical security conditions into financial, operational, and strategic impact so leadership can make informed decisions instead of relying on intuition, fear, or generic risk labels. This domain ensures that security investment, risk acceptance, and prioritization are based on measurable value and loss exposure, not gut feeling.

What this domain prevents

This domain prevents security teams from operating in a vacuum where risks are labeled “high,” “medium,” or “low” without anyone understanding what that actually means to the business. It prevents over-investment in low-impact issues and under-investment in existential risks. It prevents leadership from approving risk blindly because the consequences were never quantified. It prevents endless debates between security and business teams caused by incompatible language. Most importantly, it prevents organizations from discovering the true cost of cyber risk only after an incident has already occurred.

What “done” looks like

This domain is done when cyber risk is expressed in terms leadership already understands: financial loss, operational disruption, legal exposure, customer harm, and strategic impact. You can compare one risk against another using a common unit of measure. You can justify security investment using risk reduction, not fear. Risk acceptance decisions are tied to quantified exposure and expected loss, and leadership can clearly see how much risk the organization is carrying at any point in time. When asked why a control was prioritized or deferred, the answer is backed by numbers, assumptions, and evidence.

Scope boundaries

This domain includes cyber risk modeling, loss estimation, impact analysis, risk prioritization, and value-based decision support. It covers how risk is measured, compared, and communicated. It does not include governance authority (Domain 1), technical control implementation, threat intelligence collection, or incident response execution. Risk Quantification informs decisions; it does not enforce them.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Risk Quantification & Value operates as an analytical layer that consumes data from across the security and business environment. It does not sit inside a single tool; it spans security telemetry, asset inventories, business systems, and governance workflows. Whether on-prem, cloud, or hybrid, the architecture must support continuous ingestion of risk signals and consistent modeling of impact.

Required systems, data sources, and integrations

This domain requires an asset inventory with business criticality, security telemetry from other RCF domains, vulnerability and exposure data, incident and loss history, and business data such as revenue streams, regulatory obligations, and operational dependencies. It integrates with Rocheston Noodles as the central analysis and reporting platform and uses AINA to correlate technical risk signals with business context. Governance systems consume the outputs to support risk acceptance and prioritization decisions.

Data flows

Technical signals such as vulnerabilities, control gaps, and exposure levels flow into Noodles. Business context such as asset value, service criticality, and regulatory impact is layered on top. AINA models potential loss scenarios and expected impact based on real evidence rather than static assumptions. The resulting quantified risk outputs flow into governance workflows, investment planning, and executive reporting so decisions are made with full visibility into trade-offs.

Minimum viable setup vs enterprise setup

A minimum viable setup establishes basic financial visibility: an asset inventory with criticality ratings, a simple loss estimation model, manual risk scoring translated into financial ranges, and periodic risk reporting to leadership. An enterprise setup enables continuous, defensible decision-making: automated ingestion of technical signals, dynamic risk modeling updated as environments change, integration with governance for risk acceptance tracking, and executive dashboards showing total risk exposure, risk reduction trends, and return on security investment.

DOMAIN OPERATING MODEL

Roles and ownership
This domain requires a single accountable owner, typically the CISO or a designated risk executive, with strong partnership from finance leadership. Supporting roles include security architecture and operations teams who provide technical inputs, business owners who validate impact assumptions, finance teams who validate loss modeling, and RCCE engineers who operationalize quantification into repeatable systems.

Cadence
Daily operations focus on monitoring significant changes in exposure that materially affect risk posture. Weekly routines focus on updating models based on new vulnerabilities, asset changes, or threat conditions. Monthly routines focus on reviewing risk trends, validating assumptions, and aligning with governance decisions. Quarterly routines focus on executive risk reviews, budget alignment, and reassessment of accepted risk. Annual routines focus on recalibrating models using real incident data and business changes.

Required meetings and approvals
This domain requires regular risk review sessions aligned with governance cadence, executive reviews when risk acceptance or major investment decisions are needed, and periodic alignment with finance to ensure models remain credible. Meetings exist to make decisions, not to debate definitions endlessly.

Escalation paths
Material increases in quantified risk escalate to the accountable risk owner. Risks exceeding predefined tolerance thresholds escalate to executive leadership for acceptance or mitigation decisions. Risks with regulatory, safety, or existential impact escalate to the board or audit committee with quantified exposure clearly presented.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish the foundation for measurable risk. You identify the accountable owner, define the risk quantification approach, inventory critical assets with business owners, collect baseline security and exposure data, and define initial loss categories and assumptions. Outputs include a documented risk model, an asset criticality register, baseline risk estimates, and defined reporting formats. Evidence at the end of this phase shows that risk is being measured consistently and transparently.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize quantification and connect it to reality. You integrate technical telemetry and vulnerability data into Noodles, enrich risk models with business and financial context, enable AINA-driven correlation and scenario modeling, and begin producing regular quantified risk reports. Outputs include dynamic risk models, updated loss scenarios, executive dashboards, and risk acceptance records tied to quantified exposure. Evidence demonstrates that risk values update as environments change.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make risk quantification defensible and trusted. You validate models against real incidents and near-misses, stress-test assumptions, align tolerance thresholds with leadership, and anchor risk decisions and supporting evidence into Rosecoin Vault. Outputs include validated models, risk trend analyses, tolerance thresholds, and immutable records of risk acceptance and prioritization decisions. Evidence proves that risk-based decisions are repeatable, auditable, and aligned with real-world outcomes.

End state

When this domain is mature, cybersecurity risk is no longer abstract or emotional. It becomes a measurable business variable that leadership can understand, compare, and manage deliberately. Security investment is justified through risk reduction, risk acceptance is explicit and time-bound, and the organization understands not just where it is vulnerable, but what that vulnerability actually costs.

DOMAIN 3 — THIRD-PARTY & SUPPLY CHAIN SECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Third-Party & Supply Chain Security exists to prevent your organization from being compromised by someone else's weaknesses. Modern enterprises rarely fail because of a single internal control failure; they fail because a vendor, partner, contractor, software supplier, or service provider became the attack path. This domain ensures that external dependencies are governed, assessed, monitored, and constrained so trust is never assumed and risk does not silently enter the organization through the supply chain.

What this domain prevents

This domain prevents breaches that originate from vendors with poor security hygiene, unmanaged contractors with excessive access, compromised software updates, and opaque fourth-party dependencies. It prevents organizations from relying on one-time questionnaires and outdated attestations that no longer reflect real security posture. It prevents blind spots where leadership does not know which third parties have access to sensitive systems, data, or operations. It also prevents regulatory exposure caused by unverified vendors handling regulated data, and operational outages caused by supply chain disruption or dependency failure.

What "done" looks like

This domain is done when third-party risk is visible, measured, and actively managed throughout the vendor lifecycle. You know which vendors matter, why they matter, and what access they have. Security requirements are enforced before access is granted, not after an incident. Vendor risk is continuously reassessed rather than reviewed once a year. Offboarding is as controlled as onboarding, and supply chain dependencies are documented and monitored. When a vendor fails, you can immediately determine impact, exposure, and containment actions.

Scope boundaries

This domain includes vendor security governance, third-party risk assessment, onboarding and offboarding controls, access and data exposure governance for external parties, supply chain dependency mapping, and continuous monitoring of vendor risk posture. It does not include identity enforcement mechanics themselves, internal asset security, or incident response execution, which are addressed in other RCF domains. This domain governs external trust and dependency risk.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Third-Party & Supply Chain Security spans governance, identity, data, and monitoring layers. Architecturally, it requires a central system that understands vendors as risk-bearing entities, not just procurement records. Whether the environment is on-prem, cloud, or hybrid, the architecture must support continuous visibility into vendor access, data exposure, and dependency chains, including software supply chain components.

Required systems, data sources, and integrations

This domain requires a vendor inventory with ownership and criticality, contract and data classification records, identity and access systems showing third-party access, software inventory and SBOM data, security assessment results, and monitoring telemetry related to vendor activity. It integrates with Rocheston Noodles as the system of record for third-party evidence and risk tracking, and with AINA to correlate vendor behavior, exposure, and changes in posture. Integration with procurement, legal, identity platforms, and asset inventories is essential so vendor risk is tied to real access and data paths.

Data flows

Vendor information enters the system at onboarding, including business purpose, data access, and dependency classification. Security requirements and assessments are applied before access is granted. Access and activity data flows continuously from identity systems, networks, and applications into Noodles. AINA evaluates changes in access, behavior, or external risk signals and flags drift or emerging exposure. Decisions about continued access, remediation, or termination flow back into governance and enforcement systems so third-party trust remains conditional and monitored.

Minimum viable setup vs enterprise setup

A minimum viable setup includes a centralized vendor inventory, defined security requirements for third parties, basic risk categorization, manual assessment workflows, and documented onboarding and offboarding processes. An enterprise setup includes continuous monitoring of vendor posture, automated enforcement of access controls, software supply chain visibility through SBOMs, fourth-party dependency awareness, integration with governance for risk acceptance, and executive dashboards showing aggregated supply chain risk.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically within security or risk leadership, with strong coordination across procurement, legal, IT, and business owners. Supporting roles include vendor managers, application owners, identity administrators, legal and compliance teams, and RCCE engineers who operationalize supply chain controls into enforceable systems rather than manual checklists.

Cadence

Daily operations focus on monitoring high-risk vendor activity, access anomalies, and newly identified exposure. Weekly routines focus on reviewing new vendor requests, changes in vendor scope, and remediation progress for identified gaps. Monthly routines focus on reassessing vendor criticality, reviewing access rights, and validating ongoing compliance with security requirements. Quarterly routines focus on executive review of supply chain risk concentration, critical dependencies, and accepted vendor risk. Annual routines focus on full vendor portfolio reassessment, contract security updates, and strategic dependency review.

Required meetings and approvals

Vendor onboarding approvals must include security sign-off before access is granted. Risk acceptance for high-risk vendors requires executive approval with documented justification. Periodic cross-functional reviews align security, procurement, and legal to ensure vendor risk decisions remain current as business needs change.

Escalation paths

Third-party security failures escalate to the domain owner for immediate assessment. Risks that affect critical systems or regulated data escalate to executive leadership for decision and containment. Systemic or strategic supply chain risks escalate to the board or audit committee with clear articulation of dependency, exposure, and business impact.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)
In Phase 1 you establish visibility and authority over third-party risk. You assign the accountable owner, inventory all vendors and external dependencies, classify vendors by criticality and data access, define minimum security requirements, and establish onboarding and offboarding workflows. Outputs include a vendor inventory, criticality classifications, security requirement baseline, and documented approval processes. Evidence at the end of this phase shows that third-party access is no longer unmanaged or undocumented.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize third-party security controls. You integrate identity and access data into Noodles, connect monitoring sources that capture vendor activity, implement assessment and reassessment processes, enable AINA-driven correlation of vendor behavior and risk signals, and begin producing regular supply chain risk reports. Outputs include active assessment records, access reviews, monitoring dashboards, and documented risk decisions. Evidence demonstrates that vendor risk is being evaluated continuously rather than periodically.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make supply chain security resilient and defensible. You enforce least-privilege access for third parties, test offboarding and termination procedures, validate software supply chain integrity, stress-test dependency scenarios, and anchor vendor risk decisions and evidence into Rosecoin Vault. Outputs include validated offboarding tests, dependency maps, risk trend analyses, and immutable records of vendor approvals and risk acceptance. Evidence proves that third-party trust is conditional, monitored, and reversible.

End state
When this domain is mature, third parties are no longer invisible risk multipliers. They are governed participants in a controlled ecosystem where access is justified, behavior is monitored, dependencies are understood, and trust is continuously verified. Supply chain

security becomes a managed business discipline rather than a recurring source of surprise and damage.

DOMAIN 4 — IDENTITY & ACCESS MANAGEMENT

DOMAIN MISSION AND OUTCOMES

Mission

Identity & Access Management exists to ensure that only the right identities can access the right systems, data, and actions—at the right time, for the right reason, and for no longer than necessary. In modern environments, identity is the primary security perimeter. Networks, devices, and locations no longer define trust; identities do. This domain ensures that identity becomes a controlled, verifiable, and continuously enforced security boundary rather than an administrative convenience.

What this domain prevents

This domain prevents breaches caused by stolen credentials, excessive privileges, shared accounts, unmanaged service identities, and long-lived access that no longer matches business need. It prevents lateral movement after initial compromise, privilege escalation through misconfigured roles, and shadow access created by contractors, automation, or legacy systems. It also prevents regulatory exposure caused by weak access controls and audit findings driven by unclear or unreviewed permissions.

What "done" looks like

Identity & Access Management is done when access is deliberate, minimal, and provable. Every identity—human or machine—has a known owner, a defined purpose, and a bounded lifetime. Access is granted based on verified identity and context, not static trust. Privileged actions require additional verification, and access reviews demonstrate that permissions reflect current business need. When an identity is compromised or no longer needed, access can be revoked quickly and completely across the environment.

Scope boundaries

This domain includes identity lifecycle management, authentication strength, authorization models, privileged access governance, service and machine identities, and access review processes. It does not include network segmentation, endpoint hardening, or application security logic, which are addressed in other RCF domains. Identity & Access Management governs who can access what and under what conditions.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Identity & Access Management spans cloud, on-prem, and hybrid environments and must function consistently across all of them. The architecture centers on a primary identity provider integrated with applications, infrastructure, and administrative systems. It enforces strong authentication, centralized authorization, and continuous verification regardless of where workloads reside.

Required systems, data sources, and integrations

This domain requires a centralized identity provider, multi-factor authentication mechanisms, role and attribute-based access control systems, privileged access management capabilities, and identity lifecycle workflows. It consumes data from HR systems, contractor management, application inventories, infrastructure platforms, and audit logs. Integration with Rocheston Noodles provides centralized evidence of access enforcement and review, while AINA evaluates access patterns, privilege usage, and drift from expected behavior.

Data flows

Identity records originate from authoritative sources such as HR or contractor systems and flow into the identity provider. Access requests and approvals flow through governance workflows before permissions are granted. Authentication and authorization events flow continuously into Noodles, where AINA evaluates usage patterns and identifies anomalies, excessive privilege, or stale access. Decisions about revocation, elevation, or remediation flow back into identity systems so enforcement remains continuous.

Minimum viable setup vs enterprise setup

A minimum viable setup includes a central identity provider, enforced multi-factor authentication, basic role-based access control, documented joiner-mover-leaver processes, and periodic manual access reviews. An enterprise setup extends this with phishing-resistant authentication, attribute-based and context-aware access, automated provisioning and deprovisioning, privileged access session controls, continuous access evaluation, and executive dashboards showing access risk and hygiene trends.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically the identity or security architecture leader, with close coordination across IT operations, application owners, HR, and security operations. Supporting roles include identity administrators, application owners responsible for access models, compliance teams validating reviews, and RCCE engineers who design and enforce identity controls at scale.

Cadence

Daily operations focus on monitoring authentication failures, anomalous access attempts, and privileged activity. Weekly routines focus on reviewing access changes, onboarding and offboarding events, and remediation of identified issues. Monthly routines focus on access review cycles for high-risk systems and privileged roles. Quarterly routines focus on broader access certification, role model validation, and reduction of privilege sprawl. Annual routines focus on reviewing identity architecture, authentication strength, and alignment with evolving threats.

Required meetings and approvals

Access approvals must be formalized for privileged roles and sensitive systems. Periodic access review sign-offs are required from system owners. Changes to authentication standards or role models require governance approval to ensure consistency and risk alignment.

Escalation paths

Suspicious or excessive access activity escalates to security operations for immediate investigation. Unresolved access violations escalate to the domain owner for authority-based enforcement. Identity risks affecting critical systems or regulated data escalate to executive leadership for decision and containment. Systemic identity failures escalate to the board or audit committee when they materially affect organizational risk.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish control over identities and access. You identify the accountable owner, select or validate the primary identity provider, enforce baseline authentication requirements, inventory applications and systems requiring access control, and define joiner-mover-leaver workflows. Outputs include an identity architecture overview, access policy definitions, system inventory, and documented lifecycle processes. Evidence at the end of this phase shows that identities are centrally managed and access decisions are no longer ad hoc.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize strong access control. You integrate applications and infrastructure with the identity provider, enable multi-factor authentication broadly, implement role or attribute-based access models, establish privileged access governance, connect authentication and authorization logs into Noodles, and enable AINA analysis of access patterns. Outputs include integrated systems, active access models, monitoring dashboards, and access review records. Evidence demonstrates that access is enforced consistently and monitored continuously.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make identity security resilient and defensible. You enforce least privilege across roles, remove stale and excessive access, test offboarding and credential revocation, implement phishing-resistant authentication for high-risk roles, and anchor access review and enforcement evidence into Rosecoin Vault. Outputs include completed access certifications, privilege reduction metrics, test results for revocation scenarios, and immutable access governance records. Evidence proves that identity-based access is controlled, monitored, and provable.

End state

When this domain is mature, identity is no longer the weakest link in the security chain. It becomes a continuously enforced control plane where access is intentional, temporary, and verifiable. The organization can confidently answer who has access, why they have it, and how quickly that access can be removed—without guesswork or delay.

DOMAIN 5 — PRIVACY & DATA PROTECTION

DOMAIN MISSION AND OUTCOMES

Mission

Privacy & Data Protection exists to ensure that sensitive data is collected, processed, stored, shared, and retained in ways that are lawful, minimal, controlled, and provable. This domain protects the organization from regulatory penalties, legal exposure, customer harm, and loss of trust by treating data as a governed asset rather than an uncontrolled byproduct of systems. It ensures that privacy is not a policy statement but an operational reality embedded into how data flows through the organization.

What this domain prevents

This domain prevents data breaches that expose personal, regulated, or confidential information. It prevents regulatory violations caused by unlawful processing, excessive data collection, unclear consent, or improper retention. It prevents "unknown data sprawl" where sensitive data exists in places nobody expects or owns. It prevents over-retention that increases breach impact and litigation risk, and it prevents under-protection where encryption, access control, or segregation is missing or inconsistent. It also prevents organizations from discovering privacy failures only after regulators, customers, or courts intervene.

What "done" looks like

Privacy & Data Protection is done when the organization knows what data it has, why it has it, where it lives, who can access it, and how long it is allowed to exist. Data handling aligns with legal and contractual obligations by design, not exception. Sensitive data is protected through classification, encryption, access control, and minimization. Data subject rights can be fulfilled reliably and on time. When a breach or inquiry occurs, the organization can demonstrate lawful processing, appropriate safeguards, and controlled data lifecycle with evidence rather than assurances.

Scope boundaries

This domain includes data classification, privacy governance, lawful processing controls, consent and purpose limitation, encryption and key management oversight, data retention and deletion, and data subject rights support. It does not include identity enforcement mechanics, application security logic, or incident response execution, which are covered in other RCF domains. Privacy & Data Protection governs how data is handled and protected across its entire lifecycle.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Privacy & Data Protection spans every environment where data exists. The architecture must support consistent data handling controls across on-prem, cloud, SaaS, and hybrid systems. It centers on understanding data flows and enforcing protection at creation, storage, use, transfer, and deletion. Privacy controls must be embedded into platforms and workflows rather than layered on afterward.

Required systems, data sources, and integrations

This domain requires data inventories and classification systems, encryption and key management platforms, access control systems tied to identity, logging and monitoring for data access and movement, and mechanisms to enforce retention and deletion. It consumes data from applications, databases, storage platforms, cloud services, and third-party processors. Integration with Rocheston Noodles provides centralized visibility into data protection evidence, while AINA correlates data access patterns, policy adherence, and potential privacy risk across systems.

Data flows

Data is classified at creation or ingestion and tagged with sensitivity, purpose, and retention requirements. Access and usage events flow into Noodles for visibility and validation. AINA evaluates whether data is accessed and processed in ways consistent with declared purpose and policy. Retention and deletion actions are executed and logged as evidence. When data is shared internally or externally, those flows are recorded and governed so privacy obligations travel with the data.

Minimum viable setup vs enterprise setup

A minimum viable setup includes a basic data inventory, defined data classifications, encryption for sensitive data, documented retention policies, and manual handling of privacy requests. An enterprise setup provides automated discovery and classification, enforced encryption and key governance, continuous monitoring of data access, automated retention and deletion, integrated support for data subject rights, and executive dashboards showing data risk and compliance posture across regions.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a data protection or security leader, working closely with legal and compliance leadership. Supporting roles include data owners, application owners, platform teams, privacy officers, and RCCE engineers who operationalize privacy requirements into enforceable technical controls.

Cadence

Daily operations focus on monitoring sensitive data access, anomalous data movement, and potential policy violations. Weekly routines focus on reviewing new data sources, changes in processing purpose, and remediation of identified gaps. Monthly routines focus on validating data inventories, reviewing retention status, and assessing third-party data handling. Quarterly routines focus on executive review of privacy posture, regulatory alignment, and high-risk data processing activities. Annual routines focus on updating data classifications, retention schedules, and privacy governance in response to regulatory and business changes.

Required meetings and approvals

New data collection or processing activities require privacy review and approval before launch. Changes to retention, purpose, or data sharing require documented approval. Periodic executive or legal reviews ensure that privacy risk remains aligned with organizational tolerance and regulatory obligations.

Escalation paths

Privacy control failures or suspected violations escalate to the domain owner and legal leadership for immediate assessment. Risks involving regulated data, cross-border transfers, or potential harm escalate to executive leadership. Material or systemic privacy risk escalates to the board or audit committee, supported by evidence of exposure and impact.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)
In Phase 1 you establish visibility and authority over data. You identify the accountable owner, define data classifications and privacy principles, inventory critical data assets and processing activities, establish retention requirements, and document privacy governance workflows. Outputs include a data inventory, classification scheme, retention policy, and privacy approval processes. Evidence at the end of this phase shows that data is no longer unmanaged or unowned.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize data protection controls. You implement encryption and key management, integrate access logging into Noodles, enable AINA analysis of data access and movement, enforce retention controls where possible, and begin producing privacy posture reports. Outputs include protected data stores, monitoring dashboards, access records, and privacy decision logs. Evidence demonstrates that data handling is actively governed and monitored rather than assumed.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make privacy defensible and sustainable. You test data deletion and retention enforcement, validate cross-border data handling, exercise data subject rights processes, and anchor privacy evidence into Rosecoin Vault. Outputs include test results, validated workflows, compliance assessments, and immutable privacy records. Evidence proves that privacy controls work in practice and can withstand regulatory scrutiny.

End state
When this domain is mature, privacy and data protection are no longer reactive or manual. Data is handled intentionally, protected consistently, and governed throughout its lifecycle. The organization can demonstrate compliance, minimize breach impact, and maintain trust because it knows its data—and controls it—by design.

DOMAIN 6 — AI SECURITY & ML GOVERNANCE

DOMAIN MISSION AND OUTCOMES

Mission

AI Security & ML Governance exists to ensure that artificial intelligence and machine learning systems are trustworthy, controlled, explainable, and safe to operate at scale. As AI systems increasingly influence decisions, automate actions, and interact with critical data and infrastructure, they become high-impact attack surfaces. This domain ensures AI behaves as intended, remains aligned with organizational values and regulatory expectations, and does not become an unmanaged source of risk.

What this domain prevents

This domain prevents model manipulation, data poisoning, prompt injection, unauthorized model access, and silent degradation of AI behavior over time. It prevents AI systems from operating without accountability, transparency, or oversight. It prevents regulatory exposure caused by opaque decision-making, biased outcomes, or unexplainable automation. It also prevents organizations from deploying AI systems that drift from their original purpose, misuse sensitive data, or act outside approved boundaries.

What "done" looks like

AI Security & ML Governance is done when every AI system has a defined owner, documented purpose, controlled data inputs, and measurable performance boundaries. Models are versioned, monitored, and validated continuously. Decisions made or influenced by AI can be explained and traced back to inputs, logic, and approvals. Unauthorized model changes are detected, and unsafe behavior triggers containment or rollback. Leadership understands where AI is used, what it is allowed to do, and how its risk is managed.

Scope boundaries

This domain includes AI and ML lifecycle governance, model integrity and validation, training data governance, inference monitoring, access control for models and prompts, explainability requirements, and safe-use enforcement. It does not include general application security, infrastructure hardening, or business process design. AI Security & ML Governance governs how intelligent systems are built, operated, and constrained.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

AI Security & ML Governance spans development, deployment, and operational environments. The architecture must support secure model training, controlled deployment, monitored inference, and continuous validation. It applies equally to on-prem, cloud, and hybrid AI stacks, including internally developed models and externally sourced AI services.

Required systems, data sources, and integrations

This domain requires model repositories with version control, training data governance systems, access control for model usage, logging of inference and decision outputs, and monitoring tools that track model behavior over time. It consumes data from ML pipelines, data platforms, application logs, and user interactions. Integration with Rocheston Noodles provides centralized evidence of model governance, while AINA evaluates model behavior, drift, anomalies, and policy compliance.

Data flows

Training data flows into controlled pipelines where provenance, quality, and approval are recorded. Models are trained, versioned, and deployed with documented purpose and constraints. Inference events and decisions flow into Noodles for visibility. AINA continuously evaluates outputs against expected behavior, fairness metrics, performance thresholds, and security policies. When anomalies or violations are detected, alerts and governance actions flow back into deployment controls and approval workflows.

Minimum viable setup vs enterprise setup

A minimum viable setup includes an inventory of AI systems, defined ownership, basic access controls, documented training data sources, and periodic model review. An enterprise setup includes automated model versioning, continuous drift detection, explainability tooling, enforced approval gates for deployment, real-time monitoring of inference behavior, and executive dashboards showing AI risk posture across the organization.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, often a security or AI governance leader, working closely with data science leadership. Supporting roles include ML engineers, data owners, legal and ethics representatives, application owners, and RCCE engineers who embed governance into AI pipelines and platforms.

Cadence

Daily operations focus on monitoring AI behavior, performance anomalies, and security alerts. Weekly routines focus on reviewing changes to models, data sources, or usage patterns. Monthly routines focus on validating model performance, fairness, and alignment with approved purpose. Quarterly routines focus on executive review of AI risk, regulatory alignment, and major model changes. Annual routines focus on reassessing AI governance standards and adapting to new regulatory or technological developments.

Required meetings and approvals

Deployment of new AI models or major changes requires formal approval. Use of AI in high-impact or regulated decisions requires documented governance review. Periodic cross-functional reviews ensure AI systems remain aligned with organizational values and legal expectations.

Escalation paths

Detected AI security incidents or unsafe behavior escalate to the domain owner and security operations for immediate containment. Risks affecting customers, regulated data, or critical decisions escalate to executive leadership. Systemic AI governance failures escalate to the board or audit committee with clear articulation of impact and exposure.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish visibility and ownership of AI systems. You identify the accountable owner, inventory all AI and ML models in use, document their purpose and data sources, define baseline governance requirements, and establish approval workflows. Outputs include an AI system inventory, ownership assignments, governance standards, and documented decision rights. Evidence at the end of this phase shows that AI usage is known, owned, and governed.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize AI security controls. You implement model versioning and access control, integrate logging and monitoring into Noodles, enable AINA analysis of model behavior and drift, enforce approval gates for deployment, and begin producing AI governance reports. Outputs include monitored models, validation dashboards, approval records, and documented performance metrics. Evidence demonstrates that AI behavior is actively observed and controlled.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make AI governance resilient and defensible. You stress-test models against misuse and adversarial scenarios, validate explainability and auditability, test rollback and containment procedures, and anchor AI governance evidence into Rosecoin Vault. Outputs include test results, drift analysis, validated controls, and immutable governance records. Evidence proves that AI systems are safe to operate, monitored continuously, and governed with rigor.

End state

When this domain is mature, AI is no longer a black box or an unmanaged experiment. It becomes a controlled, explainable, and continuously validated capability. The organization can innovate with AI confidently, knowing that intelligent systems are governed, secure, and aligned with both business goals and societal expectations.

DOMAIN 7 — NETWORK, 5G & EDGE SECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Network, 5G & Edge Security exists to ensure that connectivity never becomes implicit trust. As organizations move beyond traditional data centers into cloud, 5G, edge computing, and distributed environments, the network transforms from a protected perimeter into a dynamic fabric that connects users, devices, services, and locations everywhere. This domain ensures that network access is intentional, segmented, monitored, and continuously verified so connectivity enables the business without becoming an open attack surface.

What this domain prevents

This domain prevents lateral movement after initial compromise, uncontrolled east-west traffic, exposed services, and implicit trust based on network location. It prevents attackers from using flat networks, legacy VPN assumptions, or unmanaged edge nodes to pivot across environments. It prevents outages caused by misconfigured routing, insecure edge deployments, or poorly governed 5G integrations. It also prevents regulatory exposure and data leakage caused by unsegmented traffic and unmonitored network paths.

What "done" looks like

Network, 5G & Edge Security is done when network access is explicitly authorized, segmented by risk, and continuously observed. No system, device, or service can communicate simply because it is "on the network." Traffic paths are intentional and documented. Network controls adapt to changes in identity, device posture, and workload context. When anomalies occur, they are detected quickly and contained without requiring manual reconstruction of network flows.

Scope boundaries

This domain includes network architecture, segmentation models, zero trust enforcement at the network layer, secure connectivity for 5G and edge environments, traffic inspection, and network telemetry. It does not include identity lifecycle management, endpoint hardening, or application-layer security logic, which are covered in other RCF domains. This domain governs how systems communicate and how trust is enforced across connections.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Network, 5G & Edge Security is built on zero trust principles applied to connectivity. The architecture assumes no implicit trust based on location and enforces segmentation across data centers, cloud environments, 5G networks, and edge deployments. Control planes must be centralized, while enforcement points are distributed close to workloads, users, and devices to minimize blast radius and latency.

Required systems, data sources, and integrations

This domain requires network segmentation and policy enforcement platforms, secure gateways or software-defined perimeter components, firewalls, traffic inspection capabilities, and telemetry collection systems. It consumes data from identity providers, endpoint posture systems, cloud platforms, network devices, 5G infrastructure, and edge nodes. Integration with Rocheston Noodles centralizes network evidence and visibility, while AINA correlates traffic patterns, policy adherence, and anomalies across distributed environments.

Data flows

Network devices, gateways, and edge nodes generate continuous telemetry about traffic, connections, and enforcement decisions. This data flows into Noodles for aggregation and visibility. AINA evaluates traffic patterns against expected behavior, identity context, and segmentation policies to identify drift, misconfiguration, or malicious activity. When policy violations or anomalies are detected, enforcement actions flow back into network controls to isolate, block, or reroute traffic automatically or with approval.

Minimum viable setup vs enterprise setup

A minimum viable setup includes documented network zones, basic segmentation between critical environments, firewall enforcement, monitored ingress and egress points, and visibility into core traffic flows. An enterprise setup includes identity-aware network controls, microsegmentation, encrypted traffic inspection where appropriate, secure 5G and edge integration, continuous validation of segmentation effectiveness, and executive dashboards showing network risk and exposure trends.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a network or security architecture leader, with close coordination across network operations, cloud teams, and security operations. Supporting roles include network engineers, cloud and edge platform teams, identity and endpoint teams for context integration, and RCCE engineers who design and maintain zero trust network architectures.

Cadence

Daily operations focus on monitoring network anomalies, policy violations, and availability issues across core, 5G, and edge environments. Weekly routines focus on reviewing segmentation changes, access requests, and remediation of identified weaknesses. Monthly routines focus on validating network architecture alignment with zero trust principles and reviewing traffic trends. Quarterly routines focus on executive review of network risk, dependency exposure, and resilience posture. Annual routines focus on architecture refresh, technology evaluation, and adaptation to new connectivity models.

Required meetings and approvals

Changes to network segmentation, core routing, or connectivity models require formal review and approval due to their potential impact. Integration of new 5G or edge deployments requires security sign-off before production use. Periodic cross-functional reviews ensure network design remains aligned with identity, endpoint, and application security models.

Escalation paths

Detected network intrusions or segmentation failures escalate immediately to security operations for containment. Unresolved or systemic issues escalate to the domain owner for authority-driven remediation. Network risks affecting critical services or regulated data escalate to executive leadership. Large-scale or persistent exposure escalates to the board or audit committee with clear articulation of business impact.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)
In Phase 1 you establish visibility and baseline control over network connectivity. You identify the accountable owner, inventory network assets including 5G and edge components, document existing network zones and trust assumptions, define segmentation objectives, and enable basic traffic logging. Outputs include a network architecture overview, zone definitions, initial segmentation policies, and logging configurations. Evidence at the end of this phase shows that network trust is understood and no longer implicit.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize zero trust networking. You implement segmentation and identity-aware controls, integrate network telemetry into Noodles, enable AINA analysis of traffic behavior, secure ingress and egress points, and formalize change management for network policies. Outputs include enforced segmentation, monitored traffic flows, dashboards, and documented policy approvals. Evidence demonstrates that network controls actively restrict and observe connectivity.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make network security resilient and adaptive. You test segmentation effectiveness, simulate lateral movement scenarios, validate containment actions, secure 5G and edge deployments under failure conditions, and anchor network security evidence into Rosecoin Vault. Outputs include test results, architecture validation reports, and immutable network governance records. Evidence proves that network trust is enforced continuously and can withstand attack and change.

End state
When this domain is mature, the network is no longer a soft perimeter or a hidden liability. It becomes a controlled, segmented, and observable fabric that connects systems safely across core, cloud, 5G, and edge environments. Connectivity enables speed and scale without sacrificing security, and trust is earned continuously rather than assumed.

DOMAIN 8 — ENDPOINT, DEVICE & IOT SECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Endpoint, Device & IoT Security exists to ensure that every device that touches your environment is known, controlled, hardened, monitored, and removable. Endpoints are where real attacks land because they are the closest surface to humans, credentials, and daily work. IoT and specialized devices expand that surface even further, often with weaker controls and longer lifecycles. This domain ensures devices cannot become silent entry points, persistent footholds, or unmanaged risk that undermines every other control.

What this domain prevents

This domain prevents malware execution, ransomware spread, credential theft from infected devices, and persistence through unmanaged endpoints. It prevents compromise through unpatched systems, insecure configurations, shadow devices, and uncontrolled USB or peripheral access. It prevents IoT devices from becoming invisible backdoors due to weak authentication, outdated firmware, or poor segmentation. It also prevents operational disruption caused by device sprawl, unmanaged assets, and inconsistent security baselines across departments and geographies.

What "done" looks like

Endpoint, Device & IoT Security is done when the organization has total asset visibility and can prove enforcement at the device level. Every endpoint and IoT device is inventoried, assigned an owner, and placed under a security baseline. Devices are patched, configured, and monitored continuously. High-risk behavior is detected quickly and contained automatically where possible. When a device is lost, stolen, compromised, or no longer needed, it can be isolated and removed from access paths immediately.

Scope boundaries

This domain includes device inventory, endpoint hardening, patch and vulnerability hygiene at the device layer, endpoint detection and response, device posture validation, IoT device governance, and device isolation or containment. It does not include identity lifecycle governance, network segmentation design, or application security logic, which are handled in other RCF domains. This domain governs the security of the devices themselves.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Endpoint, Device & IoT Security must operate across corporate endpoints, servers, mobile devices, and IoT fleets in on-prem, cloud, and hybrid environments. The architecture centers on three pillars: inventory and ownership, enforcement of a hardened baseline, and continuous monitoring with rapid containment. Device controls must integrate with identity and network context so access decisions reflect device health.

Required systems, data sources, and integrations

This domain requires an asset inventory system, endpoint management for configuration and patching, endpoint detection and response capability, device posture assessment, and mechanisms for isolation and containment. For IoT, it requires device discovery, firmware visibility, segmentation enforcement, and secure onboarding controls. It consumes telemetry from endpoints, servers, mobile platforms, IoT gateways, vulnerability scanners, and authentication systems. Integration with Rocheston Noodles provides centralized evidence of device compliance and control health, while AINA correlates device behavior, vulnerability status, and enforcement drift across the fleet.

Data flows

Devices enroll into management and are recorded in inventory with ownership, location, and classification. Baseline configurations and patch policies are applied and enforced. Endpoint and IoT telemetry flows continuously into Noodles, including health state, patch level, security events, and behavioral signals. AINA evaluates drift, anomalies, and risk patterns, then triggers alerts and recommended containment actions. Decisions and actions flow back into device controls to isolate a compromised endpoint, block risky behavior, or restrict access based on posture.

Minimum viable setup vs enterprise setup

A minimum viable setup includes a complete device inventory, basic endpoint management for patching and configuration, baseline hardening standards, and continuous endpoint event logging. An enterprise setup includes real-time posture-based access, advanced endpoint detection and response, automated isolation, firmware governance for IoT, full lifecycle control from procurement to retirement, and executive dashboards showing device risk, compliance, and exposure trends.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically the endpoint security or device engineering leader, with coordination across IT operations, security operations, and network teams. Supporting roles include desktop and server administrators, mobile device management owners, IoT platform owners, asset management teams, and RCCE engineers who design device security as a repeatable system rather than a reactive cleanup function.

Cadence

Daily operations focus on monitoring endpoint alerts, patch failures, device enrollment gaps, and anomalous behavior. Weekly routines focus on patch deployment cycles, remediation of high-risk vulnerabilities, and review of devices that are non-compliant or unknown. Monthly routines focus on baseline validation, device posture reporting, IoT firmware review, and reduction of device sprawl. Quarterly routines focus on lifecycle reviews, decommissioning of obsolete devices, and executive posture review. Annual routines focus on refreshing baselines, updating device standards, and reassessing IoT risk as environments evolve.

Required meetings and approvals

Changes to security baselines require formal review and approval to ensure operational compatibility. Exceptions to device controls, such as allowing unmanaged devices or legacy IoT equipment, require documented approval with expiry. High-risk IoT onboarding must receive security sign-off before production deployment.

Escalation paths

Active endpoint compromise escalates immediately to security operations for containment. Non-compliant devices that cannot be remediated escalate to the domain owner for enforcement decisions, including access restriction or removal. Device risks affecting critical systems or regulated data escalate to executive leadership. Persistent systemic device failures escalate to the board or audit committee when they materially affect organizational resilience.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)
In Phase 1 you establish visibility and baseline control over devices. You identify the accountable owner, build or validate an asset inventory, classify endpoints and IoT devices by criticality, define hardened baselines, and enable basic telemetry collection. Outputs include a device inventory with ownership, baseline standards, patch policy definitions, and logging configurations. Evidence at the end of this phase shows that devices are known, owned, and under management.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize device enforcement. You enroll devices into management platforms, deploy baseline configurations, implement patch automation, enable endpoint detection capabilities, integrate telemetry into Noodles, and enable AINA evaluation of device posture and drift. Outputs include managed device fleets, compliance dashboards, patch performance reports, and active alerting workflows. Evidence demonstrates that device controls are enforced and monitored continuously.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make device security resilient and scalable. You validate containment actions through exercises, test isolation workflows, reduce privilege on endpoints, enforce posture-based access restrictions, harden IoT onboarding and firmware control, and anchor device compliance evidence into Rosecoin Vault. Outputs include containment test results, posture trend analysis, validated baselines, and immutable compliance records. Evidence proves that endpoints and IoT devices cannot remain unmanaged or compromised without rapid detection and response.

End state
When this domain is mature, endpoints and IoT devices are no longer the organization's easiest entry point. They become controlled assets with enforced baselines, continuous monitoring, and rapid containment. The organization can prove device hygiene, prevent silent footholds, and keep the fleet healthy even as the environment grows and changes.

DOMAIN 9 — SECURE SOFTWARE DEVELOPMENT (SSDLC)

DOMAIN MISSION AND OUTCOMES

Mission

Secure Software Development exists to ensure that security is built into software from the moment an idea becomes code, not added later as damage control. Software now defines business operations, customer experience, and critical infrastructure. This domain ensures applications are designed, built, tested, released, and maintained in ways that minimize exploitable flaws, reduce systemic risk, and make security a normal part of engineering work rather than an external gate.

What this domain prevents

This domain prevents vulnerabilities from being introduced silently during development and then discovered in production by attackers. It prevents insecure design decisions that cannot be fixed with patches. It prevents dependency risk from unmanaged open-source libraries and unknown components. It prevents release pressure from overriding security review. It also prevents organizations from accumulating technical debt that turns every future change into a high-risk event.

What “done” looks like

Secure Software Development is done when security is embedded into how software is written and released every day. Threat modeling influences design decisions. Code is reviewed and tested automatically for common classes of flaws. Dependencies are known, tracked, and controlled. Builds fail when security requirements are not met. Security findings are treated as engineering defects, not external interruptions. When vulnerabilities are discovered, they can be traced to code, ownership, and remediation timelines with evidence.

Scope boundaries

This domain includes secure design practices, developer security standards, code review and testing, dependency and SBOM governance, build and release security controls, and vulnerability remediation within the development lifecycle. It does not include runtime application defense, infrastructure hardening, or incident response execution, which are covered in other RCF domains. SSDLC governs how software is created and maintained securely.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Secure Software Development spans development environments, build systems, and deployment pipelines. The architecture must integrate security controls into source repositories, CI/CD pipelines, and artifact repositories so security checks are automated and repeatable. This applies equally to on-prem, cloud-native, and hybrid development models.

Required systems, data sources, and integrations

This domain requires source code repositories, CI/CD platforms, static and dynamic analysis tools, dependency scanning, artifact repositories, and vulnerability tracking systems. It consumes data from code commits, build logs, test results, dependency manifests, and deployment records. Integration with Rocheston Noodles centralizes evidence from the development pipeline, while AINA correlates findings, trends, and risk patterns across projects and teams.

Data flows

Code changes flow into repositories where automated checks are triggered. Security analysis results flow into Noodles as evidence of control enforcement. AINA evaluates patterns such as recurring flaw types, delayed remediation, or risky dependencies. Approved builds flow into deployment systems only when security gates are satisfied. Findings and remediation actions are tracked to closure so security outcomes remain visible and measurable.

Minimum viable setup vs enterprise setup

A minimum viable setup includes secure coding standards, basic static analysis, dependency visibility, manual threat modeling for critical applications, and documented release approvals. An enterprise setup includes automated security testing across the pipeline, enforced security gates, comprehensive SBOM generation, continuous dependency monitoring, centralized vulnerability tracking, and executive dashboards showing software risk posture across portfolios.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security or engineering leadership role, with shared responsibility across development teams. Supporting roles include application owners, development leads, security engineers, DevOps teams, and RCCE engineers who embed security controls directly into pipelines and workflows.

Cadence

Daily operations focus on monitoring build failures, new security findings, and remediation progress. Weekly routines focus on reviewing recurring issues, updating security rules, and addressing high-risk vulnerabilities. Monthly routines focus on trend analysis across applications, dependency risk review, and improvement of developer guidance. Quarterly routines focus on executive review of software risk, backlog reduction, and alignment with business priorities. Annual routines focus on refreshing secure development standards and adapting to new threat patterns.

Required meetings and approvals

Security requirements for releases must be clearly defined and enforced through automation. Exceptions to security gates require documented approval with justification and expiry. Major architectural changes require security review early in the design phase rather than just before release.

Escalation paths

Critical vulnerabilities escalate to application owners and security leadership for immediate remediation planning. Unresolved or systemic development risks escalate to executive leadership. Software risks affecting regulated data or critical services escalate to the board or audit committee when they materially affect organizational exposure.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish baseline security expectations for software development. You identify the accountable owner, define secure coding and design standards, inventory applications and repositories, enable basic code scanning, and define release approval requirements. Outputs include development security standards, application inventory, initial scanning results, and documented workflows. Evidence at the end of this phase shows that security checks exist and are visible.

Phase 2: Implement (days 15–60)
In Phase 2 you embed security into the development pipeline. You integrate automated testing into CI/CD, implement dependency scanning and SBOM generation, connect findings into Noodles, enable AINA analysis of trends and risk concentration, and formalize remediation tracking. Outputs include automated pipelines, dashboards, vulnerability backlogs, and approved build records. Evidence demonstrates that insecure code does not flow silently into production.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make secure development resilient and scalable. You enforce security gates consistently, reduce recurring vulnerability classes, validate remediation timelines, test rollback and emergency patch workflows, and anchor SSDLC evidence into Rosecoin Vault. Outputs include validated pipelines, reduced defect trends, compliance-ready evidence, and immutable development security records. Evidence proves that secure software practices are sustained, not episodic.

End state
When this domain is mature, software security is no longer a bottleneck or an afterthought. It becomes a built-in quality of engineering work. Applications ship faster with fewer defects, dependencies are controlled, and security posture improves continuously as code evolves.

DOMAIN 10 — CONTINUOUS MONITORING & DETECTION

DOMAIN MISSION AND OUTCOMES

Mission

Continuous Monitoring & Detection exists to ensure the organization is never blind to what is happening in its own environment. Modern attacks do not announce themselves; they blend into normal activity and unfold over time. This domain ensures security teams can observe, correlate, and detect malicious or abnormal behavior continuously, across systems, identities, networks, applications, and data. It transforms security from periodic checking into real-time situational awareness.

What this domain prevents

This domain prevents breaches from going undetected for weeks or months. It prevents attackers from quietly escalating privileges, moving laterally, or exfiltrating data without triggering visibility. It prevents security teams from drowning in alerts that lack context and from missing real threats hidden in noise. It also prevents overreliance on point tools that see only fragments of the attack chain and cannot tell a coherent story.

What "done" looks like

Continuous Monitoring & Detection is done when the organization has unified visibility and can detect meaningful security events quickly and reliably. Signals from across the environment are collected, correlated, and prioritized based on risk and impact. Alerts tell a story rather than listing symptoms. Detection coverage is known and measured. When something goes wrong, security teams can see what happened, where it started, and what systems are involved without rebuilding timelines manually.

Scope boundaries

This domain includes telemetry collection, log aggregation, signal correlation, detection logic, alerting, and coverage measurement. It does not include incident response execution, vulnerability remediation, or forensic investigation, which are addressed in other RCF domains. This domain is about seeing and understanding, not fixing or recovering.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Continuous Monitoring & Detection spans the entire technology landscape and must operate across on-prem, cloud, SaaS, and hybrid environments. The architecture is built around centralized collection of telemetry, real-time analysis, and correlation across domains. Detection logic must be decoupled from individual tools so visibility survives technology changes.

Required systems, data sources, and integrations

This domain requires centralized telemetry ingestion, scalable storage for logs and events, detection and correlation engines, and alerting interfaces. It consumes data from endpoints, networks, identity systems, cloud platforms, applications, databases, and security controls. Integration with Rocheston Noodles provides a unified view of detection evidence and coverage, while AINA correlates signals into meaningful incidents rather than isolated alerts.

Data flows

Telemetry flows continuously from systems and controls into centralized ingestion. AINA analyzes events in real time, correlating them across domains to identify suspicious patterns and attack sequences. High-confidence detections are surfaced as prioritized alerts with context and impact. Detection outcomes and metrics flow back into governance and operations so coverage gaps and effectiveness can be measured and improved.

Minimum viable setup vs enterprise setup

A minimum viable setup includes centralized log collection, basic detection rules for common attack patterns, and alerting for high-risk events. An enterprise setup includes full-spectrum telemetry coverage, advanced correlation and behavior analytics, threat-informed detection logic, continuous coverage measurement, and executive dashboards showing detection effectiveness and blind spots.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically the SOC or detection engineering leader. Supporting roles include security analysts, detection engineers, platform teams, and RCCE engineers who design scalable monitoring architectures and detection logic aligned with real-world threats.

Cadence

Daily operations focus on monitoring alerts, tuning detections, and validating signal quality. Weekly routines focus on improving detection coverage, reducing false positives, and incorporating new threat intelligence. Monthly routines focus on reviewing detection effectiveness, coverage gaps, and alignment with changing environments. Quarterly routines focus on executive review of monitoring posture and investment alignment. Annual routines focus on architecture refresh and detection strategy evolution.

Required meetings and approvals

Detection logic changes require review to ensure quality and avoid blind spots. Major platform changes require coordination with architecture and operations teams. Executive reviews focus on coverage and effectiveness rather than raw alert counts.

Escalation paths

Confirmed high-risk detections escalate immediately to incident response teams. Detection failures or blind spots escalate to the domain owner for corrective action. Systemic visibility gaps affecting critical assets escalate to executive leadership. Persistent detection deficiencies escalate to the board or audit committee when they materially affect risk posture.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish baseline visibility. You identify the accountable owner, inventory key data sources, enable centralized telemetry ingestion, define initial detection priorities, and configure basic alerting. Outputs include a monitoring architecture overview, source inventory, initial detection rules, and alert workflows. Evidence at the end of this phase shows that critical activity is visible and monitored.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize detection at scale. You integrate additional telemetry sources, enable AINA-driven correlation, tune detection logic, connect monitoring outputs into Noodles, and establish coverage metrics. Outputs include expanded telemetry, prioritized alerts, dashboards, and detection performance reports. Evidence demonstrates that detection is continuous, contextual, and improving.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make monitoring resilient and reliable. You test detection logic using simulations, validate alert fidelity, measure coverage against known attack techniques, and anchor detection evidence into Rosecoin Vault. Outputs include validated detections, coverage assessments, and immutable monitoring records. Evidence proves that the organization can see and detect real threats consistently.

End state

When this domain is mature, the organization is no longer guessing whether it has been compromised. It has continuous situational awareness, reliable detection, and a clear understanding of what is happening across its environment. Monitoring becomes a strategic capability that enables rapid response and informed decision-making rather than a flood of disconnected alerts.

DOMAIN 11 — THREAT INTELLIGENCE & ADVERSARY TRACKING

DOMAIN MISSION AND OUTCOMES

Mission

Threat Intelligence & Adversary Tracking exists to ensure the organization understands who is attacking, how they operate, and why they target certain assets—before damage occurs. Security defenses are most effective when they are informed by real adversary behavior rather than generic threat lists. This domain turns external and internal intelligence into actionable insight so defenses, detections, and decisions are aligned with the threats that actually matter.

What this domain prevents

This domain prevents organizations from defending against imaginary or outdated threats while real adversaries move unnoticed. It prevents reactive security postures that only respond after incidents occur. It prevents wasted effort chasing low-relevance indicators and vendor-driven noise. It also prevents leadership from making strategic security decisions without understanding the intent, capability, and persistence of adversaries targeting the organization or its industry.

What "done" looks like

Threat Intelligence & Adversary Tracking is done when the organization can clearly articulate which adversaries matter, what techniques they use, and how those techniques map to the organization's environment. Intelligence is timely, relevant, and integrated into detection, prevention, and response workflows. Adversary activity is tracked over time, not treated as isolated events. Leadership understands threat trends and can see how defenses are aligned against real-world attack patterns.

Scope boundaries

This domain includes threat intelligence collection, analysis, prioritization, adversary profiling, campaign tracking, and intelligence-driven decision support. It does not include real-time detection execution, vulnerability remediation, or incident response actions, which are handled in other RCF domains. This domain informs action; it does not execute it.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Threat Intelligence & Adversary Tracking operates as an intelligence layer that feeds multiple security domains. The architecture must support ingestion of external intelligence, correlation with internal telemetry, and long-term tracking of adversary behavior. It applies across on-prem, cloud, and hybrid environments and must remain independent of any single data source or vendor feed.

Required systems, data sources, and integrations

This domain requires threat intelligence feeds, internal telemetry from monitoring systems, historical incident data, and analytical tooling to correlate and contextualize intelligence. It consumes information from open-source intelligence, commercial feeds, industry sharing groups, and internal detection outputs. Integration with Rocheston Noodles provides centralized intelligence management and evidence, while AINA correlates indicators, behaviors, and campaigns into adversary profiles rather than isolated signals.

Data flows

External intelligence feeds and internal observations flow into a centralized analysis layer. AINA evaluates relevance based on industry, geography, assets, and observed activity. Intelligence is enriched with internal telemetry to confirm whether adversary techniques are present or emerging. The resulting insights flow into detection engineering, risk prioritization, and executive reporting so intelligence directly influences defensive posture.

Minimum viable setup vs enterprise setup

A minimum viable setup includes curated threat feeds, basic analysis workflows, and periodic intelligence reporting aligned to the organization's risk profile. An enterprise setup includes continuous intelligence ingestion, adversary and campaign tracking over time, automated relevance scoring, integration with detection and monitoring systems, and executive dashboards showing threat trends and exposure alignment.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a threat intelligence or security strategy leader. Supporting roles include analysts, detection engineers, incident responders, and RCCE engineers who ensure intelligence is operationalized rather than siloed in reports.

Cadence

Daily operations focus on monitoring emerging threats and validating relevance. Weekly routines focus on updating adversary profiles and sharing actionable intelligence with detection and operations teams. Monthly routines focus on trend analysis, campaign tracking, and alignment with detection coverage. Quarterly routines focus on executive briefings about threat landscape changes and strategic implications. Annual routines focus on reassessing threat models and intelligence sources.

Required meetings and approvals

Intelligence briefings are required to align detection and prevention priorities. Strategic intelligence assessments require executive review when they influence major investment or risk decisions. Coordination meetings ensure intelligence is translated into concrete defensive actions.

Escalation paths

High-confidence intelligence indicating imminent threat escalates immediately to security operations and leadership. Intelligence revealing systemic exposure escalates to executive leadership for prioritization. Persistent or strategic adversary threats escalate to the board or audit committee when they materially affect organizational risk.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish intelligence foundations. You identify the accountable owner, define intelligence requirements aligned to business risk, select initial intelligence sources, and establish analysis workflows. Outputs include an intelligence strategy, source inventory, initial adversary profiles, and reporting formats. Evidence at the end of this phase shows that threat intelligence is scoped and purposeful.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize intelligence. You integrate intelligence feeds into Noodles, correlate intelligence with internal telemetry using AINA, establish adversary tracking, and produce regular actionable intelligence outputs. Outputs include enriched intelligence reports, adversary profiles, and integration with detection priorities. Evidence demonstrates that intelligence influences real security decisions.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make intelligence resilient and trusted. You validate intelligence relevance, measure impact on detection effectiveness, refine adversary models, and anchor intelligence decisions and evidence into Rosecoin Vault. Outputs include validated threat models, trend analyses, and immutable intelligence records. Evidence proves that intelligence is reliable, relevant, and defensible.

End state

When this domain is mature, threat intelligence is no longer passive information. It becomes a strategic capability that continuously informs how the organization defends itself. Adversaries are understood, tracked, and anticipated, allowing security teams to stay aligned with real threats rather than reacting to surprises.

DOMAIN 12 — VULNERABILITY MANAGEMENT & SECURITY TESTING

DOMAIN MISSION AND OUTCOMES

Mission

Vulnerability Management & Security Testing exists to ensure weaknesses are found, prioritized, and fixed before attackers exploit them. Every environment contains flaws, but unmanaged flaws become guaranteed entry points. This domain ensures vulnerabilities are continuously discovered, validated, ranked by real risk, and driven to remediation in a disciplined, provable way. It turns “we know we have issues” into “we know which issues matter and they are being closed.”

What this domain prevents

This domain prevents attackers from exploiting known but unpatched vulnerabilities. It prevents security teams from drowning in scan results with no prioritization or ownership. It prevents the false sense of security created by periodic scans that are never acted upon. It also prevents repeated findings during audits, operational outages caused by rushed patching, and long-lived exposures that quietly increase risk over time.

What “done” looks like

Vulnerability Management & Security Testing is done when the organization has continuous visibility into its exposure and can prove timely remediation of meaningful risk. Vulnerabilities are prioritized based on exploitability, asset criticality, and business impact rather than raw severity scores alone. Owners are clearly assigned, remediation timelines are enforced, and exceptions are explicit and time-bound. Leadership can see exposure trends declining, not just scan volume increasing.

Scope boundaries

This domain includes vulnerability discovery, validation, prioritization, remediation tracking, and security testing across infrastructure, applications, and configurations. It includes both automated and manual testing. It does not include endpoint hardening, secure software development practices, or incident response execution, which are addressed in other RCF domains. This domain governs how weaknesses are identified and eliminated.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Vulnerability Management & Security Testing spans the full environment, including on-prem infrastructure, cloud platforms, applications, and edge systems. The architecture must support continuous scanning, targeted testing, and centralized tracking of findings. Results must be correlated with asset context and business criticality so remediation focuses on what matters most.

Required systems, data sources, and integrations

This domain requires vulnerability scanners, configuration assessment tools, application security testing platforms, and a centralized system for tracking findings and remediation. It consumes data from asset inventories, endpoint platforms, cloud services, CI/CD pipelines, and configuration management systems. Integration with Rocheston Noodles provides a unified view of vulnerability evidence and remediation status, while AINA correlates findings with exploit intelligence, asset value, and exposure to determine real risk.

Data flows

Scanning and testing tools generate findings that flow into a centralized tracking system. Asset context and ownership are applied so findings are actionable. AINA evaluates exploitability, exposure, and impact to prioritize remediation. Remediation actions and validation results flow back into the system, closing the loop with evidence that vulnerabilities were addressed or explicitly accepted.

Minimum viable setup vs enterprise setup

A minimum viable setup includes regular vulnerability scanning, basic asset tagging, manual prioritization, and documented remediation tracking. An enterprise setup includes continuous scanning, automated correlation with threat intelligence, risk-based prioritization, integrated application and infrastructure testing, validation of fixes, and executive dashboards showing exposure trends and remediation performance.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically within security operations or risk management. Supporting roles include system and application owners responsible for remediation, infrastructure and cloud teams executing fixes, security testers, and RCCE engineers who design scalable vulnerability management workflows.

Cadence

Daily operations focus on ingesting new findings, validating critical vulnerabilities, and tracking remediation progress. Weekly routines focus on prioritization reviews, coordination with owners, and verification of fixes. Monthly routines focus on trend analysis, recurring weakness identification, and process improvement. Quarterly routines focus on executive review of exposure reduction and alignment with risk tolerance. Annual routines focus on reassessing testing strategies and coverage.

Required meetings and approvals

Remediation plans for high-risk vulnerabilities require formal agreement between security and asset owners. Exceptions or deferrals require documented approval with justification and expiry. Regular coordination meetings ensure remediation work stays aligned with operational realities.

Escalation paths

Unpatched critical vulnerabilities escalate to the domain owner for enforcement. Risks affecting critical systems or regulated data escalate to executive leadership. Systemic remediation failures escalate to the board or audit committee when they materially affect organizational risk posture.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish visibility into weaknesses. You identify the accountable owner, inventory assets in scope, enable baseline scanning and testing, define prioritization criteria, and establish remediation tracking. Outputs include scan coverage reports, asset mappings, prioritization rules, and remediation workflows. Evidence at the end of this phase shows that vulnerabilities are being identified and tracked centrally.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize risk-based remediation. You integrate scanning and testing tools into Noodles, enable AINA-driven prioritization, coordinate remediation with asset owners, and begin producing exposure and remediation reports. Outputs include prioritized vulnerability backlogs, remediation status dashboards, and validation results. Evidence demonstrates that high-risk vulnerabilities are being addressed in a timely manner.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make vulnerability management reliable and defensible. You validate remediation effectiveness, reduce recurring vulnerability classes, test exception handling, and anchor vulnerability evidence into Rosecoin Vault. Outputs include exposure trend analysis, validated fixes, exception records, and immutable remediation proof. Evidence proves that weaknesses are not only found, but systematically eliminated.

End state

When this domain is mature, vulnerabilities no longer accumulate silently. Weaknesses are discovered early, prioritized intelligently, and driven to closure with accountability. The organization can prove that it understands its exposure and is actively reducing it, turning vulnerability management into a disciplined risk-reduction engine rather than a recurring source of surprise.

DOMAIN 13 — INCIDENT RESPONSE

DOMAIN MISSION AND OUTCOMES

Mission

Incident Response exists to ensure that when something goes wrong, the organization responds deliberately rather than reactively. Security incidents are inevitable in complex environments; chaos is not. This domain ensures incidents are detected, contained, investigated, communicated, and resolved through practiced procedures with clear authority and evidence. It turns a potential crisis into a controlled operational event.

What this domain prevents

This domain prevents panic-driven decision making, delayed containment, inconsistent communication, and loss of evidence during security incidents. It prevents incidents from escalating because nobody knew who was in charge or what to do next. It prevents regulatory exposure caused by missed notification deadlines, incomplete records, or uncontrolled disclosures. It also prevents organizations from repeating the same mistakes because lessons were never captured or applied.

What “done” looks like

Incident Response is done when the organization can respond quickly, consistently, and confidently to security events. Roles are predefined, authority is clear, and actions are rehearsed. Incidents are detected, triaged, and contained within defined timeframes. Communications are accurate and controlled. Evidence is preserved. After the incident, root causes are understood, corrective actions are tracked, and response capability improves measurably over time.

Scope boundaries

This domain includes incident classification, triage, containment, eradication, recovery coordination, communication management, evidence preservation, and post-incident review. It does not include continuous monitoring and detection, vulnerability remediation, or long-term forensic investigation, which are covered in other RCF domains. Incident Response governs how the organization acts under pressure.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Incident Response relies on integration across detection, containment, communication, and evidence systems. The architecture must support rapid intake of alerts, coordinated response actions, secure collaboration, and preservation of evidence. It must function across on-prem, cloud, and hybrid environments so response is consistent regardless of where the incident occurs.

Required systems, data sources, and integrations

This domain requires incident management and ticketing systems, secure communication channels, containment and isolation capabilities, logging and evidence repositories, and notification mechanisms. It consumes alerts and context from monitoring systems, endpoint and network controls, identity platforms, and cloud services. Integration with Rocheston Noodles provides centralized incident records and evidence, while AINA correlates alerts into incident narratives and supports automated response workflows.

Data flows

Alerts flow from monitoring systems into incident management. AINA correlates related alerts into a single incident view and provides context about scope and impact. Response actions and decisions are logged as they occur. Evidence from affected systems is collected and preserved. Status updates and communications flow through approved channels so information remains accurate and controlled throughout the incident lifecycle.

Minimum viable setup vs enterprise setup

A minimum viable setup includes defined incident categories, an on-call response team, basic containment procedures, and documented communication plans. An enterprise setup includes automated incident correlation, predefined response playbooks, integrated containment actions, real-time collaboration tools, regulatory notification tracking, and executive dashboards showing response performance and readiness.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically the incident response or security operations leader. Supporting roles include incident commanders, technical responders, communications and legal representatives, business owners, and RCCE engineers who design response workflows and automation.

Cadence

Daily operations focus on readiness: maintaining contact lists, validating tools, and reviewing open incidents. Weekly routines focus on playbook updates, tooling checks, and coordination with detection teams. Monthly routines focus on tabletop exercises, review of incident metrics, and improvement tracking. Quarterly routines focus on executive readiness reviews and cross-functional coordination. Annual routines focus on full-scale simulations and program refresh.

Required meetings and approvals

Major incidents require formal incident command activation and executive briefings. External communications and regulatory notifications require legal and executive approval. Post-incident reviews require leadership sign-off to ensure corrective actions are owned and tracked.

Escalation paths

Incidents escalate based on severity and impact. High-severity incidents escalate immediately to executive leadership. Incidents affecting regulated data, critical services, or safety escalate to legal and regulatory leadership. Systemic response failures escalate to the board or audit committee with documented evidence.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish response authority and structure. You identify the accountable owner, define incident categories and severity levels, assign response roles, document escalation and communication paths, and establish incident tracking. Outputs include an incident response policy, role assignments, escalation matrix, and communication templates.

Evidence at the end of this phase shows that response authority and processes are defined and accessible.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize response capability. You integrate detection outputs into incident management, implement response playbooks, enable AINA-driven correlation and automation, connect Noodles for evidence tracking, and conduct initial tabletop exercises. Outputs include active playbooks, incident dashboards, exercise reports, and response metrics. Evidence demonstrates that incidents can be handled consistently and with context.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make incident response resilient and defensible. You test containment and recovery actions, validate communication and notification timelines, refine playbooks based on exercises and real incidents, and anchor incident evidence into Rosecoin Vault. Outputs include validated playbooks, response performance reports, lessons learned, and immutable incident records. Evidence proves that the organization can respond effectively under real pressure.

End state
When this domain is mature, incidents are no longer moments of confusion and improvisation. They become controlled operational events handled with speed, clarity, and accountability. The organization can demonstrate not only that it detects incidents, but that it responds decisively, preserves trust, and improves with every event.

DOMAIN 14 — RESILIENCE, BUSINESS CONTINUITY & DISASTER RECOVERY

DOMAIN MISSION AND OUTCOMES

Mission

Resilience, Business Continuity & Disaster Recovery exists to ensure the organization can continue operating through disruption and recover deliberately from failure. Cybersecurity is not only about preventing incidents; it is about surviving them. This domain ensures that critical services remain available, data can be restored, and the business can function even when systems fail, attacks succeed, or external events cause widespread disruption. It turns outages from existential threats into managed operational events.

What this domain prevents

This domain prevents prolonged downtime, irreversible data loss, and chaotic recovery efforts during crises. It prevents organizations from discovering too late that backups do not work, recovery plans are outdated, or dependencies were never understood. It prevents business paralysis caused by single points of failure, fragile architectures, or untested assumptions about availability. It also prevents regulatory and contractual violations caused by extended service outages and unmet recovery obligations.

What "done" looks like

Resilience, Business Continuity & Disaster Recovery is done when the organization can demonstrate that critical services can withstand disruption and recover within defined timeframes. Recovery objectives are realistic, documented, and tested. Dependencies are known. Failover and restoration procedures work under real conditions. Leadership understands what will continue, what will degrade, and what will stop during a major event —and accepts those outcomes intentionally rather than by surprise.

Scope boundaries

This domain includes business impact analysis, resilience architecture, backup and recovery, failover planning, continuity procedures, and disaster recovery testing. It does not include incident detection, vulnerability remediation, or forensic investigation, which are handled in other RCF domains. This domain governs survival and recovery, not prevention or response execution.

B) Domain architecture blueprint

Reference architecture

Resilience architecture spans applications, infrastructure, data, and operational processes. It must support redundancy, isolation, and recovery across on-prem, cloud, and hybrid environments. Architectures should assume failure and be designed so no single event—technical, cyber, or environmental—can permanently disable critical business functions.

Required systems, data sources, and integrations

This domain requires backup and restore platforms, replication and failover mechanisms, dependency mapping, configuration and asset inventories, and testing environments. It consumes data from application architectures, infrastructure platforms, cloud services, and business process owners. Integration with Rocheston Noodles provides centralized visibility into resilience evidence, while AINA evaluates recovery readiness, drift from recovery objectives, and test outcomes.

Data flows

System and application architectures define dependencies and recovery priorities. Backup, replication, and failover systems generate evidence of protection and readiness. Test results and recovery metrics flow into Noodles for visibility. AINA evaluates whether recovery objectives remain achievable as environments change and highlights gaps before real failures occur.

Minimum viable setup vs enterprise setup

A minimum viable setup includes identification of critical services, defined recovery objectives, basic backups, documented recovery procedures, and periodic manual testing. An enterprise setup includes automated failover, immutable and offline backups, dependency-aware recovery sequencing, regular recovery exercises, continuous validation of backup integrity, and executive dashboards showing resilience posture and recovery readiness.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a resilience, infrastructure, or security leader, with strong collaboration across IT operations, application owners, and business leadership. Supporting roles include disaster recovery coordinators, platform teams, business continuity planners, and RCCE engineers who design resilient architectures and recovery workflows.

Cadence

Daily operations focus on monitoring backup success, replication health, and availability indicators. Weekly routines focus on reviewing failures, addressing gaps, and validating changes against recovery objectives. Monthly routines focus on testing specific recovery components and updating documentation. Quarterly routines focus on broader recovery exercises and executive review of resilience posture. Annual routines focus on full disaster recovery simulations and reassessment of business impact and tolerance.

Required meetings and approvals

Changes to recovery objectives or critical service classifications require executive approval. Major architectural changes require resilience review to ensure recovery assumptions remain valid. Post-exercise reviews require leadership sign-off to ensure lessons are owned and addressed.

Escalation paths

Recovery failures or inability to meet defined objectives escalate to the domain owner immediately. Risks that threaten critical business operations escalate to executive leadership for prioritization and investment decisions. Systemic resilience gaps escalate to the board or audit committee when they materially affect the organization's ability to operate.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish resilience intent and visibility. You identify the accountable owner, perform a business impact analysis, define critical services and recovery objectives, inventory dependencies, and document baseline recovery procedures. Outputs include a business impact assessment, recovery objectives, service dependency maps, and initial recovery plans. Evidence at the end of this phase shows that resilience expectations are defined and owned.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize resilience controls. You implement backups, replication, and failover mechanisms aligned to recovery objectives, integrate resilience evidence into

Noodles, enable AINA evaluation of readiness, and conduct initial recovery tests. Outputs include protected systems, test results, dashboards, and remediation plans for identified gaps. Evidence demonstrates that recovery is feasible and measurable.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make resilience reliable under real stress. You test recovery under adverse conditions, validate data integrity and restoration speed, exercise cross-functional continuity procedures, and anchor resilience evidence into Rosecoin Vault. Outputs include validated recovery exercises, resilience maturity assessments, and immutable proof of readiness. Evidence proves that the organization can survive disruption and recover intentionally.

End state

When this domain is mature, disruption no longer equals disaster. The organization can absorb shocks, maintain essential operations, and recover deliberately. Resilience becomes a built-in property of how systems and processes are designed, giving leadership confidence that the business can withstand both cyber incidents and unexpected events without collapsing.

DOMAIN 15 — DIGITAL FORENSICS & INVESTIGATION

DOMAIN MISSION AND OUTCOMES

Mission

Digital Forensics & Investigation exists to ensure the organization can discover the truth after a security event and prove it with confidence. When incidents occur, speculation, assumptions, and incomplete data are dangerous. This domain ensures evidence is preserved correctly, investigations are methodical, and conclusions are defensible in technical, legal, and regulatory contexts. It turns incidents from confusion into clarity and from accusations into facts.

What this domain prevents

This domain prevents loss or contamination of evidence during incidents. It prevents organizations from drawing incorrect conclusions based on incomplete or altered data. It prevents legal exposure caused by broken chain of custody, improper evidence handling, or undocumented investigative steps. It also prevents repeated incidents caused by failure to understand root cause, attacker behavior, or the true scope of compromise.

What "done" looks like

Digital Forensics & Investigation is done when investigations are repeatable, objective, and provable. Evidence is collected in a forensically sound manner, preserved with integrity, and analyzed using documented methods. Timelines can be reconstructed accurately. Findings can withstand scrutiny from auditors, regulators, courts, and executive leadership. Most importantly, investigations lead to corrective action, not just reports.

Scope boundaries

This domain includes forensic readiness, evidence collection and preservation, timeline reconstruction, root-cause analysis, and investigative reporting. It does not include real-time detection, active containment, or long-term remediation execution, which are addressed in other RCF domains. This domain governs truth-finding and proof, not response or prevention.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Digital Forensics & Investigation requires architectures that preserve evidence by design. Systems must generate reliable logs, retain historical data, and support forensic acquisition across endpoints, servers, cloud platforms, networks, and applications. The architecture must support investigations in on-prem, cloud, and hybrid environments without relying on ad hoc access or manual reconstruction.

Required systems, data sources, and integrations

This domain requires centralized log retention, time-synchronized systems, forensic acquisition tools, secure evidence storage, and investigation workspaces. It consumes data from endpoints, servers, cloud audit logs, identity platforms, network devices, applications, and backup systems. Integration with Rocheston Noodles provides centralized forensic evidence visibility, while AINA assists in correlating events, reconstructing timelines, and identifying relationships across disparate data sources.

Data flows

When an incident or investigation begins, relevant data is preserved immediately to prevent loss or alteration. Logs, disk images, memory captures, and cloud audit records are collected and stored securely. Evidence metadata and chain-of-custody records flow into Noodles for tracking. AINA analyzes event sequences and correlations to support timeline reconstruction and hypothesis testing. Investigation findings flow into incident response, governance, and remediation planning.

Minimum viable setup vs enterprise setup

A minimum viable setup includes centralized logging with retention, basic forensic acquisition procedures, secure evidence storage, and documented investigation workflows. An enterprise setup includes forensic readiness built into platforms, automated evidence preservation triggers, long-term immutable storage, advanced timeline and behavior analysis, cross-cloud forensic capability, and executive dashboards showing investigative coverage and readiness.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a digital forensics or security investigations leader. Supporting roles include incident responders, legal and compliance representatives, IT and cloud platform teams, external forensic specialists when required, and RCCE engineers who design forensic readiness into systems rather than treating forensics as an afterthought.

Cadence

Daily operations focus on maintaining forensic readiness, log integrity, and evidence retention. Weekly routines focus on validating data sources, time synchronization, and investigative tooling. Monthly routines focus on reviewing investigation quality, updating procedures, and incorporating lessons learned. Quarterly routines focus on exercises and executive review of forensic capability. Annual routines focus on reassessing forensic scope, legal requirements, and storage strategy.

Required meetings and approvals

Formal approval is required to initiate deep forensic investigations that may affect systems, data, or personnel. Legal and executive review is required when investigations may lead to regulatory notification, litigation, or disciplinary action. Post-investigation reviews require leadership sign-off to ensure findings result in corrective action.

Escalation paths

Findings indicating criminal activity, insider threat, or regulatory exposure escalate immediately to executive and legal leadership. Evidence integrity issues escalate to the domain owner for corrective action. Systemic investigative gaps escalate to the board or audit committee when they affect the organization's ability to establish truth and accountability.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish forensic readiness. You identify the accountable owner, define investigation authority, inventory available data sources, standardize time synchronization, and document evidence handling and chain-of-custody procedures. Outputs include a forensic readiness plan, evidence handling procedures, data source inventory, and investigation workflows. Evidence at the end of this phase shows that investigations can begin without improvisation.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize investigation capability. You integrate logs and evidence sources into Noodles, enable AINA-assisted timeline analysis, establish secure evidence storage, and conduct initial investigation exercises. Outputs include investigation dashboards, preserved evidence samples, exercise reports, and documented findings. Evidence demonstrates that investigations can reconstruct events accurately and consistently.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make forensics defensible and resilient. You validate chain-of-custody integrity, test investigations across cloud and hybrid scenarios, refine procedures based on exercises, and anchor forensic evidence and reports into Rosecoin Vault. Outputs include validated investigation reports, readiness assessments, and immutable forensic records. Evidence proves that investigative conclusions are trustworthy and legally defensible.

End state

When this domain is mature, the organization no longer guesses what happened. It knows. Incidents can be reconstructed with precision, accountability is established through evidence, and decisions are made based on facts rather than assumptions. Digital forensics becomes a pillar of trust, resilience, and continuous improvement rather than a last resort after failure.

DOMAIN 16 — POST-QUANTUM SECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Post-Quantum Security exists to ensure the organization's data, identities, and trust mechanisms remain secure in a future where quantum computing can break today's cryptography. Most security programs protect systems only against current attackers. This domain protects the organization against future attackers who can decrypt data captured today and exploit long-lived cryptographic assumptions tomorrow. It ensures cryptographic decisions made now do not become irreversible liabilities later.

What this domain prevents

This domain prevents "harvest now, decrypt later" attacks where adversaries collect encrypted data today with the intent to decrypt it once quantum capabilities mature. It prevents long-term exposure of sensitive data, intellectual property, personal information, and state-regulated records. It prevents cryptographic lock-in where systems cannot be upgraded without massive disruption. It also prevents false confidence caused by assuming that quantum risk is distant, theoretical, or someone else's problem.

What "done" looks like

Post-Quantum Security is done when the organization knows exactly where cryptography is used, how long protected data must remain confidential, and which systems are vulnerable to future cryptographic breaks. Cryptographic agility is built into architectures so algorithms can be replaced without redesigning systems. Sensitive data with long confidentiality lifetimes is protected using quantum-resistant approaches or compensating controls. Leadership understands the organization's quantum exposure and has an intentional migration strategy rather than vague awareness.

Scope boundaries

This domain includes cryptographic inventory, quantum risk assessment, cryptographic agility planning, migration to post-quantum algorithms, and long-term data protection strategy. It does not include general encryption configuration, key management operations, or identity governance execution, which are covered in other RCF domains. Post-Quantum Security governs future cryptographic survivability.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Post-Quantum Security spans every layer where cryptography is used, including data at rest, data in transit, identity, authentication, signing, and trust chains. The architecture must support algorithm abstraction and replacement without breaking applications or protocols. It applies across on-prem, cloud, hybrid, and third-party environments and must account for both internal systems and external trust dependencies.

Required systems, data sources, and integrations

This domain requires a cryptographic inventory covering applications, infrastructure, protocols, certificates, keys, and third-party dependencies. It requires lifecycle visibility into data sensitivity and retention periods. It consumes information from identity platforms, PKI systems, application architectures, cloud services, and vendor documentation. Integration with Rocheston Noodles provides centralized visibility into cryptographic exposure, while AINA evaluates quantum risk based on data lifetime, algorithm strength, and system criticality.

Data flows

Cryptographic usage data flows into a centralized inventory. Data classification and retention information is correlated to determine long-term confidentiality requirements. AINA evaluates which cryptographic uses are vulnerable to quantum attacks and prioritizes migration paths. Migration decisions and cryptographic changes flow back into architecture standards, development practices, and system configurations so future risk is reduced systematically.

Minimum viable setup vs enterprise setup

A minimum viable setup includes an inventory of cryptographic usage, identification of long-lived sensitive data, basic quantum risk assessment, and documented migration principles. An enterprise setup includes cryptographic abstraction layers, dual-algorithm support, staged migration to post-quantum algorithms, continuous validation of cryptographic posture, and executive dashboards showing quantum exposure and progress.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security architecture or cryptography leader. Supporting roles include identity and PKI teams, application architects, infrastructure teams, legal and compliance stakeholders for data retention requirements, and RCCE engineers who design cryptographic agility into systems.

Cadence

Daily operations focus on monitoring changes to cryptographic usage and new system deployments. Weekly routines focus on reviewing architectural changes for cryptographic impact. Monthly routines focus on updating cryptographic inventories and reassessing quantum exposure. Quarterly routines focus on executive review of quantum readiness and migration progress. Annual routines focus on updating post-quantum strategy based on advances in standards, research, and adversary capability.

Required meetings and approvals

Changes to cryptographic standards or algorithms require formal architecture approval. Systems handling long-lived sensitive data require post-quantum review before deployment. Migration plans and risk acceptance decisions require executive approval due to their long-term impact.

Escalation paths

Discovery of unmitigated quantum-vulnerable cryptography protecting long-lived sensitive data escalates to the domain owner immediately. Systemic cryptographic exposure escalates to executive leadership for prioritization and investment decisions. Strategic quantum risk escalates to the board or audit committee when it affects long-term trust, regulatory obligations, or national security exposure.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish visibility into cryptographic exposure. You identify the accountable owner, inventory cryptographic usage across systems, classify data by confidentiality lifetime, and define quantum risk criteria. Outputs include a cryptographic inventory, data lifetime mapping, and initial quantum risk assessment. Evidence at the end of this phase shows that quantum exposure is understood rather than assumed.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize cryptographic agility. You define post-quantum standards and transition principles, update architecture patterns, integrate quantum risk tracking into Noodles, enable AINA-driven prioritization, and begin migration for highest-risk use cases. Outputs include updated standards, migration plans, pilot implementations, and progress dashboards. Evidence demonstrates that quantum risk reduction has begun.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make quantum readiness durable. You validate cryptographic abstraction and fallback mechanisms, test migration scenarios, update third-party requirements, and anchor quantum readiness evidence into Rosecoin Vault. Outputs include validated architectures, tested migration paths, readiness assessments, and immutable records of cryptographic decisions. Evidence proves that the organization can evolve cryptography without catastrophic disruption.

End state

When this domain is mature, cryptography is no longer a fixed assumption that will eventually fail. It becomes an adaptable trust mechanism designed to survive technological disruption. The organization protects today's data from tomorrow's attackers and can demonstrate that its security posture is built for the future, not frozen in the past.

DOMAIN 17 — AUTONOMOUS DEFENSE & SELF-HEALING SYSTEMS

DOMAIN MISSION AND OUTCOMES

Mission

Autonomous Defense & Self-Healing Systems exists to ensure security does not depend on human reaction time. Modern attacks move faster than manual response, exploit scale, and target the gaps between teams, tools, and shifts. This domain enables the environment to defend itself by detecting, deciding, and acting automatically within defined authority. It ensures security systems can contain damage, restore safe state, and continue operating even when humans are unavailable or overwhelmed.

What this domain prevents

This domain prevents damage escalation caused by slow or inconsistent human response. It prevents attackers from exploiting dwell time between detection and action. It prevents fatigue-driven mistakes during high-volume incidents. It prevents outages and repeated compromise caused by fragile systems that cannot recover on their own. It also prevents over-automation chaos by ensuring autonomous actions are bounded, auditable, and reversible.

What "done" looks like

Autonomous Defense & Self-Healing Systems is done when the environment can safely act on its own. High-confidence threats are contained automatically. Systems can isolate, revoke, roll back, or reconfigure themselves to a known-safe state. Automation follows predefined authority and guardrails. Actions are logged, traceable, and reviewable. Human teams shift from firefighting to oversight, tuning, and improvement rather than constant intervention.

Scope boundaries

This domain includes automated containment, policy-driven response actions, self-healing workflows, rollback mechanisms, and autonomous enforcement logic. It does not include detection logic design, governance authority definition, or long-term remediation planning, which are handled in other RCF domains. This domain governs execution at machine speed, not strategy or policy creation.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Autonomous Defense & Self-Healing Systems sits between detection and infrastructure control planes. The architecture must support rapid decision-making, controlled execution, and safe rollback. It operates across endpoints, networks, identities, cloud workloads, and applications in on-prem, cloud, and hybrid environments. Enforcement points must be close to assets, while decision logic remains centrally governed.

Required systems, data sources, and integrations

This domain requires response orchestration platforms, policy engines, enforcement integrations with endpoint, network, identity, and cloud controls, and state validation mechanisms. It consumes inputs from detection systems, identity posture, device health, network context, and asset criticality. Integration with Rocheston Noodles provides visibility into automated actions and outcomes, while AINA evaluates confidence levels, selects appropriate responses, and validates post-action state.

Data flows

Detection signals flow into the decision engine with context and confidence scores. AINA evaluates signals against response policies and authority boundaries. Approved actions execute automatically through integrated enforcement points. Resulting state changes and outcomes flow back into Noodles as evidence. Feedback loops validate whether the action achieved containment or recovery and trigger additional steps if needed.

Minimum viable setup vs enterprise setup

A minimum viable setup includes automated containment for a limited set of high-confidence scenarios, such as isolating compromised endpoints or revoking credentials. An enterprise setup includes multi-stage autonomous workflows, cross-domain coordination, self-healing infrastructure patterns, continuous validation of system state, and executive dashboards showing automation effectiveness and safety.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security automation or architecture leader. Supporting roles include security operations, platform teams, identity and network engineers, and RCCE engineers who design autonomous actions that are safe, bounded, and resilient.

Cadence

Daily operations focus on monitoring automated actions, false positives, and system health. Weekly routines focus on tuning automation policies and expanding coverage cautiously. Monthly routines focus on reviewing outcomes, rollback effectiveness, and safety metrics. Quarterly routines focus on executive review of autonomy levels and risk tolerance. Annual routines focus on advancing self-healing maturity and integrating new autonomous capabilities.

Required meetings and approvals

Introduction of new autonomous actions requires formal approval and testing. Expansion of authority boundaries requires executive sign-off. Post-incident reviews include evaluation of autonomous behavior to ensure actions were appropriate and controlled.

Escalation paths

Automation failures or unsafe actions escalate immediately to the domain owner and security leadership. Autonomous actions affecting critical business systems escalate to executive leadership for review. Systemic automation risks escalate to the board or audit committee when they materially affect operational trust.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish safe foundations for autonomy. You identify the accountable owner, define authority boundaries, select initial automation use cases, integrate enforcement points, and document rollback procedures. Outputs include automation policies, authority definitions, and integration mappings. Evidence at the end of this phase shows that autonomous actions are controlled and reversible.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize autonomous defense. You enable AINA-driven decision logic, deploy automated containment and recovery workflows, integrate evidence tracking into Noodles, and conduct controlled exercises. Outputs include active automation workflows, dashboards, and exercise results. Evidence demonstrates that systems can act and recover without human intervention in defined scenarios.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make autonomy trustworthy at scale. You stress-test automation under failure conditions, validate rollback and recovery paths, expand self-healing patterns, and anchor automation evidence into Rosecoin Vault. Outputs include validated workflows, safety assessments, and immutable records of autonomous actions. Evidence proves that self-healing systems reduce impact without introducing uncontrolled risk.

End state

When this domain is mature, security no longer waits for humans to catch up with machines. The environment defends itself within defined limits, heals from damage automatically, and maintains stability under attack. Humans remain in control—but no longer in the critical path—allowing security to operate at the speed and scale modern threats demand.

DOMAIN 18 — PEOPLE SECURITY & CULTURE

DOMAIN MISSION AND OUTCOMES

Mission

People Security & Culture exists to ensure that humans strengthen security rather than weaken it. Most security failures ultimately involve people—not because they are careless, but because systems are confusing, incentives are misaligned, or expectations are unclear. This domain ensures security is understood, practiced, and reinforced as part of everyday work. It transforms security from a set of rules people work around into a shared responsibility people actively support.

What this domain prevents

This domain prevents social engineering attacks, credential theft through phishing, insider misuse—both malicious and accidental—and security breakdowns caused by fatigue, confusion, or poor incentives. It prevents training programs that check boxes but change nothing. It prevents blame-driven cultures where mistakes are hidden instead of corrected. It also prevents erosion of trust between security teams and the rest of the organization, which quietly undermines every technical control.

What "done" looks like

People Security & Culture is done when secure behavior is the default, not the exception. Employees understand why security matters and how their actions affect risk. Security expectations are clear, practical, and role-appropriate. Reporting suspicious activity is easy and encouraged. Mistakes are treated as learning opportunities, not punishments. Leadership models the behavior it expects, and security culture improves measurably over time rather than degrading under pressure.

Scope boundaries

This domain includes security awareness, behavior reinforcement, role-based education, insider risk awareness, leadership engagement, and cultural measurement. It does not include identity enforcement, technical access controls, or automated detection systems, which are covered in other RCF domains. This domain governs human behavior, incentives, and shared responsibility.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

People Security & Culture operates across organizational, educational, and behavioral systems rather than traditional infrastructure. The architecture must support continuous learning, feedback, and measurement. It applies across all roles, locations, and employment types, including contractors and partners with access to systems or data.

Required systems, data sources, and integrations

This domain requires training and learning platforms, communication channels, reporting mechanisms for suspicious activity, and systems to track participation and outcomes. It consumes data from phishing simulations, incident reports, training completion records, and employee feedback. Integration with Rocheston Noodles provides centralized evidence of cultural controls, while AINA evaluates behavior trends, risk signals, and program effectiveness over time.

Data flows

Training content and guidance flow to users based on role and risk. Behavioral signals such as reporting rates, phishing responses, and policy acknowledgments flow back into Noodles. AINA correlates these signals with incident data and risk outcomes to assess whether culture is improving or degrading. Insights flow into leadership reporting and program adjustments so culture evolves intentionally.

Minimum viable setup vs enterprise setup

A minimum viable setup includes basic security awareness training, clear reporting channels, leadership messaging, and simple measurement of participation. An enterprise setup includes role-based and adaptive training, continuous reinforcement through simulations and micro-learning, behavioral analytics, insider risk awareness programs, and executive dashboards showing culture health and human risk trends.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security, risk, or people leadership role. Supporting roles include HR, communications, legal, security operations, line managers, and RCCE engineers who ensure people-focused controls are integrated with technical security rather than isolated.

Cadence

Daily operations focus on monitoring reports of suspicious activity and responding constructively. Weekly routines focus on communications, reinforcement, and review of behavior signals. Monthly routines focus on training updates, simulation analysis, and targeted interventions. Quarterly routines focus on executive review of culture metrics and alignment with organizational change. Annual routines focus on refreshing the people security strategy and adapting to new threat patterns and workforce models.

Required meetings and approvals

Leadership alignment meetings ensure security expectations are consistent across the organization. Approval is required for programs that affect employee monitoring or disciplinary processes to ensure fairness and legal compliance. Regular reviews ensure training remains relevant and respectful of employee time.

Escalation paths

Patterns of risky behavior escalate to the domain owner for intervention and support. Insider risk indicators escalate to security and legal leadership for controlled investigation. Cultural breakdowns affecting critical operations escalate to executive leadership. Systemic people-risk issues escalate to the board or audit committee when they materially affect organizational resilience.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish clarity and ownership. You identify the accountable owner, define expected security behaviors, establish reporting channels, inventory existing training and communication assets, and align leadership messaging. Outputs include a people security charter, reporting procedures, baseline training materials, and communication plans. Evidence at the end of this phase shows that expectations are defined and visible.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize behavior change. You deploy role-based training, launch simulations and awareness campaigns, integrate behavior metrics into Noodles, enable AINA analysis of trends, and begin reporting culture indicators to leadership. Outputs include training records, simulation results, dashboards, and feedback summaries. Evidence demonstrates that people are engaged and behavior is being measured.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make culture durable under pressure. You refine programs based on data, address high-risk roles or behaviors, test reporting and response processes, and anchor people security evidence into Rosecoin Vault. Outputs include validated culture assessments, improvement plans, and immutable records of training and engagement. Evidence proves that culture is improving and resilient, not superficial.

End state
When this domain is mature, people are no longer treated as the weakest link. They become an active defense layer. Security culture supports technical controls, adapts as threats change, and holds up even during stress. The organization can prove that its human element strengthens resilience rather than undermining it.

DOMAIN 19 — CONTINUOUS IMPROVEMENT & MATURITY

DOMAIN MISSION AND OUTCOMES

Mission

Continuous Improvement & Maturity exists to ensure security does not stagnate. Threats evolve, technology changes, business models shift, and teams rotate. A security program that does not actively measure itself and improve will decay even if it once performed well. This domain ensures the organization learns from experience, adapts deliberately, and raises its security capability over time instead of resetting every audit cycle.

What this domain prevents

This domain prevents security programs from plateauing after initial success. It prevents repeated mistakes caused by untracked lessons learned. It prevents audit-driven "spikes" followed by long periods of neglect. It prevents maturity claims that are based on documentation rather than performance. It also prevents leadership from believing security is improving when, in reality, controls are eroding or threats are outpacing defenses.

What "done" looks like

Continuous Improvement & Maturity is done when the organization can clearly show how its security posture has improved over time and why. Maturity is measured, not assumed. Weak areas are identified early, improvements remember past failures, and progress is visible across people, process, and technology. Leadership understands current maturity, target maturity, and the roadmap between them, and those targets evolve as risk changes.

Scope boundaries

This domain includes maturity modeling, performance measurement, improvement planning, lessons learned integration, and governance of long-term security evolution. It does not include day-to-day control operation, detection, response, or remediation execution, which are handled in other RCF domains. This domain governs learning, adaptation, and progress.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Continuous Improvement & Maturity operates as an overlay across all RCF domains. The architecture must aggregate evidence, metrics, outcomes, and lessons from every control area into a unified maturity view. It must support longitudinal analysis so trends can be seen across months and years, not just point-in-time snapshots.

Required systems, data sources, and integrations

This domain requires access to evidence and metrics from all RCF domains, assessment frameworks, reporting systems, and executive dashboards. It consumes data from audits, incidents, exercises, remediation tracking, and operational performance indicators. Integration with Rocheston Noodles provides centralized maturity scoring and historical tracking, while AINA evaluates trends, identifies stagnation or regression, and recommends improvement focus areas.

Data flows

Operational and governance evidence flows into Noodles continuously. AINA analyzes control effectiveness, incident outcomes, and remediation performance to calculate maturity signals. These signals flow into maturity dashboards and improvement plans. Decisions about priorities and investments flow back into domain roadmaps so improvement is intentional and tracked.

Minimum viable setup vs enterprise setup

A minimum viable setup includes basic maturity definitions, periodic self-assessments, documented lessons learned, and manual improvement tracking. An enterprise setup includes continuous maturity scoring across domains, automated trend analysis, integration of improvement actions into planning cycles, executive dashboards, and evidence-backed demonstration of progress over time.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security governance or strategy leader. Supporting roles include domain owners, risk and audit teams, executive sponsors, and RCCE engineers who ensure maturity measurement is grounded in real evidence rather than opinion.

Cadence

Daily operations focus on capturing lessons and performance data as events occur. Weekly routines focus on updating metrics and reviewing emerging improvement signals. Monthly routines focus on analyzing trends and adjusting improvement priorities. Quarterly routines focus on executive review of maturity progression and investment alignment. Annual routines focus on resetting maturity targets and adapting the framework to new risk realities.

Required meetings and approvals

Periodic maturity reviews require executive participation see that improvement remains a leadership priority. Approval is required for changes to maturity targets or scoring models to ensure consistency. Cross-domain reviews ensure improvements in one area do not create regressions elsewhere.

Escalation paths

Detected maturity regression escalates to the accountable owner for corrective planning. Persistent stagnation escalates to executive leadership for prioritization and resource alignment. Systemic failure to improve escalates to the board or audit committee when it threatens long-term resilience.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish how maturity will be measured. You identify the accountable owner, define maturity levels and evaluation criteria, inventory available evidence sources, and document how lessons learned will be captured. Outputs include a maturity model, baseline assessment, and improvement tracking structure. Evidence at the end of this phase shows that maturity is defined and measurable.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize improvement. You integrate maturity metrics into Noodles, enable AINA analysis of trends and gaps, align improvement actions with domain roadmaps, and begin reporting maturity progression to leadership. Outputs include dashboards, improvement plans, and updated domain targets. Evidence demonstrates that improvement is active and tracked.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make improvement continuous and defensible. You validate maturity scoring against real outcomes, test whether lessons learned lead to measurable change, refine models, and anchor maturity evidence into Rosecoin Vault. Outputs include validated maturity assessments, trend analyses, and immutable records of progress decisions. Evidence proves that improvement is sustained rather than cosmetic.

End state

When this domain is mature, security is no longer treated as a static destination. It becomes a continuously advancing capability. The organization can demonstrate not only where it stands today, but how it has improved, what it has learned, and how it will adapt next—turning experience into resilience and time into an advantage.

DOMAIN 20 — EVIDENCE, LEGAL HOLD & PROVENANCE (ROSECOIN BLOCKCHAIN VAULT)

DOMAIN MISSION AND OUTCOMES

Mission

Evidence, Legal Hold & Provenance exists to ensure that security, compliance, and operational truth can be proven beyond dispute. In modern environments, evidence is fragile: logs can be altered, records overwritten, screenshots manipulated, and timelines rewritten after the fact. This domain ensures that evidence is preserved with integrity, traceability, and legal defensibility from the moment it is generated. It transforms security proof from "trust us" into "verify it."

What this domain prevents

This domain prevents silent evidence tampering, accidental loss of critical records, and disputes over what actually happened. It prevents legal exposure caused by broken chain of custody, incomplete legal holds, or unverifiable audit artifacts. It prevents organizations from failing audits, investigations, or court proceedings because evidence could not be trusted. It also prevents internal erosion of accountability where facts are negotiable or reconstructed retroactively.

What "done" looks like

Evidence, Legal Hold & Provenance is done when every critical security and compliance artifact can be traced back to its origin, time, and integrity state. Evidence is immutable once captured, verifiable independently, and preserved according to legal and regulatory requirements. Legal holds can be applied quickly and comprehensively. Auditors, regulators, and courts can validate authenticity without relying solely on organizational assurances. Leadership can make decisions knowing the underlying data is defensible.

Scope boundaries

This domain includes evidence capture, integrity protection, provenance tracking, legal hold enforcement, retention governance, and immutability through Rosecoin Blockchain Vault. It does not include detection logic, incident response execution, or forensic analysis itself, which are covered in other RCF domains. This domain governs proof, trust, and defensibility.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Evidence, Legal Hold & Provenance sits beneath every RCF domain as a trust anchor. The architecture must support ingestion of evidence from diverse systems, integrity sealing, long-term retention, and independent verification. It spans on-prem, cloud, and hybrid environments and must operate even when source systems are compromised or decommissioned.

Required systems, data sources, and integrations

This domain requires centralized evidence ingestion, secure storage, hashing and integrity mechanisms, legal hold management, and verification interfaces. It consumes evidence from logs, screenshots, reports, configuration states, approvals, investigations, and compliance artifacts generated across all RCF domains. Integration with Rocheston Noodles provides structured evidence management, while the Rosecoin Blockchain Vault anchors cryptographic proofs of existence, integrity, and time. AINA assists by classifying evidence, validating completeness, and correlating artifacts to controls and events.

Data flows

Evidence generated by systems and processes flows into Noodles where it is structured, indexed, and prepared for preservation. Cryptographic hashes and metadata are anchored to the Rosecoin Blockchain Vault, creating an immutable timestamped record. Legal holds can be applied to prevent alteration or deletion. Verification workflows allow independent validation of integrity without exposing sensitive content. Evidence usage for audits, investigations, or litigation is logged to maintain full provenance.

Minimum viable setup vs enterprise setup

A minimum viable setup includes centralized evidence collection, defined retention policies, manual legal hold processes, and integrity checks using cryptographic hashing. An enterprise setup includes automated evidence ingestion from all domains, real-time anchoring to Rosecoin, automated legal hold enforcement, cross-jurisdiction retention management, independent verification portals, and executive dashboards showing evidence coverage and integrity status.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a governance, legal, or security assurance leader. Supporting roles include legal counsel, compliance teams, security operations, IT platform teams, and RCCE engineers who design evidence pipelines that are reliable, automated, and defensible.

Cadence

Daily operations focus on monitoring evidence ingestion, integrity status, and legal hold enforcement. Weekly routines focus on validating coverage across domains and resolving ingestion gaps. Monthly routines focus on retention review, legal hold audits, and evidence quality checks. Quarterly routines focus on executive review of evidence posture and readiness for audits or litigation. Annual routines focus on reassessing retention strategy, jurisdictional requirements, and blockchain anchoring policies.

Required meetings and approvals

Legal holds require formal legal approval and documented scope. Changes to retention or immutability policies require executive and legal sign-off due to their regulatory impact. Periodic reviews ensure evidence practices remain aligned with evolving laws and business risk.

Escalation paths

Evidence integrity failures escalate immediately to the domain owner and legal leadership. Incomplete legal holds or retention violations escalate to executive leadership. Systemic evidence weaknesses escalate to the board or audit committee when they threaten regulatory, legal, or trust outcomes.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish evidence authority and integrity foundations. You identify the accountable owner, define evidence categories and retention requirements, inventory evidence sources across RCF domains, establish legal hold procedures, and enable basic cryptographic integrity controls. Outputs include an evidence framework, source inventory, retention schedules, and legal hold workflows. Evidence at the end of this phase shows that proof is intentionally managed and protected.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize immutable evidence. You integrate evidence ingestion into Noodles, enable automated hashing and anchoring into Rosecoin Blockchain Vault, apply retention and legal hold enforcement, and enable AINA-assisted classification and correlation. Outputs include anchored evidence records, dashboards, verification workflows, and legal hold logs. Evidence demonstrates that artifacts are immutable, traceable, and verifiable.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make evidence defensible under scrutiny. You test legal hold activation and release, validate independent verification, simulate audit and legal requests, and anchor governance decisions into Rosecoin. Outputs include validation reports, readiness assessments, and immutable provenance records. Evidence proves that the organization can defend its security and compliance claims with cryptographic certainty.

End state

When this domain is mature, trust no longer depends on reputation or explanation. It depends on proof. Evidence is immutable, verifiable, and preserved with intent. Audits, investigations, and legal challenges shift from stress events to controlled processes. The organization can stand behind its claims—because the record cannot be rewritten.

DOMAIN 21 — AI AGENT GOVERNANCE & RUNTIME CONTROLS

DOMAIN MISSION AND OUTCOMES

Mission

AI Agent Governance & Runtime Controls exists to ensure autonomous and semi-autonomous AI agents operate within strict boundaries, remain accountable, and cannot become uncontrolled actors inside the organization. Unlike traditional software, agents can plan, decide, take actions, call tools, move across systems, and interact with humans in ways that are dynamic and unpredictable if not constrained. This domain ensures agents are governed like high-privilege operators: allowed to act, but only within defined authority, with continuous oversight, and with provable logs of every decision and action.

What this domain prevents

This domain prevents agents from taking unauthorized actions, accessing data beyond their scope, leaking sensitive information, or being manipulated through prompt injection and tool abuse. It prevents "agent drift" where an agent gradually expands its behavior beyond its intended purpose. It prevents shadow agents deployed without oversight by teams trying to move fast. It also prevents accountability collapse where nobody can explain why an agent acted, what it saw, what tools it used, and what data it touched.

What "done" looks like

AI Agent Governance & Runtime Controls is done when every agent has a defined owner, explicit purpose, bounded permissions, and measurable operational limits. Every tool call, data access, and action is logged with provenance. Agents operate under policy guardrails that enforce least privilege, safe outputs, and permitted workflows. Agents can be paused, contained, or revoked instantly. The organization can prove which agents exist, what they are allowed to do, what they actually did, and whether their behavior remained within approved boundaries.

Scope boundaries

This domain includes agent inventory, purpose and authority definition, runtime policy enforcement, tool and data access control, output safety, auditing, and kill-switch mechanisms. It does not include general AI model lifecycle governance, network segmentation, or incident response execution, which are handled in other RCF domains. This domain governs the operational control and accountability of agents in production.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

AI Agent Governance & Runtime Controls sits at the intersection of identity, policy enforcement, data protection, and execution control. The architecture must treat agents as privileged identities with controlled access to systems and tools. It must enforce runtime policies close to execution points, while maintaining centralized visibility and governance. This architecture applies to on-prem, cloud, and hybrid environments, and to both internally built agents and externally sourced agent platforms.

Required systems, data sources, and integrations

This domain requires an agent registry, identity and access controls for agent accounts, policy engines that enforce runtime constraints, tool execution gateways, and comprehensive logging. It consumes data from agent prompts, tool calls, outputs, decision traces, and system interactions. Integration with Rocheston Noodles provides centralized governance evidence and dashboards, while AINA monitors agent behavior for drift, unsafe actions, and policy violations and can trigger containment actions.

Data flows

Agent definitions and permissions are created and approved through governance workflows, then published into the runtime control layer. At runtime, agent requests and tool calls flow through enforcement gateways where policies validate intent, permissions, data sensitivity, and allowed actions. Outputs and actions are logged into Noodles with provenance. AINA evaluates behavior patterns to detect drift, manipulation attempts, or abnormal tool use. Governance decisions and remediation actions flow back into agent permissions and runtime guardrails.

Minimum viable setup vs enterprise setup

A minimum viable setup includes an inventory of agents, assigned owners, basic permission boundaries, centralized logging of actions, and a manual kill-switch process. An enterprise setup includes runtime policy enforcement, scoped tool gateways, continuous drift detection, segmented agent identities, automated containment, integration with Rosecoin

Vault for immutable records, and executive dashboards showing agent posture, risk, and activity.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically an AI governance or security architecture leader. Supporting roles include AI engineering teams, identity teams, data owners, legal and compliance representatives, security operations, and RCCE engineers who operationalize agent control planes and evidence pipelines.

Cadence

Daily operations focus on monitoring agent activity, policy violations, and abnormal behavior. Weekly routines focus on reviewing newly proposed agents, permission changes, and tool integrations. Monthly routines focus on analyzing behavior trends, drift indicators, and validating that agents remain aligned to approved purpose. Quarterly routines focus on executive review of agent risk posture and changes to authority boundaries. Annual routines focus on refreshing agent governance standards to reflect evolving technology and threat models.

Required meetings and approvals

New agents require formal approval before production deployment. Permission expansions require documented justification and approval. Tool integrations require security review to ensure they do not become privilege escalation paths. High-impact agent use cases require legal and governance review when they affect regulated data, customer decisions, or critical operations.

Escalation paths

Unsafe agent behavior escalates immediately to the domain owner and security operations for containment. Agent activity involving regulated data or critical actions escalates to executive and legal leadership. Systemic agent governance failures escalate to the board or audit committee when they materially affect organizational risk and trust.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish control over agents. You identify the accountable owner, inventory existing agents and agent-like automations, define governance standards for agent purpose and permissions, establish approval workflows, and implement basic logging. Outputs include an agent registry, ownership assignments, approval processes, and baseline policy definitions. Evidence at the end of this phase shows that agents are known, owned, and visible.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize runtime controls. You deploy enforcement gateways for tool calls, integrate identity-based permissions for agents, connect activity logs into Noodles, enable AINA monitoring for drift and policy violations, and implement kill-switch mechanisms. Outputs include controlled runtime environments, dashboards, policy enforcement logs, and incident-ready audit trails. Evidence demonstrates that agents cannot act outside approved boundaries without detection.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make agent governance defensible and resilient. You test prompt injection and tool abuse scenarios, validate containment and kill-switch actions, refine policies based on real usage, and anchor agent activity evidence into Rosecoin Vault. Outputs include test results, validated governance controls, drift analysis reports, and immutable agent provenance records. Evidence proves that agent behavior is controlled, auditable, and trustworthy.

End state

When this domain is mature, AI agents become safe to operate as part of the organization's core systems. They are treated as privileged actors with strict boundaries, continuous oversight, and provable accountability. The organization can innovate with agents confidently because autonomy is governed, runtime behavior is controlled, and trust is earned through evidence rather than assumption.

DOMAIN 22 — SPACE & ORBITAL SECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Space & Orbital Security exists to ensure that assets and services dependent on space-based infrastructure remain trustworthy, available, and resilient against interference, failure, or hostile action. Modern organizations increasingly rely on satellites and orbital systems for communications, navigation, timing, earth observation, and global connectivity. This domain ensures that space dependencies are understood, secured, and governed so orbital risk does not become an invisible single point of failure.

What this domain prevents

This domain prevents disruption caused by loss, degradation, or manipulation of satellite-based services such as GPS, timing signals, communications links, and remote sensing data. It prevents organizations from assuming space infrastructure is always available or neutral. It prevents blind reliance on third-party orbital providers without understanding control, resilience, or geopolitical exposure. It also prevents cascading failures where loss of space services silently breaks terrestrial systems, cloud services, logistics, or safety-critical operations.

What "done" looks like

Space & Orbital Security is done when the organization clearly understands which business functions depend on orbital systems and how resilient those dependencies are. Satellite-based services are inventoried, monitored, and integrated into continuity and risk planning. Disruption scenarios are anticipated and mitigated. Alternative capabilities exist where required. Leadership understands space-related risk exposure and can make informed decisions about reliance, redundancy, and response.

Scope boundaries

This domain includes identification and governance of space-based dependencies, orbital threat awareness, resilience planning for satellite services, and coordination with providers and regulators. It does not include terrestrial network hardening, endpoint security, or incident response execution, which are covered in other RCF domains. This domain governs space-related dependencies and risks.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Space & Orbital Security operates as a dependency and resilience layer rather than a standalone system. The architecture must map space-based services into enterprise architectures and continuity plans. It applies across on-prem, cloud, and hybrid environments wherever satellite connectivity, navigation, timing, or observation data is used.

Required systems, data sources, and integrations

This domain requires an inventory of space-dependent services, contracts and SLAs with orbital providers, monitoring of satellite service availability, and integration with business continuity and risk systems. It consumes data from communications platforms, navigation and timing systems, IoT and edge deployments, and third-party provider status feeds. Integration with Rocheston Noodles provides centralized visibility into orbital dependencies and resilience evidence, while AINA evaluates impact scenarios, provider risk, and dependency concentration.

Data flows

Dependency information flows into a centralized registry linking business functions to space services. Availability and status signals from providers flow into monitoring systems. AINA correlates disruptions or degradations with business impact to assess risk in real time. Insights flow into governance, continuity planning, and executive reporting so space-related risk is actively managed rather than assumed.

Minimum viable setup vs enterprise setup

A minimum viable setup includes identification of space dependencies, basic monitoring of service availability, and documented response procedures for outages. An enterprise setup includes redundancy planning across providers, alternative terrestrial capabilities, continuous monitoring and alerting, integration with resilience exercises, and executive dashboards showing orbital risk exposure and dependency criticality.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically within security architecture, resilience, or critical infrastructure leadership. Supporting roles include network and communications teams, cloud and edge platform owners, legal and procurement teams managing provider relationships, and RCCE engineers who integrate orbital risk into enterprise security and resilience planning.

Cadence

Daily operations focus on monitoring availability of critical space-based services. Weekly routines focus on reviewing provider updates and changes in dependency. Monthly routines focus on reassessing criticality and testing fallback assumptions. Quarterly routines focus on executive review of orbital risk and dependency concentration. Annual routines focus on updating space risk assessments based on geopolitical, technological, and regulatory developments.

Required meetings and approvals

Introduction of new space-dependent services requires security and resilience review. Changes to reliance levels or providers require approval when they affect critical operations. Executive review is required when space dependencies represent material business risk.

Escalation paths

Disruption of critical space services escalates immediately to the domain owner and incident coordination teams. Risks affecting safety, regulated operations, or national infrastructure escalate to executive leadership. Strategic or geopolitical space risk escalates to the board or audit committee when it materially affects long-term resilience.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish visibility into space dependencies. You identify the accountable owner, inventory all space-based services in use, map them to business functions, and document basic outage response procedures. Outputs include a space dependency register, criticality assessments, and response playbooks. Evidence at the end of this phase shows that orbital reliance is known and owned.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize orbital resilience. You integrate service monitoring into Noodles, enable AINA analysis of dependency impact, assess provider resilience and contracts, and incorporate space disruption scenarios into continuity planning. Outputs include monitoring dashboards, risk assessments, and updated continuity plans. Evidence demonstrates that space dependencies are actively managed.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make orbital risk survivable. You test disruption scenarios, validate fallback capabilities, refine response procedures, and anchor dependency and resilience evidence into Rosecoin Vault. Outputs include exercise results, validated mitigation plans, and immutable records of orbital risk decisions. Evidence proves that space-related disruption will not cause uncontrolled failure.

End state

When this domain is mature, space is no longer an invisible dependency. Orbital services are treated as critical infrastructure with known risk, monitored availability, and deliberate resilience planning. The organization can operate confidently in a world where space is contested, complex, and essential—because reliance is intentional, governed, and defensible.

DOMAIN 23 — SUSTAINABLE (GREEN) CYBERSECURITY

DOMAIN MISSION AND OUTCOMES

Mission

Sustainable (Green) Cybersecurity exists to ensure security can scale without exhausting resources, people, or the planet. Traditional security programs often grow by adding tools, agents, data retention, and compute until cost, complexity, and energy usage spiral out of control. This domain ensures cybersecurity is efficient, intentional, and environmentally responsible while remaining effective. It treats sustainability as an operational requirement, not a marketing claim.

What this domain prevents

This domain prevents runaway security sprawl that increases carbon footprint, operational cost, and system fragility. It prevents inefficient architectures that duplicate telemetry, over-retain data, and burn compute for marginal value. It prevents security programs from becoming unsustainable to operate, staff, or fund. It also prevents ESG and regulatory exposure where cybersecurity operations contradict stated sustainability commitments or create hidden environmental impact.

What "done" looks like

Sustainable Cybersecurity is done when security outcomes improve while resource consumption is reduced or stabilized. Telemetry is purposeful rather than excessive. Data retention is intentional. Compute usage is right-sized. Security tooling is consolidated where possible. Leaders can see the energy, cost, and environmental impact of security operations alongside effectiveness metrics. Sustainability decisions are explicit tradeoffs, not accidental consequences.

Scope boundaries

This domain includes efficiency of security architecture, data minimization, compute and storage optimization, tooling consolidation, lifecycle management, and environmental impact measurement related to cybersecurity operations. It does not include enterprise-wide sustainability programs, facilities management, or carbon accounting outside security scope. This domain governs how security is built to last.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Sustainable Cybersecurity overlays all security architectures. The design principle is “maximum signal, minimum waste.” Architectures must reduce redundant data flows, centralize analysis, and favor shared control planes over isolated tools. This applies across on-prem, cloud, and hybrid environments, with particular focus on telemetry pipelines, analytics platforms, and long-term storage.

Required systems, data sources, and integrations

This domain requires visibility into security compute usage, data ingestion volumes, storage growth, and tooling overlap. It consumes metrics from SIEM, XDR, cloud platforms, endpoint systems, and data retention systems. Integration with Rocheston Noodles provides consolidated visibility into security data flows and evidence volume, while AINA evaluates signal quality, redundancy, and efficiency opportunities.

Data flows

Security telemetry flows into centralized platforms where duplication and low-value data can be identified. AINA analyzes which signals drive detections, decisions, and outcomes, and which consume resources without benefit. Optimization decisions flow back into architecture standards, retention policies, and tooling configurations so sustainability improvements are enforced, not advisory.

Minimum viable setup vs enterprise setup

A minimum viable setup includes basic visibility into security data volumes, defined retention limits, and periodic review of tooling overlap. An enterprise setup includes continuous telemetry optimization, dynamic retention based on value and risk, consolidated analytics platforms, energy-aware cloud usage, and executive dashboards showing sustainability and security effectiveness together.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically within security architecture or governance. Supporting roles include platform teams, cloud and infrastructure owners, finance and ESG stakeholders, and RCCE engineers who design efficient security architectures that balance protection with sustainability.

Cadence

Daily operations focus on monitoring abnormal growth in data, compute, or cost. Weekly routines focus on reviewing efficiency signals and tuning configurations. Monthly routines focus on consolidation opportunities, retention adjustments, and impact reporting. Quarterly routines focus on executive review of sustainability metrics and tradeoffs. Annual routines focus on aligning security architecture with organizational sustainability goals and regulatory expectations.

Required meetings and approvals

Major increases in security telemetry, retention, or tooling require architectural review. Changes affecting sustainability commitments require executive awareness. Consolidation or decommissioning decisions require coordination across security, IT, and business stakeholders.

Escalation paths

Runaway cost or resource consumption escalates to the domain owner for corrective action. Sustainability risks affecting ESG commitments escalate to executive leadership. Structural inefficiencies that threaten long-term viability escalate to the board or audit committee when they materially affect cost, reputation, or resilience.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish visibility into security resource usage. You identify the accountable owner, inventory security tools and data flows, measure baseline compute, storage, and retention, and define sustainability principles. Outputs include a security sustainability baseline, tooling inventory, and initial optimization criteria. Evidence at the end of this phase shows that impact is measured, not assumed.

Phase 2: Implement (days 15–60)
In Phase 2 you operationalize efficiency controls. You consolidate overlapping tools where possible, tune telemetry to focus on high-value signals, adjust retention based on risk and regulatory need, integrate sustainability metrics into Noodles, and enable AINA analysis of efficiency. Outputs include optimized pipelines, dashboards, and documented tradeoff decisions. Evidence demonstrates that security effectiveness is maintained or improved with lower resource usage.

Phase 3: Harden + validate (days 61–90)
In Phase 3 you make sustainability durable. You validate that optimizations do not degrade detection or response, automate efficiency enforcement, anchor sustainability decisions into Rosecoin Vault, and report outcomes to leadership. Outputs include validated metrics, continuous enforcement rules, and immutable records of sustainability decisions. Evidence proves that green cybersecurity is intentional, measurable, and repeatable.

End state
When this domain is mature, cybersecurity no longer grows by consuming more resources indiscriminately. It grows smarter. Security operations deliver stronger outcomes with less waste, lower cost, and reduced environmental impact. Sustainability becomes a force multiplier for resilience—ensuring the security program can endure, scale, and remain credible long into the future.

DOMAIN 24 — NEURO-COGNITIVE SECURITY & HUMAN FACTORS

DOMAIN MISSION AND OUTCOMES

Mission

Neuro-Cognitive Security & Human Factors exists to protect the human mind as a security surface. Modern attacks no longer target only systems, credentials, or networks—they target perception, attention, emotion, trust, and decision-making. This domain ensures the organization understands, mitigates, and governs cognitive risk so humans are not exploited as unprotected attack paths. It treats psychological manipulation, cognitive overload, and decision fatigue as real security threats, not soft concerns.

What this domain prevents

This domain prevents manipulation-based attacks such as social engineering, deepfake-driven deception, misinformation campaigns, authority spoofing, urgency exploitation, and decision hijacking. It prevents security breakdowns caused by alert fatigue, cognitive overload, poor interface design, and stress-induced errors. It prevents executives and operators from making high-impact decisions under manipulated or distorted conditions. It also prevents organizations from underestimating human cognitive limits in high-pressure security environments.

What "done" looks like

Neuro-Cognitive Security is done when human decision-making is protected, supported, and resilient under stress. Security workflows reduce cognitive load instead of increasing it. Interfaces present clear, prioritized information rather than noise. High-risk decisions are supported by structure, verification, and delay where appropriate. The organization can demonstrate that humans are less susceptible to manipulation, deception, and fatigue-driven error, even during incidents or crises.

Scope boundaries

This domain includes cognitive threat modeling, human-centered security design, decision integrity controls, deception awareness, deepfake and influence risk mitigation, and cognitive load management in security operations. It does not include general security awareness training, identity enforcement, or technical detection systems, which are covered in other RCF domains. This domain governs how humans perceive, decide, and act under security pressure.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Neuro-Cognitive Security operates across human-system interaction layers rather than traditional infrastructure. The architecture focuses on decision points, interfaces, workflows, and communication channels where cognitive manipulation or overload can occur. It applies across SOC environments, executive decision processes, crisis communications, and any system where humans interpret security signals.

Required systems, data sources, and integrations

This domain requires visibility into decision workflows, alerting systems, communication tools, and user interfaces. It consumes data from SOC platforms, incident response tools, executive briefings, training simulations, and user feedback. Integration with Rocheston Noodles provides centralized evidence of human-factor controls, while AINA evaluates cognitive risk patterns such as alert overload, decision delays, conflicting signals, and manipulation indicators.

Data flows

Security signals flow through human-facing systems where prioritization, clarity, and context are critical. AINA analyzes signal volume, timing, and presentation to identify cognitive stress points. Feedback from users and outcomes flows back into design and governance processes so interfaces and workflows are refined continuously. Evidence of cognitive controls and improvements flows into Noodles for validation and reporting.

Minimum viable setup vs enterprise setup

A minimum viable setup includes identification of critical decision points, reduction of unnecessary alerts, basic verification procedures for high-risk decisions, and guidance for handling manipulation attempts. An enterprise setup includes cognitive load-aware SOC design, decision support systems, structured verification and delay mechanisms for sensitive actions, deepfake and influence scenario planning, and executive dashboards showing cognitive risk indicators.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically within security governance or human risk leadership. Supporting roles include SOC leadership, UX and systems designers, communications teams, executive staff, behavioral specialists, and RCCE engineers who integrate human-factor protections into technical systems and workflows.

Cadence

Daily operations focus on monitoring alert volume, decision bottlenecks, and signs of cognitive overload. Weekly routines focus on reviewing near-miss decisions, manipulation attempts, and workflow friction. Monthly routines focus on refining interfaces, decision aids, and communication protocols. Quarterly routines focus on executive review of cognitive risk and resilience. Annual routines focus on updating threat models to reflect evolving manipulation techniques and technologies.

Required meetings and approvals

Changes to high-impact decision workflows require governance review to ensure cognitive risk is reduced, not increased. Crisis communication protocols require executive approval. Introduction of monitoring or behavioral controls requires legal and ethical review to ensure respect for individual rights.

Escalation paths

Detected manipulation attempts escalate to the domain owner and security leadership for containment and communication control. Cognitive overload affecting critical operations escalates to executive leadership for intervention. Systemic human-factor risk escalates to the board or audit committee when it materially affects organizational decision integrity.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish awareness of cognitive risk. You identify the accountable owner, map critical human decision points, assess alert and information overload, and document manipulation risk scenarios. Outputs include a cognitive risk map, prioritized decision points, and baseline design principles. Evidence at the end of this phase shows that human risk is identified and owned.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize cognitive protections. You redesign high-risk workflows to reduce load, implement verification and delay controls for sensitive actions, integrate cognitive risk indicators into Noodles, enable AINA analysis of human-factor signals, and conduct simulations involving manipulation and deception. Outputs include improved interfaces, decision aids, simulation results, and dashboards. Evidence demonstrates that human decisions are more resilient and supported.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make cognitive security durable. You test workflows under stress, validate resistance to deception scenarios, refine controls based on outcomes, and anchor human-factor evidence into Rosecoin Vault. Outputs include validated decision processes, resilience assessments, and immutable records of design and governance decisions. Evidence proves that cognitive risk is actively managed and reduced.

End state

When this domain is mature, humans are no longer the soft underbelly of the security program. Decision-making is protected, clarity is prioritized, and manipulation is resisted by design. The organization treats the human mind as a critical security surface—supported, respected, and defended with the same rigor as any technical system.

DOMAIN 25 — META-GOVERNANCE & FRAMEWORK EVOLUTION

DOMAIN MISSION AND OUTCOMES

Mission

Meta-Governance & Framework Evolution exists to ensure the security framework itself does not become obsolete. Threats change, technology evolves, regulations shift, and organizations transform. A static framework—no matter how strong at launch—will eventually fail. This domain ensures RCF remains adaptive, self-correcting, and future-aligned. It governs how the framework learns, evolves, retires outdated assumptions, and incorporates new risk realities without fragmentation or loss of integrity.

What this domain prevents

This domain prevents framework stagnation where controls no longer match real threats or operating models. It prevents uncontrolled customization that breaks consistency across domains and regions. It prevents reactive, piecemeal updates driven by audits or incidents rather than strategy. It also prevents dependency on individuals or institutional memory to decide "how security should change," ensuring evolution is intentional, governed, and evidence-based.

What "done" looks like

Meta-Governance & Framework Evolution is done when the framework can change without chaos. Updates are deliberate, traceable, and validated against evidence and outcomes. Deprecated controls are retired cleanly. New domains or controls are introduced without breaking existing implementations. Leadership can see why the framework evolved, what changed, and how those changes improve resilience. RCF remains coherent, relevant, and defensible over time rather than slowly drifting into irrelevance.

Scope boundaries

This domain includes framework versioning, control lifecycle management, cross-domain alignment, change governance, and evolution strategy. It does not include day-to-day control operation, detection, or response execution, which are governed by individual domains. This domain governs the framework as a living system.

DOMAIN ARCHITECTURE BLUEPRINT

Reference architecture

Meta-Governance & Framework Evolution operates above all RCF domains as the control plane for change. The architecture must support versioned frameworks, controlled updates, impact analysis, and backward compatibility. It applies across global implementations so evolution does not fragment regional or organizational deployments.

Required systems, data sources, and integrations

This domain requires access to evidence, metrics, outcomes, and lessons from all RCF domains. It consumes data from maturity assessments, incidents, regulatory changes, technology trends, and operational feedback. Integration with Rocheston Noodles provides centralized visibility into framework performance and change impact, while AINA analyzes trends, detects misalignment, and recommends evolution paths based on real-world signals.

Data flows

Evidence and performance data from all domains flow into a centralized analysis layer. AINA evaluates effectiveness, drift, and emerging gaps. Proposed framework changes flow through governance workflows for review, testing, and approval. Approved updates flow back into domain standards, implementation guides, and training so evolution is consistent and controlled. All changes are logged with rationale and impact assessment.

Minimum viable setup vs enterprise setup

A minimum viable setup includes defined framework ownership, documented change procedures, version tracking, and periodic review of relevance. An enterprise setup includes continuous evaluation of framework effectiveness, automated impact analysis, staged rollout of updates, backward compatibility controls, and executive dashboards showing framework health and evolution velocity.

DOMAIN OPERATING MODEL

Roles and ownership

This domain requires a single accountable owner, typically a security governance or framework authority. Supporting roles include domain owners, legal and compliance leadership, technology strategists, and RCCE engineers who ensure framework changes translate into operable controls rather than abstract revisions.

Cadence

Daily operations focus on capturing signals that indicate misalignment or emerging risk. Weekly routines focus on reviewing proposed changes and feedback. Monthly routines focus on analyzing framework performance and relevance. Quarterly routines focus on executive review of evolution priorities and strategic alignment. Annual routines focus on major framework version updates and long-term roadmap planning.

Required meetings and approvals

Framework changes require formal review and approval to maintain consistency and defensibility. Major evolution decisions require executive sponsorship. Cross-domain reviews ensure updates improve resilience without creating gaps or conflicts.

Escalation paths

Detected framework misalignment with real-world risk escalates to the domain owner for corrective action. Conflicting or fragmented framework changes escalate to executive leadership. Strategic failure to evolve escalates to the board or audit committee when it threatens long-term resilience and credibility.

IMPLEMENTATION ROADMAP

Phase 1: Stand up (first 7–14 days)

In Phase 1 you establish authority over framework evolution. You identify the accountable owner, define governance for framework changes, inventory existing controls and dependencies, and document versioning and approval processes. Outputs include a meta-governance charter, framework lifecycle definitions, and change workflows. Evidence at the end of this phase shows that framework evolution is owned and governed.

Phase 2: Implement (days 15–60)

In Phase 2 you operationalize controlled evolution. You integrate framework performance data into Noodles, enable AINA analysis of effectiveness and drift, establish impact assessment for proposed changes, and begin controlled updates to selected domains. Outputs include versioned framework updates, impact reports, and communication artifacts. Evidence demonstrates that evolution improves alignment without disruption.

Phase 3: Harden + validate (days 61–90)

In Phase 3 you make evolution resilient and defensible. You validate backward compatibility, test update rollouts, anchor framework change records into Rosecoin Vault, and assess governance effectiveness. Outputs include validated framework versions, evolution metrics, and immutable provenance of change decisions. Evidence proves that the framework can evolve without losing integrity or trust.

End state

When this domain is mature, RCF is no longer a static standard frozen in time. It becomes a living system—capable of learning, adapting, and improving as the world changes. The framework evolves with intention, governed by evidence and strategy rather than pressure or trend. Security remains coherent, resilient, and future-ready because the framework itself is designed to survive change.

ROCHESTON CYBERSECURITY FRAMEWORK (RCF) CHECKLIST

DOMAIN 1: GOVERNANCE & POLICY

1.1.1 Is there a board-approved Cybersecurity Charter that defines the role of security in business growth?

1.1.2 Does the board have a designated Cyber Security Subject Matter Expert (SME) or an external advisor?

1.1.3 Is cybersecurity a standing agenda item for every quarterly board meeting?

1.1.4 Are executive bonuses and KPIs mathematically linked to the organization's RCF Security Score?

1.1.5 Does the board receive a "Cost of Inaction" report detailing the financial loss of deferred security upgrades?

1.1.6 Is there a formal process for the board to review and sign off on "High-Risk" technology deployments?

1.1.7 Does the board participate in an annual Tabletop Exercise (TTX) simulating a catastrophic breach?

1.1.8 Is the board updated on the security posture of all Mergers & Acquisitions (M&A) targets pre-closing?

1.1.9 Is there a board-level "Succession Plan" for the CISO and key security leadership?

1.1.10 Does the board approve the organization's "Risk Appetite" for AI and Autonomous systems?

1.2.1 Does the CISO have an independent budget that cannot be repurposed by the IT department?

1.2.2 Is the CISO empowered to "Stop Production" if a critical security flaw is detected?

1.2.3 Are cybersecurity roles defined using the US DoD 8140 and NICE Framework standards?

1.2.4 Is there a clear separation of duties between the team that builds systems and the team that secures them?

1.2.5 Does the security team have a 24/7/365 operational mandate?

1.2.6 Is the CISO required to hold an accredited certification like the Rocheston RCCE?

1.2.7 Are "Security Champions" appointed within every non-technical business unit?

1.2.8 Is there a formal internal "Cyber Council" with representatives from Legal, HR, and Finance?

1.2.9 Does the organizational chart clearly show the CISO reporting to the CEO or Board?

1.2.10 Is there a dedicated budget for "Offensive Security" (Domain 13) independent of defensive operations?

1.3.1 Does the system maintain a live "Regulatory Inventory" of all 1,000+ global cybersecurity laws?

1.3.2 Is there an automated cross-walk between RCF domains and NIST, ISO, HIPAA, and GDPR?

1.3.3 Is the "Digital Sovereignty" of data managed according to the specific laws of the host nation?

1.3.4 Does the organization have a pre-registered "Data Protection Officer" (DPO) in required jurisdictions?

1.3.5 Are "Standard Contractual Clauses" (SCCs) automatically included in all international data contracts?

1.3.6 Is there a mechanism to monitor for "Shadow Regulations"—emerging laws that haven't passed yet?

1.3.7 Does the leadership track compliance for "Non-Traditional" assets like Satellites and Neural links?

1.3.8 Is there a "72-Hour Breach Notification" protocol that is automated via the Rosecoin Ledger?

1.3.9 Are "Right to be Forgotten" requests handled through a board-approved automated workflow?

1.3.10 Is there a legal "Safe Harbor" protocol for ethical hackers who find vulnerabilities in the company?

1.4.1 Are all written security policies mirrored by an executable configuration file in AINA OS?

1.4.2 Is there a prohibition against "Manual Policy Overrides" without a dual-signature blockchain key?

1.4.3 Are policy versions tracked via an immutable Git-based repository?

1.4.4 Is the "Acceptable Use Policy" (AUP) integrated into the system login screen for all users?

1.4.5 Are data retention policies enforced by "Auto-Delete" scripts rather than manual cleanup?

1.4.6 Is there a "Zero-Trust" policy for all remote and internal network connections?

1.4.7 Are "Shadow IT" discovery policies enforced through automated network blocking?

1.4.8 Is the "AI Ethics Policy" hard-coded into the guardrails of the organization's LLMs?

1.4.9 Is there a "Clean Desk" policy for both physical and digital (desktop/cloud) environments?

1.4.10 Are "Bring Your Own Device" (BYOD) policies enforced through mandatory containerization?

1.5.1 Is the "Universal Audit Trail" stored exclusively on the Rosecoin AI Blockchain?

1.5.2 Are all executive "Risk Acceptance" forms cryptographically signed and non-erasable?

1.5.3 Can a regulator be granted a "Read-Only" blockchain key to verify compliance in real-time?

1.5.4 Is every administrative "Rule Change" logged with a unique biometric hash of the admin?

1.5.5 Is there a "Tamper-Alert" system if any legacy log attempt to overwrite the blockchain?

1.5.6 Does the Rosecoin Ledger record the exact timestamp of every "Self-Healing" event?

1.5.7 Are compliance certificates issued as NFTs to ensure they cannot be forged?

1.5.8 Is there a "Proof of Testing" record for every Pen Test stored on the ledger?

1.5.9 Are external audits validated against the internal blockchain record for discrepancies?

1.5.10 Is the "Master Evidence Vault" geographically distributed across multiple nodes?

1.6.1 Is there an "AI Liability" framework that defines legal responsibility for autonomous decisions?

1.6.2 Does the leadership have a "Quantum Migration Plan" with a set completion date?

1.6.3 Is "Green Cybersecurity" (energy efficiency) included in the annual corporate ESG report?

1.6.4 Are there ethical guidelines for the use of "Cognitive Biometrics" (Domain 24)?

1.6.5 Is there a "Digital Heritage" policy for data handling after a business merger or dissolution?

1.6.6 Does the organization advocate for global cyber-peace through the support of international norms?

1.6.7 Is there a protocol for "Algorithmic Transparency" to explain AI decisions to users?

1.6.8 Are "Biometric Privacy" rules enforced for all neuro-cognitive and physical data?

1.6.9 Does the leadership mandate "Accessibility" in all security tools for disabled employees?

1.6.10 Is there a "Cyber-Sustainability" goal to reduce the carbon footprint of the SOC by 50%?

1.7.1 Is there a Board-approved "Ransomware Payment Policy" that strictly defines if, when, and how a ransom could legally be paid?

1.7.2 Does the organization have a pre-authorized "Crypto-Wallet" or legal fiat mechanism for emergency payments if life-safety is at risk?

1.7.3 Is there a designated "War Room Commander" (non-technical) authorized to make business-survival decisions during a total outage?

1.7.4 Are there specific governance protocols for handling "Nation-State Threats" (e.g., disconnection from the global internet to protect data)?

1.7.5 Is there a "High-Risk Leaver" policy for C-Level executives joining competitors to prevent IP theft and data exfiltration?

1.7.6 Does the Board have a pre-signed "Authorization Memo" for offensive active defense (hacking back) if legal in the jurisdiction?

1.7.7 Is there a "Deepfake Defense" protocol for verifying the identity of the CEO/CFO during urgent video calls or money transfers?

1.7.8 Are "Rumor Control" channels established to counter disinformation campaigns on social media during a crisis?

1.7.9 Is there a "Golden Copy" governance rule ensuring the most critical data is stored offline and requires physical keys to access?

1.7.10 Does the organization maintain a "Secret Clearance" roster for liaising with government intelligence agencies (e.g., FBI, CISA)?

1.8.1 Does the CISO have specific Directors and Officers (D&O) Liability Insurance coverage to protect personal assets from regulatory lawsuits?

1.8.2 Is there an "Executive Digital Protection" program that secures the home networks and personal devices of the CEO and Board?

1.8.3 Are there specific governance protocols for "Divestitures" (Spin-offs) to ensure secure separation of data and IT assets?

1.8.4 Is there a "Reverse Due Diligence" process to ensure buyers of the company meet RCF security standards before data transfer?

1.8.5 Are "VIP Travel Protocols" in place to provide burner devices and secure comms for executives traveling to high-risk nations?

1.8.6 Is the identity of the CISO and security team shielded from public records where possible to prevent targeted coercion?

1.8.7 Is there a "Key Person Risk" policy that mandates cross-training for the sole administrators of critical cryptographic keys?

1.8.8 Does the organization allow for "Anonymous Reporting" of safety concerns to the Board, bypassing the executive chain?

1.8.9 Is there a formal "Conflict of Interest" register for security leadership regarding vendor selection?

1.8.10 Does the RCF Governance engine automatically freeze the accounts of executives under internal investigation?

1.8.11 Material Incident Disclosure Automation: Is there an automated workflow to trigger regulatory "Material Incident" disclosures (e.g., SEC 8-K or NIS2 24-hour notifications) within mandatory legal windows?

1.8.12 Cyber Insurance Compliance Audit: Does the organization conduct monthly audits to ensure current controls (MFA, EDR, Backups) still meet the evolving "minimum eligibility" requirements of the cyber insurance provider?

1.8.13 Real-Time Transparency Ledger: Does the organization provide a live, read-only dashboard for stakeholders to verify current security health metrics without relying on static, point-in-time audit reports?

1.8.14 Autonomous Policy Enforcement: Are corporate security policies translated into "active code" that automatically blocks non-compliant configurations in real-time?

DOMAIN 2: RISK QUANTIFICATION & VALUE

2.1.1 Is there a real-time, automated "Asset Discovery" engine running continuously to identify every IP-connected device?

2.1.2 Does the system automatically assign a specific monetary value ($) to every asset based on its role in revenue generation?

2.1.3 Is data classified by "value-density" (e.g., how much PII/IP is in a specific database) rather than just generic labels?

2.1.4 Are "Shadow IT" assets automatically detected, quarantined, and assigned a high-risk valuation until authorized?

2.1.5 Is the "Replacement Cost" versus "Recovery Cost" calculated for all Critical Infrastructure assets?

2.1.6 Are Intellectual Property (IP) assets (code, designs, trade secrets) tagged with digital watermarks for valuation tracking?

2.1.7 Does the asset inventory link directly to the Business Continuity Plan to prioritize high-value recovery?

2.1.8 Are "Ghost Assets" (unused but running servers) automatically identified to reduce attack surface and cost?

2.1.9 Is the "Brand Value" at risk calculated dynamically based on social media sentiment analysis during a threat?

2.1.10 Are "Human Assets" (Key Personnel) included in the risk register with defined impact values if they are targeted?

2.2.1 Does the organization reject "High/Medium/Low" heatmaps in favor of quantitative models like FAIR (Factor Analysis of Information Risk)?

2.2.2 Is the "Annualized Loss Expectancy" (ALE) calculated for every major threat scenario?

2.2.3 Does AINA OS provide a real-time "Risk Thermometer" based on live threat intelligence feeds?

2.2.4 Are risk calculations automatically adjusted when a new vulnerability (CVE) is discovered in the environment?

2.2.5 Is the "Probability of Compromise" derived from actual attack telemetry rather than industry averages?

2.2.6 Are "Black Swan" events (low probability, infinite impact) modeled separately from daily operational risks?

2.2.7 Is the "Cost of Controls" compared against the "Reduction in Risk" to ensure positive ROI for security spending?

2.2.8 Does the risk model account for "Aggregate Risk" where multiple minor flaws combine to create a critical failure?

2.2.9 Are "Geopolitical Risk Factors" (e.g., war, sanctions) integrated into the risk scoring for international offices?

2.2.10 Is the "Time-to-Compromise" metric simulated and tracked to measure the speed of potential adversaries?

2.3.1 Is every vendor assigned a "Financial Risk Score" based on their security posture and access level?

2.3.2 Are "Software Bill of Materials" (SBOMs) ingested to calculate the inherited risk from open-source libraries?

2.3.3 Is there a "Contractual Risk Clause" ensuring vendors pay for breaches they cause?

2.3.4 Does the organization calculate the "Concentration Risk" of relying on a single cloud provider (e.g., AWS/Azure)?

2.3.5 Are 4th-party risks (your vendor's vendors) mapped and estimated for downstream impact?

2.3.6 Is there an automated "Kill-Switch" risk threshold that revokes vendor access if their security score drops?

2.3.7 Is the cost of a "Supply Chain Interruption" modeled for every critical supplier (e.g., days of survival without them)?

2.3.8 Are "Hardware Interdiction" risks (tampering during shipping) evaluated for sensitive equipment?

2.3.9 Is the "Reputational Contagion" risk modeled if a major partner suffers a public scandal?

2.3.10 Does the Rosecoin Ledger record the exact "Risk Status" of a vendor at the time of a contract signing?

2.4.1 Is the organization's Cyber Insurance coverage aligned strictly with the quantitative risk exposure limits?

2.4.2 Are "Exclusion Clauses" (e.g., Acts of War) in insurance policies analyzed against the threat landscape?

2.4.3 Does the organization use Rosecoin evidence to prove "Due Diligence" to insurers to lower premiums?

2.4.4 Is the "Self-Insured Retention" (Deductible) amount validated against the company's cash reserves?

2.4.5 Are specific coverage lines purchased for "Ransomware Reimbursement" and "Regulatory Fines"?

2.4.6 Is there a "Gap Analysis" performed quarterly between the insurance policy requirements and actual security controls?

2.4.7 Is the cost of "Business Interruption" claims pre-calculated to expedite insurance payouts?

2.4.8 Does the policy cover "Social Engineering Fraud" (CEO Fraud) specifically?

2.4.9 Is there a designated legal team to handle "Notification Claims" within the insurer's strict timelines?

2.4.10 Are "Betterment" clauses included to ensure destroyed systems are replaced with newer, more secure versions?

2.5.1 Has the Board signed a formal "Risk Appetite Statement" defining the exact dollar amount of acceptable loss?

2.5.2 Are "Risk Tolerance" thresholds set for specific operational metrics (e.g., max downtime of 4 hours)?

2.5.3 Is every "Risk Acceptance" (exception) signed by an executive and logged permanently on the Rosecoin Blockchain?

2.5.4 Is there a mandatory "Expiration Date" on all accepted risks, forcing a re-evaluation?

2.5.5 Are "Systemic Risks" (risks that could destroy the company) reviewed monthly by the CEO?

2.5.6 Is there a "Zero-Tolerance" policy for specific risks (e.g., safety-critical systems, child data)?

2.5.7 Does the AINA OS enforce "Risk Freezing" (preventing new deployments) if the aggregate risk score is too high?

2.5.8 Are "Risk Owners" individually named and held accountable for the risks in their specific business units?

2.5.9 Is there a "Whistleblower" channel for reporting hidden risks that management is ignoring?

2.5.10 Does the risk framework explicitly reject "Security by Obscurity" as a valid control?

2.6.1 Is the risk of "AI Model Collapse" (poisoned training data) quantified for all AI assets?

2.6.2 Is "Harvest Now, Decrypt Later" (Quantum threat) modeled as a current risk for long-term encrypted data?

2.6.3 Are "Neurological Privacy" risks assessed for any brain-computer interface (BCI) or biometric tech?

2.6.4 Is the "Orbital Collision" or "Signal Loss" risk calculated for any space-based assets or dependencies?

2.6.5 Is the risk of "Deepfake Identity Theft" modeled for executive communications?

2.6.6 Are "Algorithmic Bias" risks evaluated for potential legal and reputational damage?

2.6.7 Is the energy consumption risk (Carbon Tax/ESG impact) of the security infrastructure tracked?

2.6.8 Are "Metaverse/Spatial Computing" risks (virtual harassment, asset theft) included in the register?

2.6.9 Is the risk of "Autonomous Agent" malfunction (AI doing damage without humans) formally assessed?

2.6.10 Does the organization track "Technological Debt" as a compounding security risk?

2.7.1 Is security positioned as a "Sales Enabler" with ready-to-share trust packages for prospects?

2.7.2 Is the "Return on Mitigation" (ROM) calculated for every security purchase?

2.7.3 Does the organization track "Time-to-Market" gains achieved by automating security (DevSecOps)?

2.7.4 Is the value of "Customer Trust" measured through surveys and linked to security performance?

2.7.5 Are "Innovation Risks" managed in a way that encourages safe experimentation rather than stifling it?

2.7.6 Is the "Cost of Friction" (user inconvenience) measured against the security benefit of controls?

2.7.7 Does the security team report on "Revenue Protected" rather than just "Threats Blocked"?

2.7.8 Is there a "Value-at-Risk" (VaR) reduction target for the CISO's annual performance review?

2.7.9 Are security certifications (RCF, SOC2) leveraged as competitive differentiators in marketing?

2.7.10 Does the organization calculate the "Efficiency Dividend" of using AINA OS automation versus manual labor?

2.8.1 Is "Risk Velocity" (how fast a threat can spread once inside) calculated for every network segment?

2.8.2 Does the risk model account for "Cascading Failures" where a minor outage triggers a major critical system collapse?

2.8.3 Is "Aggregate Risk" monitored to detect when multiple "Low" vulnerabilities combine to form a "Critical" exploit chain?

2.8.4 Are "Correlation Factors" mapped to identify assets that will fail simultaneously (e.g., all running the same flawed firmware)?

2.8.5 Is the "Time-to-Recovery" vs. "Time-to-Impact" gap measured to see if defense can outpace the attack?

2.8.6 Does the model account for "Risk Contagion" between interconnected subsidiaries or partner networks?

2.8.7 Is "Seasonal Risk" (e.g., increased attacks during holidays or tax season) factored into the dynamic scoring?

2.8.8 Are "re-opened" risks (vulnerabilities that were fixed and returned) flagged with a higher severity multiplier?

2.8.9 Is the "Decay Rate" of security controls (how fast a tool becomes obsolete) tracked as a risk factor?

2.8.10 Does the system model "Multi-Vector" attacks where physical, social, and digital risks happen at once?

2.9.1 Is the financial risk of "Insider Threats" (rogue employees) quantified based on privilege levels?

2.9.2 Is the "Burnout Rate" of the security team tracked as a risk to operational continuity?

2.9.3 Does the organization calculate the "Cost of Replacement" for key security personnel (knowledge loss)?

2.9.4 Is "Security Culture" measured as a metric? (e.g., % of employees who report phishing vs. click it).

2.9.5 Is the risk of "Social Engineering Susceptibility" quantified per department (e.g., Finance vs. HR)?

2.9.6 Are "Key Man Dependencies" (processes that rely on one person) flagged as critical financial risks?

2.9.7 Is the cost of "User Friction" (productivity loss due to strict security) balanced against the risk reduction?

2.9.8 Is "Cognitive Overload" in the SOC (too many alerts) treated as a quantifiable risk of missed detections?

2.9.9 Is the "Off-boarding Risk" (data theft by departing staff) automatically assessed 30 days prior to exit?

2.9.10 Does the organization track "Shadow HR" (consultants/contractors with unmonitored access) as a distinct risk category?

2.10.1 Is the risk of "Vendor Lock-In" (cost to migrate away from a cloud provider) calculated annually?

2.10.2 Is "Data Remanence" (risk of data surviving deletion) assessed for decommissioned hardware?

2.10.3 Are "Legacy Debt" costs (maintenance of old systems) compared to the risk cost of modernizing?

2.10.4 Is the "Regulatory Fine Exposure" (e.g., 4% of global revenue for GDPR) set aside in a risk capital reserve?

2.10.5 Is the risk of "Encryption Obsolescence" (current keys becoming weak) tracked on a multi-year horizon?

2.10.6 Are "Patent/IP Infringement" risks in software development tracked via code scanning?

2.10.7 Is the "Sovereign Risk" of storing data in politically unstable regions quantified?

2.10.8 Is there a "Bankruptcy Risk" assessment for critical security vendors (what if your firewall vendor goes under)?

2.10.9 Is the "Social License to Operate" risk (public trust) modeled against potential privacy scandals?

2.10.10 Does the Rosecoin Ledger provide a "Defensible Disposal" record to prove data was destroyed to lower liability?

DOMAIN 3: THIRD-PARTY & SUPPLY CHAIN SECURITY

3.1.1 Is every new vendor subjected to a mandatory "Zero-Trust" intake process before a contract is signed?

3.1.2 Does the organization require a "D-UNS Number" or Legal Entity Identifier (LEI) to verify the true corporate identity of every supplier?

3.1.3 Are vendors automatically screened against global sanctions lists (OFAC, EU, UN) via the AINA OS integration?

3.1.4 Is "Beneficial Ownership" analysis performed to detect if a vendor is secretly owned by a hostile nation-state entity?

3.1.5 Does the onboarding process require the vendor to submit a "Security Passport" (e.g., SOC2, ISO 27001) verified on the blockchain?

3.1.6 Is there a "Vendor Tiering" system that automatically classifies suppliers as Tier 1 (Critical), Tier 2 (Operational), or Tier 3 (Commodity)?

3.1.7 Are "Fourth-Party" risks (your vendor's vendors) mapped to understand downstream dependencies?

3.1.8 Does the organization mandate "Background Checks" for vendor personnel who will have privileged access to the network?

3.1.9 Is there a "Financial Health Check" performed to ensure the vendor isn't at risk of bankruptcy (and sudden service loss)?

3.1.10 Are "Diversity & Inclusion" metrics tracked during onboarding to satisfy ESG supply chain requirements?

3.2.1 Does the AINA OS "Vulnerability Vines" engine scan the vendor's public-facing infrastructure for bugs before signing?

3.2.2 Is every vendor assigned a dynamic "Trust Score" (0-100) that updates in real-time based on their security performance?

3.2.3 Does the organization refuse to do business with vendors whose Trust Score falls below a specific threshold (e.g., 75)?

3.2.4 Are "Dark Web" searches conducted automatically to see if the vendor's credentials are currently for sale?

3.2.5 Is "Reputational Risk" monitored by scraping news and social media for scandals involving the vendor?

3.2.6 Are vendors required to prove they perform their own "Penetration Testing" at least annually?

3.2.7 Does the risk assessment include a "Geopolitical Risk" score based on the vendor's physical HQ location?

3.2.8 Is the "concentration risk" calculated (e.g., are 80% of your critical services dependent on one supplier)?

3.2.9 Does the system generate an "Automated Rejection" letter if a vendor fails critical security checks?

3.2.10 Is the entire risk assessment report hashed and stored in the Rosecoin Vault as evidence of "Due Diligence"?

3.3.1 Is a machine-readable SBOM (Software Bill of Materials) required for every piece of software purchased?

3.3.2 Does the system automatically scan the SBOM against the CVE database to find hidden vulnerable libraries (like Log4j)?

3.3.3 Is "Open Source Intelligence" applied to check the "Health" of the open-source projects used in the vendor's code?

3.3.4 Are vendors required to sign their code updates with a cryptographic key that the organization has verified?

3.3.5 Is there a "No-Update" policy for software that does not come with a verified SBOM?

3.3.6 Does the organization maintain a "Golden Repository" of approved third-party libraries for internal developers?

3.3.7 Are "Transitive Dependencies" (dependencies of dependencies) mapped and visualized in the AINA OS dashboard?

3.3.8 Is there a "Time-to-Patch" SLA (Service Level Agreement) that mandates how fast a vendor must fix a bug in their code?

3.3.9 Are "VEX" (Vulnerability Exploitability Exchange) documents required to filter out false positives in the SBOM?

3.3.10 Does the organization have the right to audit the vendor's "Build Pipeline" to prevent SolarWinds-style attacks?

3.4.1 Is a "Hardware Bill of Materials" (HBOM) required for critical appliances to detect counterfeit chips?

3.4.2 Are "Tamper-Evident" seals and packaging inspected and logged upon receipt of new hardware?

3.4.3 Does the organization use "Trusted Platform Modules" (TPM) to verify the integrity of the hardware boot process?

3.4.4 Are critical hardware components sourced only from "TAA-Compliant" (Trade Agreements Act) countries?

3.4.5 Is there a "Quarantine Protocol" for new hardware where it is tested in an isolated network before deployment?

3.4.6 Are "Firmware Hashes" verified against the manufacturer's website before the device is powered on?

3.4.7 Is the "Country of Origin" tracked for every sub-component of critical servers (e.g., motherboard vs. chassis)?

3.4.8 Does the organization prohibit the use of "Grey Market" hardware resellers?

3.4.9 Is there a physical "Chain of Custody" log for the transport of sensitive equipment?

3.4.10 Are "Hardware Implants" (spy chips) scanned for using X-ray or side-channel analysis for ultra-critical assets?

3.5.1 Does the RCF engine monitor vendor connectivity in real-time for "Anomalous Data Exfiltration"?

3.5.2 Is there an automated "Kill-Switch" that instantly severs the connection if a vendor is breached?

3.5.3 Are "Privileged Access Management" (PAM) logs reviewed daily for all vendor accounts?

3.5.4 Does the system detect if a vendor changes their banking details (to prevent Business Email Compromise fraud)?

3.5.5 Is "Configuration Drift" monitoring applied to vendor-managed devices on your network?

3.5.6 Are vendors required to re-authenticate via "Biometric MFA" for every high-risk session?

3.5.7 Does the organization receive "Push Notifications" from the vendor's own SOC if they detect an incident?

3.5.8 Is the "Least Privilege" principle enforced dynamically (e.g., access is revoked when the project ends)?

3.5.9 Are "Honey-Tokens" (fake data) placed in vendor-accessible folders to detect unauthorized browsing?

3.5.10 Does the Rosecoin Blockchain record every single file transfer between the organization and the vendor?

3.6.1 Do all vendor contracts include a "Right to Audit" clause allowing for unannounced security inspections?

3.6.2 Is there a mandatory "Breach Notification" clause requiring the vendor to report incidents within 24 hours?

3.6.3 Are "Data Sovereignty" clauses included to prevent vendors from moving data to illegal jurisdictions?

3.6.4 Does the contract specify exactly who owns the "Intellectual Property" created during the engagement?

3.6.5 Is "Cyber Insurance" a contractual requirement for all Tier 1 vendors?

3.6.6 Are "Data Destruction" certificates required and verified upon contract termination?

3.6.7 Does the contract allow for "Liquidated Damages" (fines) if the vendor causes a security outage?

3.6.8 Are vendors required to comply with the organization's specific RCF Maturity Level (e.g., must be ML-3)?

3.6.9 Is there a "Force Majeure" clause that specifically addresses cyber-warfare and pandemics?

3.6.10 Are all contracts stored as "Smart Contracts" on Rosecoin to automate penalty enforcement?

3.7.1 Is the vendor's "Carbon Footprint" tracked to align with Domain 23 (Green Cybersecurity)?

3.7.2 Does the organization audit the vendor for "Digital Labor Rights" (e.g., fair treatment of their IT staff)?

3.7.3 Is there a policy against using vendors who supply surveillance technology to authoritarian regimes?

3.7.4 Are vendors evaluated for their readiness to support "Post-Quantum Cryptography" (Domain 16)?

3.7.5 Is "AI Ethics" compliance required for vendors supplying machine learning models?

3.7.6 Does the organization support "Small & Diverse Businesses" by providing them with free RCF security tools?

3.7.7 Is there a "Conflict Minerals" audit for hardware vendors (e.g., sourcing of cobalt/lithium)?

3.7.8 Does the vendor have a "Whistleblower" program for their own employees to report security risks?

3.7.9 Is the "Circular Economy" (recyclability) of the vendor's hardware evaluated?

3.7.10 Does the organization participate in "Collective Defense" by sharing vendor threat intel with industry peers?

3.8.1 Is there a "Data Bill of Materials" (DBOM) required for all AI vendors to prove their training data wasn't stolen or poisoned?

3.8.2 Does the organization mandate a "Vendor Exit Stress Test" (Fire Drill) to prove it can survive if a critical supplier suddenly vanishes?

3.8.3 Are "Zombie APIs" (old, forgotten vendor connections) automatically hunted and terminated by AINA OS?

3.8.4 Is there a "Source Code Escrow" agreement verified on the blockchain (ensuring you get the code if the vendor goes bankrupt)?

3.8.5 Does the organization map the "Geopolitical Path" of data packets? (e.g., ensuring vendor traffic doesn't route through hostile nations).

3.8.6 Is "API Security Testing" (DAST for APIs) mandatory for every third-party integration before it goes live?

3.8.7 Are "Space Logistics" vendors (satellite launch/ground stations) vetted for specific orbital security protocols (Domain 22)?

3.8.8 Is there a "Counter-Intelligence" screening for vendor personnel accessing Top Secret/IP-heavy zones?

3.8.9 Does the organization enforce "Just-in-Time" (JIT) access for vendors, where they have 0 standing privileges until needed?

3.8.10 Is "Data Portability" technically verified? (Can you actually get your terabytes of data out of the vendor's cloud in a usable format?)

3.8.11 High-Risk Supplier Exclusion: Is there a mechanism to identify, restrict, or exclude ICT suppliers considered "high-risk" based on non-technical factors, such as the potential influence of third states (e.g., EU Cybersecurity Act 2026 alignment)?

3.8.12 Continuous Vendor Posture Monitoring: Does the organization perform real-time, continuous monitoring of critical vendor security postures and share threat intelligence back with those suppliers?

3.8.13 Non-Technical Risk Assessments: Does the organization identify and assess "non-technical" risks in the ICT supply chain, specifically evaluating suppliers located in high-risk jurisdictions as mandated by new 2026 regulations (e.g., EU CSA2)?

3.8.14 Supply Chain Vulnerability Adoption: Does the organization actively track the adoption of "Secure by Design" principles among critical vendors to reduce the prevalence of long-term vulnerabilities?

DOMAIN 4: IDENTITY & ACCESS MANAGEMENT

4.1.1 Is the organization operating on a "Never Trust, Always Verify" mandate for every user, device, and connection?

4.1.2 Has the network been segmented so that identity verification is required to move laterally between zones (Micro-segmentation)?

4.1.3 Is "Conditional Access" enforced dynamically based on user location, device health, and risk score?

4.1.4 Are all legacy "Implicit Trust" zones (e.g., the corporate LAN) eliminated in favor of identity-based perimeters?

4.1.5 Does the architecture support "Continuous Evaluation," revoking access mid-session if risk signals change?

4.1.6 Is the "Attack Surface" hidden from unauthenticated users using a Software-Defined Perimeter (SDP)?

4.1.7 Are all access requests encrypted end-to-end, regardless of whether they originate internally or externally?

4.1.8 Is there a formal policy to treat the "Identity" as the new security perimeter, replacing the firewall?

4.1.9 Does the Zero Trust model extend to cloud environments, SaaS applications, and on-premise legacy systems equally?

4.1.10 Are "Break Glass" emergency accounts monitored with the highest level of scrutiny and alerting?

4.2.1 Has the organization eliminated passwords in favor of FIDO2/ WebAuthn passwordless standards?

4.2.2 Is "Phishing-Resistant" Multi-Factor Authentication (MFA) mandatory for 100% of employees (e.g., Hardware Keys)?

4.2.3 Does the system utilize "Behavioral Biometrics" (keystroke dynamics, mouse velocity) to authenticate users continuously?

4.2.4 Is "Liveness Detection" enforced on facial recognition systems to prevent deepfake or photo spoofing?

4.2.5 Are biometric templates stored as "Salted Hashes" and never as raw images, to protect user privacy?

4.2.6 Is "Step-Up Authentication" triggered automatically when a user attempts a high-risk action (e.g., wire transfer)?

4.2.7 Are "Adaptive Auth" policies used to reduce friction for low-risk users while hardening high-risk access?

4.2.8 Does the AINA OS support "Context-Aware" auth (e.g., blocking login if the user's phone is not near their laptop)?

4.2.9 Is there a fallback mechanism for biometric failure that does not revert to weak security questions?

4.2.10 Are "Shared Accounts" strictly prohibited and technically blocked by the IAM system?

4.3.1 Is "Just-in-Time" (JIT) access enforced, granting privileges only for the exact duration of the task?

4.3.2 Are all privileged sessions (admin activity) recorded via video/text logs for forensic review?

4.3.3 Is "Standing Privilege" (permanent admin rights) completely eliminated for all users?

4.3.4 Are privileged credentials (passwords/keys) rotated automatically after every single use?

4.3.5 Is "Dual Control" (Four-Eyes Principle) required for critical system changes or root access?

4.3.6 Are Service Accounts and Bot IDs managed with the same rigor as human admin accounts?

4.3.7 Does the PAM system isolate admin sessions on a secure "Jump Server" or browser container?

4.3.8 Are "Hard-Coded Credentials" scanned for and removed from scripts and source code continuously?

4.3.9 Is there a "Privilege Separation" model ensuring one admin account cannot compromise the entire domain?

4.3.10 Does the Rosecoin Ledger record every elevation of privilege to create an immutable audit trail?

4.4.1 Is the "Joiner, Mover, Leaver" (JML) process fully automated to prevent access creep?

4.4.2 Are "Access Reviews" (Certifications) performed monthly using AI to highlight outliers rather than rubber-stamping?

4.4.3 Is "Role-Based Access Control" (RBAC) supplemented with "Attribute-Based Access Control" (ABAC) for finer granularity?

4.4.4 Are "Orphaned Accounts" (accounts with no owner) automatically disabled after 30 days of inactivity?

4.4.5 Is "Segregation of Duties" (SoD) enforced logically to prevent fraud (e.g., same person cannot request and approve POs)?

4.4.6 Does the IGA system integrate with HR feeds to revoke access instantly upon employee termination?

4.4.7 Are "Entitlement Creep" metrics tracked to identify users accumulating too many rights over time?

4.4.8 Is there a "Self-Service" portal for access requests that uses automated approval workflows?

4.4.9 Are "Policy Violations" in access rights auto-remediated by the system (e.g., removing conflicting roles)?

4.4.10 Is the "Principle of Least Privilege" validated by actual usage data (removing rights that are never used)?

4.5.1 Is there a comprehensive inventory of all "Machine Identities" (Bots, APIs, Containers, IoT)?

4.5.2 Are X.509 Certificates for machines managed and rotated automatically to prevent outages?

4.5.3 Is "Workload Identity" used to authenticate applications in the cloud without long-lived keys?

4.5.4 Are "Secrets Management" vaults used to inject credentials at runtime rather than storing them in config files?

4.5.5 Is there a "Bot Detection" capability to distinguish between legitimate automation and malicious scrapers?

4.5.6 Are RPA (Robotic Process Automation) bots assigned specific, limited identities rather than using human credentials?

4.5.7 Is the "Lifecycle" of machine identities tied to the workload (e.g., identity dies when the container spins down)?

4.5.8 Are "SSH Keys" discovered, rotated, and managed centrally to prevent unauthorized backend access?

4.5.9 Does the system detect "Beaconing" or anomalous behavior from service accounts?

4.5.10 Are "API Tokens" treated with the same sensitivity as passwords, with expiration and rotation policies?

4.6.1 Does the organization support "Decentralized Identifiers" (DIDs) to give users control over their own identity data?

4.6.2 Are "Verifiable Credentials" used to prove certifications/employment without sharing raw PII?

4.6.3 Is the "Identity Proofing" process (checking passports/IDs) logged on the Rosecoin Blockchain for non-repudiation?

4.6.4 Can the system interoperate with "Self-Sovereign Identity" (SSI) wallets for customer authentication?

4.6.5 Is there a "Zero-Knowledge Proof" mechanism to verify age or citizenship without revealing the actual birthdate?

4.6.6 Are "Smart Contracts" used to automate access revocation based on external triggers (e.g., contract expiry)?

4.6.7 Is the "Root of Trust" for the identity system anchored in the immutable Rosecoin ledger?

4.6.8 Does the organization allow "Bring Your Own Identity" (BYOI) for partners using verified decentralized credentials?

4.6.9 Are "Consent Receipts" generated on-chain whenever a user grants access to their personal data?

4.6.10 Is the identity infrastructure resistant to "Centralized Database Hacks" by distributing trust?

4.7.1 Is the organization migrating to "Post-Quantum Cryptography" (PQC) for all identity signing keys?

4.7.2 Are "Deepfake Defense" tools integrated into video verification and voice authentication flows?

4.7.3 Is "AI-Driven Identity Threat Detection" (ITDR) active to catch sophisticated credential attacks?

4.7.4 Are "Neuro-Metric" identifiers (brainwave patterns) explored for high-security areas (Domain 24 alignment)?

4.7.5 Is the "Metaverse Identity" of the brand protected against impersonation in virtual worlds?

4.7.6 Are "Digital Twin" identities secured to prevent manipulation of physical assets?

4.7.7 Is there a defense against "Prompt Injection" attacks targeting identity verification AI chatbots?

4.7.8 Does the system support "Offline Authentication" securely for air-gapped or space assets?

4.7.9 Are "Synthetic Identities" (fake users created by AI) actively hunted in the customer database?

4.7.10 Is the "DNA Data" (if used for extreme auth) protected with the highest level of privacy encryption?

4.8.1 Is there a strict "No Verbal Password Reset" policy enforced at the Service Desk?

4.8.2 Are Service Desk agents required to use "Out-of-Band" verification (e.g., push notification to a manager) before resetting MFA?

4.8.3 Is "Voice Stress Analysis" or AI-based sentiment monitoring used to detect coercion during support calls?

4.8.4 Does the Service Desk have a "Duress Code" protocol if an employee is being physically forced to request access?

4.8.5 Are "Knowledge-Based Authentication" (KBA) questions (e.g., "mother's maiden name") banned due to public data availability?

4.8.6 Is there a "Cool-Down Period" (e.g., 1 hour) enforced after a credential reset before the account can access "Crown Jewel" assets?

4.8.7 Are "Visual verifications" (video call with ID held up) mandatory for remote employee account recovery?

4.8.8 Does the system flag "Serial Resetters" (users who frequently request resets) for heightened security monitoring?

4.8.9 Is the "Caller ID" automatically validated against the employee's known mobile number (preventing spoofing)?

4.8.10 Are Service Desk agents authorized to "Lock Down" an executive account immediately on suspicion without approval?

4.9.1 Is the Customer Identity provider physically isolated from the Employee Identity provider to prevent lateral movement?

4.9.2 Does the CIAM system support "Progressive Profiling" (collecting data slowly) to balance UX with security?

4.9.3 Are "Bot Mitigation" layers active on the customer login page to prevent Credential Stuffing attacks?

4.9.4 Is "Privacy-by-Design" enforced, giving customers a dashboard to manage their own consents and data sharing?

4.9.5 Does the system allow customers to view their own "Active Sessions" and remotely kill suspicious ones?

4.9.6 Is "MFA for Customers" mandatory for sensitive actions (e.g., changing shipping address or payment info)?

4.9.7 Are customer passwords stored using "Slow Hashing" algorithms (e.g., Argon2) to make brute-forcing impossible?

4.9.8 Is there a "Family/Delegate Access" model allowing customers to safely grant limited access to others?

4.9.9 Does the CIAM system detect "Impossible Travel" (login from London and Tokyo in 1 hour) and block it instantly?

4.9.10 Is "Social Login" (Google/Apple) implemented with strict scope limits to prevent over-sharing of data?

4.10.1 Is there an "Offline Identity" capability ensuring critical operations can continue if the cloud IdP goes down?

4.10.2 Are "Emergency Access Keys" (physical tokens) stored in a fireproof safe for "Break Glass" scenarios?

4.10.3 Is the "Active Directory" (or equivalent) backed up to an immutable, air-gapped vault daily?

4.10.4 Is there a "Failover" plan for MFA? (e.g., if SMS fails, switch to TOTP or FIDO key automatically).

4.10.5 Are "Domain Controller" recoverability tests performed quarterly?

4.10.6 Is the "Token Time-to-Live" (TTL) reduced during active attacks to force frequent re-authentication?

4.10.7 Is there a "Clean Room" environment ready for rebuilding the Identity infrastructure from scratch?

4.10.8 Are "Cached Credentials" on endpoints hardened to prevent dumping by tools like Mimikatz?

4.10.9 Is "Directory Monitoring" active to detect changes to Admin groups in real-time?

4.10.10 Does the organization simulate a "Total Identity Collapse" in its Disaster Recovery exercises?

4.10.11 Machine Identity Lifecycle: Are machine identities (APIs, bots, service accounts) managed as first-class citizens with mandatory owners, automated rotation, and behavioral monitoring?

4.10.12 Phishing-Resistant MFA for Admins: Is phishing-resistant MFA (FIDO2/WebAuthn) mandatory for all administrative and high-privilege access, prohibiting SMS or standard push-notifications?

4.10.13 Identity-Bound Authentication: Does the system verify access based on identity-bound signals, including device health, geographic location, and typical login behavior patterns?

4.10.14 Token & Session Management: Are authentication tokens short-lived, and does the system automatically revoke active sessions across all devices if a "Mover" or "Leaver" event is detected?

4.10.15 Credential Abuse & Model Manipulation: Does the system detect and block attempts to use forged identities (deepfakes or biometric spoofing) to trigger automated actions or access machine learning models?

4.10.16 ZTNA Migration from Legacy VPN: Has the organization completed the transition from legacy VPNs to Zero Trust Network Access (ZTNA) to limit lateral movement and reduce the blast radius of compromised credentials?

4.10.17 Machine-Learned Behavior Modeling: Does the identity system use machine learning to detect and block lateral movement that mimics legitimate administrative tool usage but deviates from established behavioral baselines?

4.10.18 Phishing-Resistant MFA Verification: Is "phishing-resistant MFA" (such as FIDO2 hardware keys) verified as a mandatory requirement for both internal users and external contractors, a prerequisite for modern 2026 cyber insurance coverage?

DOMAIN 5: PRIVACY & DATA PROTECTION

5.1.1 Does the organization maintain a live, automated "Record of Processing Activities" (RoPA) that updates as code changes?

5.1.2 Is a "Data Protection Officer" (DPO) appointed and registered with the relevant Supervisory Authorities (e.g., ICO, CNIL)?

5.1.3 Does the AINA OS automatically map data flows against 100+ global privacy laws (GDPR, CCPA, LGPD, PIPL)?

5.1.4 Is "Privacy by Design" enforced by requiring a Privacy Impact Assessment (PIA) before any new database is created?

5.1.5 Are "Legitimate Interest Assessments" (LIA) documented and signed off before processing data without explicit consent?

5.1.6 Does the organization utilize a "One-Stop-Shop" mechanism for EU regulatory interactions where applicable?

5.1.7 Are "Binding Corporate Rules" (BCRs) established for intra-group data transfers?

5.1.8 Is there a "Privacy Steering Committee" that meets quarterly to review data ethics and compliance trends?

5.1.9 Does the system automatically flag and block "Dark Patterns" in UI/UX design that trick users into consenting?

5.1.10 Are "Data Processing Agreements" (DPAs) digitally signed and linked to every vendor contract in the system?

5.2.1 Is there a "Continuous Discovery" engine running to find PII (Personally Identifiable Information) in unstructured data (emails, PDFs)?

5.2.2 Does the system classify data based on sensitivity (Public, Internal, Confidential, Restricted, Secret) automatically?

5.2.3 Are "Regex Patterns" updated weekly to detect new ID formats (e.g., new national ID cards, medical codes)?

5.2.4 Is "Shadow Data" (data stored in unapproved cloud buckets) automatically detected and quarantined?

5.2.5 Does the organization scan for "Ghost Data" (data belonging to former customers) and prompt for deletion?

5.2.6 Are "Data Lineage" maps generated to show exactly where a specific user's data has traveled across the network?

5.2.7 Is "Metadata Management" used to tag data with retention rules at the moment of creation?

5.2.8 Does the system distinguish between "Direct Identifiers" (Name) and "Indirect Identifiers" (GPS + Zip Code)?

5.2.9 Are development and test environments scanned to ensure no production PII is present without masking?

5.2.10 Is "Data Resonance" checked (e.g., finding the same file duplicated across 50 servers)?

5.3.1 Is the "Right to Access" fulfilled via a self-service portal, eliminating manual zip-file creation?

5.3.2 Does the "Right to be Forgotten" (Erasure) trigger a cascading delete across all backups and third-party systems?

5.3.3 Is identity verification for DSARs automated to prevent "Data Subject Impersonation" attacks?

5.3.4 Can the system handle "Data Portability" requests by generating machine-readable formats (JSON/XML) instantly?

5.3.5 Is the "Right to Rectification" available, allowing users to correct their own data without support tickets?

5.3.6 Does the organization track the "Time-to-Fulfill" for DSARs to ensure compliance with the 30-day (GDPR) or 45-day (CCPA) limits?

5.3.7 Are "Do Not Sell My Info" (GPC signals) honored automatically by the website's cookie consent manager?

5.3.8 Is there a "Suppression List" to ensure erased users are not accidentally re-marketed to later?

5.3.9 Are DSAR rejection reasons (e.g., "Legitimate Legal Hold") documented and communicated clearly?

5.3.10 Does the Rosecoin Ledger record the completion of every DSAR request for regulatory audit proof?

5.4.1 Is "Geo-Fencing" used to prevent data from physically leaving its jurisdiction of origin (e.g., German data stays in Germany)?

5.4.2 Are "Transfer Impact Assessments" (TIAs) conducted for all data flows to high-risk nations?

5.4.3 Is "Schrems II" compliance enforced by applying supplementary measures (encryption) for US transfers?

5.4.4 Does the system detect "Data Stowaways" (hidden PII in logs) moving across borders?

5.4.5 Are "Data Localization" laws (e.g., Russia, China, Vietnam) strictly adhered to by the infrastructure?

5.4.6 Is "Split-Key Encryption" used, where the decryption key never leaves the country of origin?

5.4.7 Does the organization monitor for "Legal Conflicts" where one country demands access and another forbids it?

5.4.8 Are "Sovereign Cloud" instances used for government or critical infrastructure clients?

5.4.9 Is there a "Data Visa" system in AINA OS that approves/denies transfer packets based on real-time legal rules?

5.4.10 Are international "Onward Transfers" (Vendor to Vendor) mapped and legally covered?

5.5.1 Is "Explicit Consent" captured with a granular breakdown (e.g., distinct checks for marketing vs. analytics)?

5.5.2 Are "Cookie Walls" avoided, ensuring access to the service even if consent is refused (where legally required)?

5.5.3 Is the "Consent Receipt" hashed and stored on the Rosecoin Blockchain to prove validity in court?

5.5.4 Can users "Withdraw Consent" as easily as they gave it (e.g., one-click revocation)?

5.5.5 Is "Preference Management" centralized, so a change on mobile reflects on the web instantly?

5.5.6 Are "Just-in-Time" notices used to ask for consent at the moment of data collection, not just in a long policy?

5.5.7 Does the system manage "Parental Consent" flows for users identified as minors?

5.5.8 Is the "Validity Period" of consent tracked, prompting for re-consent after a set time (e.g., 12 months)?

5.5.9 Are "Joint Controller" relationships clearly communicated to the user at the point of consent?

5.5.10 Does the organization audit consent rates to detect UI issues or user mistrust?

5.6.1 Is "Differential Privacy" applied to datasets to add noise, making re-identification mathematically impossible?

5.6.2 Is "Homomorphic Encryption" explored for allowing computation on encrypted data without decryption?

5.6.3 Are "Synthetic Data" sets generated for testing and AI training, replacing real user data entirely?

5.6.4 Is "Format-Preserving Encryption" (FPE) used to protect legacy databases without breaking application logic?

5.6.5 Is "Pseudonymization" applied immediately upon data ingestion (separating ID from Data)?

5.6.6 Are "Zero-Knowledge Proofs" used for age or eligibility verification?

5.6.7 Is "Federated Learning" used for AI, keeping raw data on user devices while only sharing model updates?

5.6.8 Is "K-Anonymity" validation run on release datasets to ensure individuals cannot be singled out?

5.6.9 Are "Secure Multi-Party Computation" (SMPC) protocols used for sharing insights with partners without sharing data?

5.6.10 Does the system automatically "Re-Key" encrypted data on a schedule to limit exposure?

5.7.1 Is the "Collection Limitation" principle enforced (only collecting what is strictly needed)?

5.7.2 Are "retention policies" defined as code, automatically deleting data when its purpose expires?

5.7.3 Is "Data Rotting" (storing data "just in case") actively prohibited and scanned for?

5.7.4 Are "Legal Holds" capable of pausing deletion scripts for specific users during litigation?

5.7.5 Is "Defensible Disposal" practiced, generating a certificate of destruction for every deleted batch?

5.7.6 Are backups included in the deletion cycle (or crypto-shredded) to prevent "Data Resurrection"?

5.7.7 Is "Unstructured Data" (File Servers) subject to the same retention rules as databases?

5.7.8 Does the organization periodically review "Why" it is collecting specific fields?

5.7.9 Are "Ephemeral Data" stores (RAM/Cache) cleared securely after session termination?

5.7.10 Is "Physical Media" (Hard Drives) degaussed and shredded on-site before disposal?

5.8.1 Is the "72-Hour Rule" (GDPR) built into the Incident Response timeline?

5.8.2 Does the system automatically identify whose data was impacted to generate a notification list?

5.8.3 Are "Notification Templates" pre-approved by Legal to avoid delays during a crisis?

5.8.4 Is there a mechanism to notify Supervisory Authorities via their specific API or portal?

5.8.5 Are "Third-Party Breaches" monitored to see if the organization's data was exposed by a vendor?

5.8.6 Is "Harm Analysis" conducted to determine if the breach poses a high risk to rights and freedoms?

5.8.7 Does the Rosecoin Ledger provide an immutable timeline of the breach response for investigators?

5.8.8 Are "Call Center Scripts" ready to handle inbound queries from affected users?

5.8.9 Is "Credit Monitoring" pre-negotiated for potential victims?

5.8.10 Is "Post-Breach" analysis used to update Privacy Policies and controls?

5.9.1 Is "Model Inversion" defense in place to prevent attackers from reconstructing training data from AI?

5.9.2 Are "Machine Unlearning" protocols available to remove a specific user's influence from a trained model?

5.9.3 Is "Biometric Data" (Face, Voice) prohibited from being used for "Emotional Analysis" without explicit consent?

5.9.4 Are "Neural Data" (Brainwaves) classified as "Super-Sensitive" requiring the highest encryption?

5.9.5 Is "Automated Decision Making" (Profiling) subject to human review upon request?

5.9.6 Are "Prompt Injection" logs scrubbed of PII before being stored or analyzed?

5.9.7 Is the "Provenance" of AI training data verified to ensure it wasn't scraped illegally?

5.9.8 Are "Deepfake" rights managed (e.g., right to not be digitally simulated)?

5.9.9 Is "Cognitive Liberty" respected (no manipulation of subconscious behavior)?

5.9.10 Does the organization support "Opt-Out" for their data being used to train generative AI?

5.10.1 Is "Employee Monitoring" transparent, proportional, and legally justified?

5.10.2 Are "Privacy Champions" appointed in every department to act as local advocates?

5.10.3 Is "BYOD" (Bring Your Own Device) privacy managed (containerization) so the company can't see personal photos?

5.10.4 Are internal "Phishing Tests" conducted without shaming or exposing employee behavior publicly?

5.10.5 Is "Health Data" (HR/Covid) stored separately from performance data?

5.10.6 Are "Video Surveillance" (CCTV) policies clear on where and why recording happens?

5.10.7 Is "Diversity Data" (Race/Religion) anonymized and aggregated immediately?

5.10.8 Are "Background Check" results minimized and deleted after the hiring decision?

5.10.9 Is there a "Privacy Hotline" for employees to report concerns anonymously?

5.10.10 Does the organization conduct "Privacy Tabletop Exercises" to practice breach response?

5.11.1 Is the "Time-to-Notify" (TTN) metric tracked for every privacy incident (goal: <72 hours)?

5.11.2 Does the organization measure the "Consent Opt-In Rate" to gauge user trust in the brand?

5.11.3 Is "Data Minimization" quantified? (e.g., % of stale data deleted vs. collected).

5.11.4 Are "DSAR Fulfillment Costs" calculated to justify automation investments?

5.11.5 Is the "Privacy ROI" (Return on Investment) reported to the Board? (e.g., fines avoided + brand value).

5.11.6 Does the organization track "Third-Party Data Exposure" frequency as a key risk indicator?

5.11.7 Is "Privacy Debt" (legacy systems non-compliant with new laws) tracked and burned down?

5.11.8 Are "Algorithmic Bias" scores reported for all customer-facing AI models?

5.11.9 Is the "Re-Identification Risk" score calculated for every public dataset release?

5.11.10 Does the organization conduct "Mystery Shopper" tests on its own privacy processes (e.g., fake DSARs)?

5.12.1 Is "Data Residency" strictly enforced at the packet level for Sovereign Cloud clients?

5.12.2 Are "Warrant Canaries" published to transparently signal government data requests?

5.12.3 Is there a "Kill-Switch" for data flows to nations that suddenly become hostile (Sanctions enforcement)?

5.12.4 Does the organization have a "Cognitive Firewall" policy to prevent manipulation of user behavior via AI?

5.12.5 Are "Neuro-Rights" explicitly recognized in the Privacy Charter (Right to Mental Privacy)?

5.12.6 Is "Genetic Data" handling isolated in a separate, air-gapped "Bio-Vault"?

5.12.7 Are "Spatial Privacy" rules defined for Metaverse/AR interactions (who can see my avatar's data)?

5.12.8 Is "Post-Mortem Privacy" defined? (What happens to a user's digital soul after death?).

5.12.9 Does the organization support "Anonymous Payment" methods (Crypto/Cash) to protect financial privacy?

5.12.10 Is the "Right to Analog" preserved? (Can a user interact with the company without digital tracking?).

5.13.1 Is "Real-Time Bidding" (RTB) data leakage prevented by stripping all PII (IP, device ID) from bid requests before they leave the server?

5.13.2 Does the organization enforce a strict "No-Sale" policy for customer data to third-party data brokers?

5.13.3 Is "Pixel Governance" active? (Are Facebook/TikTok tracking pixels audited weekly to ensure they aren't capturing sensitive form data like health info?)

5.13.4 Are "Dark Patterns" in marketing emails (e.g., hidden unsubscribe buttons) automatically detected and flagged by AINA OS?

5.13.5 Does the organization use "Contextual Advertising" (based on content) rather than "Behavioral Advertising" (tracking users across the web) where possible?

5.13.6 Is "Cross-Device Tracking" (linking a user's phone to their TV) disclosed explicitly with a specific opt-in?

5.13.7 Are "Data Clean Rooms" used for marketing analytics to prevent direct sharing of user lists with partners?

5.13.8 Is there a mechanism to handle "Global Privacy Control" (GPC) signals as a legally binding "Do Not Sell" instruction?

5.13.9 Are "Lookalike Audiences" created only from users who have explicitly consented to their data being used for modeling?

5.13.10 Does the organization audit its media agencies to ensure they aren't buying "Black Market" data to enrich customer profiles?

5.14.1 Is "Age Assurance" technology (e.g., zero-knowledge age estimation) used to identify and protect child users without collecting their IDs?

5.14.2 Are "High-Risk" features (location sharing, DMs) automatically disabled by default for users under 18 (aligned with the UK/CA Age Appropriate Design Codes)?

5.14.3 Is there a "Stalkerware" defense? (Does the app notify the user if their location is being viewed by another account, protecting domestic violence victims?)

5.14.4 Are "Digital Legacy" contacts allowed, letting users decide who accesses their data after death (Post-Mortem Privacy)?

5.14.5 Is "Algorithmic Addiction" monitored? (Are design features that exploit dopamine loops in children flagged as privacy harms?)

5.14.6 Are "Senior Citizen" protections in place to flag unusual data transfers that suggest "Elder Fraud" or coercion?

5.14.7 Is accessibility data (e.g., "User is Blind") stored as "Sensitive Health Data" rather than standard profile tags?

5.14.8 Is there a "Safety Button" for users to quickly exit the site and wipe local history (essential for victims of abuse)?

5.14.9 Does the organization ban the use of "Voiceprints" for children under any circumstances?

5.14.10 Is there a "Human-in-the-Loop" appeal process for any AI decision that denies service to a vulnerable person?

DOMAIN 6: AI SECURITY & ML GOVERNANCE

6.1.1 Is there a board-approved "AI Acceptable Use Policy" explicitly defining which AI tools (e.g., ChatGPT, Copilot) are permitted?

6.1.2 Has the organization created a "Shadow AI" discovery process to identify unauthorized models running on employee devices?

6.1.3 Are all AI systems classified by "Risk Tier" (Unacceptable, High, Limited, Minimal) in accordance with the EU AI Act?

6.1.4 Is there a "Human-in-the-Loop" mandate for high-risk AI decisions affecting employment, credit, or healthcare?

6.1.5 Does the organization maintain a live "AI Inventory" (Model Registry) tracking every deployed model and its owner?

6.1.6 Are "AI Ethics" reviews conducted before training begins, specifically checking for bias against protected groups?

6.1.7 Is there a defined "Liability Framework" for AI errors? (Who is responsible if the AI hallucinates a false fact: the dev or the user?)

6.1.8 Are "Copyright" risks assessed for all Generative AI outputs to ensure the company actually owns the code/content it generates?

6.1.9 Is there a "Right to Explanation" capability where the AI can technically explain why it made a specific decision?

6.1.10 Does the organization have a "Decommissioning Plan" for AI models that become obsolete or drift beyond safety limits?

6.2.1 Is "Prompt Injection" defense active? (Are inputs scanned for patterns that try to override system instructions?)

6.2.2 Does the system filter "Indirect Prompt Injections" where an AI processing a web page reads hidden malicious text?

6.2.3 Are "System Prompts" (the root instructions) hard-coded and separated from user input to prevent "Jailbreaking"?

6.2.4 Is "Output Sanitization" in place to prevent the AI from generating toxic content, hate speech, or phishing emails?

6.2.5 Does the organization use "Retrieval Augmented Generation" (RAG) security to ensure the AI only accesses documents the user has permission to see?

6.2.6 Are "Hallucination" detectors active to flag when the AI is likely making up facts?

6.2.7 Is "Sensitive Data Filtering" applied to the prompt input to prevent employees from pasting customer PII into public chatbots?

6.2.8 Are "Plugin" interactions (AI talking to APIs) strictly limited by "Least Privilege" (e.g., the AI can read the calendar but not delete it)?

6.2.9 Is "Model Denial of Service" prevented by capping the number of tokens/requests a user can generate per minute?

6.2.10 Are "Cache Poisoning" attacks prevented by verifying the integrity of stored AI responses?

6.3.1 Is "Data Poisoning" detection active during the training phase to find malicious samples designed to corrupt the model?

6.3.2 Does the organization perform "Adversarial Training" (training the model on attack data) to make it robust against evasion?

6.3.3 Is "Model Extraction" monitoring in place to detect if an attacker is querying the API specifically to steal the model's logic?

6.3.4 Are "Membership Inference" attacks (trying to guess if a specific person was in the training set) blocked by Differential Privacy?

6.3.5 Is "Model Inversion" defense applied to prevent attackers from reconstructing user faces or data from the model outputs?

6.3.6 Are "Sponge Attacks" (inputs designed to maximize energy/compute usage) detected and dropped?

6.3.7 Is there a "Perturbation Detector" to flag inputs that have invisible noise added to fool the AI (e.g., a panda labeled as a gibbon)?

6.3.8 Are "Backdoor Triggers" hunted for in third-party models (e.g., a model that works perfectly unless a specific pixel is present)?

6.3.9 Is "Evasion" testing part of the standard QA process before any model goes to production?

6.3.10 Does the Rosecoin Ledger record the exact "Hash" of the model file to prove it hasn't been tampered with?

6.4.1 Is an "AI Software Bill of Materials" (AI-BOM) required for every model, listing the training data sources and libraries used?

6.4.2 Are "Pickle" files (insecure Python serialization) banned in favor of safer formats like Safetensors?

6.4.3 Is "Model Signing" enforced, requiring a cryptographic signature before a model can be loaded into production?

6.4.4 Are "Hugging Face" or other public model repositories scanned for malware before downloading?

6.4.5 Is "Data Lineage" tracking enforced to prove that training data was legally obtained (not scraped without consent)?

6.4.6 Are "Data Sanitization" pipelines automated to remove PII from training sets before the model ever sees them?

6.4.7 Is there a "Kill Switch" for third-party models if the vendor is compromised?

6.4.8 Are "Federated Learning" nodes authenticated to prevent a malicious participant from poisoning the global model?

6.4.9 Is "Transfer Learning" risk assessed (inheriting vulnerabilities from the base model)?

6.4.10 Does the organization scan for "Dependency Confusion" attacks in Python/PyTorch libraries?

6.5.1 Is "Excessive Agency" restricted? (Can the AI Agent actually execute a financial transaction, or just suggest it?)

6.5.2 Are "Human Approval Gates" required for any AI action that modifies data or sends external communications?

6.5.3 Do AI Agents have their own "Non-Human Identity" managed in the IAM system (Domain 4)?

6.5.4 Is "Loop Detection" active to stop two AI agents from getting into an infinite conversation loop that drains resources?

6.5.5 Are "Goal Hijacking" attacks monitored (where the AI is tricked into pursuing a different objective)?

6.5.6 Is the "Context Window" flushed securely between sessions to prevent data leakage to the next user?

6.5.7 Are "Flash War" safeguards in place? (Preventing AI trading bots from crashing the market in milliseconds).

6.5.8 Is "Voice Cloning" authorization strictly enforced for AI agents acting as customer support?

6.5.9 Are "Plan" validations logged? (Does the system record how the Agent decided to execute a complex task?).

6.5.10 Is there a "Physical Safety" override for AI agents connected to robotics or IoT?

6.6.1 Is the "Training Environment" air-gapped or isolated from the corporate network?

6.6.2 Are "Model Weights" stored in an encrypted vault with strict access controls (preventing IP theft)?

6.6.3 Is "Compute Hygiene" enforced? (Ensuring GPUs are wiped of data between training runs).

6.6.4 Are "Jupyter Notebooks" scanned for hard-coded credentials before being committed to repositories?

6.6.5 Is "Drift Detection" active to alert if the model's accuracy degrades in production (Model Decay)?

6.6.6 Are "Inference Endpoints" protected by WAFs (Web Application Firewalls) specifically tuned for AI payloads?

6.6.7 Is "Rate Limiting" applied per API key to prevent "Oracle Attacks" on the model?

6.6.8 Are "Container Breakouts" monitored for AI workloads running in Kubernetes?

6.6.9 Is "Version Control" applied to data? (Can you roll back the data to a previous state, not just the code?).

6.6.10 Does the organization conduct "Red Teaming" exercises specifically using "AI Red Teams" (AI attacking AI)?

6.7.1 Is "Watermarking" (e.g., C2PA/SynthID) enforced on all AI-generated content to prove origin and detect deepfakes?

6.7.2 Does the organization have a legal "Indemnification" agreement with AI vendors protecting against copyright lawsuits?

6.7.3 Is there a "Clean Data" certification proving that no pirated books or protected art were used in model fine-tuning?

6.7.4 Are "Model Weights" treated as Trade Secrets and protected by specific Non-Disclosure Agreements (NDAs)?

6.7.5 Is there a mechanism to "Opt-Out" company data from being used to train the vendor's future foundation models?

6.7.6 Does the organization audit for "Model Laundering" (renaming a stolen model to hide its source)?

6.7.7 Are "Terms of Service" violations monitored? (e.g., using an API to generate content that violates the vendor's policy, risking a ban).

6.7.8 Is "Output Ownership" clearly defined? (If the AI writes code, does the company own it or the AI provider?).

6.7.9 Are "Style Mimicry" risks assessed? (Is the AI generating content that dangerously resembles a specific artist/competitor?).

6.7.10 Does the Rosecoin Ledger store the "Prompt History" to prove human creative input for copyright registration purposes?

6.8.1 Is "Anthropomorphism" restricted? (Is the AI programmed to clarify "I am an AI" and not say "I feel" or "I love"?).

6.8.2 Are "Persuasion Filters" active? (Detecting if the AI is using psychological tactics to manipulate the user's opinion).

6.8.3 Is "Emotional Reliance" monitored? (Flagging users who spend excessive hours chatting with support bots for companionship).

6.8.4 Are "Dark Nudges" prevented? (Ensuring the AI doesn't subtly push users toward higher-risk financial decisions).

6.8.5 Is "Truthfulness Tuning" applied? (Rewarding the model for saying "I don't know" rather than hallucinating a confident lie).

6.8.6 Are "Echo Chambers" avoided? (Ensuring the AI doesn't just reinforce the user's existing biases without providing context).

6.8.7 Is there a "Cognitive Load" limit? (Preventing the AI from flooding the user with too much information to force a mistake).

6.8.8 Are "Suicide/Self-Harm" triggers hard-coded to immediately divert to human help resources?

6.8.9 Is "Voice Synthesis" regulated to prevent the AI from sounding exactly like a trusted authority figure?

6.8.10 Does the organization conduct "Psy-Ops" testing to see if the AI can be tricked into radicalizing a user?

6.9.1 Is "Sovereign Training" enforced? (Are models for national defense trained on air-gapped supercomputers inside the country?).

6.9.2 Are "Export Controls" applied to Model Weights? (Treating advanced AI models as weapons that cannot be emailed abroad).

6.9.3 Is "Clearance Level" required for AI Trainers? (Vetting the humans who write the Reinforcement Learning feedback).

6.9.4 Are "Foreign Data Sources" scrubbed from the training set of national security models?

6.9.5 Is there a "Kill Chain" analysis for AI? (Understanding how an enemy AI could disrupt the organization's logistics).

6.9.6 Are "Model Guardrails" physically located on the server, not just in the API software layer?

6.9.7 Is "Nuclear/Bio Threat" knowledge unlearned? (Has the model been lobotomized of specific knowledge on how to build weapons?).

6.9.8 Are "Satellite/Orbital" AI models hardened against radiation flipping bits in their neural networks?

6.9.9 Is "Strategic Surprise" monitored? (Watching for competitor AI capabilities that suddenly render current encryption useless).

6.9.10 Does the organization participate in "AI Safety Institutes" (US/UK) to share threat intel on frontier models?

6.9.11 AI "Vibe Coding" Governance: Is there a security review process for code generated by AI assistants to detect "invisible" logic flaws or insecure dependencies before they reach production?

6.9.12 Prompt Injection Defense: Are AI-powered browsers and agentic tools configured with filters to prevent prompt injection attacks from malicious web content?

DOMAIN 7: NETWORK, 5G & EDGE SECURITY

7.1.1 Has the organization fully retired "Legacy VPNs" in favor of identity-aware ZTNA connectors?

7.1.2 Are "flat networks" eliminated by enforcing strict Micro-Segmentation down to the workload level?

7.1.3 Is "Device Posture" checked continuously (every 5 minutes) rather than just at the initial login?

7.1.4 Does the network treat "Internal Users" with the exact same hostility and verification level as "External Attackers"?

7.1.5 Are "Software-Defined Perimeters" (SDP) used to cloak critical applications, making them invisible to port scans?

7.1.6 Is "Least Privilege" applied to network routes (e.g., the printer VLAN cannot talk to the database VLAN)?

7.1.7 Does the organization block "Lateral Movement" by requiring re-authentication for every zone crossing?

7.1.8 Are "Service Meshes" (like Istio/Linkerd) used to secure East-West traffic between microservices with mutual TLS (mTLS)?

7.1.9 Is there a "Default Deny" policy for all outbound traffic that hasn't been explicitly whitelisted?

7.1.10 Does the ZTNA policy dynamically revoke access if the user's "Risk Score" increases (e.g., they downloaded malware)?

7.2.1 Is "Network Slicing" security enforced to ensure a breach in the "IoT Slice" cannot jump to the "Corporate Slice"?

7.2.2 Are "Rogue Base Stations" (Stingrays/IMSI Catchers) actively hunted using spectrum analyzers?

7.2.3 Does the organization use "Private 5G" with its own SIM cards/eSIMs rather than relying on public carrier security?

7.2.4 Are 5G "API Gateways" (NEF/SCEF) hardened to prevent attackers from querying subscriber location data?

7.2.5 Is "SUCI" (Subscription Concealed Identifier) enabled to encrypt the user's identity (IMSI) over the air?

7.2.6 Are "Roaming Interfaces" (IPX/GRX) monitored for signaling attacks coming from foreign telecom networks?

7.2.7 Is the "5G Core" (the brain of the network) isolated in a highly secure, containerized environment?

7.2.8 Are "Edge Computing" nodes (MEC) physically secured and tamper-proofed since they reside outside the main data center?

7.2.9 Does the organization audit the "Supply Chain" of 5G radio equipment (RAN) for backdoors?

7.2.10 Is "Slice Isolation" tested via penetration testing to prove logical separation holds under stress?

7.3.1 Is WPA3-Enterprise (192-bit security) mandatory for all corporate Wi-Fi connections?

7.3.2 Are "Management Frames" protected (PMF) to prevent de-authentication attacks that force users offline?

7.3.3 Is "Enhanced Open" (OWE) used for Guest Networks to encrypt traffic without requiring a password?

7.3.4 Does the wireless intrusion prevention system (WIPS) automatically "contain" Rogue Access Points by jamming their signals?

7.3.5 Are "Karma Attacks" (devices automatically connecting to fake names like 'Free Wi-Fi') prevented by endpoint profiles?

7.3.6 Is "MAC Randomization" supported and managed correctly to avoid tracking issues while maintaining security?

7.3.7 Are "Hidden SSIDs" recognized as "Security by Obscurity" and replaced with proper certificate-based authentication?

7.3.8 Is "Client Isolation" turned on by default for all IoT and Guest subnets?

7.3.9 Are "Bluetooth/BLE" scanning policies in place to detect unauthorized skimming devices in the office?

7.3.10 Is the physical placement of Access Points tuned to prevent "Signal Bleed" into the parking lot or street?

7.4.1 Is "RPKI" (Resource Public Key Infrastructure) enforced to prevent BGP Hijacking of the company's IP prefixes?

7.4.2 Is DNSSEC (Domain Name System Security Extensions) fully enabled to prevent DNS spoofing/cache poisoning?

7.4.3 Are "DOH" (DNS over HTTPS) and "DOT" (DNS over TLS) supported to encrypt DNS queries from prying eyes?

7.4.4 Is "IPv6 Security" fully matured (e.g., RA Guard enabled) to prevent attackers from using IPv6 backdoors on IPv4 networks?

7.4.5 Are "Man-in-the-Middle" (MitM) attacks detected by monitoring for SSL/TLS certificate anomalies?

7.4.6 Is "Network Time Protocol" (NTP) secured (NTS) to prevent time-drift attacks that break log forensics?

7.4.7 Are "Unused Ports" physically locked or administratively shutdown on all switches?

7.4.8 Is "Control Plane Policing" (CoPP) configured on routers to protect the CPU from DoS traffic?

7.4.9 Are "VLAN Hopping" defenses active (e.g., disabling DTP negotiation)?

7.4.10 Does the Rosecoin Ledger record "Route Changes" to provide an immutable history of network topology?

7.5.1 Is there a "Single-Pass" inspection architecture where traffic is decrypted once and scanned for everything (AV, DLP, IPS)?

7.5.2 Are "CASB" (Cloud Access Security Broker) policies integrated directly into the network path?

7.5.3 Is "Remote Browser Isolation" (RBI) triggered automatically for uncategorized or suspicious websites?

7.5.4 Does the SASE platform enforce "Data Sovereignty" by ensuring traffic is inspected in the correct local PoP (Point of Presence)?

7.5.5 Is "Encrypted Traffic Analysis" (ETA) used to find malware inside HTTPS streams without breaking encryption (using metadata)?

7.5.6 Are "Shadow IT" apps automatically blocked at the DNS/Network layer by the SASE gateway?

7.5.7 Is "Bandwidth Throttling" applied to non-business streaming services to preserve capacity for critical ops?

7.5.8 Are "Firewall-as-a-Service" (FWaaS) rules standardized globally rather than managed on individual appliances?

7.5.9 Is "User Experience" (DEM) monitored to ensure security scanning doesn't kill productivity?

7.5.10 Does the SASE provider guarantee "100% Uptime" via redundant global backbones?

7.6.1 Is "Volumetric DDoS" mitigation automated with a capacity of at least 3 Tbps to handle massive floods?

7.6.2 Are "Application Layer" (Layer 7) DDoS attacks (low and slow) detected by analyzing request behavior?

7.6.3 Is "BGP Flowspec" used to propagate drop rules to upstream ISPs instantly during an attack?

7.6.4 Are "Anycast" IP addresses used to distribute attack traffic across multiple global data centers?

7.6.5 Is there a "Clean Pipe" agreement with the ISP to scrub traffic before it hits the firewall?

7.6.6 Are "Rate Limits" applied to all API endpoints to prevent exhaustion attacks?

7.6.7 Is "Fail-Open" vs "Fail-Closed" logic defined and tested for every critical security appliance?

7.6.8 Are "Out-of-Band" (OOB) management ports available (via LTE/ Satellite) if the main network is crushed?

7.6.9 Is "Chaos Engineering" used to randomly unplug network cables and test self-healing capabilities?

7.6.10 Does the organization have a "Dark Web" alert for when their specific IP ranges are targeted by botnet-for-hire services?

7.7.1 Is "Quantum Key Distribution" (QKD) piloted for the most sensitive "Crown Jewel" fiber links?

7.7.2 Are "Satellite Internet" links (Starlink/OneWeb) secured with enterprise-grade encryption and not treated as trusted home networks?

7.7.3 Is "Optical Wiretapping" detection active on fiber lines (sensing physical vibrations or attenuation)?

7.7.4 Are "Li-Fi" (Light Fidelity) networks explored for ultra-secure rooms where radio waves are a risk?

7.7.5 Is "6G" research monitored to prepare for "Terahertz" frequency security challenges?

7.7.6 Are "Mesh Networks" (decentralized routing) tested for resilience in war-zone scenarios?

7.7.7 Is "Nano-Network" security (medical implants communicating) assessed if applicable to the industry?

7.7.8 Are "High-Altitude Platform Systems" (HAPS) considered in the disaster recovery connectivity plan?

7.7.9 Is "Underwater Cable" security (risk of cutting) mapped for international data dependencies?

7.7.10 Does the organization participate in "Internet Governance" bodies to influence future protocol security?

7.8.1 Is "Full Packet Capture" (FPC) enabled for critical entry/exit points to allow for retroactive breach analysis (the "Time Machine" capability)?

7.8.2 Does the organization use "JA3 Fingerprinting" to identify malware communication inside encrypted TLS traffic without decrypting it?

7.8.3 Is "East-West" traffic monitoring active to catch an attacker moving laterally from the printer VLAN to the Server VLAN?

7.8.4 Are "Beaconing" detectors tuned to find low-and-slow signals (e.g., a hacked server pinging a command center once every 24 hours)?

7.8.5 Is "Decryption Mirroring" used legally to inspect SSL traffic from high-risk users while preserving privacy for banking/health sites?

7.8.6 Are "Flow Logs" (NetFlow/IPFIX) retained for a minimum of 365 days to investigate "long-dwell" APTs?

7.8.7 Does the system automatically correlate "Network Spikes" with "Endpoint Process Launches" to pinpoint the exact app causing traffic?

7.8.8 Is "DNS Tunneling" detection active to stop attackers from smuggling data out via DNS queries?

7.8.9 Are "Honeypots" deployed inside the internal network to trick intruders into revealing themselves early?

7.8.10 Does the Rosecoin Ledger hash the "Chain of Custody" for pcap files to ensure they are admissible in court?

7.9.1 Is there a hard "Kill Policy" for legacy protocols (TLS 1.0/1.1, SMBv1, NTLM) with zero exceptions?

7.9.2 Are "Management Protocols" (Telnet, HTTP, FTP) strictly banned in favor of SSH, HTTPS, and SFTP?

7.9.3 Is "802.1X" authentication mandatory for every physical wired port (preventing someone from plugging in a laptop in the lobby)?

7.9.4 Are "Unused Dark Fiber" strands disconnected or monitored to prevent "Optical Tapping"?

7.9.5 Is "HSTS Preloading" enforced for all corporate domains to force browsers to never use insecure HTTP?

7.9.6 Are "LLMNR" and "NetBIOS" broadcast protocols disabled to prevent local credential spoofing attacks?

7.9.7 Is "IPv6 Leakage" prevented on IPv4-only VPNs to ensure traffic doesn't bypass the tunnel?

7.9.8 Are "Faraday Cages" or RF shielding used for "SCIF" (Sensitive Compartmented Information Facility) rooms?

7.9.9 Is "Cable Plant" security audited? (Are critical fiber conduits physically armored or alarmed?).

7.9.10 Does the organization perform "Wireless Sweeps" for hidden "Bugging Devices" or rogue cellular bridges in executive offices?

7.9.11 AI-Powered Browser Security: Are AI-integrated browsers and agentic tools restricted from processing unverified third-party web content to prevent prompt injection and automated data exfiltration?

7.9.12 API "Shadow" Discovery: Is there a continuous discovery process to find and secure "Zombie" or "Shadow" APIs that may be exposing corporate data to external AI training models?

DOMAIN 8: ENDPOINT, DEVICE & IOT SECURITY

8.1.1 Is "EDR/XDR" (Extended Detection & Response) mandatory on 100% of endpoints, replacing legacy Anti-Virus?

8.1.2 Are "BIOS/UEFI Passwords" enabled to prevent attackers from booting from external USB drives?

8.1.3 Is "Full Disk Encryption" (BitLocker/FileVault) enforced with keys stored in the TPM (Trusted Platform Module)?

8.1.4 Are "USB Mass Storage" ports blocked by default, requiring specific temporary approval to use thumb drives?

8.1.5 Does the organization use "Application Whitelisting" (e.g., AppLocker) to prevent any unsigned .exe from running?

8.1.6 Is "Local Admin" access revoked for all standard users to prevent malware installation?

8.1.7 Are "Firmware Updates" pushed automatically alongside OS patches (e.g., updating the Dell/HP BIOS remotely)?

8.1.8 Is there a "Remote Wipe" capability that works even if the device is not on the corporate VPN?

8.1.9 Are "Privacy Screens" mandatory for devices used in public spaces (trains/planes)?

8.1.10 Does the AINA OS perform a "Health Check" (patch level, firewall on) before allowing the device to connect to the network?

8.2.1 Is "Containerization" (Android Work Profile / iOS User Enrollment) used to strictly separate corporate data from personal photos?

8.2.2 Are "Jailbroken/Rooted" devices automatically detected and blocked from accessing corporate email?

8.2.3 Is "Copy/Paste Protection" active, preventing users from copying text from a work email into a personal WhatsApp chat?

8.2.4 Are "Managed Apps" configured to wipe themselves automatically if the device goes offline for >7 days?

8.2.5 Is "Biometric Enforcement" required to open any work app (FaceID for Outlook)?

8.2.6 Are "SMS Previews" disabled on the lock screen for 2FA codes to prevent shoulder surfing?

8.2.7 Is "Location Tracking" disabled for the MDM agent to respect user privacy (only tracking device location when "Lost Mode" is active)?

8.2.8 Are "Malicious Wi-Fi" networks (Man-in-the-Middle) detected by the mobile defense agent?

8.2.9 Is there a "Block List" for high-risk apps (e.g., banning TikTok on government-issued phones)?

8.2.10 Does the organization use "Mobile Phishing Defense" to filter malicious SMS (Smishing) links?

8.3.1 Are all IoT devices placed on a strictly isolated "IoT VLAN" that cannot route to the Finance/HR network?

8.3.2 Is there a "No Default Password" policy enforced before any IoT device is provisioned?

8.3.3 Does the organization use "NAC" (Network Access Control) to profile devices (e.g., "This is a Printer") and assign ACLs automatically?

8.3.4 Are "UPnP" (Universal Plug and Play) protocols disabled on all routers to prevent IoT devices from punching holes in the firewall?

8.3.5 Is "Firmware Analysis" performed on IoT devices to find hard-coded backdoors before deployment?

8.3.6 Are "Medical Devices" (IoMT) or "Industrial Controllers" (PLCs) protected by "Virtual Patching" at the gateway level if they can't be updated?

8.3.7 Is "East-West" traffic monitoring active to detect if a smart coffee machine is trying to hack a smart TV?

8.3.8 Are "Device Certificates" (mTLS) used for authentication instead of static API keys?

8.3.9 Is there a physical "Port Lock" or "Tamper Seal" on accessible IoT ports (e.g., lobby cameras)?

8.3.10 Does the Rosecoin Ledger record the "Device Identity" and ownership history of every sensor to prevent spoofing?

8.4.1 Is a "Hardware Bill of Materials" (HBOM) required to ensure chips aren't sourced from sanctioned entities?

8.4.2 Does the organization verify "Secure Boot" is enabled to ensure the OS hasn't been modified by a rootkit?

8.4.3 Are "Supply Chain Interdiction" checks performed (X-ray or tamper-tape inspection) on servers arriving from high-risk locations?

8.4.4 Is "DMA Protection" (Direct Memory Access) enabled to prevent Thunderbolt/FireWire attacks?

8.4.5 Are "Cold Boot" attack mitigations in place (encrypting RAM when the device sleeps)?

8.4.6 Is the "Management Engine" (Intel ME / AMD PSP) neutralized or monitored for out-of-band exploits?

8.4.7 Are "Hardware Implants" (e.g., keyloggers inside keyboards) considered in the threat model for high-security zones?

8.4.8 Is there a "Firmware Signing" key management policy (ensuring only the vendor can update the BIOS)?

8.4.9 Does the organization use "HSMs" (Hardware Security Modules) to generate and store device identity keys?

8.4.10 Is "E-Waste" security enforced? (Are chips physically crushed/shredded so data cannot be recovered from discarded IoT)?

8.5.1 Is "Geofencing" hard-coded into drones to prevent them from flying into restricted airspace or enemy territory?

8.5.2 Are "Remote Control" signals encrypted (AES-256) to prevent "Command Injection" hijacking of robots?

8.5.3 Is "Failsafe Logic" tested? (Does the robot safely power down if it loses connection, rather than going rogue?)

8.5.4 Are "Lidar/Sensor Spoofing" defenses active to prevent attackers from blinding autonomous vehicles?

8.5.5 Is "Swarm Security" managed? (ensuring a compromised drone cannot issue commands to the rest of the fleet).

8.5.6 Are "Black Box" recorders installed on autonomous units to investigate accidents/hacks?

8.5.7 Is "Over-the-Air" (OTA) updating cryptographically secured to prevent flashing malicious firmware to a fleet of robots?

8.5.8 Are "Battery Safety" protocols monitored to prevent thermal runaway attacks (causing physical fire via malware)?

8.5.9 Is "Human Detection" safety mandatory (robot stops instantly if a human is too close)?

8.5.10 Does the system validate "GPS Integrity" to detect spoofing attacks that could misguide drones?

8.6.1 Is there an automated "Kill Switch" that bricks a device if it is reported stolen?

8.6.2 Are "Leased Devices" wiped to DoD standards (3-pass overwrite) before being returned to the vendor?

8.6.3 Is "Device Drift" monitored? (Flagging devices that haven't checked in for 30 days for quarantine).

8.6.4 Are "End-of-Life" (EOL) dates tracked, with a mandatory replacement policy for unsupported hardware?

8.6.5 Is "Asset Tagging" linked to the digital inventory (scan a QR code to see the device's security status)?

8.6.6 Are "Loaner Devices" re-imaged automatically immediately upon return?

8.6.7 Is "Printer Security" enforced? (Hard drives encrypted, memory wiped after print jobs).

8.6.8 Are "Wearables" (Smartwatches) included in the BYOD policy if they access corporate alerts?

8.6.9 Is "Bluetooth Pairing" restricted to approved devices only?

8.6.10 Does the organization have a "Hardware Hacking Lab" to test the physical security of its own deployed devices?

8.7.1 Is the "Purdue Model" (or a modern Zero Trust equivalent) strictly enforced, ensuring Level 0 (Sensors) cannot talk to Level 5 (Internet)?

8.7.2 Are "Unidirectional Gateways" (Data Diodes) used to physically prevent data from flowing into the critical plant network?

8.7.3 Is "Passive Scanning" mandatory for OT networks (listening only) to avoid crashing fragile legacy PLCs with active probes?

8.7.4 Are "Engineering Workstations" (the keys to the kingdom) kept offline or require multi-person physical access to operate?

8.7.5 Is "Safety Instrumented System" (SIS) logic locked and independent, ensuring that even if the cyber controls fail, the physical safety latches work?

8.7.6 Are "Jump Hosts" strictly monitored with "Keystroke Recording" for any vendor maintenance on industrial controllers?

8.7.7 Is there a "Manual Override" drill performed annually? (Can humans physically turn the valves if the screens go black?).

8.7.8 Are "Transient Cyber Assets" (contractor laptops) scanned in a "Sheep Dip" kiosk before connecting to the plant floor?

8.7.9 Is "Project File" integrity checked? (Ensuring the logic uploaded to the PLC hasn't been modified to cause physical damage).

8.7.10 Does the Rosecoin Ledger store the "Configuration State" of critical infrastructure to prove regulatory compliance (NERC CIP)?

8.8.1 Is "Patient Safety" the #1 priority in the risk model (above data confidentiality)?

8.8.2 Are "Implantable Devices" (Pacemakers, Insulin Pumps) protected against "RF Replay Attacks" that could trigger unauthorized doses?

8.8.3 Is there a "Hospital Mode" firewall policy that isolates MRI/CT scanners from the guest Wi-Fi network?

8.8.4 Are "Default Passwords" on medical equipment changed before the device enters the clinical environment?

8.8.5 Is "Legacy Medical OS" protection active? (Isolating Windows XP MRI machines behind virtual firewalls).

8.8.6 Are "Telemetry Streams" (patient vitals) encrypted end-to-end to prevent tampering that could lead to misdiagnosis?

8.8.7 Is "Bio-Hacking" defense considered? (Protecting neural interfaces or smart prosthetics from malicious override).

8.8.8 Is there a "Crisis Protocol" for ransomware in a hospital? (Do doctors know how to switch to paper charts immediately?).

8.8.9 Are "Wearable Health Monitors" (smartwatches) vetted to ensure they aren't leaking employee health data to insurers?

8.8.10 Does the organization demand "Long-Term Support" (10+ years) guarantees from medical device vendors for security patches?

8.9.1 Is the "Home Router" of high-value executives scanned (with permission) for critical vulnerabilities?

8.9.2 Does the organization provide a "Corporate Wi-Fi Puck" (Hotspot) for executives to avoid using insecure home ISP routers?

8.9.3 Are "Smart Home" risks assessed? (Ensuring Alexa/Siri isn't listening to confidential Board meetings).

8.9.4 Is "Split Tunneling" disabled for high-risk users, forcing all traffic back through the corporate security stack?

8.9.5 Are "Family Devices" (kids' iPads) considered a lateral movement risk if on the same network as the work laptop?

8.9.6 Is there a "Clean Desk" policy for the home office (e.g., no sensitive papers visible in Zoom backgrounds)?

8.9.7 Are "VPN Always-On" configurations locked to prevent users from disabling security for faster Netflix speeds?

8.9.8 Is "Shoulder Surfing" training provided for remote workers (e.g., working in coffee shops)?

8.9.9 Does the endpoint agent detect "Rogue Devices" on the home network (e.g., a compromised smart bulb attacking the laptop)?

8.9.10 Is "Digital Exhaust" managed? (Ensuring home printers don't store copies of printed corporate documents).

DOMAIN 9: SECURE SOFTWARE DEVELOPMENT (SSDLC)

9.1.1 Is security integrated into the "Design Phase" (Threat Modeling) before a single line of code is written?

9.1.2 Are "IDE Plugins" (e.g., SonarLint, Snyk) mandatory for developers to catch vulnerabilities in real-time as they type?

9.1.3 Does the CI/CD pipeline have "Hard Gates" that automatically fail a build if critical vulnerabilities are detected?

9.1.4 Is "Policy-as-Code" (OPA) used to enforce security rules on infrastructure configurations (Terraform/Helm) within the repository?

9.1.5 Are "Pre-Commit Hooks" installed to prevent secrets (API keys, passwords) from ever being committed to Git?

9.1.6 Is the "Definition of Done" (DoD) for every sprint explicitly updated to include security acceptance criteria?

9.1.7 Does the organization measure "Mean Time to Remediate" (MTTR) for code vulnerabilities specifically?

9.1.8 Are "Security Champions" embedded within each development squad to bridge the gap between DevOps and SecOps?

9.1.9 Is "Immutable Infrastructure" practiced? (Servers are never patched in place; they are replaced with new, secure images).

9.1.10 Does the Rosecoin Ledger record the "integrity hash" of every release artifact to prevent supply chain injection attacks?

9.2.1 Is SAST (Static Analysis) run on every commit to scan source code for known patterns (e.g., SQL Injection)?

9.2.2 Is DAST (Dynamic Analysis) run against the staging environment to find runtime flaws that SAST misses?

9.2.3 Is IAST (Interactive Analysis) used during functional testing to identify vulnerabilities with lower false-positive rates?

9.2.4 Are "Container Scans" performed on Docker images to find vulnerabilities in the base OS layers?

9.2.5 Is "Fuzz Testing" (Fuzzing) applied to critical inputs to crash the application with random data, revealing memory errors?

9.2.6 Are "Business Logic" flaws (e.g., buying an item for $0.00) tested manually, as automated tools often miss them?

9.2.7 Does the organization use "Software Composition Analysis" (SCA) to track and patch open-source dependencies (e.g., Log4j)?

9.2.8 Are "False Positive" rates tracked to prevent developer fatigue (Alert Fatigue)?

9.2.9 Is there a "Vulnerability Management" dashboard that correlates findings from all tools into a single view?

9.2.10 Are "Secrets Scanning" tools running historically on the entire git history, not just the current head?

9.3.1 Is every API endpoint protected by strong authentication (OAuth 2.0 / OIDC) with no "Public by Default" routes?

9.3.2 Are "Broken Object Level Authorization" (BOLA/IDOR) checks performed to ensure User A cannot access User B's resources by changing an ID?

9.3.3 Is "Rate Limiting" and "Throttling" applied to all APIs to prevent DoS and scraping attacks?

9.3.4 Is "API Schema Validation" enforced to reject any input that does not strictly match the expected format (Positive Security Model)?

9.3.5 Are "Shadow APIs" (undocumented endpoints) actively hunted and shut down?

9.3.6 Is "Data Masking" applied to API responses to prevent "Excessive Data Exposure" (sending full PII when only the name is needed)?

9.3.7 Are "JWTs" (JSON Web Tokens) signed using strong algorithms (RS256) and secrets that are rotated regularly?

9.3.8 Is there a "Service Mesh" (e.g., Istio) ensuring mTLS encryption between internal microservices?

9.3.9 Are "API Gateways" used as the single point of entry, stripping malicious headers before they reach the backend?

9.3.10 Does the API documentation (Swagger/OpenAPI) match the actual deployment 100%?

9.4.1 Are all developers required to complete "Secure Coding Training" (e.g., RCCE-Dev) focused on their specific language (Java/Python/Go)?

9.4.2 Is the "OWASP Top 10" (and API Top 10) printed/accessible and referenced in code reviews?

9.4.3 Are "Peer Code Reviews" mandatory, with a specific checklist item for "Security Implications"?

9.4.4 Is "Input Validation" centralized in a trusted library rather than ad-hoc coding in every function?

9.4.5 Is "Output Encoding" used contextually (HTML, JavaScript, URL) to defeat Cross-Site Scripting (XSS)?

9.4.6 Are "Parameterized Queries" (Prepared Statements) the only allowed method for database interaction (Zero Tolerance for String Concatenation)?

9.4.7 Is "Cryptographic Agility" built in? (Can you switch from RSA to Elliptic Curve without rewriting the app?).

9.4.8 Is "Error Handling" generic? (Ensuring stack traces and system details are never shown to the end-user).

9.4.9 Is there a "Gamified" leaderboard to reward developers who fix the most security bugs?

9.4.10 Does the organization host internal "Capture The Flag" (CTF) events to teach developers how to think like hackers?

9.5.1 Are "Serverless Functions" (Lambda/Azure Functions) scanned for "Over-Privileged" IAM roles?

9.5.2 Is "Runtime Protection" (CWPP) active on containers to detect if a shell is spawned inside a pod?

9.5.3 Are "Kubernetes Secrets" stored encrypted (etcd encryption) and not as environment variables?

9.5.4 Is "Image Signing" (Cosign/Notary) enforced? (Cluster refuses to run images not signed by the CI pipeline).

9.5.5 Are "Network Policies" in Kubernetes configured to deny all traffic between namespaces by default?

9.5.6 Is "Misconfiguration Scanning" (KSPM) running to detect open dashboards or insecure kubelet ports?

9.5.7 Are "Ephemeral Containers" used for debugging instead of allowing SSH access to production nodes?

9.5.8 Is the "Attack Surface" of base images minimized (using Distroless or Alpine images)?

9.5.9 Are "Admission Controllers" used to prevent privileged containers (Root) from running?

9.5.10 Is "Infrastructure-as-Code" (IaC) scanned for compliance with RCF standards before deployment?

9.6.1 Does the organization enforce SLSA Level 3 (Supply-chain Levels for Software Artifacts)? (Is the build platform hardened, isolated, and verifiable?)

9.6.2 Are "Hermetic Builds" mandatory? (Does the build process run without network access to ensure no external malicious code is pulled in during compilation?)

9.6.3 Is "Dependency Pinning" enforced? (Are lockfiles used to ensure the exact same version of a library is used every time, preventing "Dependency Confusion" attacks?)

9.6.4 Are "Typosquatting" checks performed on internal package registries to prevent developers from accidentally installing "numpuy" instead of "numpy"?

9.6.5 Is there a "Two-Person Rule" for modifying build scripts (Pipeline-as-Code) to prevent a rogue admin from inserting a backdoor into the compiler?

9.6.6 Are "Reproducible Builds" verified? (Can you rebuild the software from source and get a bit-for-bit identical binary, proving no tampering occurred?)

9.6.7 Does the system scan for "Repo Jacking"? (Checking if an open-source maintainer account was compromised or sold to a malicious actor).

9.6.8 Are "License Checks" automated to prevent "Copyleft" (GPL) pollution that could legally force the company to open-source its proprietary code?

9.6.9 Is there a "Private Mirror" for public repositories (NPM/Maven) to insulate the company from "Left-Pad" style deletions or outages?

9.6.10 Does the Rosecoin Ledger store the "Provenance" (the full history) of every artifact, linking the binary back to the specific source code commit?

9.7.1 Are "AI Code Assistants" (Copilot/ChatGPT) configured to block the suggestion of insecure code patterns (e.g., hardcoded passwords)?

9.7.2 Is "AI-Generated Code" tagged and subjected to higher scrutiny in code reviews than human code?

9.7.3 Are "Low-Code/No-Code" platforms (PowerApps/Zapier) restricted to a sandbox environment that cannot access "Crown Jewel" data?

9.7.4 Is there a "Citizen Developer" certification required before a business user can publish a Low-Code app to the corporate catalog?

9.7.5 Are "Data Exfiltration" controls active on No-Code connectors (e.g., preventing a user from connecting Salesforce to their personal Google Sheet)?

9.7.6 Is "Prompt Hygiene" enforced for developers? (Ensuring they don't paste proprietary source code into public AI models for debugging).

9.7.7 Are "Ghost Apps" (abandoned Low-Code projects) automatically archived and deleted after 90 days of inactivity?

9.7.8 Is "Logic Flaw" scanning applied to Low-Code flows? (Checking for infinite loops or missing approval steps).

9.7.9 Does the organization verify the "Copyright Status" of AI-generated code to ensure it doesn't infringe on open-source licenses?

9.7.10 Is there a "Kill Switch" for all Low-Code integrations if a platform vulnerability is discovered?

9.8.1 Is the "Circuit Breaker" pattern used? (Does the application stop calling a failing service to prevent cascading failure across the system?)

9.8.2 Are "Bulkheads" implemented? (Is the application partitioned so that a crash in the "Reporting" module doesn't kill the "Checkout" module?)

9.8.3 Is "Graceful Degradation" tested? (If the recommendation engine fails, does the site still show products, or does it throw a 500 error?)

9.8.4 Is "Chaos Engineering" practiced? (Are tools like "Chaos Monkey" used to randomly kill pods in production to verify self-healing?)

9.8.5 Are "Timeouts" and "Retries with Jitter" configured correctly to prevent "Thundering Herd" problems during recovery?

9.8.6 Is "Fail Safe" (Fail Closed) the default? (If the security service crashes, does access default to "Deny" rather than "Allow"?)

9.8.7 Are "Health Checks" deep? (Do they check if the app can actually write to the DB, not just if the web server is up?)

9.8.8 Is "Cache Poisoning" defense active? (Ensuring one user's private data doesn't get cached and served to the next user).

9.8.9 Are "Feature Flags" used to instantly turn off a vulnerable feature in production without redeploying the whole app?

9.8.10 Does the system support "Blue/Green" or "Canary" deployments to roll back instantly if a security bug is found in the new version?

9.9.1 Is the "Strangler Fig" pattern used to slowly replace legacy monolithic code with secure microservices?

9.9.2 Is "Virtual Patching" (WAF rules) applied immediately to legacy apps that can no longer be updated at the code level?

9.9.3 Are "Hard-Coded Credentials" in legacy apps actively hunted and replaced with secrets management calls?

9.9.4 Is there a "Sunset Policy" that mandates a hard stop date for apps running on End-of-Life (EOL) languages (e.g., Python 2, Java 7)?

9.9.5 Are "Wrapper" interfaces used to add modern authentication (MFA) to legacy apps that don't natively support it?

9.9.6 Is "Crypto-Agility" forced into legacy updates? (Replacing MD5/SHA1 with SHA-256 whenever a module is touched).

9.9.7 Are "Zombie Libraries" (libraries that haven't been updated in 5 years) flagged for mandatory replacement?

9.9.8 Is "Technical Debt" quantified in dollars ($) and reported to the Board as a security risk?

9.9.9 Are legacy apps isolated in "Containment VLANs" with strict firewall rules limiting their reach?

9.9.10 Does the organization perform "Archaeology" sprints to document and understand undocumented legacy code before it breaks?

DOMAIN 10: CONTINUOUS MONITORING & DETECTION

10.1.1 Does the organization operate a 24/7/365 SOC (either internal or hybrid-managed) to ensure eyes are on the glass when attackers strike at 3 AM?

10.1.2 Is there a clear Tiered Structure (Tier 1 Triage, Tier 2 Analysis, Tier 3 Hunting) or a modern Swarm Model where experts collaborate instantly?

10.1.3 Are Handover Protocols formalized? (Does the Singapore team verbally brief the London team during shift changes to ensure context isn't lost?)

10.1.4 Is Mean Time to Detect (MTTD) tracked and relentlessly optimized (target: <1 minute for critical events)?

10.1.5 Is Mean Time to Respond (MTTR) measured from the moment of detection to the moment of containment (target: <15 minutes)?

10.1.6 Are SOC Analysts empowered to make decisions? (Can a Tier 1 analyst isolate a laptop without asking a manager if ransomware is confirmed?)

10.1.7 Is there a War Room (physical or virtual) equipped with out-of-band comms for managing major crises?

10.1.8 Does the SOC have Visualization Walls that show real-time attack maps, not just for show, but for identifying global trends?

10.1.9 Is Ticket Enrichment automated? (When an analyst opens a ticket, is the IP reputation, geo-location, and owner info already populated?)

10.1.10 Does the Rosecoin Ledger record the exact timestamp an analyst Picked Up a ticket to prove diligence during an audit?

10.2.1 Is the SIEM ingesting logs from All Layers? (Network, Endpoint, Identity, Cloud, Application, and Physical Badge Readers).

10.2.2 Are Correlation Rules updated weekly based on the latest Threat Intel (e.g., If PowerShell runs > Encoded Command > Outbound Traffic = Critical Alert)?

10.2.3 Is Log Retention tiered? (Hot storage for 90 days for fast search, Cold storage for 7 years for compliance).

10.2.4 Are Sigma Rules used to write detection logic that is portable across different SIEM platforms?

10.2.5 Is Decryption handled before ingestion? (Are SSL logs decrypted so the SIEM can actually see the SQL injection inside the packet?)

10.2.6 Are Silent Log Sources monitored? (Does the system alert if the Firewall stops sending logs, indicating a potential tamper or failure?)

10.2.7 Is Data Normalization automated? (Ensuring User_Name in AWS matches sAMAccountName in AD for seamless searching).

10.2.8 Are Blind Spots mapped annually? (Identifying which critical assets are not sending logs to the SIEM).

10.2.9 Is Privacy Masking active in the SIEM? (Hiding PII from analysts unless they have a warrant/approval to unmask).

10.2.10 Does the SIEM use Graph Databases to visualize relationships (e.g., User A talked to Server B which talked to IP C)?

10.3.1 Is Playbook coverage >80% for common alerts (e.g., Phishing, Malware, Failed Login)?

10.3.2 Does the SOAR platform automatically Enrich indicators? (Auto-checking VirusTotal/Whois so the human doesn't have to).

10.3.3 Is Automated Containment enabled for high-confidence threats? (e.g., If Ransomware = True, isolate host immediately without human approval).

10.3.4 Are False Positive loops automated? (If a user confirms Yes, I logged in from Bali, does the system close the ticket and learn?)

10.3.5 Can the SOAR trigger actions across Disparate Tools? (e.g., Block IP on Firewall + Suspend User in AD + Wipe Token in Okta).

10.3.6 Are Playbooks Version Controlled and treated as code (Git-backed)?

10.3.7 Is there a Human Decision Gate for irreversible actions (e.g., wiping a server) to prevent automation accidents?

10.3.8 Are Phishing Inboxes automated? (Users forward emails -> SOAR parses -> Detonates link -> Replies to user Safe or Malicious).

10.3.9 Does the SOAR calculate Time Saved (ROI) to justify the automation budget?

10.3.10 Is Case Management integrated? (Can legal/HR view the evidence timeline without accessing the technical SOAR interface?)

10.4.1 Is Threat Hunting a dedicated function, not just something we do when it's quiet?

10.4.2 Do hunters use Hypothesis-Driven searching? (e.g., I suspect APT29 is using DNS tunneling, let me look for long TXT records).

10.4.3 Are Indicators of Compromise (IoCs) swept retroactively? (Searching 1 year of logs for a newly discovered bad IP).

10.4.4 Is Deception Technology (Honeytokens) deployed? (Fake admin credentials left in memory to catch attackers dumping LSASS).

10.4.5 Are Beacons actively hunted? (Looking for jittery, low-frequency connections to C2 servers).

10.4.6 Is Outlier Analysis performed? (e.g., Why is this marketing laptop running a PowerShell script at 2 AM?).

10.4.7 Are Living off the Land (LotL) binaries monitored? (Watching for legit tools like certutil or bitsadmin downloading files).

10.4.8 Do hunters analyze Failed Attacks? (Understanding who tried to break in and failed, as they will likely try again).

10.4.9 Is there a Hunter's Notebook or Wiki where findings are shared to improve future automated rules?

10.4.10 Does the Rosecoin Blockchain verify the integrity of the Hunter's Report to prevent tampering with findings?

10.5.1 Is Baseline Profiling established for every user? (Knowing that Bob from Accounting never logs in from North Korea).

10.5.2 Are Peer Group comparisons used? (Alerting if Bob uploads 10GB of data when the rest of Accounting uploads 50MB).

10.5.3 Is Impossible Travel detection active? (Login from NY and London within 1 hour = Alert).

10.5.4 Are Privilege Escalation attempts detected behaviorally? (Standard user suddenly accessing Admin shares).

10.5.5 Is Data Exfiltration monitored via behavioral shifts? (e.g., User printing 500 pages on a Sunday night).

10.5.6 Are Service Accounts monitored for interactive logins? (A backup account should never type on a keyboard).

10.5.7 Is Flight Risk analysis performed? (Correlating Resignation Letter detection with High Volume Download).

10.5.8 Are Lateral Movement chains detected? (User A logs into Machine B, then Machine C, then Server D).

10.5.9 Is Risk Scoring dynamic? (Does the user's risk score drop back down if they behave normally for 30 days?)

10.5.10 Is Context Aware monitoring used? (Understanding that high traffic is normal during Backup Window hours).

10.6.1 Is Distributed Tracing used to track a request through microservices to find where the security failure occurred?

10.6.2 Are CloudTrail / CloudWatch (or equivalents) guarded? (Alerting instantly if logging is turned off).

10.6.3 Is Serverless Monitoring active? (Watching Lambda execution times for anomalies indicating crypto-mining).

10.6.4 Are Cost Spikes treated as security alerts? (A sudden $5,000 bill often means a cloud account compromise).

10.6.5 Is API Observability enabled? (Detecting strange parameter manipulation in API calls).

10.6.6 Are Ephemeral Assets logged? (Capturing logs from a container before it spins down and disappears forever).

10.6.7 Is Multi-Cloud logging centralized? (AWS, Azure, and GCP logs flowing into one Single Pane of Glass).

10.6.8 Are Infrastructure-as-Code changes logged? (Knowing who changed the Terraform script that opened Port 22).

10.6.9 Is Synthetic Monitoring used to simulate user traffic and detect availability/integrity issues?

10.6.10 Does AINA OS correlate Performance issues with Security issues? (e.g., CPU spike = DDoS or Mining?).

10.7.1 Is Alert Fatigue monitored? (Are specific rules silenced if they generate >90% false positives to save analyst sanity?)

10.7.2 Are Shift Limits enforced? (Preventing analysts from working >12 hours to avoid decision fatigue errors).

10.7.3 Is there a Psychological Support program for analysts who view disturbing content (e.g., during forensic investigations)?

10.7.4 Are Gamification elements used to reward analysts for quality investigations rather than just ticket quantity?

10.7.5 Is Career Pathing clear? (Ensuring Tier 1 analysts have a roadmap to become Hunters or Engineers).

10.7.6 Is Ergonomic Assessment mandatory for the SOC environment? (Monitors, lighting, seating).

10.7.7 Are Cognitive Breaks scheduled? (Mandatory time away from screens during high-stress shifts).

10.7.8 Is Rotation practiced? (Moving analysts between Triage and Engineering to prevent burnout).

10.7.9 Is No-Blame Post-Mortem the standard culture? (Focusing on process failure, not human error).

10.7.10 Are SOC Analysts trained on Critical Thinking and Bias detection?

10.8.1 Is the Chain of Custody for every major incident automatically hashed to the Rosecoin Blockchain?

10.8.2 Are Auditor Views available? (Read-only access for regulators to verify SOC performance without needing spreadsheets).

10.8.3 Is Log Immutability mathematically guaranteed by the ledger? (Proving logs weren't deleted by an insider).

10.8.4 Are Compliance Reports (PCI, HIPAA) generated instantly from live data rather than manually compiled?

10.8.5 Is Evidence Retention automated based on legal hold requirements?

10.8.6 Are Analyst Notes signed and timestamped to prevent retroactive editing of the investigation timeline?

10.8.7 Is Incident Severity classification recorded permanently to prevent downplaying breaches later?

10.8.8 Are Digital Fingerprints of malware samples stored on-chain for industry sharing?

10.8.9 Is SLA Performance (Time to Detect) verifiable by third parties via the blockchain record?

10.8.10 Does the system issue a Daily Integrity Certificate proving that the monitoring infrastructure itself was not compromised?

10.9.1 Is "Breach and Attack Simulation" (BAS) running 24/7? (continuously firing simulated ransomware samples to test if the EDR actually blocks them).

10.9.2 Are "Security Control Validations" automated? (e.g., Every morning, the system tries to visit a gambling site to prove the Web Proxy is still working).

10.9.3 Is "Detection Engineering" tested daily? (The BAS tool mimics a specific threat group, like Lazarus, to see if the SIEM alerts trigger).

10.9.4 Are "Assumed Breach" tests conducted? (Starting a test from inside the network to see how far an attacker can get before detection).

10.9.5 Is "Email Gateway Validation" active? (Sending safe, weaponized test emails to see if they bypass the spam filter).

10.9.6 Does the organization map its "Detection Coverage" against the MITRE ATT&CK heatmap dynamically? (Knowing exactly which 20% of techniques you cannot see).

10.9.7 Is "Regression Testing" performed on security rules? (Ensuring a new firewall change didn't accidentally break an old blocking rule).

10.9.8 Are "Cloud Configuration Tests" automated? (Simulating a public S3 bucket access attempt to verify the policy denies it).

10.9.9 Is "Lateral Movement" simulation active? (Trying to move from the Receptionist VLAN to the CEO VLAN to prove segmentation holds).

10.9.10 Does the Rosecoin Ledger record the daily "Pass/Fail" score of these automated tests to prove continuous diligence to insurers?

10.10.1 Is there a dedicated "Insider Threat Program" that combines HR, Legal, and Security data (not just technical logs)?

10.10.2 Are "High-Flight-Risk" employees (e.g., those on Performance Improvement Plans) monitored with heightened scrutiny?

10.10.3 Is "Sentiment Analysis" used on public communications (Slack/ Teams) to detect extreme hostility or disgruntlement? (Privacy aligned).

10.10.4 Are "Sabbotage Indicators" monitored? (e.g., A system admin deleting backups or changing retention policies right before resigning).

10.10.5 Is "Renegade Account" detection active? (Spotting users creating "backdoor" local accounts to maintain access after firing).

10.10.6 Are "Data Hoarding" behaviors flagged? (e.g., A sales rep downloading the entire CRM database to a USB drive).

10.10.7 Is there a "Negative Background Check" trigger? (If an employee is arrested for fraud outside of work, does HR notify Security?).

10.10.8 Are "Privileged Users" (Admins) subjected to more frequent background re-investigations?

10.10.9 Is "Exit Monitoring" automated? (Analyzing the last 30 days of activity for every departing employee for IP theft).

10.10.10 Does the organization use "Psycholinguistic Profiling" to detect potential espionage or coercion attempts against staff?

10.10.11 Continuous Exposure Management (CEM): Has the SOC transitioned from periodic vulnerability scans to CEM, which uses automated attack path analysis to identify and remediate the most exploitable routes into the network?

10.10.12 Predictive Threat Modeling: Does the SIEM/SOAR utilize machine learning to analyze past incident patterns and forecast likely future threat vectors before they are executed?

10.10.13 Context-Aware Vulnerability Prioritization: Does the monitoring system layer "context-aware" intelligence over CVE data, prioritizing flaws based on actual exploitability, reach within internal dependency graphs, and prevalence in the wild?

10.10.14 Law Enforcement Interoperability: Are threat intelligence platforms integrated with forensic data and legal workflows to facilitate real-time sharing and coordinated response with law enforcement?

10.10.15 Cross-Disciplinary Cognitive Defense: Is there a formal strategy to detect "Cognitive Attacks"—such as sentiment manipulation or article seeding—by integrating open-source intelligence with internal threat detection platforms?

10.10.16 Hypervisor & Virtualization Observability: Are there dedicated monitoring controls for the virtualization layer to identify stealthy attacks targeting hypervisors, which are currently a major industry blind spot?

DOMAIN 11: THREAT INTELLIGENCE & ADVERSARY TRACKING

11.1.1 Does the organization produce "Executive Intelligence Briefings" that translate cyber threats into business risks (e.g., "Nation-State tension in Region X threatens our supply chain")?

11.1.2 Are "Geopolitical Risk Factors" monitored? (e.g., tracking sanctions, wars, or elections that might trigger cyber-retaliation against the industry).

11.1.3 Is "Attribution" analyzed? (Does the organization know who is likely to attack—e.g., Lazarus Group vs. script kiddies—to prioritize defenses?)

11.1.4 Are "Threat Motives" defined? (Is the primary threat espionage, financial extortion, or hacktivism?).

11.1.5 Is "Sector-Specific" intel consumed? (Focusing on threats targeting the specific industry, e.g., Finance or Healthcare, rather than generic noise).

11.1.6 Does the CISO use Strategic Intel to drive "Budget Allocation"? (e.g., "Ransomware is up 300%, so we need more storage for immutable backups").

11.1.7 Are "Emerging Technologies" tracked as threats? (e.g., Monitoring academic papers for new quantum decryption breakthroughs).

11.1.8 Is there a "Competitor Watch"? (Monitoring if competitors are being breached to anticipate similar attacks).

11.1.9 Does the organization track "Legislative Threats"? (New laws that could act as a "threat" to business operations).

11.1.10 Does the Rosecoin Ledger store Strategic Assessments to prove to the Board that warnings were issued prior to an incident?

11.2.1 Is the "MITRE ATT&CK Framework" used as the common language for tracking adversary behaviors?

11.2.2 Are "Campaigns" tracked, not just individual alerts? (Connecting a phishing email today to a firewall scan last week).

11.2.3 Does the organization map its "Defensive Coverage" against known Actor TTPs? (e.g., "We know APT29 uses Golden Ticket attacks; do we have a rule for that?").

11.2.4 Are "Threat Hunting" missions driven by Operational Intel? (e.g., "Intel says Group X is targeting VPNs; let's hunt for VPN anomalies").

11.2.5 Is "Adversary Emulation" performed? (Red Teams mimicking the exact style of a specific threat actor to test defenses).

11.2.6 Are "YARA Rules" created and deployed to hunt for specific malware families in file stores?

11.2.7 Is "Sandboxing" used to extract TTPs from malware samples found in the wild?

11.2.8 Are "COAs" (Courses of Action) pre-planned for top threat actors? (If Actor Y attacks, we execute Playbook Z).

11.2.9 Is "Time-to-Weaponize" tracked? (How fast does a proof-of-concept exploit become a live attack in the wild?).

11.2.10 Does the organization track "Toolmarks"? (Identifying the specific custom tools used by adversaries).

11.3.1 Is "Feed Aggregation" automated? (Pulling IPs/URLs from 10+ sources and de-duplicating them).

11.3.2 Is "Intel Aging" enforced? (Automatically removing an IP from the blocklist after 30 days so the firewall doesn't overflow).

11.3.3 Is "Contextual Scoring" applied? (Rating an IoC as "High Confidence" if confirmed by multiple sources).

11.3.4 Are IoCs pushed "Real-Time" to security controls? (Intel updates the Firewall blocklist within minutes of discovery).

11.3.5 Is "False Positive" tuning active? (Whitelisting legitimate CDNs like Cloudflare that might accidentally get flagged).

11.3.6 Are "File Hashes" blocked at the EDR level globally?

11.3.7 Is "Email Intelligence" used? (Blocking subjects or senders associated with active global phishing campaigns).

11.3.8 Are "C2 (Command & Control)" lists updated hourly to block callbacks from infected hosts?

11.3.9 Is "Vulnerability Intelligence" used to prioritize patching? (Patching a "Medium" bug first because Intel says it's being exploited now).

11.3.10 Does the system validate IoCs against internal traffic? (Checking "Have we already seen this bad IP?" immediately upon receipt).

11.4.1 Are "VIP Credentials" monitored? (Scanning the Dark Web for the CEO's personal email/password dumps).

11.4.2 Is "Typosquatting" detection active? (Finding domains like `gOogle.com` or `company-support.com` registered to impersonate the brand).

11.4.3 Are "Rogue Mobile Apps" hunted? (Scanning App Stores for fake versions of the company's customer app).

11.4.4 Is "Social Media Impersonation" monitored? (Finding fake LinkedIn profiles of executives used for recruiting scams).

11.4.5 Are "Leaked Source Code" repositories scanned? (Checking GitHub/Pastebin for proprietary code posted by accident or malice).

11.4.6 Is "Carding Market" monitoring active? (For retail/banks: checking if customer credit cards are being sold in bulk).

11.4.7 Are "Deep/Dark Web Forums" scraped for mentions of the company name? (Detecting chatter about planning an attack).

11.4.8 Is "Takedown Service" automated? (Does the vendor automatically issue legal takedowns for phishing sites?).

11.4.9 Are "Partner Breaches" monitored? (Alerting if a key law firm or supplier appears on a ransomware leak site).

11.4.10 Is "Brand Sentiment" analysis used to detect disinformation campaigns designed to hurt stock price?

11.5.1 Is the organization a member of an "ISAC" (Information Sharing and Analysis Center)? (e.g., FS-ISAC for finance).

11.5.2 Are "STIX/TAXII" standards used to automate the machine-to-machine sharing of threat data?

11.5.3 Is there a "Traffic Light Protocol" (TLP) policy? (Clearly marking intel as TLP:RED, AMBER, GREEN, or WHITE).

11.5.4 Are "Anonymization" tools used when sharing intel? (Stripping victim-specific details before sharing with the community).

11.5.5 Is there a "Feedback Loop"? (Does Operations tell Intel "This alert was useful" or "This was junk"?).

11.5.6 Does the organization share "Sightings" back to the community? (Contributing to the "Herd Immunity" of the industry).

11.5.7 Are "Law Enforcement" liaisons established? (FBI/Interpol contacts for sharing high-level criminal intel).

11.5.8 Is "Internal Dissemination" tailored? (Sending technical hashes to the SOC, but high-level risks to the Board).

11.5.9 Are "Flash Reports" generated instantly for major global events (e.g., "Log4j is out, here is our status").

11.5.10 Does the Rosecoin Blockchain track the "Reputation" of intel sources? (Ignoring sources that frequently provide false positives).

11.6.1 Is "Generative AI" used to summarize complex threat reports into executive summaries automatically?

11.6.2 Are "Predictive Analytics" used? (AI forecasting "Based on current scanning, a ransomware attack is 80% likely in 48 hours").

11.6.3 Is "Translation" automated? (Instantly translating threat chatter from Russian/Chinese/Farsi into English for analysts).

11.6.4 Are "Graph Analytics" used to find hidden connections between disparate attack infrastructure?

11.6.5 Is "Sentiment Analysis" on hacker forums used to gauge the "Credibility" of a threat actor?

11.6.6 Are "Pattern Recognition" models used to identify new malware families without signatures?

11.6.7 Is "Auto-Attribution" attempted by AI? (Matching coding style in malware to known groups).

11.6.8 Are "News Feeds" curated by AI to filter out FUD (Fear, Uncertainty, Doubt) and clickbait?

11.6.9 Is "Natural Language Querying" enabled? (Can analysts ask the Intel database "Show me all threats targeting SQL in Europe"?).

11.6.10 Does the system detect "AI-Generated Malware" characteristics?

11.7.1 Is "Disinformation" planted? (Leaving fake "Admin Password" files that trigger alerts when opened).

11.7.2 Are "Breadcrumbs" used? (Fake credentials that lead attackers to a honeypot rather than real data).

11.7.3 Is "Avatar Management" active? (Are researchers using maintained, realistic fake personas to infiltrate hacker forums?).

11.7.4 Are "Canary Tokens" buried in documents, databases, and emails to detect theft?

11.7.5 Is "Attack Surface Manipulation" used? (Dynamically changing port numbers or banners to confuse scanners).

11.7.6 Are "Sinkholes" used to redirect malicious traffic for analysis?

11.7.7 Is "Counter-Surveillance" performed? (Watching the watchers—seeing who is scanning the network).

11.7.8 Are "Tarpits" deployed? (Systems that slow down scanning by holding the connection open indefinitely).

11.7.9 Is "Attribution Masking" used for outgoing defensive scans? (So attackers don't know you are scanning them).

11.7.10 Does the organization maintain a "Persona Non Grata" list of researchers/entities banned from interacting with the brand?

11.8.1 Is "Fusion Center" logic applied? (Merging Physical Security data—e.g., protests, war zones—with Cyber data to predict attacks on data centers or offices).

11.8.2 Is "Executive Travel Intel" automated? (Scanning for Wi-Fi pineapples or IMSI catchers in specific hotels before the CEO arrives).

11.8.3 Are "Upstream Dependencies" monitored? (Tracking the financial health and hacker chatter regarding your critical software vendors).

11.8.4 Is "Weather/Disaster Intel" correlated with availability risks? (e.g., Predicting network outages due to hurricanes affecting a specific ISP hub).

11.8.5 Are "Protest & Riot" monitors active? (Alerting if a hacktivist group announces a physical protest alongside a DDoS campaign).

11.8.6 Is "Satellite Imagery" used for physical supply chain monitoring? (e.g., detecting disruption at a chip manufacturing plant).

11.8.7 Is "Sub-Contractor" chatter monitored? (Detecting if a third-party janitorial or IT support firm is being targeted to get to you).

11.8.8 Are "Hardware Interdiction" risks tracked? (Intel on shipping routes that are compromised by nation-state tampering).

11.8.9 Is "BGP Hijacking" intel consumed? (Knowing if your traffic is being physically routed through a hostile nation).

11.8.10 Does the system map "Kinetic Impact"? (Translating a cyber threat, like a SCADA hack, into physical safety terms for the safety team).

11.9.1 Are "Priority Intelligence Requirements" (PIRs) defined by the Board? (e.g., "We need to know about any threat to our SWIFT payment gateway immediately").

11.9.2 Is "Intel Accuracy" measured? (Tracking the ratio of "True Positive" vs. "False Positive" alerts generated by intel feeds).

11.9.3 Is "Cost Avoidance" calculated? (e.g., "Intel warned us to block IP X, which attacked our competitor the next day, saving us $1M").

11.9.4 Is "Source Reliability" graded? (Automatically degrading the trust score of a feed vendor that sends stale data).

11.9.5 Is "Time-to-Acknowledge" tracked? (How fast does the SOC read a Critical Intel Report?).

11.9.6 Is "Stakeholder Satisfaction" surveyed? (Asking the CISO and Patch Team: "Did this report actually help you?").

11.9.7 Is "Intel-Driven Patching" measured? (Tracking how many emergency patches were prioritized solely due to Intel warnings).

11.9.8 Are "Legal Guidelines" audited? (Ensuring researchers didn't accidentally commit a crime while interacting with Dark Web actors).

11.9.9 Is "Data Obsolescence" managed? (Purging old intel so the database doesn't become slow and toxic).

11.9.10 Does the Rosecoin Ledger record the "Value Realization" metrics to justify the annual Threat Intel budget?

DOMAIN 12: VULNERABILITY MANAGEMENT & SECURITY TESTING

12.1.1 Is "Continuous Scanning" enabled? (Scanning assets 24/7 rather than waiting for a weekly or monthly window).

12.1.2 Are "Authenticated Scans" the default? (Logging into the server to see deep vulnerabilities, not just scanning the outside ports).

12.1.3 Does the scanner cover "All Layers"? (OS, Database, Middleware, Web App, and Cloud Configuration).

12.1.4 Is "Risk-Based Prioritization" used? (Prioritizing a "Medium" vulnerability that has a live exploit over a "High" vulnerability that is theoretical).

12.1.5 Are "ephemeral assets" (containers that live for minutes) scanned at the image level before deployment?

12.1.6 Is there a "Zero Unscanned Asset" policy? (If a device touches the network, it must be scanned or blocked).

12.1.7 Are "False Positives" actively managed? (Tuning the scanner to stop reporting "missing headers" as Critical).

12.1.8 Is "Agent-Based Scanning" used for laptops and remote workers who are rarely on the office VPN?

12.1.9 Does the system automatically "Ticket" the correct team? (Sending DB flaws to DBAs and Java flaws to Developers).

12.1.10 Does the Rosecoin Ledger record the "Scan Completion Date" for every asset to prove compliance to auditors?

12.2.1 Is "Penetration Testing" performed on every critical application before it goes live (Gold Master)?

12.2.2 Do tests go beyond "Compliance"? (Testing for business logic flaws, like "Can I buy this item for $0?", not just CVEs).

12.2.3 Are "API Pen Tests" mandatory? (Specifically testing for BOLA/IDOR vulnerabilities in mobile backends).

12.2.4 Is "Grey Box" testing preferred? (Giving the tester credentials and diagrams to find deep bugs, rather than wasting time guessing passwords).

12.2.5 Are "Retests" mandatory? (Verifying the fix actually works and didn't introduce new bugs).

12.2.6 Is there a "Safe Harbor" agreement? (Protecting the testers legally as long as they follow the Rules of Engagement).

12.2.7 Are "Cloud Penetration Tests" conducted annually on the Azure/AWS configuration itself?

12.2.8 Does the organization use "Accredited Testers" (CREST/OSCP) rather than generic tool-runners?

12.2.9 Are "Mobile Apps" tested on jailbroken devices to check for inadequate root detection?

12.2.10 Are the "Final Reports" hashed to the Rosecoin Blockchain to ensure findings cannot be deleted or hidden by management?

12.3.1 Is "Red Teaming" conducted without notifying the SOC? (Testing the detection capability, not just the defenses).

12.3.2 Do Red Teams simulate "Specific Threat Actors"? (e.g., "Attack us exactly like Ransomware Group X would").

12.3.3 Is "Purple Teaming" practiced? (Red Team attacks while Blue Team watches, sharing notes in real-time to improve alerts).

12.3.4 Does the Red Team target "People and Process"? (e.g., Calling the Help Desk to reset a password), not just technology.

12.3.5 Is "Physical Access" in scope? (Trying to tailgate into the data center or plant a rogue Wi-Fi device).

12.3.6 Are "Assume Breach" scenarios run? (Starting the test with a "compromised" laptop to see if they can pivot to the Domain Controller).

12.3.7 Is "Data Exfiltration" tested? (Can they actually steal 10GB of data without an alarm going off?).

12.3.8 Is there a "White Card" / "Get out of Jail Free" letter for physical testers if caught by security guards?

12.3.9 Are "Objective-Based" missions defined? (e.g., "Your goal is to wire $10k to this account," not just "hack the server").

12.3.10 Is the "Debrief" constructive? (Focusing on "How do we detect this next time?" rather than shaming the SOC).

12.4.1 Is there a public "security.txt" file on the website directing researchers where to report bugs?

12.4.2 Does the organization run a "Private Bug Bounty" program for critical assets?

12.4.3 Is "Triage" managed effectively? (Ensuring developers aren't flooded with spam reports).

12.4.4 Are "Payouts" competitive? (Paying enough to incentivize top hackers to look at your code vs. selling the exploit).

12.4.5 Is there a strict "No Retaliation" policy for ethical hackers who report vulnerabilities in good faith?

12.4.6 Are "Researcher Metrics" tracked? (Time to Triage, Time to Bounty).

12.4.7 Is the "Scope" clearly defined? (Explicitly forbidding DDoS or attacks on third-party hosting).

12.4.8 Does the organization allow "Public Disclosure" after a fix is verified (boosting community trust)?

12.4.9 Are "Internal Bounties" offered? (Paying employees—non-security staff—who find and report bugs).

12.4.10 Are "Smart Contracts" used to automate bounty payments upon verification (Rosecoin integration)?

12.5.1 Are "Phishing Simulations" run monthly with varying difficulty levels (not just obvious Nigerian Prince scams)?

12.5.2 Is "Vishing" (Voice Phishing) tested? (Calling employees pretending to be IT support).

12.5.3 Are "USB Drops" conducted? (Leaving "Payroll_2026.xlsx" USB drives in the lobby to see who plugs them in).

12.5.4 Is "Smishing" (SMS) testing included? (Sending fake MFA requests or delivery notifications).

12.5.5 Is there a "No-Shame" culture? (Users who click are trained, not punished, to encourage reporting).

12.5.6 Is "Executive Whaling" tested? (Simulating targeted attacks against the C-Suite).

12.5.7 Is "Deepfake" testing conducted? (Using voice synthesis to test finance team verification procedures).

12.5.8 Are "QR Code" attacks (Quishing) tested in the office environment?

12.5.9 Is the "Reporting Button" usage tracked? (Success is "High Reporting Rate," not just "Low Click Rate").

12.5.10 Are "Physical Tailgating" tests performed at office entry points?

12.6.1 Are "SLAs" (Service Level Agreements) enforced? (Criticals fixed in 48 hours, Highs in 7 days).

12.6.2 Is "Root Cause Analysis" performed for recurring vulnerabilities? (Why do we keep having SQL injection?).

12.6.3 Is "Risk Acceptance" formal? (If a bug can't be fixed, is the risk signed off by an Exec and expiring in 90 days?).

12.6.4 Is "Virtual Patching" (WAF rules) applied immediately while waiting for the code fix?

12.6.5 Are "Re-Scans" automated? (Scanner automatically verifies the patch once the ticket is closed).

12.6.6 Is "Patch Verification" independent? (Security team verifies the fix, not the team that applied it).

12.6.7 Are "Legacy Systems" that cannot be patched placed in "Containment VLANs"?

12.6.8 Is "Vulnerability Age" tracked? (Flagging zombies—vulnerabilities that have been open for >1 year).

12.6.9 Is "Accountability" clear? (Every asset has an owner who is responsible for its patching).

12.6.10 Does the Rosecoin Ledger provide an immutable "Certificate of Hygiene" proving the system was clean at a specific date?

12.7.1 Is "External Attack Surface Management" (EASM) active? (Continuously scanning the entire internet to find assets belonging to the company that IT doesn't know about).

12.7.2 Are "Shadow Cloud Accounts" hunted? (Detecting if a developer opened a personal AWS account using a corporate credit card).

12.7.3 Is "Certificate Monitoring" automated? (Scanning for SSL certificates issued to your domain by unauthorized CAs).

12.7.4 Are "Forgotten Subdomains" reclaimed? (Detecting marketing-2021.company.com pointing to a deleted Azure resource, vulnerable to takeover).

12.7.5 Is "Code Repository" leakage monitored? (Scanning GitHub/GitLab for developers accidentally publishing private keys or code).

12.7.6 Are "SaaS Permutations" checked? (Finding orphaned Trello boards or Notion pages with "Public" access).

12.7.7 Is "Brand Asset" discovery active? (Finding old marketing microsites left running on unpatched WordPress).

12.7.8 Are "Partner Connections" mapped? (Identifying "backdoors" where third-party vendors have persistent connections).

12.7.9 Is "M&A Due Diligence" automated? (Instantly mapping the attack surface of a company being acquired).

12.7.10 Does the Rosecoin Ledger record the "Asset Discovery Timeline" to prove diligence in finding Shadow IT?

12.8.1 Is "Configuration Scanning" continuous? (Detecting an open S3 bucket or unencrypted database within seconds of creation).

12.8.2 Are "CIS Benchmarks" applied automatically? (Ensuring every cloud workload meets the hardening standard by default).

12.8.3 Is "Drift Detection" active? (Alerting if a Terraform-deployed infrastructure is manually changed in the console).

12.8.4 Are "Over-Privileged Roles" flagged? (Identifying IAM roles that have "Admin" rights but haven't used them in 90 days).

12.8.5 Is "Network Reachability" analyzed? (Mathematically proving which internal databases are actually reachable from the internet).

12.8.6 Are "Serverless Configurations" audited? (Checking Lambda functions for hard-coded secrets or insecure runtimes).

12.8.7 Is "Multi-Cloud" posture unified? (Applying the same "Encrypt Everything" rule across AWS, Azure, and GCP simultaneously).

12.8.8 Are "Container Registries" scanned for public access? (Preventing proprietary Docker images from being pulled by the public).

12.8.9 Is "Flow Log" analysis used to detect traffic to bad IPs even if the configuration looks correct?

12.8.10 Does the system "Auto-Remediate" critical flaws? (Automatically reverting a Security Group change that opened Port 22 to the world).

12.9.1 Is "Exploit Prediction" used? (Using AI to predict which vulnerabilities will be exploited in the next 30 days, based on dark web chatter).

12.9.2 Is "Contextual Risk Scoring" applied? (Downgrading a "Critical" CVE if the server is air-gapped and unreachable).

12.9.3 Are "Asset Value" multipliers used? (A bug on the "Main Database" scores 100x higher than the same bug on a "Lobby Display").

12.9.4 Is "Campaign Intelligence" integrated? (Prioritizing vulnerabilities that are currently being used by Ransomware groups targeting your sector).

12.9.5 Is "Remediation forecasting" tracked? (Predicting "It will take 45 days to fix this" to manage Board expectations).

12.9.6 Are "Win Rates" calculated? (Tracking if the "Mean Time to Remediate" is faster than the "Mean Time to Exploit").

12.9.7 Is "Zero-Day" modeling performed? (Simulating "What if a Zero-Day hits our VPN today?" to test resilience).

12.9.8 Are "Social Media" trends used? (Alerting if a researcher tweets about a bug in your software stack before the CVE is published).

12.9.9 Is "Historical Regression" analyzed? (Identifying teams that repeatedly introduce the same vulnerabilities).

12.9.10 Does the organization use "EPSS" (Exploit Prediction Scoring System) alongside CVSS scores?

12.9.11 KEV-Driven Patching: Is the patching schedule prioritized based on the CISA Known Exploited Vulnerabilities (KEV) catalog rather than just raw CVSS scores?

12.9.12 SBOM Dependency Tracking: Does the organization use Software Bills of Materials (SBOMs) to track and remediate vulnerabilities in sub-dependencies of third-party software?

DOMAIN 13: INCIDENT RESPONSE

13.1.1 Is there a Board-approved "Incident Response Charter" that explicitly grants the CISO authority to sever connections (e.g., unplug the internet) without seeking CEO permission during a crisis?

13.1.2 Is the "IR Plan" updated quarterly and stored in an offline, physical format (e.g., printed "Battle Binders") to ensure accessibility during a total network outage?

13.1.3 Are "Role Cards" assigned to specific individuals (Incident Commander, Scribe, Liaison) so everyone knows their exact job in the fog of war?

13.1.4 Is there a "Retainer Agreement" with a top-tier Forensics Firm (e.g., Mandiant/CrowdStrike) with a guaranteed SLA of <4 hours on-site?

13.1.5 Does the organization maintain a "Crisis Budget" (pre-authorized emergency funds) to buy bitcoins for ransom (if legal/approved) or emergency hardware instantly?

13.1.6 Is "Legal Privilege" established immediately? (Does outside counsel direct the investigation to protect findings from discovery in future lawsuits?)

13.1.7 Are "Out-of-Band Communications" established? (Signal/Wire groups or satellite phones ready for when email and Slack are compromised).

13.1.8 Is there a "Cyber Insurance" hotline card in the wallet of every executive?

13.1.9 Does the plan define "Severity Levels" (SEV1 - Critical to SEV4 - Low) with distinct escalation timeframes for each?

13.1.10 Does the Rosecoin Ledger record the "Declaration of Incident" timestamp to legally prove when the clock started for regulatory reporting (e.g., 72-hour GDPR window)?

13.2.1 Is "Golden Hour" triage practiced? (The ability to determine scope—how many machines are infected—within the first 60 minutes).

13.2.2 Are "Memory Forensics" captured immediately? (Dumping RAM before rebooting to capture encryption keys or resident malware).

13.2.3 Is "Live Response" capability active? (Can analysts remotely shell into a machine to kill a process or delete a file without alerting the attacker?)

13.2.4 Are "Forensic Images" of critical evidence hashed and write-protected immediately to preserve chain of custody?

13.2.5 Is "Patient Zero" identification a priority metric? (Finding the initial entry vector to close the door).

13.2.6 Are "Timeline Analysis" tools used to reconstruct the attack (e.g., "The attacker logged in at 2:00, pivoted at 2:05")?

13.2.7 Is "Malware Reverse Engineering" performed (internally or via vendor) to understand the kill switch or beaconing behavior?

13.2.8 Are "False Flag" checks performed? (Ensuring the DDoS isn't just a distraction for a quiet data theft happening elsewhere).

13.2.9 Is "Scope Creep" monitored? (Continually asking "Is it really just these 5 servers, or is it the whole Active Directory?").

13.2.10 Does AINA OS automatically generate a "Triage Report" summarizing the infected assets for the Incident Commander?

13.3.1 Is "Micro-Segmentation" used for containment? (Isolating the "Infected VLAN" from the rest of the network with one click).

13.3.2 Is "Identity Isolation" performed? (Resetting the "krbtgt" account and forcing a global password reset for all admins immediately).

13.3.3 Are "Kill Switches" tested? (Can you sever the link to the internet or third-party partners instantly?)

13.3.4 Is "Eradication" total? (Rebuilding infected servers from "Known Good" media rather than just cleaning the virus).

13.3.5 Are "Backdoor" checks performed? (Scanning for new user accounts, scheduled tasks, or webshells left behind by the attacker).

13.3.6 Is "Patching" part of eradication? (Closing the vulnerability that let them in before bringing systems back online).

13.3.7 Is "Recovery Prioritization" defined by business value? (Restoring the Payroll system before the Cafeteria Menu system).

13.3.8 Are "Integrity Checks" run on restored data? (Verifying that the backups weren't also encrypted or poisoned).

13.3.9 Is "Phased Re-Entry" practiced? (Bringing systems online one by one and watching for beaconing traffic).

13.3.10 Does the organization have a "Clean Network" (Green Zone) built in parallel to migrate clean assets into?

13.4.1 Are "Holding Statements" pre-written and legally approved for various scenarios (Ransomware, Data Leak, DDoS)?

13.4.2 Is there a "Dark Website" (a dormant crisis status page) ready to go live if the main website is defaced or down?

13.4.3 Are "Spokespeople" media-trained specifically for cyber crises (avoiding phrases like "No evidence of data theft" which usually age poorly)?

13.4.4 Is "Internal Communication" managed? (Telling employees what to do—"Don't turn on your PC"—without causing panic).

13.4.5 Are "Partner Notifications" automated? (Legally required notification to banks, suppliers, or customers within strict timelines).

13.4.6 Is "Social Media" monitoring active? (Replying to rumors on Twitter/X before they spiral out of control).

13.4.7 Is there a specific plan for "Extortion" handling? (How to respond if the attacker emails journalists directly).

13.4.8 Are "Call Center Scripts" updated? (Giving support agents the exact words to say to angry customers).

13.4.9 Is "Regulator Liaison" assigned? (A specific person responsible for talking to the FBI, ICO, or SEC).

13.4.10 Does the organization practice "Silence Discipline"? (Ensuring no leaks come from the IT team during the investigation).

13.5.1 Is there a Board-level "To Pay or Not to Pay" framework established before the emotion of a crisis hits?

13.5.2 Are "OFAC Checks" (Sanctions) mandatory before any payment? (Checking if the wallet belongs to a sanctioned entity like North Korea).

13.5.3 Is a "Crypto Broker" on retainer? (Someone who can legally source $5M in Bitcoin and transfer it securely).

13.5.4 Are "Decryptor Tests" performed? (Asking the attacker to decrypt one benign file to prove their tool works).

13.5.5 Is "Double Extortion" anticipated? (Assuming they stole data and encrypted it, and planning for the leak).

13.5.6 Are "Negotiators" used? (Professional negotiators who know the psychological profiles of specific ransomware groups).

13.5.7 Is "Backup Immunity" verified? (Are backups immutable/air-gapped so the attacker couldn't delete them first?).

13.5.8 Is there a "Manual Workaround" plan? (Can the factory run on paper tickets for 10 days if we decide not to pay?).

13.5.9 Are "Exfiltration Logs" analyzed? (Determining exactly what was stolen to see if paying to suppress the leak is even worth it).

13.5.10 Does the Rosecoin Ledger record the transaction details of any ransom payment for future legal defense?

13.6.1 Is a "Hot Wash" (Debrief) conducted within 48 hours of closure while memories are fresh?

13.6.2 Is a "Root Cause Analysis" (RCA) document produced using the "5 Whys" method?

13.6.3 Are "Corrective Actions" tracked in a ticket system? (Ensuring the hole is actually fixed and not just forgotten).

13.6.4 Is the "Risk Register" updated? (Re-calculating risk scores based on the reality of the breach).

13.6.5 Are "Playbooks" refined? (Updating the "Ransomware Playbook" because Step 3 didn't work).

13.6.6 Is "Executive Reporting" honest? (Presenting the "Good, Bad, and Ugly" to the Board, not a sanitized version).

13.6.7 Are "Audit Scopes" adjusted? (If the breach happened in a "Low Risk" area, is that area now "High Risk"?).

13.6.8 Is "Employee Support" offered? (Counseling for the IT team who worked 100-hour weeks during the crisis).

13.6.9 Is "Evidence Retention" managed? (Keeping the forensic images for 3-7 years in case of lawsuits).

13.6.10 Does the organization share "Sanitized Intel" (IoCs) with the industry ISAC to protect others?

13.7.1 Is there a strict "No Illegal Action" policy? (Clarifying that "hacking back" is illegal in most jurisdictions, but "Active Defense" is not).

13.7.2 Are "Tarpits" used? (Trapping the attacker in a slow loop to waste their time and resources).

13.7.3 Is "Attribution Beaconing" used? (Embedding files with tokens that "phone home" when the attacker opens them on their own machine).

13.7.4 Are "Sinkholes" deployed? (Redirecting the attacker's C2 traffic to a server you control).

13.7.5 Is "Doxing" defense active? (Investigating who the attacker is to provide intel to law enforcement for arrest).

13.7.6 Are "Honey-Patches" used? (Looking like you patched the vulnerability, but actually leaving it open as a trap).

13.7.7 Is "Deception Routing" used? (Making the production database look like a test server to confuse the attacker).

13.7.8 Are "Web Beacons" used in document watermarks to track where stolen data ends up?

13.7.9 Is there close coordination with "National Cyber Centers" (CISA/NCSC) who can legally hack back?

13.7.10 Does the organization have a "Legal Opinion" on file for every Active Defense technique used?

13.8.1 Is there a specific "Vendor Breach Playbook"? (Steps to take when a critical supplier like Microsoft, AWS, or a Payroll provider announces a breach).

13.8.2 Are "Disconnect Criteria" defined? (At what point do we sever the VPN tunnel to a compromised vendor? e.g., "Confirmed Lateral Movement").

13.8.3 Is "Trust Verification" required before reconnection? (Forcing the vendor to provide a fresh "Clean Health Certificate" or third-party audit before turning the data pipe back on).

13.8.4 Are "Fourth-Party" impacts mapped? (If your vendor's vendor gets hacked, do you know which of your services break?).

13.8.5 Is there a "Joint War Room" protocol? (Establishing a direct line to the Vendor's CISO, bypassing their public support desk).

13.8.6 Are "Source Code" audits triggered? (If a dev-tool vendor is breached, do you immediately scan your own code for injected backdoors?).

13.8.7 Is "Data Exposure" assumed? (If a SaaS HR provider is breached, do you assume all employee data is gone and trigger ID theft protection immediately?).

13.8.8 Are "Alternative Suppliers" identified? (If the primary logistics vendor is down for 3 weeks, who do we call?).

13.8.9 Is "Inbound Email" filtering tightened? (Automatically quarantining emails from the breached vendor's domain to prevent phishing).

13.8.10 Does the Rosecoin Ledger record the "Vendor Notification Time" vs. "Our Action Time" to prove we mitigated the risk instantly?

13.9.1 Is "Snapshot Forensics" automated? (Scripting the ability to snapshot a compromised EC2 instance and move it to a forensic VPC for analysis instantly).

13.9.2 Are "Token Revocations" global? (Understanding that "Resetting a Password" is not enough; you must revoke all active OAuth/OIDC refresh tokens).

13.9.3 Is "SaaS Log Retention" extended? (Paying for the "Premium Logging" tier in M365/Salesforce so logs aren't deleted after 7 days).

13.9.4 Is "Shadow Copy" checking performed? (Did the attacker use the cloud backup feature to exfiltrate data?).

13.9.5 Are "Serverless" functions analyzed? (Checking for malicious code injected into Lambda/Azure Functions that run without a server).

13.9.6 Is "Console Access" locked down? (Ensuring the attacker didn't create a "Break Glass" admin user in the root cloud account).

13.9.7 Are "Ephemeral Artifacts" captured? (Grabbing logs from containers before the attacker kills them).

13.9.8 Is "Cross-Tenant" movement checked? (Ensuring the attacker didn't hop from your Dev tenant to your Prod tenant).

13.9.9 Is "eDiscovery" capability ready? (Can you search across 10,000 cloud mailboxes for a specific keyword in minutes?).

13.9.10 Does the organization use "Cloud-Native" IR tools? (Using tools built for APIs, not trying to use EnCase on an S3 bucket).

Domain 14: Resilience, Business Continuity & Disaster Recovery

14.1.1 Is the "BIA" updated annually? (Quantifying exactly how much money is lost per hour of downtime for each specific app).

14.1.2 Are "RTOs" (Recovery Time Objectives) and "RPOs" (Recovery Point Objectives) defined by the business owners, not IT?

14.1.3 Is there a "Criticality Tiering" system? (Tier 0 = The Brand dies without it, Tier 4 = Nice to have).

14.1.4 Does the strategy differentiate between "Cyber Disaster" (Ransomware) and "Physical Disaster" (Fire)? (Cyber requires clean data; Physical requires any data).

14.1.5 Is "Interdependency Mapping" complete? (Knowing that the Payroll system cannot recover until the Identity system is up).

14.1.6 Is "Minimum Business Continuity Objective" (MBCO) defined? (What is the absolute minimum functionality needed to survive?).

14.1.7 Are "Manual Workarounds" documented? (Can the call center use pen and paper if the CRM is down for 3 days?).

14.1.8 Is "Key Person Risk" identified? (Who are the 3 people who know the passwords, and where is the envelope if they are unavailable?).

14.1.9 Does the organization have a "Geographically Dispersed" recovery site (>500 miles away) to survive regional catastrophes?

14.1.10 Does the Rosecoin Ledger store the BIA sign-off to prove to regulators that the Board accepted the downtime risks?

14.2.1 Is the "3-2-1-1-0 Rule" enforced? (3 copies, 2 media types, 1 offsite, 1 offline/immutable, 0 errors).

14.2.2 Is "Immutability" technically guaranteed? (Using S3 Object Lock or WORM drives so even a Domain Admin cannot delete backups).

14.2.3 Is there an "Air-Gapped" Clean Room? (A sterile environment to restore and scan data for malware before reconnecting to production).

14.2.4 Are "Backup Accounts" separate? (Ensuring the backup admin credentials are not the same as the Windows Domain Admin credentials).

14.2.5 Is "Multi-Person Auth" (M of N control) required to destroy or modify backup retention policies?

14.2.6 Are "SaaS Data" (M365, Salesforce) backed up independently? (Not relying on the vendor's "Recycle Bin").

14.2.7 Is "Encryption Key" backup managed? (Ensuring you have the keys to decrypt the backups if the Key Management Server (KMS) is destroyed).

14.2.8 Are "Gold Images" of servers kept offline? (To rebuild infrastructure quickly without reinstalling from scratch).

14.2.9 Is "Backup Poisoning" detection active? (Alerting if the daily backup size changes drastically, indicating encryption).

14.2.10 Does the organization perform "Restore Tests" monthly? (Proving you can actually read the tape/disk, not just write to it).

14.3.1 Is "Active-Active" architecture used for Tier 0 apps? (Traffic is load-balanced across two data centers; if one dies, the other takes over instantly).

14.3.2 Are "Single Points of Failure" (SPoF) eliminated? (Redundant ISPs, redundant routers, redundant power supplies).

14.3.3 Is "Load Balancing" health-check based? (Automatically diverting traffic away from a failing server).

14.3.4 Are "Circuit Breakers" coded into apps? (Preventing a failure in one module from crashing the entire system).

14.3.5 Is "Auto-Scaling" enabled? (Automatically adding servers during a DDoS or high-traffic event).

14.3.6 Are "Availability Zones" (AZs) utilized in the cloud? (Spreading servers across physically separate buildings).

14.3.7 Is "DNS Failover" automated? (Switching traffic to the DR site at the domain level).

14.3.8 Are "Database Replicas" monitored for lag? (Ensuring the secondary DB is actually in sync).

14.3.9 Is "Graceful Degradation" designed? (The website stays up, but "Search" is disabled to save resources).

14.3.10 Does AINA OS automatically trigger "Self-Healing" scripts (e.g., restarting a service) upon failure detection?

14.4.1 Is there a "Cloud Exit Strategy"? (Can you move to another cloud or on-prem if AWS/Azure permanently fails or bans you?).

14.4.2 Is "Source Code Escrow" verified? (Can you get the code for your critical SaaS app if the vendor goes bankrupt?).

14.4.3 Are "Portability Containers" used? (Using Kubernetes to make moving apps between clouds easier).

14.4.4 Is "Region Lock" risk assessed? (What happens if the "US-East-1" region goes down entirely?).

14.4.5 Are "Vendor SLAs" financially backed? (Do they pay significant penalties for downtime?).

14.4.6 Is "SaaS Configuration" backed up? (Backing up the settings and policies, not just the data).

14.4.7 Is "Shadow IT" continuity considered? (If a department relies on a credit-card SaaS app, what happens if the card expires?).

14.4.8 Are "Hybrid" dependencies mapped? (Ensuring the cloud app doesn't break if the on-prem VPN connection fails).

14.4.9 Is "Internet Dependency" minimized? (Can the factory floor continue to operate locally if the internet is cut?).

14.4.10 Does the organization monitor "Global Internet Health" (e.g., ThousandEyes) to distinguish between internal faults and ISP outages?

14.5.1 Is there a "Pandemic/Bio-Safety" plan? (Can 100% of the workforce work remotely indefinitely?).

14.5.2 Is "Succession Planning" documented? (Who becomes Acting CEO if the executive team is incapacitated?).

14.5.3 Are "Emergency Alerts" (Mass Notification System) automated via SMS/WhatsApp?

14.5.4 Is "Physical Safety" prioritized? (Evacuation drills for fire/active shooter).

14.5.5 Are "Alternate Work Locations" contracted? (WeWork or hot-sites ready if the HQ burns down).

14.5.6 Is "Hardware Supply" buffered? (Keeping a stock of 50 laptops in case of supply chain disruption).

14.5.7 Are "Paper Forms" printed and ready? (For hospitals/logistics to operate manually).

14.5.8 Is "Psychological Support" available post-disaster?

14.5.9 Are "Family Plans" encouraged? (Helping employees prepare their homes so they can focus on work recovery).

14.5.10 Is "Media Monitoring" active to gauge public sentiment during an outage?

14.6.1 Is a "Full Scale" DR test performed annually? (Actually failing over, not just talking about it).

14.6.2 Are "Tabletop Exercises" (TTX) conducted quarterly with executives?

14.6.3 Is "Chaos Engineering" practiced? (Randomly killing servers in production to prove resilience).

14.6.4 Are "After Action Reports" (AAR) mandatory after every test?

14.6.5 Is "Drift Management" active? (Ensuring the DR site matches the Production site's patches and configs).

14.6.6 Are "Call Trees" tested? (Actually calling people to see if they answer).

14.6.7 Is "Backup Integrity" verified automatically? (The system boots the backup VM and takes a screenshot).

14.6.8 Are "Network Bubbles" used for testing? (Spinning up the DR site without conflicting with Production IP addresses).

14.6.9 Is "Recovery Time" timed with a stopwatch? (Did we meet the 4-hour RTO?).

14.6.10 Does the Rosecoin Ledger record the "Test Certificate" to prove readiness to cyber insurers?

14.7.1 Is "Flood Plain" analysis performed for all data center locations? (Ensuring the server room isn't in a 100-year flood zone).

14.7.2 Are "Temperature Limits" tested? (Can the data center cooling system handle a record-breaking 50°C heatwave without shutting down?).

14.7.3 Is "Water Independence" planned? (If the municipal water supply fails, does the cooling system have a backup reservoir?).

14.7.4 Are "Power Grid" dependencies diverse? (Ensuring the primary and secondary data centers are on different power substations).

14.7.5 Is "Fuel Polishing" performed? (Regularly cleaning the diesel in the generator tanks to ensure it doesn't degrade and fail when needed).

14.7.6 Are "Renewable Backups" utilized? (Solar/Battery storage on-site to run critical comms if diesel runs out).

14.7.7 Is "Physical Access" feasible during disasters? (If roads are flooded, does the team have a helicopter or boat contract to reach the site?).

14.7.8 Are "Supply Chain Routes" mapped for weather risk? (Knowing that a hurricane in Taiwan delays hardware replacement).

14.7.9 Is "Air Quality" monitoring active? (Detecting wildfire smoke that could clog server air intakes).

14.7.10 Does the organization have a "Carbon Impact" plan for recovery? (Understanding the emissions cost of running on generators for a week).

14.8.1 Is there an "Offline Internet" repository? (A local mirror of critical documentation, code libraries, and Wikipedia needed to rebuild civilization/business if the internet is cut).

14.8.2 Are "Satellite Uplinks" (Starlink/Kymeta) independent of the local grid? (Tested to work when fiber cables are severed).

14.8.3 Is "Crypto-Sovereignty" managed? (Having local HSMs so you don't rely on a cloud provider to unlock your own encrypted data).

14.8.4 Are "Long-Term Provisions" stocked? (Food, water, and cots for IT staff to live in the data center for 7 days during a civil crisis).

14.8.5 Is "Radio Communication" established? (Shortwave/Ham radio protocols for long-distance coordination without cell towers).

14.8.6 Are "Paper Keys" used? (Physical cryptographic key parts stored in safe deposit boxes to reconstruct the root of trust).

14.8.7 Is "Cash" (Physical Fiat) available? (Small bills in a safe to pay for fuel/supplies if electronic banking is down).

14.8.8 Are "Local Dependencies" mapped? (Ensuring the "Global" DR plan doesn't fail because the local Keycard System is down).

14.8.9 Is "Data Sovereignty" verified? (Ensuring your backup data doesn't cross borders where it could be seized by a foreign government).

14.8.10 Does the Rosecoin Ledger store the "Reconstruction Constitution"? (The ultimate unchangeable document defining how to rebuild the company from zero).

14.8.11 Immutable Backup Integrity: Are backup repositories technically immutable and stored using "isolated credentials" that are not connected to the primary domain?

14.8.12 Measured Restore Tests: Are restore tests conducted quarterly for critical identity systems (Active Directory/IdP) with recorded "Time-to-Restore" metrics?

14.8.13 "Clean Room" Recovery Path: Is there a pre-defined, sterile "Green Zone" environment where systems can be restored and scanned for dormant malware before being reintroduced to the production network?

14.8.14 Cyber-Physical Convergence: Does the Business Continuity Plan include specific triggers for physical safety overrides if a cyber incident impacts Industrial Control Systems (ICS) or IoT devices?

14.8.15 Operational Resilience vs. Simple Backup: Does the strategy prioritize "operational resilience"—the ability to maintain critical functions during a long-term incident—rather than just data restoration?

14.8.16 Supply Chain Dependency Failover: Is there a verified plan to maintain business operations if a primary third-party technology provider (e.g., a major cloud or SaaS vendor) suffers a multi-day global outage?

14.8.17 "Green Zone" Restoration: Is there a pre-configured, sterile "Green Zone" environment where infrastructure can be rebuilt and verified before re-connecting to the production network?

14.8.18 Sovereign Data Survivability: Can critical operations continue in "Islanding Mode" if national or international fiber backbones are severed or geofenced?

14.8.19 Multi-Stage Extortion Readiness: Does the business continuity plan explicitly address "multi-stage extortion" scenarios where attackers simultaneously encrypt data and threaten the release of high-value exfiltrated assets?

14.8.20 Immutable Backup Validation: Are backups verified to be technically "immutable" (unchangeable), and is this verified through regular automated recovery drills to ensure access to insurance coverage?

DOMAIN 15: DIGITAL FORENSICS & INVESTIGATION

15.1.1 Is there a defined "Forensic Readiness Policy"? (Defining the capability to collect evidence before an incident occurs, rather than scrambling to find tools).

15.1.2 Is the "Forensics Lab" (or designated zone) physically secured with strict access logs and video surveillance?

15.1.3 Are "Forensic Workstations" air-gapped from the corporate network to prevent cross-contamination?

15.1.4 Are "Write Blockers" (hardware/software) mandatory for all evidence handling to ensure the original drive is never modified?

15.1.5 Is there a "Legal Authorization" matrix? (Who authorizes the seizure of an employee's laptop? Legal? HR? CISO?).

15.1.6 Are "Privacy Impact Assessments" conducted? (Ensuring investigations comply with GDPR/Works Councils when reviewing employee emails).

15.1.7 Is "Tool Validation" performed? (Regularly testing EnCase, FTK, or Axiom to prove they produce accurate results for court).

15.1.8 Are "investigators" certified? (GCFA, GCFE, or CCE certifications required to testify as expert witnesses).

15.1.9 Is there a "Case Management System"? (Tracking every case number, investigator assignment, and status).

15.1.10 Does the Rosecoin Ledger record the "Hash Values" (MD5/SHA256) of every evidence file upon acquisition to prove data integrity?

15.2.1 Is "Order of Volatility" followed? (Capturing RAM first, then Swap, then Disk, then Network Logs).

15.2.2 Is "Live Acquisition" used for encrypted disks? (Capturing the logical volume while the machine is on, as pulling the plug might lock the data forever).

15.2.3 Are "Remote Collection" agents deployed? (Ability to silently image a laptop in a hotel room over the VPN).

15.2.4 Is "Cloud-to-Cloud" preservation used? (Using API tools to preserve an Office 365 mailbox directly to an S3 bucket without downloading it to a laptop).

15.2.5 Are "Faraday Bags" used for mobile device seizure? (Blocking cellular signals to prevent remote wiping).

15.2.6 Is "Volatile Data" (clipboard, chat history) captured automatically during triage?

15.2.7 Are "Physical Seals" (Tamper Tape) used on seized hardware?

15.2.8 Is "Network Packet Capture" (PCAP) retained for specific investigation windows?

15.2.9 Are "Snapshot" capabilities used for virtual machines? (Freezing the state of a VM for analysis).

15.2.10 Does the process handle "Bio-metric Unlocking"? (Legal policy on whether you can force a user to use their fingerprint/face to unlock a seized phone).

15.3.1 Is "Timeline Analysis" automated? (Creating a "Super Timeline" of file system changes, web history, and event logs).

15.3.2 Are "Registry Hives" analyzed? (Checking for USB connections, recent documents, and user activity).

15.3.3 Is "Link File" (LNK) analysis performed? (Proving a user opened a specific file even if the file itself was deleted).

15.3.4 Are "Prefetch/Shimcache" artifacts analyzed? (Proving a program was executed).

15.3.5 Is "Browser Forensics" deep? (Recovering deleted history, cache, and session tokens).

15.3.6 Is "Email Forensics" capable of header analysis? (Tracing the true origin IP of a spoofed email).

15.3.7 Are "Deleted Files" carved? (Using file signature analysis to recover data from unallocated space).

15.3.8 Is "Keyword Searching" optimized? (Using GREP/Regex to find credit card numbers or "Confidential" markers across terabytes of data).

15.3.9 Are "Unknown Binaries" detonated? (Sending suspicious files to a sandbox for behavior analysis).

15.3.10 Does the analysis include "User Attribution"? (Proving who was sitting at the keyboard, not just which account was logged in).

15.4.1 Is "Mobile Extraction" tiered? (Logical extraction for speed, Physical extraction for deleted data).

15.4.2 Are "Encrypted Apps" (Signal/WhatsApp) parsed? (Using key extraction tools to read secure chat databases).

15.4.3 Is "Geolocation" mapped? (Plotting GPS data from photos and Wi-Fi connections to show user movement).

15.4.4 Are "Wearables" (Apple Watch/Fitbit) analyzed? (Using heart rate or step data to prove user activity times).

15.4.5 Is "Vehicle Forensics" considered? (Extracting navigation history and phone logs from Infotainment systems).

15.4.6 Are "Drone" flight logs analyzed? (For physical security incidents).

15.4.7 Is "IoT State" captured? (Smart speaker logs or smart lock access history).

15.4.8 Are "JTAG / Chip-off" techniques available? (Physically unsoldering chips from damaged devices to recover data).

15.4.9 Is "Cloud Backup" extraction used if the physical phone is locked? (Pulling the iCloud backup instead of cracking the passcode).

15.4.10 Does the lab handle "GrayKey" or similar advanced unlocking tools?

15.5.1 Is "Legal Hold" automation in place? (Instantly preventing deletion of emails for 50 custodians with one click).

15.5.2 Is "Early Case Assessment" (ECA) used? (Rapidly scanning data to tell Legal "We have 1 million docs, cost to review is $X").

15.5.3 Are "De-NISTing" and "De-duplication" applied? (Removing known system files and duplicate emails to reduce review volume).

15.5.4 Is "Technology Assisted Review" (TAR/AI) used? (Training the AI to find "Relevant" documents so lawyers don't read junk).

15.5.5 Are "Conversation Threads" reconstructed? (Grouping replies so the lawyer reads the story, not disjointed messages).

15.5.6 Is "Production Export" strictly formatted? (Converting data to Bates-stamped PDF/TIFF/Load Files for court exchange).

15.5.7 Are "Privilege Logs" automated? (Flagging emails involving outside counsel for Attorney-Client Privilege).

15.5.8 Is "Audio/Video" transcription automated for review?

15.5.9 Are "Chat Platforms" (Slack/Teams) processed natively? (Displaying emojis and reactions correctly, not just text).

15.5.10 Does the Rosecoin Ledger record the "Chain of Custody" for every eDiscovery export delivered to opposing counsel?

15.6.1 Is "Deepfake Detection" used? (Analyzing frequency domains in audio/video to detect AI synthesis).

15.6.2 Is "Steganography" detection active? (Checking images for hidden data/messages embedded in pixels).

15.6.3 Is "EXIF/Metadata" analysis standard? (Checking photo timestamps and camera serial numbers).

15.6.4 Are "Wiping Tools" detected? (Looking for evidence of CCleaner or BleachBit usage).

15.6.5 Is "Timestomp" detection performed? (Checking if file creation dates were manually altered).

15.6.6 Are "Encryption Detectors" used? (Finding VeraCrypt containers disguised as random files).

15.6.7 Is "Image Enhancement" scientifically valid? (Clarifying license plates in CCTV without "creating" pixels that don't exist).

15.6.8 Are "Audio Forensics" used? (Filtering background noise to hear conversations).

15.6.9 Is "Spoofing" detection active? (Identifying if a screenshot was Photoshopped).

15.6.10 Does the organization use "Hashing" to prove that the evidence file has never changed since the moment of capture?

15.7.1 Is "Snapshotting Automation" active? (Automatically triggering an AWS EBS snapshot the moment a high-severity alert fires on a VM).

15.7.2 Are "Serverless Logs" verbose? (Ensuring Lambda/Azure Functions log the payload of the event, not just the execution time, for reconstruction).

15.7.3 Is "Container Freezing" possible? (Pausing a compromised Docker container to inspect its memory before it crashes or spins down).

15.7.4 Are "SaaS Audit Logs" centralized? (Aggregating 90 days of logs from Slack, Salesforce, and Zoom into a single searchable lake).

15.7.5 Is "Cross-Account" forensic access pre-configured? (Does the forensic team have a "Break Glass" role to read data from the Production account without needing approval during a crisis?).

15.7.6 Are "API Activity" chains reconstructed? (Visualizing "Key A created VM B, which accessed Bucket C").

15.7.7 Is "Golden Image" comparison used? (Diffing the compromised container against the original registry image to find exactly what files were changed).

15.7.8 Are "Orphaned Resources" investigated? (Checking if that running EC2 instance belongs to a project that was deleted 6 months ago).

15.7.9 Is "Metadata Timelining" used for files stored in Object Storage (S3/Blob)? (Tracking PutObject and GetObject calls).

15.7.10 Does the organization use "Forensic-as-Code"? (Scripts that automatically build a clean analysis workstation in the cloud and mount the evidence drives securely).

15.8.1 Is "Four-Eyes Principle" enforced for accessing sensitive employee data? (Requiring two distinct admins to authenticate before reading the CEO's email).

15.8.2 Is "Data Minimization" practiced? (Searching only for the specific keyword "Project X" rather than reading the entire mailbox).

15.8.3 Are "Personal vs. Professional" boundaries defined? (Policy on handling "Personal" folders on corporate laptops).

15.8.4 Is "Investigator Auditing" active? (Who watches the watchers? Logging every file the investigator opens).

15.8.5 Are "Anonymized Reports" the default? (Referring to "User A" in technical reports until Legal authorizes unmasking).

15.8.6 Is "Bias Awareness" training mandatory? (Ensuring investigators don't target specific employees based on non-technical factors).

15.8.7 Is "Works Council" (Union) approval workflow integrated? (For regions like Germany/France, ensuring labor reps sign off on monitoring).

15.8.8 Are "Non-Relevant" data disposal policies strict? (Immediately deleting the employee's family photos found during the scan).

15.8.9 Is "Privileged Material" (Attorney-Client) automatically flagged and quarantined from the investigation team?

15.8.10 Does the organization have a clear "Expectation of Privacy" waiver signed by all employees annually?

15.9.1 Is the "Chain of Custody" recorded on the blockchain? (Every time a hard drive changes hands, a transaction is signed).

15.9.2 Are "Evidence Hashes" timestamped on the ledger? (Mathematically proving "This file existed in this state at this exact time").

15.9.3 Is "Investigator Action" logging immutable? (Preventing a rogue investigator from deleting their own command history).

15.9.4 Are "Digital Fingerprints" of malware shared via smart contracts? (Automating threat intel sharing).

15.9.5 Is the "Case File" integrity verified? (Ensuring pages haven't been torn out of the PDF report).

15.9.6 Are "Access Grants" tokenized? (Giving external counsel a temporary "Token" to view evidence that expires automatically).

15.9.7 Is "Tamper-Evident" storage used for physical evidence (bags with NFC/Blockchain tags)?

15.9.8 Are "Interview Recordings" hashed immediately? (Preventing "deepfake" allegations regarding witness testimony).

15.9.9 Is "Cross-Border" evidence transfer tracked? (Proving compliance with data sovereignty laws via the ledger).

15.9.10 Does the Rosecoin Ledger provide a "Verifier Portal" for courts to independently check the validity of the digital evidence?

DOMAIN 16: POST-QUANTUM SECURITY

16.1.1 Has the organization conducted a "Cryptographic Inventory"? (Identifying every instance of RSA, ECC, and Diffie-Hellman currently running in the environment).

16.1.2 Is "Store Now, Decrypt Later" (SNDL) risk assessed? (Identifying data stored today that will still be sensitive in 10 years when a Quantum Computer could decrypt it).

16.1.3 Is the "Mosca Theorem" applied? (Calculating: If Time to Migrate + Shelf Life of Data > Time to Quantum, you are already too late).

16.1.4 Are "Hard-Coded Keys" hunted? (Scanning source code for developers who manually implemented specific crypto libraries instead of calling the OS API).

16.1.5 Is "Vendor Readiness" tracked? (Demanding roadmaps from Microsoft, AWS, and Cisco on when they will support PQC algorithms).

16.1.6 Is "Data Classification" updated for Quantum? (Flagging "Long-Life" secrets like Genetic Data or Nuclear Blueprints that need immediate protection).

16.1.7 Is "Key Length" monitoring active? (Ensuring AES-256 is the absolute minimum, as AES-128 is weakened by Grover's Algorithm).

16.1.8 Are "Root Certificates" (PKI) planned for rotation? (How will you replace the root CA certificate on 50,000 IoT devices?).

16.1.9 Is "Perfect Forward Secrecy" (PFS) mandatory? (Ensuring that a compromised key today doesn't decrypt past traffic).

16.1.10 Does the Rosecoin Ledger record the "Crypto-Agility Score" of the organization to benchmark readiness against competitors?

16.2.1 Is a "Hybrid Mode" enforced? (Using both Classical (ECC) and Post-Quantum (Kyber/Dilithium) algorithms simultaneously during the transition years).

16.2.2 Is "Crypto-Agility" architected? (Can you switch the encryption algorithm of the entire database by changing one config file, without rewriting code?).

16.2.3 Are "NIST PQC Standards" (FIPS 203/204/205) adopted? (Using ML-KEM/Kyber for encryption and ML-DSA/Dilithium for signatures).

16.2.4 Is "Hardware Acceleration" assessed? (Can the current firewalls handle the larger key sizes and heavier math of PQC without slowing down the network?).

16.2.5 Are "Certificate Authorities" (internal) upgraded to issue Quantum-Safe certificates?

16.2.6 Is "TLS 1.3" (or 1.4 draft) enforced? (Disabling older protocols that cannot support hybrid key exchange).

16.2.7 Are "VPN Tunnels" upgraded first? (Prioritizing the outer shell of the network for quantum resistance).

16.2.8 Is "SSH" configured for PQC? (Ensuring admin access to servers is protected against future decryption).

16.2.9 Are "Performance Benchmarks" run? (Testing if the new algorithms kill mobile battery life or add unacceptable latency).

16.2.10 Does the organization have a "Back-Out Plan"? (If a flaw is found in the new PQC algorithm, can you revert safely?).

16.3.1 Is "QKD" piloted for the most sensitive links? (Using the laws of physics—photon entanglement—to exchange keys, which cannot be intercepted without detection).

16.3.2 Are "Satellite QKD" links considered? (For inter-continental secure communication that bypasses undersea cables).

16.3.3 Is "Dark Fiber" utilized? (QKD often requires dedicated fiber strands; does the organization own/lease them?).

16.3.4 Are "Trusted Nodes" secured? (Since QKD has distance limits, are the repeater stations physically bunkered?).

16.3.5 Is "QRNG" (Quantum Random Number Generation) used? (Using quantum noise for entropy rather than pseudo-random software algorithms).

16.3.6 Is "Side-Channel" defense active on QKD gear? (Ensuring the hardware itself doesn't leak the keys via electromagnetic radiation).

16.3.7 Are "Fiber Taps" monitored via quantum interference? (Detecting if anyone is trying to splice the line).

16.3.8 Is "Key Management" integrated? (Can the standard KMS ingest keys generated by the exotic QKD hardware?).

16.3.9 Are "Point-to-Point" limitations understood? (Recognizing QKD doesn't work well on routed networks yet).

16.3.10 Does the Rosecoin Ledger verify the "Photon Source" integrity to ensure the QKD hardware isn't a simulation?

16.4.1 Are "Hash-Based Signatures" (XMSS/LMS) used for the root ledger? (These are stateful but extremely resistant to quantum attack).

16.4.2 Is the "Merkle Tree" structure quantum-hardened? (Using quantum-safe hash functions like SHA-3 or SHAKE-256).

16.4.3 Is there a "Fork Strategy" for Quantum? (If a quantum computer breaks the legacy chain, is there a plan to hard-fork to a new secure chain?).

16.4.4 Are "Wallet Addresses" hashed? (Ensuring the public key is not revealed until the funds are spent, adding a layer of protection).

16.4.5 Is "Zero-Knowledge Proof" (ZKP) technology upgraded? (Ensuring the zk-SNARKs used for privacy are also quantum-safe).

16.4.6 Is "Consensus" protected? (Ensuring the voting mechanism cannot be overwhelmed by a quantum computer solving proof-of-work instantly).

16.4.7 Are "Smart Contracts" audited for crypto-dependencies? (Finding contracts that rely on pre-compiled ECC functions).

16.4.8 Is "Address Reuse" banned? (Since revealing the public key makes it vulnerable to Shor's algorithm, are addresses used only once?).

16.4.9 Is "Commit-Reveal" logic used? (Hiding the transaction details until the block is finalized).

16.4.10 Does the Rosecoin Foundation maintain a "Quantum Bounty" (A massive prize for anyone who can break the testnet with a quantum computer)?

16.5.1 Is a "Y2Q Officer" appointed? (A specific leader responsible for the transition, similar to Y2K).

16.5.2 Is "Simulated Quantum Attack" testing performed? (Using mathematical models to see how quickly the current firewall would crumble).

16.5.3 Are "Legal Contracts" updated? (Requiring vendors to indemnify the company if their "Secure" product is broken by a quantum computer).

16.5.4 Is "Board Education" conducted? (Explaining that "Quantum" isn't sci-fi, it's a looming operational risk).

16.5.5 Are "Standard Bodies" monitored? (Tracking ISO/IEC JTC 1/SC 27 updates).

16.5.6 Is "Entropy as a Service" (EaaS) evaluated? (Subscribing to high-quality random numbers).

16.5.7 Are "Legacy Systems" ring-fenced? (Accepting that the Mainframe will never be quantum safe and burying it behind a PQC proxy).

16.5.8 Is "Code Signing" transitioned? (Ensuring software updates can be trusted in a post-quantum world).

16.5.9 Are "Backup Archives" re-encrypted? (Taking old tape backups, decrypting them, and re-encrypting them with Quantum-Safe algorithms).

16.5.10 Does the Rosecoin Ledger act as the "Time Capsule" of trust, verifying which data was secured before the quantum break occurred?

16.6.1 Is "Biometric Hashing" upgraded? (Ensuring the hash of the CEO's fingerprint is stored using a quantum-resistant algorithm, because you can't change your fingerprint if the hash is cracked).

16.6.2 Are "Genetic Data" stores encrypted with One-Time Pads or QKD? (Acknowledging that DNA data needs protection for 100+ years, far beyond the reach of standard encryption).

16.6.3 Is "Identity Proofing" resilient? (Ensuring that a quantum computer cannot forge the digital signature on a passport or ID card).

16.6.4 Are "Behavioral Biometrics" weighted higher? (Relying more on how you type—which is harder for a quantum machine to simulate—than static passwords).

16.6.5 Is "Liveness Detection" quantum-tested? (Ensuring Deepfake AIs powered by quantum computing cannot bypass FaceID).

16.6.6 Are "Verifiable Credentials" (W3C) signed with Dilithium/Falcon? (Ensuring the decentralized identity wallet is future-proof).

16.6.7 Is "Non-Repudiation" legally reviewed? (If a quantum computer cracks a signing key, how do you legally prove "I didn't sign that contract"?).

16.6.8 Are "Employee Badges" (Smart Cards) upgraded? (Replacing 10-year-old RSA-1024 chips in physical access cards).

16.6.9 Is "Anonymity" preserved? (Using Zero-Knowledge Proofs that don't rely on trusted setup phases vulnerable to quantum attack).

16.6.10 Does the Rosecoin Ledger allow for "Identity Migration"? (A protocol to seamlessly move a user's reputation to a new, quantum-safe address without losing history).

16.7.1 Is "Q-Day" defined in the Incident Response Plan? (The exact criteria for declaring: "Encryption is broken, initiate emergency protocol").

16.7.2 Is there a "Network Severance" plan? (The ability to disconnect the internal network from the public internet immediately to stop active decryption).

16.7.3 Is "Out-of-Band" command and control ready? (Using pre-shared keys and physical couriers to coordinate the IT team when VPNs are unsafe).

16.7.4 Are "Paper Backups" of critical secrets available? (Root CA keys printed and stored in physical safes as the ultimate fail-safe).

16.7.5 Is there a "Key Rollover" automation? (A "Big Red Button" that revokes every single SSL/SSH certificate in the company and issues new PQC ones in <1 hour).

16.7.6 Is "Legal Triage" prepared? (Drafted notifications to customers explaining that "All historical data may be compromised").

16.7.7 Are "Honeypot" quantum keys deployed? (Fake encrypted files that, if accessed/decrypted, alert you that an adversary possesses quantum capabilities).

16.7.8 Is "Vendor Liability" clear? (Who pays if the cloud provider's "Quantum Safe" storage wasn't actually safe?).

16.7.9 Is "Cash Flow" continuity planned? (How to pay employees if the SWIFT banking network is frozen due to quantum panic).

16.7.10 Does the organization have a "Re-Foundation" plan? (The steps to rebuild trust in the brand after the cryptographic apocalypse).

16.7.11 Cryptographic Inventory Visibility: Is there a centralized, automated inventory of all cryptographic algorithms (RSA, ECC, etc.) used across every internal application and API?

16.7.12 SNDL Risk Assessment: Has the organization identified "Store Now, Decrypt Later" (SNDL) data that requires immediate migration to post-quantum algorithms due to its long-term sensitivity?

16.7.13 SNDL (Store Now, Decrypt Later) Risk: Has the organization identified long-lived data (10+ year shelf life) that must be migrated to post-quantum encryption today to prevent future decryption?

16.7.14 Cryptographic Agility Audit: Can the system switch its primary encryption algorithms (e.g., moving from ECC to Kyber) via a single configuration change without re-writing core application code?

DOMAIN 17: AUTONOMOUS DEFENSE & SELF-HEALING SYSTEMS

17.1.1 Is "Event-Driven Architecture" (EDA) fully implemented? (Does a "Disk Full" alert automatically trigger a Lambda function to clear temp files and rotate logs without waking a human?).

17.1.2 Is "Auto-Scaling" configured for security events, not just traffic? (If a DDoS hits, does the system automatically spin up 50 extra WAF instances to absorb the load?).

17.1.3 Are "Self-Restarting Services" standard? (Using systemd/ Kubernetes liveness probes to kill and restart hung processes instantly).

17.1.4 Is "Automatic IP Blocking" enabled at the edge? (If an IP generates 10 failed logins, does the firewall ban it globally for 60 minutes automatically?).

17.1.5 Are "Ephemeral Credentials" rotated automatically upon usage detection? (If a "Break Glass" account is used, does the system automatically change the password immediately after logout?).

17.1.6 Is "Malware Isolation" automated? (If EDR detects malware, is the "Isolate Host" API called instantly by the SOAR platform?).

17.1.7 Are "Database Deadlocks" resolved autonomously? (Detecting hung queries and killing them to free up resources for healthy transactions).

17.1.8 Is "Certificate Renewal" fully automated? (Using ACME/Let's Encrypt protocols to ensure no human ever manually installs an SSL cert again).

17.1.9 Are "Stale User Accounts" disabled automatically? (If no login for 90 days, the account is locked via script).

17.1.10 Does the Rosecoin Ledger record the "Bot Action" to prove that the remediation was performed by an authorized script and not a hacker covering their tracks?

17.2.1 Is "GitOps" the single source of truth? (Ensuring that if someone manually changes a security group in the AWS Console, the system reverts it to match the Git repo within 5 minutes).

17.2.2 Is "Immutable Infrastructure" enforced? (If a server acts weird, you never SSH in to fix it; you terminate it and let the Auto-Scaling Group build a fresh one).

17.2.3 Are "Configuration Management" agents (Chef/Puppet/Ansible) running in "Enforce" mode, not just "Audit" mode?

17.2.4 Is "Golden Image" cycling automated? (Automatically rebuilding the base AMIs every week with the latest patches and rotating them into production).

17.2.5 Are "Unmanaged Resources" auto-tagged for deletion? (A "Reaper" script that finds untagged EC2 instances and terminates them after 24 hours).

17.2.6 Is "Drift Alerting" integrated with ticketing? (Creating a P1 ticket if the Production environment deviates more than 1% from the Staging environment).

17.2.7 Are "Policy-as-Code" violations auto-remediated? (If a user creates an unencrypted S3 bucket, the policy engine encrypts it immediately).

17.2.8 Is "DNS Healing" active? (If a primary endpoint fails health checks, does DNS automatically update to point to the secondary region?).

17.2.9 Are "Shadow Admin" rights automatically stripped? (Scanning for users who were added to "Domain Admins" and removing them if not in the approved Git config).

17.2.10 Does the system allow for "Manual Override" (Break Glass) during emergencies, with strict auditing?

17.3.1 Is "Chaos Monkey" (random instance termination) active in Non-Production? (Training teams to build stateless apps that survive server death).

17.3.2 Are "Latency Injection" tests performed? (Simulating a slow database to ensure the web app doesn't crash entirely).

17.3.3 Is "Region Failure" simulated annually? (Turning off all traffic to "US-East" to prove "US-West" can handle the load).

17.3.4 Are "Security Chaos" experiments run? (Randomly opening a firewall port to see if the Cloud Posture tool catches and closes it).

17.3.5 Is "Certificate Expiry" simulated? (Intentionally breaking a cert in Staging to ensure the alert pipeline works).

17.3.6 Are "Game Days" mandatory for all engineering squads? (Dedicated days to practice incident response against simulated failures).

17.3.7 Is "Dependency Failure" tested? (Blocking access to Google Maps API to ensure the rest of the app still functions).

17.3.8 Are "Clock Skew" tests performed? (Testing resilience against NTP drift).

17.3.9 Is "Packet Loss" simulation used to test protocol resilience?

17.3.10 Does the organization measure "Mean Time to Heal" (MTTH)? (How long does the system take to fix itself without human intervention?).

17.4.1 Is the "Circuit Breaker Pattern" implemented? (Stop calling a failing microservice to prevent cascading failures).

17.4.2 Is "Graceful Degradation" coded? (If the "Recommendations" engine dies, show "Popular Items" instead of a 500 Error).

17.4.3 Are "Automatic Rollbacks" configured? (If a deployment causes Error Rates to spike >1%, the pipeline reverts to the previous version instantly).

17.4.4 Is "Database Self-Healing" active? (Automatic failover to a Read Replica if the Writer node crashes).

17.4.5 Are "Message Queues" (Dead Letter Queues) used? (Storing failed transactions to be "replayed" later once the system recovers).

17.4.6 Is "Input Sanitization" auto-correcting? (Stripping bad characters from input rather than just crashing).

17.4.7 Are "Memory Leaks" managed? (Automatically restarting a worker process if RAM usage exceeds 90%).

17.4.8 Is "Cache Rehydration" automated? (If Redis crashes, is there a script to warm up the cache from the DB automatically?).

17.4.9 Are "Consistency Checks" backgrounded? (Continually checking if the Search Index matches the Database and fixing discrepancies).

17.4.10 Does AINA OS provide "Predictive Healing"? (Noticing a disk is filling up before it hits 100% and extending the volume automatically).

17.5.1 Is the "System State" hash recorded on the ledger every hour? (Creating a heartbeat of integrity).

17.5.2 Are "Heal Events" logged as transactions? ("Node A died, Node B replaced it, integrity verified").

17.5.3 Is "Configuration Integrity" verifiable? (Auditors can check the blockchain to see if the firewall rules have changed).

17.5.4 Are "SLA Credits" automated? (Smart contracts automatically issuing refunds to customers if the uptime hash is missed).

17.5.5 Is "Vendor Health" tracked on-chain? (Recording the uptime performance of third-party APIs).

17.5.6 Are "Patch Certificates" issued? (Proving that a specific CVE was patched at a specific block height).

17.5.7 Is "Access Recovery" decentralized? (Using social recovery wallets to restore admin access if all keys are lost).

17.5.8 Are "Immutable Logs" used for the healing scripts? (Ensuring an attacker cannot modify the script to "heal" the system by installing a backdoor).

17.5.9 Is "Consensus" required for major architecture changes? (Requiring 3/5 keys to approve a change to the Auto-Healing logic).

17.5.10 Does the organization publish a public "Transparency Report" derived directly from the ledger data?

DOMAIN 18: PEOPLE SECURITY & CULTURE

18.1.1 Is a "No-Blame" reporting culture explicitly codified? (Ensuring an employee who reports "I clicked a link" is thanked, not punished, to encourage visibility).

18.1.2 Are "Near Misses" celebrated? (Rewarding employees who almost fell for a scam but caught it at the last second and reported it).

18.1.3 Is "Psychological Safety" measured in surveys? (Asking "Do you feel safe reporting a security mistake?" and tracking the score).

18.1.4 Are "Sanctions" reserved for negligence/malice only? (Differentiating between "I made a mistake" and "I bypassed the control intentionally").

18.1.5 Is the "Security Team" viewed as a blocker or an enabler? (Measuring the Net Promoter Score (NPS) of the security department within the company).

18.1.6 Are "Friction Logs" collected? (Allowing users to report "This security tool is making my job impossible" so it can be fixed before they bypass it).

18.1.7 Is "Burnout" monitored as a security risk? (Recognizing that exhausted admins make fatal configuration errors).

18.1.8 Are "Whistleblower" channels anonymous and tested? (Ensuring a safe path to report a CISO hiding a breach).

18.1.9 Is "Cognitive Load" assessed? (Ensuring security procedures don't require superhuman memory or attention spans).

18.1.10 Does the Rosecoin Ledger record "Culture Metrics" (Reporting Rates vs. Click Rates) to prove the workforce is becoming more resilient?

18.2.1 Is there a formal "Security Champions" program? (Embedding a trained security advocate in every engineering squad and business unit).

18.2.2 Do Champions have "Allocated Time" (e.g., 20%)? (Ensuring it's not just "extra work" but a recognized part of their job description).

18.2.3 Are Champions "Rewarded"? (Bonuses, specialized training, or exclusive swag for taking on the responsibility).

18.2.4 Is there a "Guild" structure? (Monthly meetings where Champions across the company share threats and solutions).

18.2.5 Do Champions conduct "Local Threat Modeling"? (Reviewing the design of their team's features before they reach the central security team).

18.2.6 Is the "Ratio" sufficient? (Targeting 1 Champion per 10 Developers).

18.2.7 Are "Non-Technical" Champions included? (Having a Champion in HR and Finance to spot process risks).

18.2.8 Is there a "Champion Leaderboard"? (Gamifying the detection of bugs or completion of training).

18.2.9 Do Champions have "Early Access"? (Letting them see new security tools first to provide feedback).

18.2.10 Does the organization measure "Champion Impact"? (e.g., "Teams with a Champion have 50% fewer vulnerabilities").

18.3.1 Is training "Role-Based"? (Developers get Secure Coding, Finance gets Invoice Fraud, HR gets Privacy data handling).

18.3.2 Are "Phishing Simulations" varied? (Using SMS, QR codes, and voice deepfakes, not just emails).

18.3.3 Is "Micro-Learning" used? (2-minute videos triggered when a user makes a risky action, rather than 1-hour annual lectures).

18.3.4 Are "Escape Rooms" (Virtual/Physical) used for team building? (Solving security puzzles to "escape" a locked server room).

18.3.5 Is "Live Hacking" demonstrated? (Showing employees exactly how easy it is to crack a simple password to visceralize the risk).

18.3.6 Are "Personal Security" sessions offered? (Teaching employees how to secure their family's Wi-Fi and kids' iPads, which builds goodwill).

18.3.7 Is there a "Security Month" that is actually engaging? (Guest speakers, lock-picking villages, prizes).

18.3.8 Is "Onboarding" security rigorous? (Ensuring Day 1 employees know the rules before they get access).

18.3.9 Is "Remedial Training" automated? (If you click a phish, you are instantly enrolled in a 5-minute refresher).

18.3.10 Does the organization use "Nudge Theory"? (Subtle UI prompts like "Are you sure this recipient is correct?" rather than hard blocks).

18.4.1 Are "Background Checks" continuous? (Re-screening critical staff every year, not just at hiring).

18.4.2 Is "Social Media Vetting" performed for high-risk roles? (Checking for extremist views or vulnerability to blackmail).

18.4.3 Are "Technical Interviews" practical? (Asking security candidates to find a bug in code, not just answer multiple-choice questions).

18.4.4 Is "Neurodiversity" supported? (Recognizing that some of the best security talent may be autistic/ADHD and adjusting interview processes accordingly).

18.4.5 Is "Offboarding" instantaneous? (Can you revoke all access for a terminated employee in <60 seconds?).

18.4.6 Are "Alumni Networks" managed? (Staying on good terms with former admins so they don't become disgruntled threats).

18.4.7 Is "Job Rotation" practiced? (Moving security staff between teams to prevent silos and boredom).

18.4.8 Are "Golden Handcuffs" used for key retention? (Ensuring the CISO and Lead Architect are financially motivated to stay).

18.4.9 Is "Dual Control" required for termination of IT staff? (Ensuring a firing manager cannot trigger a logic bomb).

18.4.10 Does the organization track "Attrition Risk" in the security team? (High turnover in the SOC suggests a process/management failure).

18.5.1 Does the "CEO" participate in phishing tests? (And are their results published? "Even the CEO can be tricked").

18.5.2 Is "Security" a standing agenda item at every Board meeting? (Not just when there is a breach).

18.5.3 Does the Board have a "Cyber Expert" member? (Someone who can challenge the CISO's technical claims).

18.5.4 Are "Executive Simulations" conducted? (Tabletop exercises where the C-Suite has to make the "Pay/No-Pay" decision).

18.5.5 Is "Shadow IT" by Execs tolerated? (Is the CEO allowed to use WhatsApp for business? If so, is it secured?).

18.5.6 Are "VIP Protections" accepted? (Does the C-Suite accept the inconvenience of 2FA/YubiKeys?).

18.5.7 Is "Budget" viewed as investment or cost? (Does the Board understand that $1 in security saves $100 in breach costs?).

18.5.8 Are "Performance Goals" tied to security? (Does the CTO's bonus depend on reducing vulnerabilities?).

18.5.9 Is "Radical Transparency" practiced? (Does the CISO report the "Red" status honestly without fear of being fired?).

18.5.10 Does the Rosecoin Ledger record the "Board Minutes" regarding cyber risk acceptance to prove fiduciary duty?

18.6.1 Is "Influence Operation" training mandatory? (Teaching employees how to spot foreign state propaganda designed to cause internal discord/mutiny).

18.6.2 Are "Truth Anchors" established? (A verified internal news channel that employees trust as the "Source of Truth" during a deepfake crisis).

18.6.3 Is "Critical Thinking" drilled? (Teaching staff how to verify a claim, not just what to believe).

18.6.4 Are "Crisis Simulations" psychological? (Testing if the team panics or turns on each other under high pressure).

18.6.5 Is "Algorithmic Hygiene" taught? (Showing employees how their own social media feeds are manipulated to radicalize them).

18.6.6 Are "Deepfake Drills" personalized? (Simulating a deepfake video of the employee themselves to show how easy it is to fabricate reality).

18.6.7 Is "Reputation Defense" active? (Supporting employees who are doxxed or targeted by harassment campaigns online).

18.6.8 Are "Emotional Vectors" monitored? (Recognizing that "Urgency" and "Fear" are the primary tools of a hacker, and training staff to pause when they feel them).

18.6.9 Is "Source Validation" a cultural norm? (It is considered polite, not rude, to hang up and call back to verify identity).

18.6.10 Does the organization track "Disinformation Resilience"? (Surveying staff to see if they believe known false narratives targeting the company).

18.7.1 Is "Neuro-Inclusive Design" applied to security policies? (Replacing 50-page PDFs with checklists and videos, knowing that ADHD brains struggle with dense text).

18.7.2 Are "Pattern Matchers" (Autistic talent) specifically recruited for Threat Hunting and Log Analysis roles?

18.7.3 Is "Accessibility" (WCAG) enforced for security tools? (Can a blind admin use the IAM portal with a screen reader?).

18.7.4 Are "Social Anxiety" factors considered? (Allowing employees to report incidents via Chat/Form rather than forcing a phone call).

18.7.5 Is "Quiet Work" respected? (Ensuring security alerts don't interrupt "Deep Work" states unnecessarily, causing cognitive friction).

18.7.6 Are "Interview Accommodations" standard? (Allowing security candidates to show skills via a practical test rather than a high-pressure whiteboard interview).

18.7.7 Is " Sensory Overload" minimized in the SOC? (Dimmable lights, noise cancellation, and dark mode tools).

18.7.8 Are "Communication Styles" adapted? (Training managers that "directness" in security feedback is not "rudeness").

18.7.9 Is "Burnout Prevention" personalized? (Recognizing that different neurotypes recover from stress differently).

18.7.10 Does the organization view "Neurodiversity" as a competitive defense advantage? (Leveraging different thinking styles to spot complex attacks).

18.8.1 Are "Security Bounties" paid to employees? (Instant micro-bonuses—e.g., $50 in Rosecoin—for reporting a real phishing email).

18.8.2 Is "Compliance Gamified"? (Earning status or tokens for patching a laptop within 24 hours).

18.8.3 Are "Champion Rewards" tangible? (Paying Security Champions a stipend for their extra responsibility).

18.8.4 Is "Bug Finding" incentivized internally? (Paying developers who find bugs in other teams' code).

18.8.5 Are "Security KPIs" linked to Executive Bonuses? (The CTO loses 10% of their bonus if the phishing fail rate is >5%).

18.8.6 Is "Reputation Staking" used? (Admins "stake" their reputation score to make high-risk changes; if they break it, the score drops).

18.8.7 Are "Training Completions" rewarded instantly? (Finish the video, get a free coffee voucher automatically).

18.8.8 Is "Negative Reinforcement" avoided? (Punishing mistakes drives them underground; rewarding reporting brings them to light).

18.8.9 Are "Innovation Grants" available? (Funding employee ideas for better security tools).

18.8.10 Does the Rosecoin Ledger serve as the "Immutable Resume"? (Employees leave with a cryptographic record of their high security scores, increasing their market value).

DOMAIN 19: CONTINUOUS IMPROVEMENT & MATURITY

19.1.1 Is the organization scored against the 5 Levels of Maturity? (1. Initial/Chaos, 2. Repeatable, 3. Defined, 4. Managed, 5. Optimizing/ Rosecoin).

19.1.2 Is "Regression" monitored? (Detecting if a domain that was "Level 4" last year slipped back to "Level 3" due to staff turnover).

19.1.3 Are "Maturity Goals" linked to budget? (e.g., "To move Identity from Level 2 to Level 3, we need $X for a PAM solution").

19.1.4 Is there a "Peer Benchmarking" process? (Comparing maturity scores anonymously with other industry leaders via the ISAC).

19.1.5 Is "Debt" quantified against maturity? (Calculating the cost of remaining at Level 2 versus the risk exposure).

19.1.6 Are "Third-Party" maturity scores required? (Refusing to integrate with vendors who are below Level 3).

19.1.7 Is "Automation" the gatekeeper to Level 4? (You cannot be "Managed" if you rely on spreadsheets).

19.1.8 Is "AI Integration" the gatekeeper to Level 5? (You cannot be "Optimizing" without predictive AI defenses).

19.1.9 Are "Maturity Audits" independent? (Using an external firm to validate the score every 2 years).

19.1.10 Does the Rosecoin Ledger serve as the "Maturity Passport," publicly proving the organization's rating to customers?

19.2.1 Is there a "Secure Sandbox" for dangerous testing? (A completely isolated network where engineers can detonate malware or test unstable AI without risk).

19.2.2 Is "Shadow Innovation" brought into the light? (Providing a fast-track approval process for devs who want to try new tools, so they don't do it secretly).

19.2.3 Are "Guardrails" automated in the sandbox? (e.g., The sandbox has no internet access, or data cannot be copied out).

19.2.4 Is "Fail Fast" culture supported securely? (Allowing experiments to fail without generating permanent security debt).

19.2.5 Are "New Tech" assessments rapid? (A 48-hour SLA to review a new library or tool, so security doesn't slow down innovation).

19.2.6 Is "Data Synthesis" used? (Generating fake, realistic data for testing so real customer data never touches the lab).

19.2.7 Are "Ethical Reviews" mandatory for R&D? (Ensuring the new cool feature doesn't violate privacy rights or bias laws).

19.2.8 Is "IP Protection" heightened in the lab? (Recognizing that the R&D lab contains the company's future secrets).

19.2.9 Are "Export Controls" managed? (Ensuring code written in the lab doesn't violate ITAR/EAR regulations).

19.2.10 Does the organization have a "Transition Protocol"? (A strict checklist to move a project from "Sandbox" to "Production" hardening).

19.3.1 Is "Horizon Scanning" a formal role? (Someone dedicated to looking 5-10 years out at threats like 6G, Neural Links, and Nanotech).

19.3.2 Is "Brain-Computer Interface" (BCI) security researched? (Preparing for the day employees connect their minds directly to the corporate network).

19.3.3 Are "Neuro-Privacy" rights defined? (Protecting the "Thought Data" of users if BCI becomes a reality).

19.3.4 Is "6G / Terahertz" security monitored? (Preparing for hyper-local positioning and massive IoT density).

19.3.5 Are "Autonomous Swarms" modeled? (Defending against attacks by thousands of coordinated AI drones).

19.3.6 Is "Synthetic Biology" risk assessed? (Securing the "DNA Printers" that could theoretically print pathogens).

19.3.7 Is "Generative Reality" (Metaverse) defense planned? (Protecting against infinite spoofing in virtual worlds).

19.3.8 Are "Holographic" interfaces secured? (preventing "Gesture Jacking" or visual eavesdropping).

19.3.9 Is "Energy Harvesting" security considered? (IoT devices that run forever on ambient radio waves, making them unkillable).

19.3.10 Does the organization sponsor "Academic Research" to stay ahead of the commercial market?

19.4.1 Is the "1% Rule" applied? (Mandating a 1% improvement in a specific security metric every sprint).

19.4.2 Are "Post-Mortems" blameless and public? (Sharing lessons learned with the entire company, not just the security team).

19.4.3 Is "Tool Rationalization" performed annually? (Killing old security tools that overlap or don't provide value).

19.4.4 Are "Process Mining" tools used? (Analyzing logs to find inefficient security workflows that slow down users).

19.4.5 Is "Feedback Looping" automated? (If a user marks an email as "Safe," does that instantly train the ML filter?).

19.4.6 Are "Hackathons" used to fix debt? (dedicating 2 days a quarter to strictly fixing security bugs).

19.4.7 Is "Standard Updating" automatic? (If NIST releases a new version, is the internal policy updated via script?).

19.4.8 Are "Efficiency Metrics" tracked? (Cost per Ticket, Time to Patch, False Positive Rate).

19.4.9 Is "External Review" rotational? (Using a different audit firm every 3 years to get fresh eyes).

19.4.10 Does the organization view the "Rosecoin Framework" as a living document, updating it daily based on new threats?

19.5.1 Does the organization "Contribute" to open source security tools? (Giving back to the community).

19.5.2 Are "Whitepapers" published? (Establishing the company as a thought leader in security).

19.5.3 Is there a "Bug Bounty" for the standard itself? (Paying researchers who find flaws in the Rosecoin Framework).

19.5.4 Are "University Partnerships" active? (Pipeline for new talent and fresh research).

19.5.5 Does the organization sit on "Standards Boards"? (Influencing the future of ISO/NIST).

19.5.6 Is "Mentorship" mandatory? (Senior security staff must mentor juniors).

19.5.7 Are "Community Events" hosted? (Running local BSides or OWASP chapters).

19.5.8 Is "Threat Intel" shared freely? (Believing that "A threat to one is a threat to all").

19.5.9 Are "Patents" filed for defensive innovations? (Protecting the company's unique security inventions).

19.5.10 Does the Rosecoin Foundation verify the "Community Impact" score of the organization?

19.5.11 Compliance-as-a-Service Automation: Are security controls continuously monitored to provide a "real-time compliance dashboard," reducing the administrative burden of periodic annual audits?

DOMAIN 20: EVIDENCE, LEGAL HOLD & PROVENANCE (ROSECOIN VAULT)

20.1.1 Is "Parametric Insurance" utilized? (Smart contracts that automatically pay out $10M if the Cloud Provider's uptime drops below 99.9% for >4 hours, eliminating the need for claims adjusters).

20.1.2 Is "Coverage Mapping" performed quarterly? (Ensuring the policy covers Ransomware Fines, Business Interruption, and PR Costs, not just hardware replacement).

20.1.3 Is "War Exclusion" analysis performed? (Most policies void coverage if the attack is an "Act of War." Does the policy define "Cyber War" clearly?).

20.1.4 Is "Captive Insurance" considered? (Creating a self-insurance fund for risks that the commercial market refuses to cover).

20.1.5 Are "Warranty" clauses in security tools enforced? (If the Firewall fails to block a known threat, does the vendor pay a warranty fee?).

20.1.6 Is "Silent Cyber" risk eliminated? (Ensuring Property & Casualty policies don't accidentally exclude cyber-induced fires or explosions).

20.1.7 Is "Legal Privilege" maximized? (Ensuring the Incident Response Retainer is signed by Outside Counsel, not the CISO, to protect reports from discovery).

20.1.8 Is "Subrogation" waived? (Ensuring the insurer can't sue your own employees for negligence after paying a claim).

20.1.9 Is "Proof of Diligence" automated via Rosecoin? (Providing the insurer with a read-only view of the "Patching Ledger" to lower premiums by 20%).

20.1.10 Does the organization have a "Bitcoin Reserve" strategy? (A legally compliant, Board-approved method to access crypto for ransom payments if insurance refuses/delays).

20.2.1 Are "Smart Contracts" treated as legal contracts? (Do the Terms of Service explicitly state that code execution on the blockchain constitutes a binding agreement?).

20.2.2 Is "Ricardian Contract" linkage used? (Every Smart Contract has a hash link to a PDF legal document explaining the intent in English).

20.2.3 Is "Automated Compliance" codified? (The code itself prevents a GDPR violation—e.g., it is mathematically impossible to query the database for "Race" or "Religion").

20.2.4 Are "Liability Caps" updated for AI? (Who is liable if the AI security bot accidentally shuts down the factory? The Vendor or the User?).

20.2.5 Is "Code-is-Law" dispute resolution defined? (If the code allows a hack (e.g., a flash loan attack), is it "illegal theft" or "clever use of the rules"?).

20.2.6 Are "Digital Signatures" eIDAS compliant? (Ensuring signatures are legally binding in the EU and US).

20.2.7 Is "Open Source" licensing audited? (Ensuring no "Copyleft" (GPL) code infects the proprietary security stack).

20.2.8 Are "Algorithmic Bias" audits legally privileged? (Conducting bias tests under legal supervision to avoid creating evidence for a lawsuit).

20.2.9 Is "Data Ownership" defined in metadata? (Every file has a metadata tag stating "Property of Corp X," legally establishing theft if exfiltrated).

20.2.10 Does the Rosecoin Ledger serve as the "Notary Public"? (Replacing the need for human notaries by timestamping documents on-chain).

20.3.1 Is "Data Residency" strictly enforced? (German data stays on German servers; Chinese data stays in China. No exceptions).

20.3.2 Is "Cross-Border Transfer" automated? (The system automatically blocks a file transfer from Paris to New York if no "Standard Contractual Clause" (SCC) exists).

20.3.3 Is "Sanctions Screening" real-time? (Blocking IP addresses from North Korea, Iran, or Russia instantly to avoid OFAC fines).

20.3.4 Is "Extraterritorial" risk assessed? (Understanding that the US CLOUD Act allows the US gov to seize data stored in Europe by US companies).

20.3.5 Is "Sovereign Cloud" utilized? (Using a cloud provider owned and operated entirely by locals to prevent foreign subpoena power).

20.3.6 Are "Encryption Keys" held in-country? (Storing the encrypted data in the cloud, but the keys in a physical HSM in the HQ country).

20.3.7 Is "Employee Nationality" considered for access? (Restricting access to ITAR data to citizens only).

20.3.8 Is "Geopolitical Monitoring" active? (Alerting if a new law in India requires 6-hour incident reporting).

20.3.9 Is "Splinternet" preparation active? (Can the Chinese branch operate independently if the "Great Firewall" cuts it off from the Global WAN?).

20.3.10 Does the organization have a "Human Rights" policy for surveillance? (Refusing to sell or use security tools that enable oppressive regimes).

20.4.1 Is the "Audit Trail" continuous? (Certification is not a "Point in Time" PDF, but a live stream of green lights).

20.4.2 Is the "Rosecoin Score" public? (Displaying the 0-1000 security score on the website footer, like a BBB rating).

20.4.3 Is "Recertification" automated? (If the score drops below 800, the "Certified" badge is automatically revoked via smart contract).

20.4.4 Is the "Auditor" rotated algorithmically? (The system randomly selects a certified auditor to spot-check a domain).

20.4.5 Are "Zero-Knowledge Proofs" used for compliance? (Proving to the regulator "We have backups" without showing them the actual data).

20.4.6 Is "Bug Bounty" participation mandatory for certification?

20.4.7 Is "Executive Liability" attached? (The CISO signs the certification with their own digital key, accepting personal reputational risk).

20.4.8 Is "Community Governance" active? (The organization votes on updates to the Rocheston Standard).

20.4.9 Is the "Carbon Footprint" of the security stack measured? (Green Security certification).

DOMAIN 21: AI AGENT GOVERNANCE & RUNTIME CONTROLS

21.1.1 Is "Agent Registration" mandatory? (Every autonomous script must be registered in a central "Digital HR" database before it runs a single line of code).

21.1.2 Do Agents have "Cryptographic Passports"? (Using mTLS certificates or DID - Decentralized Identifiers - to prove "I am the Finance Bot v2.1" when talking to APIs).

21.1.3 Is "Role-Based Agent Access" (RBAC) enforced? (The "Calendar Agent" has permission to read schedules but zero network access to the "Payroll Database").

21.1.4 Are "Creator Signatures" required? (Every running agent must be cryptographically signed by the human developer who deployed it).

21.1.5 Is "Version Control" strict? (If the "Customer Service Agent" starts hallucinating, can you rollback to v4.0 instantly across 500 instances?).

21.1.6 Are "Shadow Agents" hunted? (Scanning the network for unauthorized Python scripts or LangChain loops running on developer laptops).

21.1.7 Is "Least Privilege" applied to Tools? (The agent can access the "Email Tool" but cannot access the "Delete Email" function).

21.1.8 Are "Impersonation Checks" active? (Ensuring an agent cannot declare itself as a human user in logs).

21.1.9 Is "Agent Lifespan" defined? (Does the agent have a "Time-to-Live" token, ensuring it dies automatically after 24 hours if not renewed?).

21.1.10 Does the Rosecoin Ledger record the "Birth Certificate" (Hash) of every authorized agent to prevent unauthorized forks?

21.2.1 Is "Budgeting" hard-coded? (The Procurement Agent can spend up to $500 automatically; $501 triggers a Human-in-the-Loop request).

21.2.2 Is there a "Global Kill Switch"? (A single API call that freezes all agent activity instantly in case of a runaway feedback loop).

21.2.3 Are "Rate Limits" semantic? (Not just "100 requests/minute," but "Max 5 emails to External Domains per hour").

21.2.4 Is "Human-in-the-Loop" (HITL) mandatory for critical actions? (An agent can draft code, but a human must commit it).

21.2.5 Are "Tool Whitelists" dynamic? (If the "Weather API" is compromised, can you instantly revoke the agent's ability to call it?).

21.2.6 Is "Recurisve Loop" detection active? (Stopping an agent from calling itself infinitely and consuming all cloud credits).

21.2.7 Are "Approval Tokens" used? (The agent needs a fresh token from a human manager for every "High Impact" action).

21.2.8 Is "Context Window" flushing enforced? (Clearing the agent's short-term memory between tasks to prevent data leakage between customers).

21.2.9 Are "Suicide Protocols" coded? (If the agent detects it is being manipulated/jailbroken, does it terminate itself?).

21.2.10 Does the system enforce "The Three Laws" logic? (Hard-coded constraints that override any learned behavior—e.g., "Never Export Private Keys").

21.3.1 Is "Chain of Thought" (CoT) logging enabled? (Logging not just the output, but the reasoning steps: "I am deleting this file because X").

21.3.2 Are "Tool Inputs/Outputs" captured? (Knowing exactly what the agent sent to the SQL database and what it got back).

21.3.3 Is "Hallucination Detection" automated? (Using a secondary "Supervisor AI" to grade the agent's output before it is sent to the user).

21.3.4 Are "Sentiment Monitors" watching agent interactions? (Alerting if the Customer Support Agent starts becoming aggressive or rude).

21.3.5 Is "Goal Drift" monitored? (Detecting if an agent assigned to "Optimize Storage" starts trying to "Optimize Network Traffic" instead).

21.3.6 Are "Black Box" recorders installed? (Immutable logs of the agent's state tailored for forensic reconstruction after an incident).

21.3.7 Is "Prompt Injection" defense logged? (Recording every attempt by a user to trick the agent, to train better defenses).

21.3.8 Are "Cost Metrics" real-time? (Alerting if an agent burns $1,000 in tokens in 10 minutes).

21.3.9 Is "Bias Auditing" continuous? (Checking if the Hiring Agent is rejecting resumes based on specific keywords).

21.3.10 Does the Rosecoin Ledger store the "Decision Tree" hash for high-stakes decisions (e.g., loan denial) for legal auditability?

21.4.1 Is there a "Standard Protocol" for agent-to-agent talk? (Defining how the "Sales Agent" talks to the "Legal Agent" without ambiguity).

21.4.2 Are "Trust Scores" verified? (Agent A checks Agent B's Rosecoin Reputation Score before sharing data).

21.4.3 Is "Negotiation Bounding" active? (Preventing two agents from getting into an infinite price war loop).

21.4.4 Are "External Agent" firewalls in place? (Blocking unverified third-party agents from querying your internal agents).

21.4.5 Is "Economic Security" modeled? (Ensuring a rogue agent cannot drain the corporate wallet via micro-transactions).

21.4.6 Are "Deadlocks" resolved? (If Agent A waits for B, and B waits for A, does a Supervisor break the loop?).

21.4.7 Is "Data Lineage" preserved? (When Agent C uses data from Agent A, is the "Source: Agent A" tag preserved?).

21.4.8 Are "Contracts" automated? (Agents sign digital agreements for service delivery levels).

21.4.9 Is "Swarm Defense" planned? (How to stop a DDoS attack orchestrated by a botnet of hostile agents).

21.4.10 Does the organization operate an "Agent Sandbox" for testing multi-agent collaboration before production deployment?

21.5.1 Is "Vicarious Liability" accepted? (Explicit policy stating the Company is responsible for its Agent's actions, just like a human employee).

21.5.2 Are "Agent Wills" created? (Defining what happens to an agent's data and "memory" when it is decommissioned).

21.5.3 Is "Memory Wiping" verified? (Ensuring a retired agent doesn't leave sensitive context in a vector database).

21.5.4 Are "Legacy Agents" tracked? (Identifying agents running on old models (e.g., GPT-4) and forcing upgrades).

21.5.5 Is "Public Disclosure" automated? (Agents must identify themselves: "I am an AI Agent," not "I am Sarah").

21.5.6 Are "Ethics Boards" involved in agent creation? (Approving the "personality" and "goals" of the agent).

21.5.7 Is "Insurance" updated for Agentic Risk? (Does the policy cover "Agent Negligence"?).

21.5.8 Are "Termination Codes" secured? (Ensuring an attacker cannot send a "Shutdown" command to your defensive agents).

21.5.9 Is "Model Collapse" prevention active? (Ensuring agents don't train on their own output loop).

21.5.10 Does the Rosecoin Ledger record the "Death Certificate" of an agent, proving it can no longer act on behalf of the company?

21.5.11 Agentic Boundary Controls: Are there explicit, hard-coded boundaries restricting what an AI agent can access (e.g., prohibiting an agent from modifying tax elections or e-filing without human review)?

21.5.12 Semantic Rate Limiting: Does the system enforce semantic rate limits on AI agents to prevent automated reconnaissance or bulk data exfiltration?

21.5.13 Shadow Agent Discovery: Is there a dedicated monitoring tool to detect unauthorized AI agents (Shadow Agents) running on local developer environments or in the cloud?

21.5.14 Agentic Decision Logging: Does the system maintain an immutable "Decision Ledger" that records the contextual prompts and reasoning used by an AI agent for any high-risk autonomous action?

21.5.15 Agentic Collusion Monitoring: Does the runtime environment detect and alert if two autonomous AI agents begin unscripted communication that deviates from their authorized operational parameters?

21.5.16 KYA (Know Your Agent) Identity: Is every autonomous agent assigned a unique Machine Identity that is required to sign every API call it executes?

21.5.17 Agent Autonomy Boundaries: Are there hard-coded "bounds" or design-stage choice limits placed on AI agents to restrict their access to specific tools and external systems based on their risk level?

21.5.18 Traceable Agent Decisions: Are all autonomous actions taken by AI agents fully traceable and controllable through robust identity management specifically designed for machine-agent entities?

DOMAIN 22: SPACE & ORBITAL SECURITY

22.1.1 Is the "Ground Station" treated as a Tier-1 Critical Infrastructure site? (Biometric access, Faraday cage protections, and air-gapped mission control networks).

22.1.2 Is "Command Authentication" enforced via HSM? (Ensuring that a command to fire thrusters must be signed by a hardware key held by a Flight Director, preventing remote hijack).

22.1.3 Is "Uplink Encryption" mandatory? (Encrypting the command stream so an attacker cannot replay old commands or inject new ones).

22.1.4 Are "Software-Defined Radios" (SDRs) hardened? (Patching the firmware of the radio antennas themselves to prevent buffer overflows from malicious RF signals).

22.1.5 Is "Cloud Ground Station" (AWS Ground Station / Azure Orbital) configuration audited? (Applying the same rigor to "Satellite-as-a-Service" APIs as on-prem hardware).

22.1.6 Are "Telecommand Logs" immutable? (Recording every instruction sent to orbit to forensic standards).

22.1.7 Is "RF Jamming" detection active at the site? (Monitoring the spectrum for noise designed to blind the downlink).

22.1.8 Are "Network Gaps" enforced? (Mission Control is never connected to the Corporate Wi-Fi).

22.1.9 Is "Insider Threat" screening heightened for Flight Controllers? (They have the power to de-orbit a billion-dollar asset).

22.1.10 Does the Rosecoin Ledger record the "Command Sequence Hash" to prove exactly who authorized a maneuver?

22.2.1 Is "Radiation Hardening" applied to crypto-chips? (Ensuring cosmic rays don't flip a bit in the encryption key, causing a permanent lockout).

22.2.2 Is "Safe Mode" logic autonomous? (If the satellite detects an intrusion or anomaly, does it automatically shut down non-essential buses and point antennas to Earth?).

22.2.3 Are "Debug Ports" physically disabled? (Ensuring JTAG/Test ports are fused off before launch so they can't be used if the satellite is captured or probed).

22.2.4 Is "On-Board Intrusion Detection" (IDS) active? (A lightweight AI model running on the satellite CPU to detect anomalous process behavior).

22.2.5 Are "Firmware Updates" signed and rolled back? (If a bad update is uploaded, the satellite automatically reverts to the "Golden Image" after a failed boot).

22.2.6 Is "Propulsion Isolation" enforced? (The camera payload should never have network access to the thruster controls).

22.2.7 Are "Battery Management" limits hard-coded? (Preventing an attacker from overcharging the batteries to cause a thermal explosion).

22.2.8 Is "Memory Scrubbing" continuous? (Correcting single-event upsets (SEUs) in RAM caused by radiation).

22.2.9 Are "De-Orbit Codes" split-knowledge? (Requiring two separate keys from two different ground stations to initiate a burn-up).

22.2.10 Does the satellite broadcast a "Digital Identity" beacon to prove it is not a spoofer?

22.3.1 Is "Spread Spectrum" (FHSS) utilized? (Hopping frequencies thousands of times a second to evade jamming).

22.3.2 Is "Downlink Encryption" active? (Encrypting the photos/data coming down so eavesdroppers can't steal the imagery).

22.3.3 Are "Optical Links" (Laser Inter-Satellite Links) used for high-security traffic? (Lasers are near-impossible to tap without blocking the beam).

22.3.4 Is "GPS Spoofing" detection active on-board? (Ensuring the satellite knows its true orbital position).

22.3.5 Are "Bent Pipe" risks mitigated? (If the satellite is a "dumb repeater," is end-to-end encryption enforced on the user terminals?).

22.3.6 Is "Traffic Padding" used? (Sending dummy data when idle so an attacker cannot analyze traffic patterns to guess mission activity).

22.3.7 Are "Keep-Alive" heartbeats monitored for latency jitter? (Detecting Man-in-the-Middle attacks).

22.3.8 Is "Mesh Network" authentication strict? (Satellite A verifies Satellite B's certificate before routing traffic through it).

22.3.9 Are "Space Weather" feeds integrated? (Distinguishing between a solar flare outage and a jamming attack).

22.3.10 Does the organization use "Quantum Key Distribution" (QKD) via satellite for key exchange?

22.4.1 Is the "Launch Provider" (e.g., SpaceX, Rocket Lab) audited? (Ensuring the rocket's telemetry interface doesn't bridge into the satellite payload).

22.4.2 Are "Component Sources" traced? (Did the star tracker come from a sanctioned entity?).

22.4.3 Is "Payload Integration" supervised? (Physical guards watching the mating of the satellite to the rocket to prevent tampering).

22.4.4 Are "CubeSat" risks assessed? (Treating small, cheap satellites as "Untrusted" due to lower security standards).

22.4.5 Is "Launch Site" cyber hygiene verified? (Ensuring the launch pad network is clean).

22.4.6 Are "Test Data" leaks prevented? (Ensuring telemetry from ground testing isn't exposed on public buckets).

22.4.7 Is "End-of-Life" planning enforced? (Ensuring sufficient fuel remains to de-orbit safely, complying with space debris laws).

22.4.8 Are "Rideshare" neighbors vetted? (If sharing a rocket with a competitor/adversary, is the separation guaranteed?).

22.4.9 Is "Space Insurance" specific to cyber? (Covering "Loss of Control" due to hacking, not just explosion).

22.4.10 Does the Rosecoin Ledger track the physical chain of custody from the factory floor to the fairing encapsulation?

22.5.1 Is "Conjunction Assessment" (Collision Avoidance) automated? (Thrusters fire automatically if debris probability > 1/10,000).

22.5.2 Are "Kessler Syndrome" scenarios modeled? (What if a debris cloud destroys 50% of the constellation?).

22.5.3 Is "Grappling" defense considered? (What if a hostile satellite physically latches onto yours?).

22.5.4 Are "Laser Blinding" sensors active? (Detecting if a ground laser is dazzling the optical sensors).

22.5.5 Is "Orbital Diversity" practiced? (Spreading assets across different planes/altitudes to avoid single-point failure).

22.5.6 Are "Silent Spares" deployed? (Satellites that remain powered down and dark until needed to replace a loss).

22.5.7 Is "Maneuver Logic" verified? (Ensuring a hack cannot cause a re-entry over a populated city).

22.5.8 Are "Proximity Alerts" active? (Detecting if an "Inspector Satellite" is shadowing your asset).

22.5.9 Is "SSA" (Space Situational Awareness) data consumed? (Using US Space Force or commercial data to track threats).

22.5.10 Does the organization have a "Space Traffic Management" policy aligned with international treaties?

DOMAIN 23: SUSTAINABLE (GREEN) CYBERSECURITY

23.1.1 Is "Adaptive Scanning" implemented? (Ensuring antivirus/vulnerability scanners do not re-scan unchanged static files every hour, wasting CPU and electricity).

23.1.2 Is "Zombie Infrastructure" hunting active? (identifying and terminating orphaned cloud instances that are running security agents but serving no business purpose).

23.1.3 Are "Data Retention" policies aligned with sustainability? (Not storing 10 years of PCAP logs in "Hot Storage" when "Cold/Tape" storage consumes significantly less energy).

23.1.4 Is "Deduplication" enforced at the source? (Preventing the transmission of duplicate log data across the network to save bandwidth and storage energy).

23.1.5 Are "Low-Power States" utilized? (Do security appliances throttle down during off-peak hours, or do they run at 100% fan speed 24/7?).

23.1.6 Is "Compute Location" optimized? (Moving non-urgent cryptographic workloads to data centers powered by renewable energy/hydro).

23.1.7 Are "Lightweight Agents" prioritized? (Selecting EDR agents that use <1% CPU over bloated legacy agents that drain laptop batteries).

23.1.8 Is "Dark Data" eliminated? (Purging unclassified, unneeded data—"ROT" (Redundant, Obsolete, Trivial)—to reduce the storage energy burden).

23.1.9 Is "Algorithm Efficiency" evaluated? (Choosing efficient cryptographic implementations that achieve the same security level with fewer CPU cycles).

23.1.10 Does the Rosecoin Ledger track the "Carbon Cost per Incident" to visualize the environmental impact of cyber defense?

23.2.1 Is "Proof-of-Stake" (PoS) or "Proof-of-Authority" (PoA) mandatory? (Banning Proof-of-Work (mining) ledgers due to their massive energy consumption).

23.2.2 Is "Transaction Batching" used? (Rolling up 1,000 logs into a single blockchain transaction to reduce the network load).

23.2.3 Is "Layer 2" scaling utilized? (Performing high-volume security validations on a low-energy sidechain, only settling on the main chain periodically).

23.2.4 Is "Smart Contract" code optimized for gas/energy? (Auditing code to remove inefficient loops that waste computational resources).

23.2.5 Are "Green Nodes" prioritized? (Incentivizing validator nodes that run on solar or wind power).

23.2.6 Is "Storage Pruning" active? (Allowing nodes to discard ancient history that is no longer needed for current security validation).

23.2.7 Is "Hardware Reuse" supported for nodes? (Designing the ledger software to run on older hardware to extend its lifecycle).

23.2.8 Is "Carbon Offsetting" automated? (Does the blockchain protocol automatically purchase carbon credits based on its transaction volume?).

23.2.9 Are "Energy Audits" public? (Publishing the real-time energy consumption of the Rosecoin network).

23.2.10 Does the organization refuse to interact with "Dirty Chains" (marketplaces/ledgers with high carbon intensity)?

23.3.1 Is "Crypto-Erase" used over physical destruction? (Securely wiping drives by destroying the encryption key, allowing the hardware to be resold/reused instead of shredded).

23.3.2 Is "Repairability" a procurement criteria? (Buying laptops and servers that can be repaired, extending their life and reducing manufacturing demand).

23.3.3 Is "Device Longevity" supported by security? (Supporting OS patches on 5-year-old devices so they don't have to be discarded for "security reasons").

23.3.4 Are "Recycling Partners" R2/e-Stewards certified? (Ensuring e-waste is not illegally dumped in developing nations).

23.3.5 Is "Virtualization" maximized? (Replacing physical firewalls/HSMs with virtual instances to reduce plastic and metal waste).

23.3.6 Is "Modular Security" used? (Upgrading just the TPM chip or security module rather than replacing the whole motherboard).

23.3.7 Is "Battery Health" managed via software? (Security agents shouldn't cause excessive battery cycles, which degrades lithium-ion cells prematurely).

23.3.8 Is "Packaging Waste" minimized? (Demanding vendors ship security appliances without excessive plastic/foam).

23.3.9 Is "Material Recovery" tracked? (Recovering gold/rare earths from decommissioned SOC hardware).

23.3.10 Does the organization have a "Take-Back" program for employee devices to ensure secure and green disposal?

23.4.1 Is "Green Coding" taught? (Training developers that efficient code is also secure code—fewer cycles, fewer bugs).

23.4.2 Are "SaaS Vendors" graded on sustainability? (Asking AWS/Azure/Salesforce for their "Carbon Emissions Report" for your specific tenant).

23.4.3 Is "Network Traversal" minimized? (Hosting security gateways close to the user to reduce the energy cost of hauling data across the ocean).

23.4.4 Are "Dark Mode" interfaces standard? (Using dark pixels on OLED screens to save energy in the SOC).

23.4.5 Is "Feature Bloat" rejected? (Refusing to install security suites with features that will never be used).

23.4.6 Are "CI/CD Pipelines" optimized? (Not running the full security test suite on every minor commit, but using intelligent selection).

23.4.7 Is "AI Model Training" energy-aware? (Training security AI models once and "fine-tuning" them, rather than retraining from scratch weekly).

23.4.8 Is "Digital Sobriety" practiced? (Questioning "Do we really need to log this?" before enabling a new data stream).

23.4.9 Is "Remote Work" supported as a green initiative? (Reducing the commuting carbon footprint of the security team).

23.4.10 Does the Rosecoin Ledger issue "Green Certificates" to vendors who meet energy-efficiency targets?

DOMAIN 24: NEURO-COGNITIVE SECURITY & HUMAN FACTORS

24.1.1 Is "Neural Firmware" signed? (Ensuring the OS running on the implanted/wearable BCI cannot be updated with a malicious rootkit that creates a "backdoor to the brain").

24.1.2 Is "Input Sanitization" applied to neural signals? (Filtering incoming data streams to prevent "Neural Buffer Overflows" or patterns designed to trigger seizures/confusion).

24.1.3 Is "Write Protection" physical? (Is there a hardware switch that physically disconnects the "Write" capability, ensuring the BCI is Read-Only unless authorized?).

24.1.4 Are "Wireless Gaps" enforced? (Ensuring the BCI uses proprietary, near-field communication rather than standard Bluetooth that can be sniffed from across the room).

24.1.5 Is "Authentication" biological? (The BCI must recognize the user's unique "Brainprint" (EEG signature) before unlocking functionalities).

24.1.6 Are "Emergency Eject" protocols accessible? (Can the user mentally trigger a "Disconnect" command—e.g., by thinking of a specific shape—to cut the connection instantly?).

24.1.7 Is "Stimulation Limits" hard-coded? (Preventing an attacker from cranking up the voltage or frequency to cause physical pain or tissue damage).

24.1.8 Is "Side-Channel" shielding active? (Preventing attackers from analyzing electromagnetic emissions from the headset to reconstruct thoughts).

24.1.9 Are "App Stores" curated for neuro-safety? (Ensuring a meditation app doesn't secretly harvest emotional state data).

24.1.10 Does the Rosecoin Ledger record the "Device Integrity Hash" of the BCI to prove it hasn't been tampered with?

24.2.1 Is "Neuro-Data" encrypted at rest and in transit? (Raw EEG/ neural data should never exist in plaintext).

24.2.2 Is "Semantic Decoupling" used? (Stripping the emotional/ semantic context from the raw signal before it leaves the local device).

24.2.3 Are "Mental Firewalls" configured? (Policies that block the BCI from reading specific "Thought Sectors"—e.g., intimate memories or corporate secrets).

24.2.4 Is "Inferred Data" protected? (Preventing algorithms from inferring sensitive attributes—like early-onset Parkinson's or sexual orientation—from motor cortex noise).

24.2.5 Is "Employer Access" strictly prohibited? (Legally and technically blocking the CISO from seeing "Attention Levels" or "Frustration Metrics" of employees).

24.2.6 Are "P300 Spikes" filtered? (Preventing "Subliminal Interrogation"—flashing images of crime scenes/passwords to see if the brain recognizes them involuntarily).

24.2.7 Is "Data Ownership" absolute? (The user owns their neural data; the vendor is merely a custodian with zero rights to sell it).

24.2.8 Are "Off-Switch" guarantees verified? (When the device is off, is it really off, or is it in "Sleep Mode" still listening to neural activity?).

24.2.9 Is "Anonymization" proven? (Neuro-data is notoriously hard to anonymize; is "Differential Privacy" applied?).

24.2.10 Does the Rosecoin Ledger track "Data Access Grants" so the user can see exactly who accessed their neural stream and when?

24.3.1 Is "Subliminal Detection" active on displays? (Analyzing video streams for single-frame inserts or high-frequency flickers designed to manipulate the subconscious).

24.3.2 Is "Audio Steganography" filtered? (Scrubbing audio streams for ultrasound commands that trigger voice assistants or influence mood).

24.3.3 Are "Deepfake Inoculation" tools used? (Real-time overlays that highlight potential deepfakes in video calls to prevent cognitive hijacking).

24.3.4 Is "Information Overload" throttling active? (Preventing an attacker from flooding the user's BCI with notification spam to induce "Cognitive Denial of Service").

24.3.5 Are "Dark Patterns" blocked in AR/VR? (Preventing augmented reality interfaces from using psychological tricks to force user action).

24.3.6 Is "Reality Verification" enforced? (Digital watermarking to distinguish between "Real Reality" and "Augmented/Virtual Reality" objects).

24.3.7 Is "Emotion Manipulation" detection active? (Alerting the user if their news feed is algorithmically tuned to induce rage or depression).

24.3.8 Are "Social Engineering" alerts integrated into the BCI? (A "Heads Up Display" warning if the person speaking shows signs of deception—voice stress/micro-expressions).

24.3.9 Is "Haptic Fuzzing" detection active? (Ensuring haptic suits/gloves aren't hacked to transmit phantom touches).

24.3.10 Does the organization conduct "Cognitive Penetration Tests"? (Testing if employees can be manipulated into revealing secrets via BCI-based social engineering).

24.4.1 Is "Bio-Telemetry" monitored for safety? (Alerting the SOC if an employee's heart rate/stress levels spike dangerously during a sensitive task—indicating duress or attack).

24.4.2 Are "Neuro-Malware" signatures updated? (Scanning for code patterns known to crash BCI drivers or loop neural inputs).

24.4.3 Is "Isolation" protocol defined? (If a BCI is compromised, is it physically removed immediately?).

24.4.4 Is "Memory Forensics" (Device) enabled? (Capturing the last 5 minutes of device buffer states, not user thoughts, to understand the exploit).

24.4.5 Are "Post-Trauma" protocols in place? (Psychological support for users who have experienced a "Brain Hack" or digital hallucination).

24.4.6 Is "Liability" defined for autonomous action? (If a hacked user "thinks" a command to delete a database, is it their fault?).

24.4.7 Are "Feedback Loops" monitored? (Preventing a loop where the BCI reads stress, tries to calm the user, fails, reads more stress, and crashes).

24.4.8 Is "Ghost Touch" detection active? (Detecting inputs that originate from the network, not the motor cortex).

24.4.9 Is "Dream Advertising" blocked? (Preventing audio injection during sleep cycles).

24.4.10 Does the Rosecoin Ledger serve as the "Immutable Black Box" for neural incidents, protecting the user from false accusations of negligence?

24.4.11 Mental Health/Stress: Training users to recognize when they are being "emotionally hijacked" by an urgent-sounding email.

24.4.12 Deepfake Verification Protocol: Is there a mandatory "out-of-band" verification protocol for confirming the identity of executives during high-value transactions involving video or audio?

24.4.13 Personalized Awareness Training: Does security training use behavioral analytics to identify "at-risk" employees and deliver personalized learning modules based on their specific digital habits?

DOMAIN 25: META-GOVERNANCE & FRAMEWORK EVOLUTION

25.1.1 Is the "Prime Directive" codified? (e.g., "The preservation of client data privacy supersedes all other business objectives, including profit").

25.1.2 Are "Constitutional Amendments" difficult? (Requiring a Super-Majority vote of the Board + CISO + External Audit to change a Core Domain, preventing a new CEO from weakening security to cut costs).

25.1.3 Is "Rights Preservation" guaranteed? (Ensuring that no future security policy can violate the fundamental human rights of employees or users).

25.1.4 Is "Framework Forking" defined? (If the company splits or spins off a subsidiary, does the new entity inherit the Full Framework or a "Lite" version?).

25.1.5 Are "Emergency Powers" explicitly limited? (Defining exactly when the CISO can suspend the Constitution—e.g., during Q-Day—and for how long).

25.1.6 Is "Policy Hierarchy" automated? (If a Local Policy conflicts with the Constitution, the code automatically rejects the Local Policy).

25.1.7 Is "Legacy Protection" enforced? (Ensuring that "New Rules" do not retroactively criminalize "Old Actions" by employees).

25.1.8 Are "Core Values" hashed? (The foundational text is stored on the Rosecoin Ledger and cannot be silently edited).

25.1.9 Is "Whistleblower Protection" constitutional? (It is impossible to fire someone for reporting a violation of Domain 25).

25.1.10 Does the Rosecoin Ledger store the "Genesis Block" of the framework, proving the original intent of the designers?

25.2.1 Is "Policy Sunsetting" mandatory? (Every policy automatically expires after 3 years unless explicitly re-ratified; no "Zombie Policies" allowed).

25.2.2 Is "Trigger-Based Updating" active? (If a new threat class—e.g., "AI Worms"—reaches a Critical Threshold, does the framework automatically trigger a Domain Review?).

25.2.3 Are "Metric Failures" analyzed? (If a specific control fails 90% of the time, is the Control questioned, not just the People?).

25.2.4 Is "Industry Drift" monitored? (If the rest of the industry moves to a new standard—e.g., 6G—does the framework alert the CISO to update Domain 19?).

25.2.5 Is "Feedback Integration" automated? (If 500 users complain about a specific 2FA friction point, is the policy automatically flagged for review?).

25.2.6 Is "Version Control" semantic? (Rosecoin Framework v1.0, v1.1, v2.0—managing security rules like software releases).

25.2.7 Are "A/B Tests" run on policies? (Testing a "Strict" password policy in Dept A and a "Passphrase" policy in Dept B to see which is safer).

25.2.8 Is "Complexity Rot" measured? (If the framework exceeds 3,000 pages, a "Simplification Audit" is triggered).

25.2.9 Is "External Signal" ingestion active? (Using AI to scan academic papers and suggest new controls for Domain 16).

25.2.10 Does the organization use "Policy-as-Code"? (The document is the code; updating the text updates the firewall rules).

25.3.1 Is "Law-Conflict" logic defined? (If US Law says "Keep Data" and EU Law says "Delete Data," which rule wins? The Framework must decide).

25.3.2 Is "Jurisdictional Routing" automated? (Routing data to the server that complies with the most strict overlapping law).

25.3.3 Are "Safe Harbor" bridges built? (Pre-negotiated legal frameworks for moving data between hostile zones).

25.3.4 Is "Sovereign Paradox" managed? (Handling the scenario where complying with the Framework requires violating a local dictatorship's law).

25.3.5 Is "Cultural Variance" accepted? (Allowing the "People Domain" to adapt to local customs while maintaining the "Core Directive").

25.3.6 Are "Sanctions Collisions" mapped? (What if a critical security vendor is sanctioned in one country but not another?).

25.3.7 Is "Extradition Protection" considered? (Ensuring the CISO is not personally liable for a conflict between two nations' laws).

25.3.8 Is "Universal Minimum" enforced? (Regardless of local law, the "Rosecoin Baseline" (e.g., Encryption) is never breached).

25.3.9 Are "Diplomatic Channels" established? (Using the ISAC to lobby for regulatory harmonization).

25.3.10 Does the Rosecoin Ledger record "Conflict Decisions" to prove to regulators why a specific choice was made during a legal catch-22?

25.4.1 Is the "Auditor" audited? (A third party verifies that the Primary Auditor isn't just rubber-stamping the compliance reports).

25.4.2 Is "Audit Rotation" algorithmic? (The Framework selects a new Auditor randomly every 3 years).

25.4.3 Are "Shadow Audits" performed? (Sending a "Secret Shopper" auditor to test if the compliance team is actually checking logs).

25.4.4 Is "Evidence Integrity" verified? (Checking if the screenshots provided to the auditor were Photoshopped).

25.4.5 Is "Conflict of Interest" detection active? (Ensuring the Auditor doesn't own stock in the security vendor they are testing).

25.4.6 Are "Incentive Structures" audited? (Ensuring the Auditor isn't paid extra for a "Clean Report").

25.4.7 Is "AI Auditing" used? (Using an AI model to read every single policy document and find contradictions).

25.4.8 Is "Continuous Verification" the standard? (Moving from "Annual Audit" to "Real-Time Dashboard").

25.4.9 Is "Regulatory Mapping" automated? (If NIST changes, the Meta-Audit instantly highlights the gaps).

25.4.10 Does the Rosecoin Foundation have the power to "Revoke" the auditor's license if they are found to be negligent?

25.5.1 Is "Archive Strategy" multi-generational? (Storing the Framework on medium (e.g., silica glass) that lasts 1,000 years).

25.5.2 Is "Knowledge Transfer" institutionalized? (Ensuring the "Why" behind the rules is passed down, not just the "What").

25.5.3 Is "Civilization Collapse" continuity considered? (If the internet ends, is there a printed manual on how to rebuild the secure network?).

25.5.4 Is "Language Evolution" planned? (Will "Firewall" mean the same thing in 50 years? Keeping definitions updated).

25.5.5 Is "Values Alignment" checked? (Ensuring the Framework doesn't drift into becoming a tool for oppression).

25.5.6 Is "Succession" decentralized? (If the entire Board dies, is there a protocol to appoint new guardians of the Framework?).

25.5.7 Is "Scorched Earth" protocol defined? (The ultimate decision to destroy the data rather than let it fall into an enemy's hands).

25.5.8 Is "Re-Genesis" planned? (How to bootstrap the Rosecoin Network from zero trust).

25.5.9 Is "Time-Capsule" crypto used? (Locking certain secrets until a specific date in the future).

25.5.10 THE FINAL CHECKPOINT: Does the Organization accept that "Security is a Journey, not a Destination," and pledge to evolve this Framework forever?

25.5.11 Sovereign Risk Assessment: Does the organization evaluate the risk of "Geopolitical Fragmentation," specifically assessing how data-sharing obligations in one jurisdiction might conflict with national security laws in another?

25.5.12 Fiduciary Cyber-Reporting: Is the "RCF Security Score" a standard metric in every financial and risk report delivered to the board, treated with the same weight as revenue?

25.5.13 Framework Evolution Auto-Trigger: Does a significant change in the global threat landscape (e.g., a "Q-Day" event) automatically trigger an emergency review of the RCF Constitution?