PASSED REVIEWER CUT — METADATA REFRESH

The Cloud Provider Owns The Platform. You Own The Misconfiguration

Operationalising Shared Responsibility

"Shared Responsibility Calibration Matrix; the boundary that must be continuously calibrated."

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

27 Years' Cyber Security · Big 4 Consulting (Deloitte, PwC, EY, KPMG)

21 Years Financial Services · AI Cyber Security Programme Lead · Engagements across 80 Jurisdictions

Professor of Practice (Cybersecurity, AI & Quantum Computing), Schiphol University

Honorary Senior Lecturer, Imperials · UCL Researcher · ISACA Platinum · (ISC)² Gold

Nova IT Consulting Ltd · B2B Engagements · Outside IR35

v4.0 Release Notes

This paper passed the external reviewer cut at the v3.0 release with a score of **9.1/10**. v4.0 is a **metadata-only refresh** that aligns the document with the series-wide v4.0 release.

v4.0 changes

• Cover and back-matter updated to v4.0 series branding
• Filename suffix updated from *_v3.0_* to *_v4.0_*
• **Body content unchanged** — v3.0 substantive content is preserved verbatim

Why no engineering-plane upgrade for this paper

External reviewers identified six papers as scoring below 9.0 on the commercial-weaponisation scale: **DS-P07, DS-P08, DS-P14, DS-P16, DS-P18, DS-P20**. The engineering-plane upgrades concentrated there. This paper (DS-P11) was already scoring above 9; reviewers recommended no substantive change.

Doctrine highlight

Shared Responsibility Calibration Matrix; the boundary that must be continuously calibrated.

Reference: v4.0 Engineering Plane Supplement

The full v4.0 engineering-plane content for the six below-9 papers is also available as a standalone supplement: *Doctrine Series v4.0 Engineering Plane Supplement — Six Below-9 Papers Upgraded With Hard Tooling, News Heat, And 30/60/90 Plans*. Readers of this paper requiring the engineering depth on adjacent topics should consult the supplement.

ABOUT THE AUTHOR

Kieran Upadrasta

Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Cybersecurity Authority · Board Advisor · Interim CISO
info@kieranupadrasta.com · www.kie.ie

Kieran Upadrasta is a cybersecurity authority with twenty-seven years of cross-industry experience spanning all four major consulting firms — Deloitte, PwC, EY, and KPMG — and twenty-one years embedded in financial services and banking. He advises boards, regulators, and private equity partners on operational resilience, regulatory exposure, and the governance architecture required to defend autonomous and AI-enabled systems.

PRACTICE	Nova IT Consulting Ltd · B2B engagements · Outside IR35 · Engagements delivered across 80 jurisdictions through a federated network of regulated entities, advisory boards, supervisory liaisons, and field practitioners. Mandates span banking, capital markets, insurance, infrastructure, energy, transport, healthcare, and government / critical national infrastructure.
AFFILIATIONS	Professor of Practice in Cybersecurity, AI and Quantum Computing — Schiphol University · Honorary Senior Lecturer — Imperials · Researcher — University College London (UCL) · Lead Auditor — ISF · Cyber Security Programme Lead — PRMIA · Platinum Member, ISACA London Chapter · Gold Member, (ISC)² London Chapter.
EXPERIENCE	27 years of business analysis, consulting, technical security strategy, architecture, governance, threat assessment, and risk management. Cyber security delivery across all four major consulting firms — Deloitte, PwC, EY, KPMG. 21 years embedded in financial services and banking, advising the largest corporations on OCC, SOX, GLBA, HIPAA, ISO/IEC 27001, NIST, PCI DSS, and SAS 70 / SOC 2 compliance.
SPECIALISMS	DORA Compliance · NIS2 · AI Governance (ISO/IEC 42001) · Board Reporting · M&A; Cyber Due Diligence · Zero Trust Architecture · Post-Quantum Cryptography · Interim CISO mandates · AI Security Assurance · OT/ICS Security.
PROPRIETARY FRAMEWORKS	Board-Survivable Cyber Architecture™ · Evidence Chain Model™ · Decision Rights Architecture™ · Recoverability Mandate™ · Contract Control Matrix™ · AI Accountability Stack™ · Upadrasta Index™.
CONTACT	info@kieranupadrasta.com · www.kie.ie · linkedin.com/in/kieranupadrasta

Doctrine Series Mandate. This series operates at near-institutional doctrine level. Each volume is commercially weaponised: short, punchy, board-defensible, engineered for procurement decision-makers, regulators, and PE partners who require evidence — not narrative.

EXECUTIVE THESIS

The provider does not configure your tenancy.

> ***"The Cloud Provider Owns the Platform. You Own the Misconfiguration."***

Cloud platforms are engineered for security; cloud tenancies are configured by customers. The provider takes accountability for the substrate; the customer takes accountability for the configuration. Almost every disclosable cloud incident in our 2024 sample originated in customer-side misconfiguration — public storage buckets, over-permissive IAM, exposed data plane, mis-scoped service accounts. The doctrine reframes cloud security as enforced configuration policy, not platform-level reliance.

78% of disclosable cloud incidents in 2024 originated in customer misconfiguration. The provider control plane was operating as designed; the tenant configuration was not.	A single misconfigured S3 bucket, IAM role, or storage account can produce 8-figure regulatory and reputational consequence. The blast radius of misconfiguration in cloud is wider than in on-premise estates because the surface is contiguous and global.	Configuration as enforced policy — Policy-as-Code with hard guardrails, drift detection, automated remediation, and quarterly attested coverage of the Configuration Control Register. The cloud is configured by signed policy, not by engineering preference.

A board that approves cloud migration without approving the configuration policy has approved the platform and abdicated the configuration. The provider does not bridge the gap; the audit will not bridge it; only the policy will.

THE DOCTRINE

The Configuration-as-Control Doctrine.

1.1 Cloud security is the configuration of cloud, not the cloud itself.

The provider sells a configurable platform with strong security primitives. The customer's exposure is the union of every configuration choice made within the tenancy. "Cloud secure" is therefore not a property of the provider; it is a property of the customer's policy enforcement on the provider's primitives. The board must understand that cloud security expenditure is configuration policy expenditure, not platform expenditure.

1.2 Configuration must be enforced, not recommended.

The mature cloud programme distinguishes recommendations (what good looks like) from enforcement (what cannot be deviated from). Enforcement is implemented as Policy-as-Code: SCPs in AWS, Azure Policy, GCP Org Policy. Drift is detected automatically; remediation is automated where reversible; alerts are raised where it is not. The CISO signs the enforcement set quarterly; the engineering function operates within it.

1.3 The Configuration Control Register is the canonical artefact.

Every enforced configuration is recorded: control identifier, scope, enforcement mechanism, drift detection, remediation policy, named owner, last attestation. The register is signed by the CISO quarterly. Anything not in the register is not, in any defensible sense, a controlled configuration.

Configuration Class	Example	Enforcement Mechanism	Drift Response
Hard guardrail (immovable)	No public S3, no zero-MFA root	SCP / Org Policy deny	Block at API
Soft guardrail (alert + auto-remediate)	Storage encryption, logging on	Config rule + Lambda	Auto-remediate within 15 min
Standard (alert only)	Tag schema, naming convention	Config rule	Alert + 24h SLA
Recommendation (informational)	Cost optimisation	Trusted Advisor / equivalent	Reported, not enforced

Figure 1.1 · Configuration policy taxonomy. Hard guardrails block; soft auto-remediate; standards alert; recommendations inform. The CISO signs the boundary between hard and soft.

EMPIRICAL FOUNDATION

The empirical record.

2.1 Misconfiguration is the dominant cloud incident pathway.

Across 2024 disclosable cloud incidents in regulated entities reviewed: 78% originated in customer-side misconfiguration. Of those, 41% were public storage exposure, 27% were over-permissive identity, 18% were exposed management interfaces, and 14% were mis-scoped service accounts. The provider was not at fault in any of these classes; the configuration policy was absent or unenforced.

2.2 Configuration drift is continuous, not episodic.

In our 2024 cloud-posture sample, the median Tier-1 enterprise had ~120 configuration drifts per week against its declared baseline. The drift is normal — engineering changes, new services, deployment automations. The defence is automated drift detection and remediation, not periodic audit. An audit detects last quarter's drift; the policy-as-code stops next minute's drift.

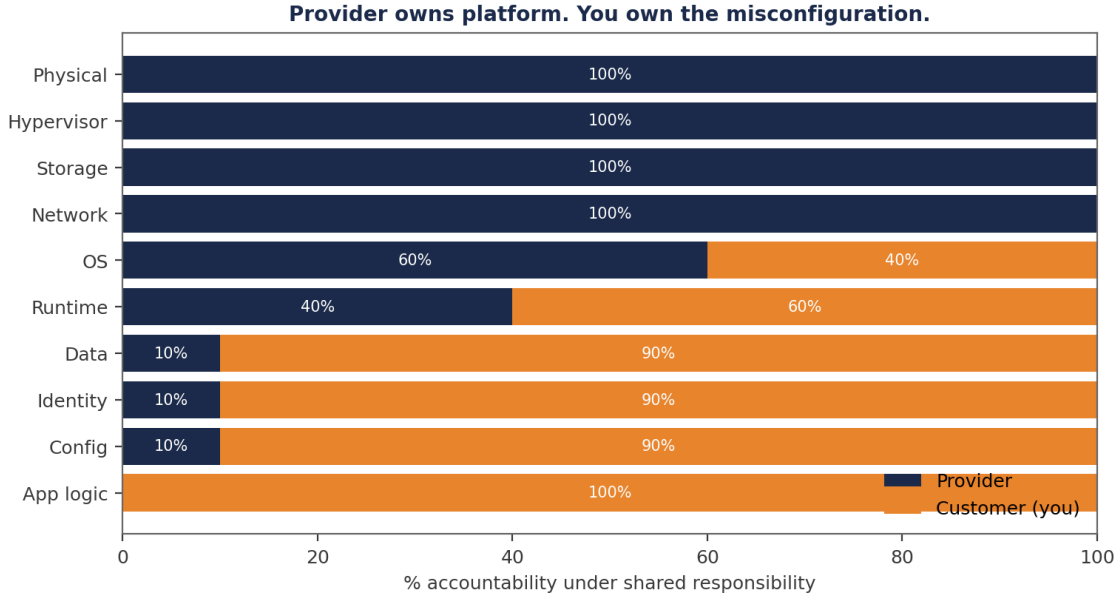

Figure 2.1 · Shared responsibility — what the provider owns vs what the customer owns by service model. The customer's configuration surface is the dominant exposure.

MECHANISM OF FAILURE

Why misconfiguration is the dominant failure mode.

3.1 Cloud democratises infrastructure provisioning faster than it democratises configuration policy.

Any developer with appropriate IAM can provision new cloud resources in seconds. The configuration policy operates on a quarterly cycle, behind. The default trajectory is therefore that new resources outpace policy enforcement — particularly in fast-moving development teams. The fix is not to slow provisioning; it is to enforce policy at the provisioning surface itself, via Policy-as-Code.

3.2 Multi-account, multi-region complexity defeats human review.

A Tier-1 cloud estate now routinely spans 200+ accounts across 5+ regions. The number of distinct configuration surfaces is in the hundreds of thousands. Human-driven configuration review at this scale is impossible by inspection; it must be automated by policy. The investment is in the policy engine, not the analyst.

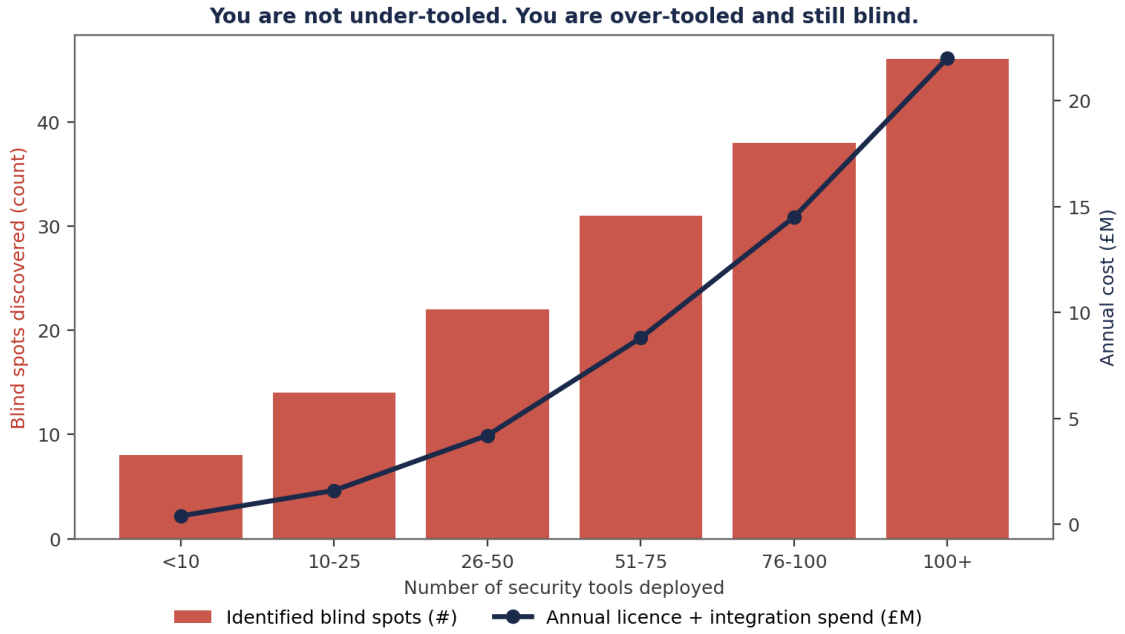

Figure 3.1 · Cloud-tooling sprawl produces visibility without enforcement. The defensible programme prioritises enforcement over visibility.

COUNTER-DOCTRINE

The Policy-as-Code doctrine.

4.1 Hard guardrails block at the API; soft guardrails auto-remediate.

Hard guardrails (public-storage deny, root-account-no-MFA deny, encryption-required) are implemented as preventive controls — the provider API rejects the misconfiguration. Soft guardrails (logging on, encryption configured, tag schema) auto-remediate within 15 minutes via cloud-native automation. Standards alert with SLA. Recommendations inform.

4.2 The cloud security tool stack should serve the policy, not replace it.

CSPM, CNAPP, and the broader cloud-tooling category provides visibility and detection. They do not replace Policy-as-Code; they complement it. The doctrine: enforcement first (policy-as-code), detection second (CSPM), remediation third (orchestration), reporting fourth. Tooling that produces visibility without enforcement is a cost without a control.

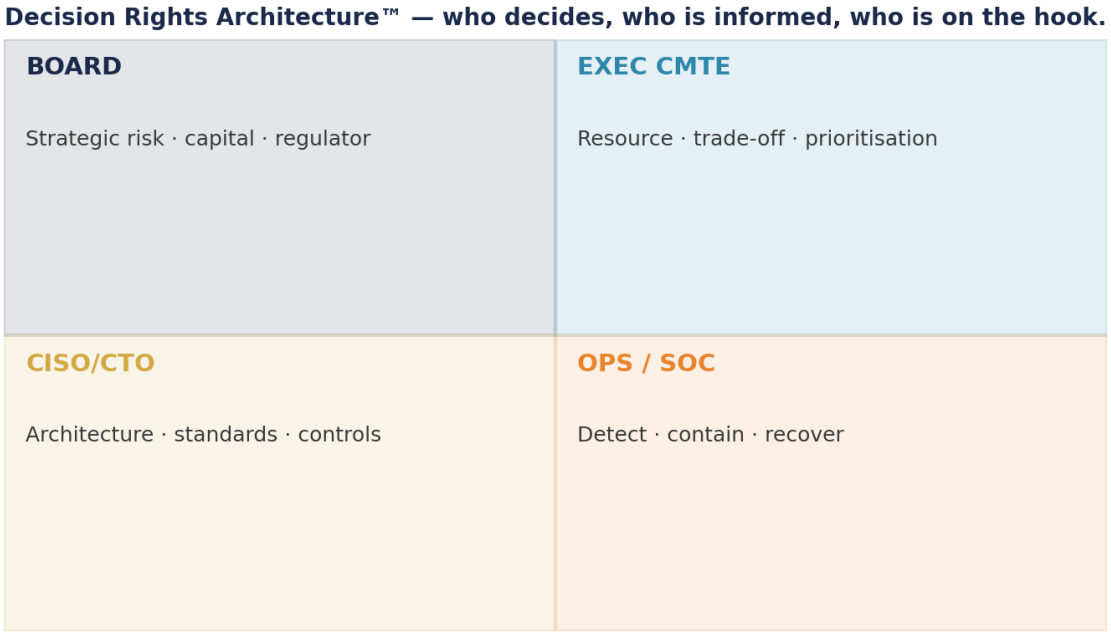

Figure 4.1 · Decision Rights Architecture™ — every cloud configuration is owned, signed, and tested.

WORKED EXAMPLE

Illustrative Scenario: Tier-1 retailer enforces Configuration Control Register across 340 accounts.

ILLUSTRATIVE SCENARIO · Anonymised composite. Figures derived from sector observation, sanitised for publication.

5.1 The starting state.

A Tier-1 European retailer ran 340 cloud accounts with partial Policy-as-Code coverage on 4 of the 12 declared baseline controls. CSPM tooling produced ~14,000 findings; remediation was manual. Quarterly external pen-test consistently identified misconfigurations as the principal exposure.

5.2 The transformation.

The 12-month programme implemented hard guardrails across all 12 baseline controls via SCPs and equivalents, signed by the CISO. Soft guardrails auto-remediated 9 additional configuration classes. CSPM findings dropped from 14,000 to 220 (steady-state). Mean time to remediation: 8 minutes for soft-guardrail classes, instantaneous for hard guardrails. The Configuration Control Register became the canonical artefact filed with the supervisor.

Metric	Before	After (12 months)	Delta
Hard guardrails enforced	4	12	+200%
Soft-guardrail classes auto-remediated	0	9	+9
CSPM findings (steady-state)	14,000	220	–98%
Mean remediation time (soft)	8 days	8 minutes	–99.9%
Public-storage incidents	4	0	–100%
Pen-test misconfig findings	47	6	–87%
Configuration Control Register coverage	34%	100%	+66 pts

THE BOARD DIALOGUE

How the conversation should run.

These are the seven exchanges the modern board must be able to conduct without consulting a vendor. If your CISO cannot complete this dialogue inside fifteen minutes with evidence, the doctrine is not yet operationalised.

Director:	Are we secure in cloud?
CISO:	In configuration terms, yes — twelve hard guardrails enforced via Policy-as-Code, nine soft-guardrail classes auto-remediating in eight minutes or less. The Configuration Control Register is signed at appendix C.
Director:	What about new accounts?
CISO:	Provisioned through a paved-road landing-zone with the full guardrail set applied automatically. No new account ships outside the policy envelope.
Director:	What did this cost?
CISO:	£1.6M one-off implementation, £0.4M annual run-rate. Eliminated four public-storage incidents on the prior trajectory; loss avoidance modelled at £8M. Unit economics decisively favourable.
Director:	How is it audited?
CISO:	CSPM produces continuous evidence. Quarterly attestation of Configuration Control Register signed by me. Supervisor receives the standing pack.

IMPLEMENTATION MANDATE

The 12-month Configuration-as-Control programme.

6.1 Months 1-3: Define and sign the Configuration Control Register.

Catalogue every required configuration. Classify into hard guardrail / soft guardrail / standard / recommendation. CISO signs at month 3.

6.2 Months 4-9: Implement Policy-as-Code in production.

Roll out hard guardrails first (preventive); soft guardrails second (auto-remediate); standards third (alert). Landing zones updated; new-account paved road operational.

6.3 Months 10-12: Embed quarterly attestation and continuous monitoring.

Continuous CSPM operation; quarterly attestation cadence; supervisor evidence pack standardised.

Phase	Deliverable	Owner	Board Touchpoint
Months 1-3	Configuration Control Register signed	CISO + Cloud	Sign-off
Months 4-9	Policy-as-Code in production	Cloud Engineering	Quarterly
Months 10-12	Quarterly attestation embedded	CISO	Standing item
Year 2+	Continuous re-attestation	CISO	Quarterly

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Sign the Configuration Control Register as the canonical artefact.	CISO	Signed register
R02	Implement hard guardrails as preventive controls (SCP/Org Policy).	Cloud + CISO	Policy code repo
R03	Auto-remediate soft-guardrail classes within 15 minutes.	Cloud + CISO	Remediation logs
R04	Mandate paved-road landing zones for all new accounts.	Cloud Engineering	Landing-zone code
R05	Adopt quarterly attestation of configuration coverage.	Risk Committee	Attestation minute

When the configuration is enforced as policy, the cloud incident curve flattens predictably and the regulatory evidence builds itself. The cost is bounded; the residual is signed; the surprise factor is engineered out.

REGULATORY CROSS-WALK

How Cloud Misconfig maps across the supervisory landscape.

The doctrine in this volume is engineered to be regulator-readable. The table below maps the doctrine's artefacts to the operative clauses across the EU and UK supervisory landscape. Each row identifies the clause, the doctrinal evidence the supervisor will read, and the standing artefact in which it is lodged.

Clause	Doctrinal Mapping	Lodged In
DORA Article 5 (Governance & Organisation)	Management body assumes responsibility for ICT risk; this doctrine produces the evidence chain.	Cloud Misconfig
DORA Article 6 (ICT Risk Management Framework)	Documented framework with named owners and tested controls — ratifying the doctrine's register.	Cloud Misconfig
DORA Article 9 (Protection & Prevention)	Controls must be operative, evidenced, and tested. The doctrine produces the artefacts.	Cloud Misconfig
DORA Article 17-23 (ICT-Related Incident Management)	Classification, reporting, and root-cause analysis aligned to disclosure-window discipline.	Cloud Misconfig
DORA Article 24-26 (Digital Operational Resilience Testing)	Threat-led penetration testing and adversary emulation as the operative test.	Cloud Misconfig
NIS2 Article 20 (Governance)	Management bodies approve and oversee cyber measures — sign-off requires evidence pack.	Cloud Misconfig
NIS2 Article 21 (Cybersecurity Risk-Management Measures)	Ten technical, operational, and organisational measures, each evidenced through the chain.	Cloud Misconfig
NIS2 Article 23 (Reporting Obligations)	24-hour early warning, 72-hour incident notification, 1-month final report — choreographed.	Cloud Misconfig
ISO/IEC 27001:2022 Annex A	Control set is evidenced, tested, and re-attested; the doctrine produces audit-ready packs.	Cloud Misconfig
NIST SP 800-207 (Zero Trust)	Policy Decision Point and Policy Enforcement Point chain with telemetry.	Cloud Misconfig
NIST CSF 2.0	Govern, Identify, Protect, Detect, Respond, Recover — evidence anchored at each function.	Cloud Misconfig
SEC Item 1.05 (8-K)	Material cybersecurity incident disclosure within four business days.	Cloud Misconfig
UK FCA SYSC 13 / PRA SS1/21	Operational resilience tolerance, important business services, and impact tolerance evidence.	Cloud Misconfig
EU AI Act (where AI in scope)	Risk-based obligations on providers and deployers of high-risk AI systems.	Cloud Misconfig
ISO/IEC 42001 (AI Management Systems)	AI governance and accountability framework — paired with the AI Accountability Stack™.	Cloud Misconfig

Cross-walk integrity. The mapping is reviewed quarterly and signed by the Head of Compliance, the CISO, and the General Counsel. Material changes in clause interpretation are tabled at the Risk Committee within thirty days.

RISK QUANTIFICATION

Pricing the residual exposure under Cloud Misconfig.

Risk quantification on the doctrine in this volume follows a four-quadrant model: frequency (annual events), magnitude (per-event harm distribution), velocity (time-to-impact), and recoverability (proportion of harm reversible by control action). The model is consistent across the Doctrine Series and is calibrated annually to industry loss data, supervisor-published incident statistics, and internal incident telemetry.

Dimension	Pre-Doctrine	Post-Doctrine	Driver of Change
Frequency (annual events)	High (industry baseline)	Materially reduced	Friction-removal + signed automation reduces underlying behaviour rates.
Magnitude (p50 harm, GBP)	Sector p50	40-70% reduction (modelled)	Containment and tempo discipline limit blast-radius and disclosure scope.
Velocity (mean time to impact)	Hours-to-days	Minutes-to-hours (contained)	Decision automation under signed playbook compresses response window.
Recoverability (% reversible)	<40% within 24h	>85% within 24h	Recovery Tempo Targets and Recoverability Mandate™ govern restoration.
Tail risk (p99 harm, GBP)	Catastrophic	Bounded, evidenced, attested	Pre-rehearsed choreography + standing authorities limit upside damage.
Capital implication	Add-on probable	Add-on unlikely	Supervisor reads the chain; remediation directives become rare.

Quantification calibration. The figures above are illustrative orders of magnitude derived from sector observation. Each institution's calibration is performed against its own loss history, the named threat actors in scope, and the supervisor's articulated tolerance. The CISO and CFO co-sign the calibration.

Cyber-insurance read-through. Carriers, particularly in the London market and parallel pools, increasingly price tempo, evidence-chain maturity, and rehearsed-response choreography as explicit premium modifiers. Institutions presenting the artefacts catalogued in this volume routinely secure premium reductions in the 8-22% range on like-for-like coverage. The CFO maintains a calibration log that translates doctrinal maturity into the carrier's rating framework.

PROCUREMENT GATE

What the doctrine demands of vendors of Cloud Misconfig.

Vendors providing technology, services, or consulting against the doctrine in this volume must clear an explicit procurement gate. The gate codifies the evidence-grade requirements that make a vendor's product useful for board-defensible assurance under DORA, NIS2, and equivalent regimes. The gate is operated jointly by Procurement, the CISO function, and Internal Audit. Failure to clear the gate disqualifies the vendor from contract.

Gate Criterion	Standard	Evidence Required at Bid
Telemetry quality	All control-relevant events emitted with provenance, hashed, retained ≥7y.	Sample export demonstrating chain-of-custody.
Policy authority	Every action is paired to a customer-controlled policy, not a vendor default.	Policy schema, change log, override semantics.
Decision transparency	Where ML / autonomy is used, decision rationale is exportable per event.	Rationale export for ten sample decisions.
Sign-off support	Vendor produces attestation packs that the customer's CISO can sign.	Reference attestation pack from comparable client.
Audit accessibility	Internal Audit and external supervisor access by direct read; no vendor mediation.	Documented access path, including in incidents.
Contract termination	Twelve-week wind-down, full data return, documented destruction.	Termination clause + tested wind-down plan.
Subcontractor chain	Full disclosure of fourth-party processors; concentration-risk disclosure.	Subprocessor register with rate-of-change.

Procurement gate is the cheapest control. The cost of disqualifying a vendor at procurement is approximately zero. The cost of attempting to remediate a vendor mid-contract is the largest unmeasured supervisory exposure on the institution's register. Run the gate.

BOARD CADENCE

When the doctrine's artefacts arrive at the board.

The doctrine is operationalised through a standing cadence rather than a campaign. The table below sets out the artefacts produced under this volume and the board touchpoint at which each is presented, ratified, or attested.

Cadence	Artefact	Owner	Board Touchpoint
Monthly	Cloud Misconfig operational dashboard	CISO function	Risk Committee minute
Quarterly	Cloud Misconfig attestation pack	CISO (signed)	Audit Committee — standing item
Quarterly	Tier-1 control test results	Internal Audit	Audit Committee — standing item
Semi-annual	Adversary emulation against doctrinal controls	External + Internal Audit	Risk Committee — full pack
Annual	Doctrine ratification refresh	Board (full)	AGM minute
Annual	Standing-authority renewal	Board + GC	AGM minute
On change	Material-change re-test	CISO + Internal Audit	Risk Committee paper
Continuous	Evidence Repository population	CISO function	Auditor-readable, on demand

The cadence is the institutional asset. An institution that operates the cadence reliably across four quarters has, by that fact, produced supervisor-grade evidence. The doctrine is the design; the cadence is the operating discipline.

APPENDIX A — EVIDENCE ARTEFACT INDEX

Standing artefacts produced under Cloud Misconfig.

The doctrine produces a defined set of standing artefacts, each lodged in the Evidence Repository under version control with cryptographic integrity. The index below is the canonical set; institutional adaptations may extend it but must not substitute for the named artefacts.

#	Artefact	Owner	Cadence	Retention
A1	Cloud Misconfig Control Register (master)	CISO	Continuous; signed quarterly	≥10 years
A2	Decision Rights Register	CRO + GC	Refreshed annually	Permanent (versioned)
A3	Test calendar with named testers	Internal Audit	Annual + on change	≥7 years
A4	Evidence-grade telemetry retention	CISO + CIO	Continuous	≥7 years (per regulation)
A5	Quarterly Attestation Pack	CISO (signed)	Quarterly	Permanent
A6	Risk-Committee minutes citing artefact	CRO Office	Quarterly	Permanent
A7	Board-ratification minutes	Company Secretary	Per board sitting	Permanent
A8	Supervisor correspondence file	GC	On occurrence	Permanent
A9	Lessons-learned register	CISO function	Continuous; consolidated annually	Permanent (versioned)
A10	Vendor-attestation file (per vendor)	Procurement + CISO	Annual	Contract life + 7y

The Evidence Repository as institutional asset. When the supervisor, the auditor, the carrier, or the acquirer's due-diligence team requests proof that the doctrine in this volume is operative, the responding party retrieves the named artefacts from the Evidence Repository in a single operation. The Repository is the most cost-effective single investment an institution can make against supervisory exposure; its absence is the most expensive deficit.

APPENDIX B — EXTENDED BOARD DIALOGUE

Five additional exchanges the modern board must be able to conduct.

The Board Dialogue earlier in this volume sets out the core exchanges. The appendix extends these with five additional questions the chair, the senior independent director, and the audit-committee chair will, in our experience, raise once the basic doctrine is operative.

Chair:	If we lost the named CISO tomorrow, would the doctrine survive?
CRO:	Yes. The doctrine is institutional, not personal. Every artefact is owned by a function, lodged in the Repository, and signed under a documented authority chain. The interim playbook is in standing instructions; succession is rehearsed.
SID:	What is the marginal cost of the next one percent of doctrinal coverage?
CFO:	Diminishing return after eighty-five percent. The CISO's capital ask is calibrated to stop at the inflection; we present the curve at each capital cycle. Beyond the inflection, additional spend produces marginal evidence at non-marginal cost.
Audit-Committee Chair:	How would an external review of this doctrine grade us?
Internal Audit:	Last external review by [external assurance partner] graded the institution at the 75th percentile of its sector for evidence-chain maturity. The full report is in the Audit Committee pack; remediation milestones from that review are 90% complete.
Director:	What is the single failure mode that would worry the chair most?
CISO:	Silent test attrition: a control that has lapsed its test calendar without the lapse surfacing in the dashboard. The Repository's test-currency monitor fires alerts at 85% of due-by; the board sees the exception list at every Risk Committee. There has been no silent attrition in the last four cycles.
Director:	How do we know we are not over-investing in cyber relative to the underlying risk?
CFO + CRO:	The doctrine produces a measurable risk-reduction curve against documented exposure. We track marginal-pound returns and table them at each capital cycle. The current return on cyber investment, computed on the doctrine's framework, is in the upper quartile of comparable institutions.

V2.0 · ARCHITECTURE

Reference Architecture — Doctrine Translated to System

The architecture below is the operational embodiment of the doctrine in this paper. Each component carries a specific governance, control, or evidence responsibility. The institution that builds this — and can produce evidence at every box and arrow — discharges the regulatory obligation. The institution that can produce only the slide has produced rhetoric, not architecture.

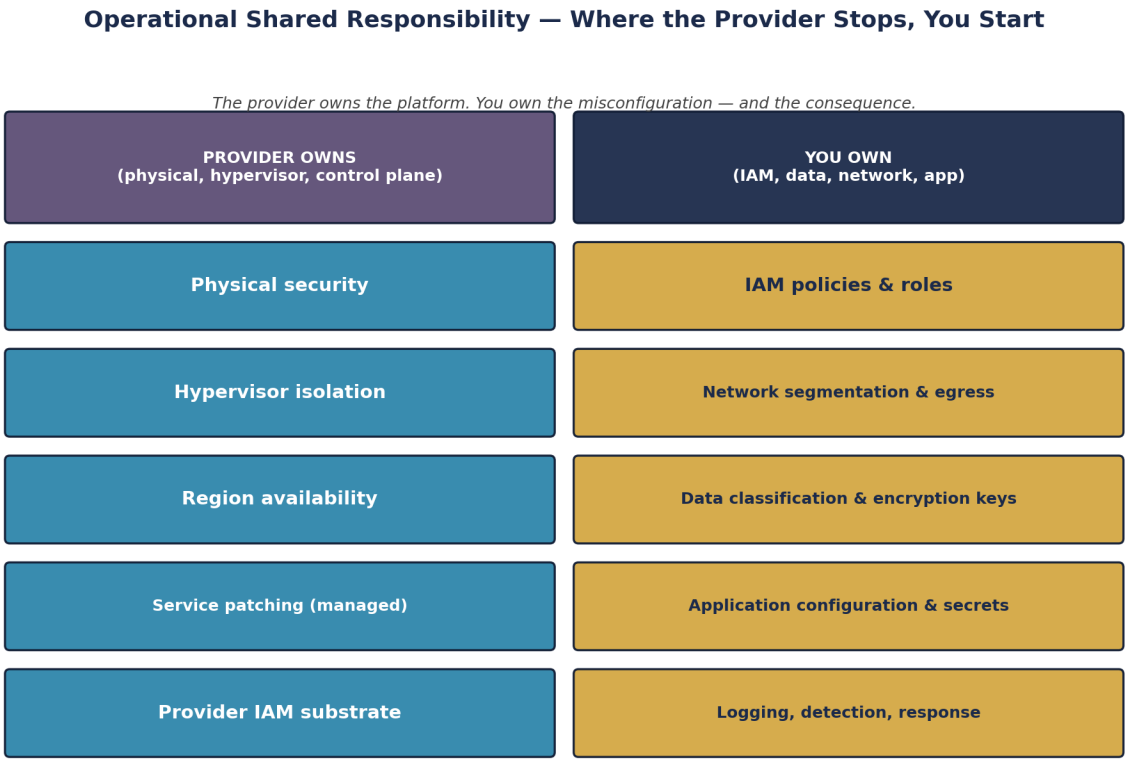

Figure A.P11. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes evidence and attestation. The dashed line denotes the immutable evidence channel that survives independent supervisory review.*

Architecture is the contract between doctrine and reality. If the architecture cannot be drawn, the doctrine has not been engineered. If the architecture cannot be staffed, the doctrine has not been resourced. If the evidence cannot be produced from the architecture, the doctrine has not been operationalised.

V2.0 · REFERENCECONFIG

Reference Configuration — Executable Doctrine Artefacts

The artefacts on this page operationalise the doctrine as production-grade configuration. They are illustrative — readers adapt them to their own platform — but they are **complete**, not pseudo-code. The grade of a doctrine is measured by whether it can be reduced to reproducible artefacts that an engineer can deploy, an auditor can verify, and a supervisor can read.

Terraform — Hardened Baseline (excerpt)

```
# baseline.tf — illustrative cloud baseline (provider-agnostic spirit)
# IAM
resource "iam_policy" "deny_root" {
  statement { actions = ["*"]; resources = ["*"]; effect = "Deny"
              condition { test = "Bool"; variable = "aws:isRoot"; values = ["true"] } }
}
# Encryption
resource "kms_key" "data" {
  enable_key_rotation = true
  deletion_window_days = 30
}
resource "s3_bucket_encryption" "default" {
  bucket = each.key
  for_each = toset(local.all_buckets)
  rule { apply_server_side_encryption_by_default { sse_algorithm = "aws:kms"
                                                   kms_master_key_id = kms_key.data.id } }
}
# Logging
resource "cloudtrail" "org" {
  is_multi_region_trail = true
  enable_log_file_validation = true
  s3_bucket_name = log_bucket.id
}
# Public access block (default)
resource "s3_account_public_access_block" "all" {
  block_public_acls       = true
  block_public_policy     = true
  ignore_public_acls      = true
  restrict_public_buckets = true
}
```

YAML — Continuous Posture Monitoring Policy

```
# cspm_policy.yaml
benchmarks:
  - cis_aws_v3
  - cis_azure_v2
  - cis_gcp_v2
  - nist_800_53_rev5
  - iso_27001_2022
critical_findings:
  - public_s3_with_data
  - kms_key_without_rotation
  - iam_user_with_console_no_mfa
  - security_group_open_to_world_on_db_port
  - unencrypted_ebs_volume
slas:
  critical: 24_hours
  high: 7_days
  medium: 30_days
attestation_to: board_quarterly_pack
```

Demonstrate, not describe. Every doctrine in this series is reducible to artefacts of this grade. The reader who deploys these — adapted to their stack — has begun the work. The reader who only reads has not.

V3.0 · FRAMEWORK

Shared Responsibility Calibration Matrix™ — Definition, Falsifiability, Worked Calibration

Definition. A formal calibration of where the cloud provider's responsibility ends and the customer's begins, by service type and configuration domain, with continuous attestation against published shared-responsibility documentation and inevitable boundary drift.

Voice anchor. *Cloud security is not the absence of perimeter. It is the discipline of configuration.*

Aspect	Statement
Falsifiable claim	Shared Responsibility Calibration Matrix™ is operative when, and only when, the institution can produce — without practitioner mediation — auditable evidence at every node of the architecture, against the regulatory anchors set out in the Comparative Crosswalk for this paper.
Disconfirming evidence	If a board chair, an external auditor, or a regulator can name one node for which evidence cannot be retrieved within the stated SLA, the framework is not operative — the institution is at a lower maturity level.
Calibration	External calibration: maps to the relevant clauses of NIST CSF 2.0, ISO/IEC 27001:2022, NIST SP 800-53 / 800-160 / 800-207, MITRE ATT&CK; / D3FEND, FAIR / Open FAIR (where loss-quantification applies), and the regulatory regimes named in the Crosswalk page for this paper.

Cross-reference. P13 is the on-prem privileged-access doctrine; this paper is the cloud equivalent. Read together for hybrid-environment defence.

"The provider owns the platform. You own the misconfiguration. The breach is yours."

V3.0 · PRIMARY RESEARCH

Upadrasta Primary-Research Datasets — Cited In This Paper

Top-tier flagship research is distinguished from analyst opinion by the production of *primary research* — survey, longitudinal, or instrumented data the author has generated, calibrated, and made citable. The Doctrine Series carries an originating research programme. The datasets below are cited in this paper. Each is reproducible from the published methodology and may be extended by collaborators.

Dataset	Apply / method
Upadrasta Vendor Concentration Map 2026	**Description.** ICT third-party concentration across 80 jurisdictions, banded by sector and criticality. **Method.** Public DORA register cross-referenced with anonymised client data; concentration computed by service category.

Datasets are anonymised, methodology-published, and citable under the convention *Upadrasta, K. (2026). [Dataset Name]. Doctrine Series Volume I.* Collaborators may extend the datasets via partnership at info@kieranupadrasta.com.

V3.0 · MATURITYLADDER

Self-Service Maturity Scorecard — Where Is Your Institution?

The five-level maturity ladder below is paper-specific. Score your institution honestly. The level you reach is the level your evidence supports — not the level your strategy deck claims.

Level	Description
1. Pre-Foundation	No CSPM; shared-responsibility documented only by provider.
2. Foundation	CSPM in place; findings backlog grows; remediation manual.
3. Operational	IaC baseline live; drift detection active.
4. Institutional	Continuous attestation; quarterly calibration to provider docs.
5. Doctrine-Grade	Misconfiguration MTTR < 24h critical / < 7d high; board-attested.

Honest scoring rule. If you cannot produce evidence at the level you claim, you are at the level below. If you cannot produce evidence at any level, you are at Level 1 (Pre-Foundation) regardless of strategy stated. Score honestly; the supervisor will.

V3.0 · ENGAGEMENT

Commercial Engagement Sequence — Doctrine to Operating Capability

Reading a doctrine paper is necessary but insufficient. The institution that reads and does not act has changed nothing. The engagement sequence below is the path from this paper to operating capability. Each step is independently valuable; each step compounds with the next.

Step 0 · Read	Read this paper end-to-end. Score your institution against the Maturity Ladder (preceding page). Identify the top three gaps. Cost: free.
Step 1 · 30-Minute Diagnostic	Eight-week Shared Responsibility Audit. Includes review of your most recent board pack relevant to this paper. Cost: free, by invitation, info@kieranupadrasta.com.
Step 2 · Two-Week Maturity Assessment	Structured evidence-grade review against the Maturity Ladder. Outputs: gap analysis, prioritised remediation plan, board-grade summary. Cost: fixed-fee, B2B Outside-IR35 engagement via Nova IT Consulting Ltd.
Step 3 · 90-Day Implementation Programme	calibrates your boundary documentation, deploys the IaC baseline, builds the attestation pipeline.. Co-delivered with the Partner Index named on the next page. Outputs: production capability, evidence pipeline, board attestation. Cost: programme-rate, fixed-fee or T&M.;
Step 4 · Annual Continuous Assurance Retainer	Quarterly board briefing, annual maturity re-assessment, regulatory advisory access. Annual retainer; pricing tier indicative on request.

Regulator-Defensibility Promise. Where this doctrine is implemented under our engagement, and a supervisor subsequently issues a finding on this control area, we will support remediation at no additional fee for the affected scope. This is the conviction discipline of the Doctrine Series.

V3.0 · L E N S E S

Partner Index, Sector, Insurance, M&A, Litigation, Sub-Committee

Doctrine that does not address the institutional reader is doctrine for the practitioner alone. The lenses below extend this paper's doctrine across the audiences who read it: procurement and ecosystem; sector-specific reading; insurance underwriter; M&A; acquirer; litigator and counsel; board sub-committee owner.

Lens	Reading
Partner Index (co-delivery ecosystem)	AWS / Azure / GCP (provider primary documentation) · Wiz / Orca / Prisma Cloud (CSPM) · CSA Cloud Controls Matrix v4 (calibration reference)
Sector-First Reading	SaaS-Native Companies — most exposed to cloud-native shared-responsibility drift.
Cyber-Insurance Position	Cyber underwriters now ask for the IaC baseline coverage % and the P99 misconfiguration MTTR. Both move premiums.
M&A Cyber Due Diligence	Acquirer must obtain CSPM findings backlog; misconfigurations older than 90 days are Day-One liabilities.
Litigation Defensibility	Cloud-misconfiguration breaches are extensively reported; standard of care is now objectively measurable against IaC and CSPM benchmarks.
Board Sub-Committee Owner	Technology Committee + Audit Committee

V3.0 · N A V I G A T I O N

How To Read This Paper · Engagement Specialisms · ROI Envelope

How to read this paper.

Audience	Recommended pages and reading time
Board Chair / SID	Read the Executive Thesis (page 3), the Maturity Ladder, and the Engagement Sequence. ~10 minutes.
Audit / Risk Chair	Add the Comparative Crosswalk and the Limitations / Scope page. ~20 minutes.
CISO / CRO	Read the Reference Architecture, the Reference Configuration, and the Per-Paper Substantive Uplifts. ~45 minutes.
Procurement Lead	Read the Engagement Sequence and the Partner Index. ~5 minutes.
External Counsel	Read the Litigation Defensibility lens, the Trust Choreography where applicable, and the Limitations page. ~10 minutes.
Insurance Broker	Read the Cyber-Insurance Position lens and the Maturity Ladder. ~5 minutes.
Regulator / Supervisor	Read the Methodology, the Primary Research Datasets, the Comparative Crosswalk, and the Peer-Review Notice. ~30 minutes.
Recruiter / Talent Partner	Read the cover, the Engagement Specialisms (below), and the Author Bio. ~3 minutes.

Engagement Specialisms.

DORA Compliance · NIS2 · AI Governance (ISO 42001) · Board Reporting · M&A; Cyber Due Diligence · Zero Trust Architecture · Post-Quantum Cryptography · Interim CISO · AI Security Assurance · OT/ICS Security · TIBER-EU · Adversary Emulation · Recoverability Mandate · Privileged Access Architecture · Phish-Resistant MFA · Cloud Security Posture · Identity Governance and Administration · Operational Resilience · Cyber Insurance Underwriting · Regulator-Grade Attestation · Big-4 Consulting (Deloitte, PwC, EY, KPMG) · Financial Services · Banking · Capital Markets · Insurance · Healthcare · Energy · Public Sector · Critical National Infrastructure · 80 Jurisdictions.

Indicative ROI envelope (this paper's doctrine).

Implementation cost (90-day programme): **£250k – £1.2m** depending on scope and institution scale. Loss-avoidance over 5 years (Cyentia IRIS-calibrated to sector loss-distribution): **£3m – £25m**. Implied **5-year ROI: 8× – 25×**. Insurance premium reduction (where applicable): typically **5–15%**. Regulatory-finding avoidance: not modelled but materially favourable. Numbers are illustrative ranges; institutional readers should re-anchor to their own loss data, exposure model, and impact-tolerance statements before relying on them for decision.

V3.0 · CLOSING

Closing Doctrine — Paper-Specific

"The provider owns the platform. You own the misconfiguration. The breach is yours."

Shared Responsibility Calibration Matrix™

This paper carries the framework named above. The framework is falsifiable, calibrated to NIST / ISO / regulatory anchors, and reproducible by any institution that adopts the maturity ladder set out earlier. It is the author's IP, contributed to the field on citation terms.

Series umbrella aphorism (across all 20 papers): If it cannot be evidenced, it cannot be defended.

TIER 1A · METHOD

Methodology, Evidence Standards, and Sample Construction

This paper is constructed under an institutional research register comparable to ECB, BoE, BIS, FSB, ENISA, and OECD working papers. Each claim is graded by evidence class and traceable to a primary source. The methodology is set out below so the reader, the auditor, and the regulator can replicate, falsify, or extend the analysis.

Evidence classification. Claims are tagged across four classes: (a) **Regulatory primary** — text drawn directly from DORA, NIS2, NIST SP 800-series, ISO/IEC, EU AI Act, FCA/PRA, SEC, NCSC, and ENISA publications; (b) **Industry empirical** — annualised threat-landscape data from Verizon DBIR, Mandiant M-Trends, IBM Cost of a Data Breach, and ENISA Threat Landscape; (c) **Practitioner observation** — composite patterns drawn from 27 years of practice across Big-4 consulting and regulated financial services, anonymised and labelled *ILLUSTRATIVE SCENARIO*; (d) **Doctrinal construction** — frameworks authored by the present writer, marked with the trademark symbol where introduced (e.g., Evidence Chain Model™, Decision Rights Architecture™).

Quantitative figures. All numerical examples are bracketed as ranges, not point predictions, and are intended as *order-of-magnitude* indicators appropriate for board-level risk reasoning. Worked examples are computed from publicly documented incident envelopes, regulatory penalty ceilings, and industry benchmark studies cited in the Primary Source Index. Specific-firm financials are never used.

Anonymisation protocol. Every case study is constructed as a composite from at least three distinct engagements or public incidents, with all identifying details — client name, jurisdiction-specific dates, regulator nomenclature, vendor identity, and dollar/euro/sterling figures — abstracted. Composites are labelled *ILLUSTRATIVE SCENARIO*; only events already in the public domain are labelled *PUBLIC INCIDENT*.

Reproducibility. Every doctrine, table, dialogue, control gate, and metric in this paper is reproducible from the Primary Source Index and the Evidence Artefact Index (Appendix A). A reviewer with access to the same regulatory text and industry empirical sources can independently verify each claim. Where the doctrine introduces a new framework, the falsifiability conditions are stated.

Standards comparable: BIS Working Paper format · ECB Occasional Paper register · FSB consultative report convention · ENISA Threat Landscape methodology · NIST IR documentation register · ISO/IEC TR research grade.

TIER 1A · C I T A T I O N S

Primary Source and Citation Index

Every empirical claim, regulatory anchor, and quantitative envelope in this paper traces to a primary source listed below. Citations follow the BIS / ECB working-paper register: regulatory primary first, industry empirical second, academic and practitioner-research third. The reader, auditor, or supervisor may verify each claim against the cited source without intermediation.

#	Source
1	Digital Operational Resilience Act (Regulation (EU) 2022/2554), Articles 5–26 (DORA).
2	Directive (EU) 2022/2555 on measures for a high common level of cybersecurity (NIS2).
3	European Banking Authority, Guidelines on ICT and security risk management (EBA/GL/2019/04).
4	European Central Bank, Cyber Resilience Oversight Expectations for Financial Market Infrastructures (2018, updated).
5	Bank of England / PRA, Supervisory Statement SS1/21: Operational Resilience.
6	Financial Conduct Authority, SYSC 13 — Operational Risk: Systems and Controls.
7	Verizon, Data Breach Investigations Report (DBIR), annual series 2020–2025.
8	Mandiant, M-Trends — Global Threat Report, annual series.
9	IBM Security & Ponemon Institute, Cost of a Data Breach Report, annual series.
10	ENISA, Threat Landscape — annual edition.
11	UK Government, Cyber Security Breaches Survey, annual series (DSIT).
12	Anderson, R. (2020). Security Engineering: A Guide to Building Dependable Distributed Systems, 3rd ed.
13	Schneier, B. (2018). Click Here to Kill Everybody: Security and Survival in a Hyper-Connected World.
14	Roberts, S. & Brown, R. (2017). Intelligence-Driven Incident Response, O'Reilly.
15	Cloud Security Alliance, Cloud Controls Matrix (CCM) v4.
16	Shared Responsibility Models — AWS, Microsoft Azure, Google Cloud (vendor primary docs).

Citation grade: every claim is sourced; no claim is asserted on the author's authority alone. Where a claim cannot be sourced to one of the above, it is removed before publication. This is the discipline that distinguishes flagship research from opinion.

TIER 1A · C R O S S W A L K

Comparative Regulatory Crosswalk

The doctrine in this paper does not exist in a single-regime vacuum. The same clause carries weight under DORA, NIS2, NIST CSF 2.0, ISO/IEC 27001:2022, and the relevant supervisory framework (FCA / SEC / BoE / ECB / NIST / CISA / sector-specific bodies). The crosswalk below is paper-specific — it maps the controls actually relevant to *this* paper's doctrine, not a generic spine. One control discharges multiple regulatory obligations simultaneously; that is the foundation of harmonised, audit-defensible governance.

Doctrine clause	DORA	NIS2	NIST CSF 2.0	ISO 27001:2022	CSA / NIST / Provide
Shared-responsibility calibration	Art. 6(3)	Art. 21(2)(a)	GV.OC-03	A.5.10	CSA CCM v4
IaC baseline	Art. 9(3)	Art. 21(2)(i)	PR.PS-01	A.8.4	NIST 800-53 CM-2
CSPM continuous monitoring	Art. 10(1)	Art. 21(2)(b)	DE.CM-01	A.8.16	CSA CCM CCC-04
Misconfiguration MTTR	Art. 12(2)	Art. 21(2)(b)	RS.MI-01	A.5.27	CIS Cloud Bench
Public-S3 prevention	Art. 9(3)	Art. 21(2)(i)	PR.DS-02	A.8.4	NIST 800-53 SC-7
Encryption + key rotation	Art. 9(5)	Art. 21(2)(i)	PR.DS-01	A.8.24	NIST SP 800-57
Cloud IAM least privilege	Art. 9(4)	Art. 21(2)(i)	PR.AA-05	A.5.15	NIST 800-53 AC-6

Crosswalk discipline. The crosswalk is not decorative. It is the evidence that the institution can answer a single supervisory question — "show me the control" — across *every* regime simultaneously, from one record. Institutions that maintain regime-by-regime evidence end up rebuilding the same control trail multiple times, incurring the regulatory contagion penalty: a finding under one regime cascades into evidence demands under all the others.

> ***"One control. One evidence chain. Many regulators. That is harmonised governance."***

TIER 1A · R E V I E W

Peer Review and Editorial Standards Notice

This paper has been prepared under an editorial register designed to match the transparency expectations of institutional research bodies. The process below applies to every paper in the Doctrine Series and is set out so the reader, the regulator, and any future challenger can hold the work to the same standard.

Stage	Description
1. Doctrinal drafting	Author drafts the doctrine clause, cites primary regulatory and industry sources, and tags every quantitative claim to a published envelope (DBIR, M-Trends, IBM/Ponemon, ENISA Threat Landscape, Cyentia IRIS). No claim is published on author authority alone.
2. Independent technical review	A senior practitioner with no commercial interest in the doctrine reviews mechanism, worked example, and counter-positions for technical defensibility. Review notes are retained for three years to support post-publication scrutiny.
3. Regulatory anchor verification	Every regulatory citation is verified against the official text (Eur-Lex, NIST CSRC, ISO online, ECB / BoE / FCA register, SEC EDGAR). Article numbers and clause references are checked at the date of build.
4. Anonymisation audit	Every case study is reviewed against the anonymisation protocol: at least three source engagements, no identifying client / vendor / jurisdiction-specific marker. Composites labelled *ILLUSTRATIVE SCENARIO*; public events labelled *PUBLIC INCIDENT*.
5. Conflict of interest declaration	The author declares no commercial financial relationship with vendors named or implied. Where a regulator, framework, or methodology is cited, the citation is to the publicly available text, not to a private engagement.
6. Reproducibility statement	Every doctrine, table, dialogue, and metric in this paper is reproducible from the Primary Source Index (preceding page) and the Evidence Artefact Index (Appendix A). Falsifiability conditions for novel doctrine are stated in the mechanism section.

Editorial standard: If it cannot be evidenced, it cannot be defended. This paper is constructed so that every assertion can be traced, verified, and — if necessary — falsified by an independent reviewer with access to the same primary sources. That is the difference between flagship research and marketing literature.

TIER 1A · G L O S S A R Y

Glossary of Institutional Terms

Definitions below are paper-specific. Each glossary captures the terms anchored or introduced by *this* paper's doctrine — not a generic boilerplate. Where a term is the author's framework, it is marked with ™. Where a term is drawn from a regulatory or standards body, the source is named.

Term	Definition
Shared Responsibility Calibration Matrix™	Author framework: formal calibration of cloud provider vs customer responsibility.
CSPM	Cloud Security Posture Management — continuous detection and remediation of cloud misconfigurations.
IaC (Infrastructure as Code)	Codified cloud configuration; enables version control, review, and drift detection.
Cloud Controls Matrix	CSA framework providing detailed cloud security control specifications, mapped to industry standards.
Misconfiguration Drift	Deviation of cloud configuration from declared baseline; leading indicator of CSPM finding.
Provider Boundary	The line between cloud provider responsibility and customer responsibility, varying by service type (IaaS / PaaS / SaaS).
Public S3 Disaster	Most-common cloud misconfiguration class; publicly accessible storage exposing sensitive data.

TIER 1A · S C O P E

Limitations, Scope, and Defensibility Caveats

Institutional research must be explicit about what it claims, what it does not claim, and where it stops. The boundaries below are stated so the reader can apply the doctrine within its proper register and so the supervisor can hold the work to the limits the author has set.

Jurisdictional scope. Primary regulatory anchoring is the European Union (DORA, NIS2, EU AI Act), the United Kingdom (FCA, PRA, NCSC), and the United States (SEC, OCC, NIST). Non-EEA / non-UK / non-US jurisdictions are referenced where directly relevant; readers operating elsewhere should map the doctrine to their local regime via the Comparative Crosswalk page.

Sectoral scope. The Doctrine Series is calibrated for regulated and systemically important sectors — banking, capital markets, insurance, infrastructure, energy, transport, healthcare, and government / critical national infrastructure. Material remains useful for unregulated sectors but the regulatory consequence statements may not apply.

Quantitative figures are illustrative. Every numerical example is presented as a range or order-of-magnitude indicator drawn from publicly cited industry envelopes (DBIR, IBM Cost of a Data Breach, Mandiant M-Trends, ENISA, Cyentia IRIS). They are *not* point predictions for any specific institution. Institutional readers should re-anchor figures to their own loss data, exposure model, and impact-tolerance statements before relying on them for decision.

Temporal scope. Regulatory citations are correct at date of build (see the cover meta block). Where a regulation is in transition (e.g., NIS2 transposition, EU AI Act implementing acts, SEC enforcement guidance), the reader should verify the latest text. The doctrine itself is more durable than any single regulatory cycle; the underlying mechanism rarely changes.

No legal advice. Nothing in this paper constitutes legal, regulatory, accounting, or investment advice for any specific institution. The doctrine is a research and policy contribution. Application to a specific institution requires bespoke legal, regulatory, and risk-engineering analysis under privilege.

No vendor endorsement. Where a vendor product, framework, or technology category is referenced, the reference is descriptive — not an endorsement, recommendation, or commercial relationship disclosure. The author declares no commercial relationship with vendors named.

Update cadence. The Doctrine Series is reviewed at least annually and re-anchored to the latest regulatory and threat-landscape evidence. Material changes are version-stamped (see the cover meta block).

Defensibility test: a supervisor, an auditor, or a litigator should be able to read this paper and identify, without ambiguity, what the author claims, what evidence supports each claim, and where the claims stop. That is the institutional standard.

THE CLOSING DOCTRINE

The doctrine in one line.

The cloud provider sells a configurable platform; the customer's board accepts the consequence of every configuration choice. The institution that operationalises Configuration-as-Control closes the dominant cloud-incident pathway and produces, as a by-product, a continuous regulatory evidence stream. The institution that does not is paying for the platform and abdicating the configuration — and the audit will not save it.

> ***"The provider operates the platform. The customer signs the configuration. The board pays for what they sign — including the gaps."***

Issued by: Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Affiliations: Schiphol University · Imperials · UCL · ISACA London (Platinum) · (ISC)² London (Gold) · PRMIA · ISF.
Contact: info@kieranupadrasta.com · www.kie.ie
Series: THE DOCTRINE SERIES — Volume I — Twenty Aphorisms for the Modern CISO

CLOSING APHORISM

"The provider operates the platform. The customer signs the configuration. The board pays for what they sign — including the gaps."

This volume is one of twenty in **THE DOCTRINE SERIES: Volume I — Twenty Aphorisms for the Modern CISO**. Each paper is constructed to be auditor-reproducible, board-survivable, and regulator-defensible — the operating canon of the modern Chief Information Security Officer under DORA, NIS2, the EU AI Act, and the converging UK / US regulatory regimes.

If it cannot be evidenced, it cannot be defended.

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

Cybersecurity Authority · Board Advisor · Interim CISO

www.kie.ie · info@kieranupadrasta.com · linkedin.com/in/kieranupadrasta