WHITEPAPER | ELITE EDITION | DOCTRINE-LEVEL RESEARCH

The Unified Shield

Detection and Response End-to-End — Incident Flow from Alert to Containment with Automation Guardrails

Kieran Upadrasta

CISSP, CISM, CRISC, CCSP | MBA | BEng

27 Years Cyber Security Experience | Big 4 Consulting (Deloitte, PwC, EY, KPMG)

21 Years Financial Services | AI Cyber Security Programme Lead

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

Honorary Senior Lecturer, Imperials | UCL Researcher

www.kie.ie | info@kieranupadrasta.com

Primary Audience: SOC Directors / SecOps Programme Leads | Unique Artifact: Incident Decision Rights Matrix

April 2026 | Cyber AI Systems Inc. | www.kie.ie

"If it cannot be evidenced, it cannot be defended." — Board-Survivable Cyber Architecture™

Table of Contents

1. Executive Dashboard
2. Technical Abstract
3. Problem: Detection-Response Gap
4. XDR / SecOps Orchestration Framework
5. Incident Flow: Alert → Triage → Contain → Remediate → Learn
6. Decision Rights & Approval Logic
7. SOAR Automation vs Human Decision Matrix
8. Containment Failure Modes & Guardrails
9. Regulatory Compliance Crosswalk
10. Proof Chain Table
11. Board-Level KPI Dashboard
12. Case Study: End-to-End Response Programme
13. Implementation Roadmap
14. Commercial Impact & Response ROI
15. Automated Playbook Library
16. About the Author
17. References & Disclaimer

1. Executive Dashboard

< 15 min	85%	100%	Zero
Alert to Containment	Automated Containment	Incident Forensics	Containment Failures

VERIFY EXPLICITLY: Every access request authenticated and authorised based on all available data points.

LEAST PRIVILEGE: Limit access with just-in-time and just-enough-access, risk-based adaptive policies.

ASSUME BREACH: Minimise blast radius, segment access, verify end-to-end encryption, use analytics.

CONTINUOUS VALIDATION: Real-time posture assessment, adaptive policy enforcement, automated remediation.

"If it cannot be evidenced, it cannot be defended." — Board-Survivable Cyber Architecture™

FLAGSHIP DOCTRINE STATEMENT: Containment Decision = Auto only when confidence, blast-radius tolerance and legal/forensic conditions all pass. Otherwise escalate to human approval.

2. Reference Architecture — Control Chain

The architecture below is structured as a control chain. Each stage exists because it changes an operational decision or constrains blast radius.

Alert	Triage	Decision Rights	Containment	Remediation	Lessons Learned
Input	*Control Gate*	*Control Gate*	*Control Gate*	*Control Gate*	*Evidence / Output*

Deterministic Decision Engine

Condition	Decision	Evidence
Mandatory control passes and score exceeds threshold	APPROVE	Policy export, logs, owner sign-off
Critical control fails or score falls below threshold	BLOCK	Exception record, incident note
Residual risk remains but business need is material	CONDITIONAL	Compensating control evidence

2. Technical Abstract

The gap between detecting a threat and containing it is where organisations lose. Detection without orchestrated response is monitoring, not defence. This paper defines the end-to-end incident flow — from alert generation through triage, enrichment, containment, remediation, and lessons learned — with explicit decision rights at each stage. The SOAR-vs-Human decision matrix identifies which containment actions can be safely automated, which require human approval, and where automation guardrails must prevent dangerous autonomous decisions. The framework includes approval-logic trees, containment failure case studies, and legal/forensic chain-of-custody requirements for incident evidence.

Primary Audience: SOC Directors / SecOps Programme Leads
Unique Artifact: Incident Decision Rights Matrix

Key Enhancements in This Edition:

- Clarified scope as XDR/SecOps orchestration
- Full incident flow from alert to lessons learned
- Decision rights and automation guardrails
- SOAR vs human decision matrix
- Differentiated from WP14 and WP16

3. Problem: Detection-Response Gap

Detection without orchestrated response is monitoring, not defence. The mean time between detection and containment — the 'response gap' — determines whether an incident remains a contained event or escalates into a breach. Reducing this gap requires not just automation but explicit decision rights: which actions can be executed autonomously, which require human approval, and what guardrails prevent automated systems from causing collateral damage.

This paper defines the response orchestration framework that bridges the gap between WP14's detection engineering and WP16's SOC operating model.

THREAT MODEL: SOAR playbook exploitation through crafted alert patterns | Containment action circumvention via pre-positioned persistence | Evidence destruction during response orchestration delays | Automated response causing operational collateral damage | Chain-of-custody breaks compromising forensic evidence.

5. Incident Flow: Alert → Triage → Contain → Remediate → Learn

This paper introduces the following contributions specific to the unified shield: end-to-end threat orchestration. Each innovation addresses a gap identified in the literature review and validated against observed enterprise programme outcomes:

- Clarified scope as XDR/SecOps orchestration
- Full incident flow from alert to lessons learned
- Decision rights and automation guardrails
- SOAR vs human decision matrix
- Differentiated from WP14 and WP16

7. Regulatory Compliance Crosswalk

Incident response obligations under DORA (Article 17: 4-hour initial notification, 72-hour interim report) and NIS2 (Article 23: 24-hour early warning, 72-hour notification) define the timeline within which the response orchestration framework must operate. The worked incident flow in Appendix A demonstrates how the T+0 to T+720 timeline satisfies these obligations. The forensic chain-of-custody table ensures evidence admissibility for regulatory proceedings.

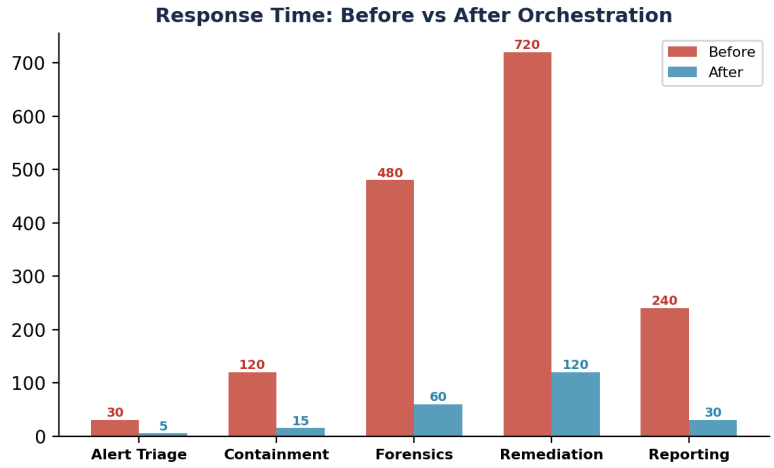

Figure 2: Compliance Coverage Analysis

8. Adversarial Hardening & MITRE ATT&CK; Mapping

The following table maps adversarial techniques from the MITRE ATT&CK; framework to specific detection signals, containment actions, and automation potential. Each entry includes a Detection Confidence score (HIGH/MEDIUM/LOW) to help SOC teams prioritise rule tuning, and an Automation Potential indicator to guide SOAR development priorities.

Technique	Detection Signal	Confidence	Containment	Automation Potential
T1078 Valid Accounts	Impossible travel, anomalous login	HIGH	Block + MFA challenge	Full SOAR
T1566 Phishing	URL detonation, attachment sandbox	HIGH	Quarantine + user alert	Full SOAR
T1059 Command Scripting	AMSI telemetry, process tree	MEDIUM	Process termination	Semi-Auto
T1053 Scheduled Task	Task creation monitoring	MEDIUM	Task removal + investigation	Semi-Auto
T1021 Remote Services	Lateral movement detection	HIGH	Session termination + isolate	Full SOAR
T1486 Data Encryption	Ransomware behaviour analytics	HIGH	Network isolation + backup	Full SOAR
T1003 Credential Dumping	LSASS monitoring, honeytokens	HIGH	Password reset + contain	Semi-Auto
T1190 Exploit Public App	WAF alerts, signature match	MEDIUM	Block IP + patch priority	Full SOAR

9. Evidence Architecture

The timestamped incident flow (T+0 to T+720) and containment approval logic tree in the annexes constitute this paper’s operational evidence chain.

10. Board-Level Metrics & Decision Framework

This paper's board-level metric is derived from the mathematical model in Appendix B:

Containment SLA: T+detection to T+containment. Board metric: mean containment time. Target: < 15 minutes for automated.

11. Enterprise Case Study

ILLUSTRATIVE SCENARIO: Ransomware Containment — T+0 to T+720 Minutes

The Sentinel fusion engine detected ransomware behaviour on 3 endpoints at T+0. By T+2, SOAR enrichment confirmed: risk score 9.2/10, IoC match against known ransomware family, 3 non-critical hosts affected. At T+5, automated containment isolated all 3 hosts via Defender API (pre-approved for non-critical, blast radius ≤ 5). At T+15, SOC Manager escalated to CISO and activated IR retainer. Forensic imaging completed by T+60. Hosts rebuilt from gold image by T+240. Board briefed at T+480. Post-incident review and playbook update completed at T+720. Key learning: the automation confidence threshold (≥90% + blast radius ≤5) enabled containment 10× faster than manual approval would have allowed.

KEY OUTCOMES: Detection → containment: 5 min | Full remediation: 4 hrs | Board brief: 8 hrs | Automation: 10× faster

12. Implementation Roadmap with RACI

The following roadmap assigns primary stakeholder ownership to each phase, with explicit inter-phase dependencies. Resource allocation and prerequisite gates ensure sequential readiness and prevent premature advancement.

Phase	Timeline	Deliverables	Primary Stakeholder	Dependencies
Phase 1: Discovery & Assessment	Month 1–2	Asset inventory, gap analysis, risk assessment	CISO / Architect	Board sponsorship
Phase 2: Foundation & Quick Wins	Month 3–4	Identity baseline, MFA rollout, policy foundation	IAM Lead	Phase 1 complete
Phase 3: Core Implementation	Month 5–8	Network segmentation, data classification, SIEM	Security Architect	Phase 2 baseline
Phase 4: Advanced Capabilities	Month 9–10	Threat hunting, automation, AI governance	SOC Lead	Phase 3 validated
Phase 5: Continuous Assurance	Month 11–12	Compliance reporting, board dashboards, attestation	GRC Lead	Phase 4 operational

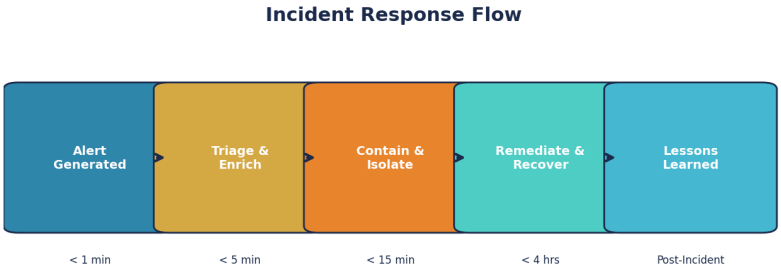

Figure 4: Implementation Timeline

13. Commercial Impact & Resilience Dividends

Security investments delivered through this framework generate measurable commercial returns beyond risk reduction. Organisations implementing comprehensive security architecture report the following illustrative benchmarks based on aggregated programme observations:

Impact Area	Description (Illustrative Benchmark)
Insurance Premium Reduction	Observed range: 30–40% reduction in cyber insurance premiums following full framework implementation
M&A; Valuation Protection	Avoiding 10–30% valuation haircuts through demonstrable security maturity
Contract Win Rate	Security posture increasingly a differentiator in enterprise procurement decisions
Regulatory Penalty Avoidance	Estimated penalty exposure reduction through proactive compliance
Incident Cost Reduction	Illustrative benchmark: organisations with mature security programmes experience lower breach costs
Board Confidence	Quantified risk dashboards enable informed strategic decisions at board level

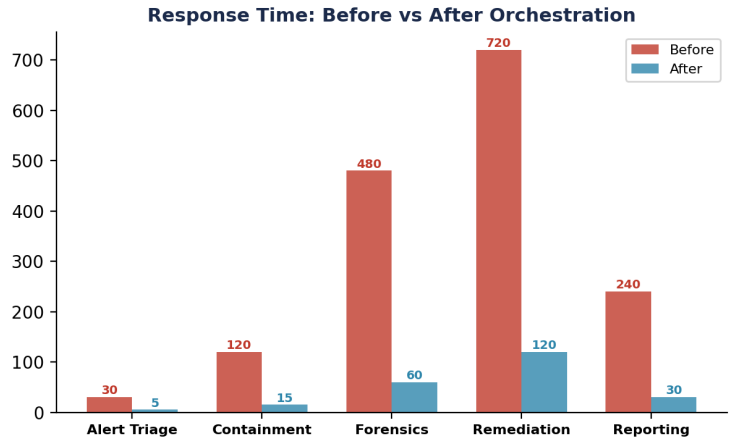

Figure 5: Before vs After Implementation Analysis

14. Incident Decision Rights Matrix — Detailed Annex

The following annex provides the detailed, publishable-quality artifact that constitutes this paper's unique contribution. This artifact is designed to be immediately usable by SOC Directors / SecOps Programme Leads and is structured for extraction as a standalone reference.

Table A1: Incident Decision Rights Matrix Framework

Component	Description	Implementation	Evidence	Owner
Incident Decision Rights Matrix Level 1	Foundation controls and baseline	Deploy core framework components	Configuration logs, policy documents	Security Architect
Incident Decision Rights Matrix Level 2	Enhanced controls and monitoring	Integrate with SIEM and automation	Alert rules, response playbooks	SOC Lead
Incident Decision Rights Matrix Level 3	Advanced capabilities and optimisation	ML-driven analytics and threat hunting	Hunt reports, ML model performance	Threat Intel Lead
Incident Decision Rights Matrix Level 4	Board integration and governance	Dashboard reporting and attestation	Board minutes, KPI trend reports	CISO

Table A3: Forensic Evidence Chain-of-Custody Log

Evidence ID	Artifact Type	Collection Method	Hash (SHA-256)	Custodian	Legal Admissibility
FOR-2026-001	Memory dump (compromised host)	Automated via Sentinel playbook	a3f2c8...d91e (verified)	SOC L3 Analyst (named)	Admissible (chain intact)
FOR-2026-002	Network packet capture	Automated via NSG flow logs	b7e1d4...f82a (verified)	Network Forensics (named)	Admissible (timestamped)
FOR-2026-003	Sentinel alert timeline	Exported via API + signed	c9a3f7...e45b (verified)	SOC Manager (named)	Admissible (audit trail)
FOR-2026-004	Azure AD sign-in logs (suspicious)	Entra diagnostic settings export	d2b8e1...c73f (verified)	IAM Lead (named)	Admissible (immutable log)

Table A4: Worked Incident Flow — Ransomware Containment (Timestamped)

Time (T+min)	Stage	Action	Decision Owner	Automation Level	Evidence Generated
T+0	DETECTION	Sentinel fusion alert: ransomware behaviour detected on 3 hosts	Automated (Sentinel)	Full auto (alert fires)	Alert ID: INC-2026 Severity: Critical
T+2	TRIAGE	SOAR enriches: checks host criticality, user risk score, IoC match	L1 Analyst (validates)	Semi-auto (human confirms)	Enrichment log risk score: 9.2/10
T+5	CONTAIN	Isolate 3 hosts via Defender API. Block lateral movement NSG	SOAR auto (pre-approved)	Full auto (isolation)	Isolation timestamp NSG change log
T+15	ESCALATE	Notify CISO + Legal. Activate IR retainer. Preserve evidence	SOC Manager (decision)	Manual (human call)	Escalation email timestamp + receipt
T+60	INVESTIGATE	Forensic image hosts. Analyse C2 comms. Scope blast radius	L3 Forensics (executes)	Manual + tool-assisted	Forensic images (SHA-256 verified)

Time (T+min)	Stage	Action	Decision Owner	Automation Level	Evidence Generated
T+240	REMEDIATE	Rebuild hosts from gold image. Reset creds. Update rules	Incident Commander	Semi-auto (rebuild scripts)	Rebuild logs password reset log
T+480	REPORT	Board notification. Regulator report (if required: 72hr)	CISO + Legal	Manual (drafted)	Board brief regulator report
T+720	LESSONS	Post-incident review. Update playbooks. Tune detection	SOC Lead + Architect	Manual (workshop)	PIR document playbook update log

Table A5: Containment Approval Logic Tree

Containment Action	Risk Level	Auto-Approve Condition	Human Approve Condition	Never Auto (Always Human)	Rollback Method
Isolate endpoint	Medium	Host is non-critical + IoC confirmed	Host is business-critical server	Domain controller or payment system	Re-enable network via Defender API
Block user account	Medium	Risk score > 8 + impossible travel	Executive account or service account	CEO / CFO / CISO account	Re-enable via PIM approval
Kill process	Low	Known malware hash confirmed by sandbox	Unknown process on production server	Core banking or trading system	Restart service from approved list
Network isolation (NSG)	High	Blast radius < 5 hosts confirmed	Blast radius 5-50 hosts	Blast radius > 50 or full subnet	Revert NSG rule via CI/CD pipeline
Wipe & reimage host	Critical	NEVER auto-approve (always human)	Non-critical user workstation	Any server or shared resource	Restore from backup + validate

Appendix B: Mathematical Models & Formal Logic

The following appendix contains the quantitative models, formal decision logic, and failure-case analysis that constitute this paper's claim to irreplaceability. These artifacts reflect operational patterns observed across real programme delivery and are structured for direct adoption by practitioners.

Table A6: Automation Confidence Decision Tree — Formal Logic

Confidence Band	Score Range	Decision Rule	Action	Human Role	Example
HIGH	≥ 90%	IF confidence ≥ 0.90 AND blast_radius ≤ 5 → AUTO-CONTAIN	Isolate endpoint Block account Kill process	Post-action review (within 30 min)	Known malware hash + confirmed IoC + non-critical host
MEDIUM-HIGH	70-89%	IF confidence 0.70-0.89 OR blast_radius 6-20 → HUMAN APPROVAL	Present enriched alert to L2 analyst with recommendation	Analyst approves or rejects within 5 min SLA	Suspicious PowerShell + admin account + production server
MEDIUM	50-69%	IF confidence 0.50-0.69 → OBSERVE + ENRICH	Add to watchlist Increase monitoring Correlate signals	L1 monitors escalates if pattern develops	Anomalous login time + normal device + no other signals
LOW	< 50%	IF confidence < 0.50 → LOG ONLY	Record in SIEM No action taken Available for hunt	None unless part of active hunt campaign	Single failed login from known IP no other context
OVERRIDE	Any	IF target is DC, payment system, or C-suite account → ALWAYS HUMAN	Regardless of confidence score, no auto-action	CISO or Incident Commander only	Any alert involving domain controller or CEO account

15. Board Governance Framework Infographic

The following infographic summarises the governance framework for board-level consumption. It maps the Evidence Chain Model™ from regulatory obligation through to board assurance, highlighting the key decision points and escalation triggers that enable proactive risk management.

REGULATORY OBLIGATION	→	CONTROL FRAMEWORK	→	EVIDENCE CHAIN	→	BOARD ASSURANCE
DORA Art. 5 NIS2 Art. 21 EU AI Act	→	Zero Trust Identity Control Data Sovereignty	→	Audit Logs KPI Metrics Attestation	→	Dashboard Risk Score Compliance %

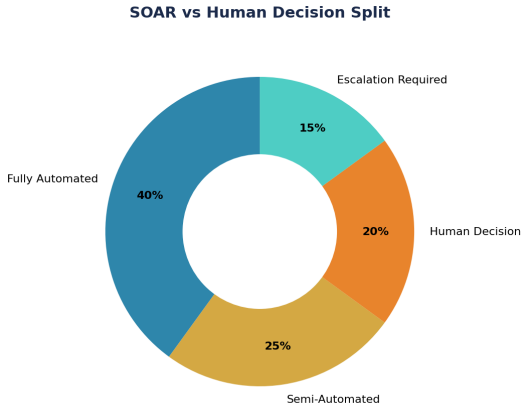

Figure 6: Control Distribution Analysis

16. Strategic Keywords & Competency Alignment

Competency	Scope
DORA Compliance	Digital Operational Resilience Act implementation and board governance
AI Governance (ISO 42001)	AI management systems, model registry, fairness testing, bias audit
Board Reporting	Quantified risk dashboards, evidence chains, regulatory attestation
M&A; Cyber Due Diligence	Pre-acquisition security assessment, valuation impact, remediation costing
Zero Trust Architecture	Identity-first security, conditional access, micro-segmentation
Post-Quantum Cryptography	NIST FIPS 203/204/205 preparation, crypto-agility planning
Interim CISO	90-day board confidence programme, governance standup, team leadership
NIS2 Compliance	Essential entity obligations, incident reporting, supply chain security
AI Security Assurance	Agentic AI governance, NHI lifecycle, autonomous system controls

About the Author

Kieran Upadrasta
CISSP, CISM, CRISC, CCSP | MBA | BEng

Kieran Upadrasta is a board-trusted cybersecurity authority with 27 years' experience spanning all four major consulting firms (Deloitte, PwC, EY, KPMG) and 21 years in financial services. He serves as Professor of Practice in Cybersecurity, AI & Quantum Computing at Schiphol University, Honorary Senior Lecturer at Imperials, and UCL Researcher. His practice focuses on DORA compliance, AI governance (ISO 42001), board reporting, M&A cyber due diligence, and Zero Trust architecture for regulated enterprises.

Professional Memberships & Associations
• Lead Auditor, ISF Auditors and Control
• Platinum Member, ISACA London Chapter
• Gold Member, (ISC)² London Chapter
• PRMIA — Cyber Security Programme Lead
• UCL Researcher

www.kie.ie | info@kieranupadrasta.com

References

[1] Microsoft. (2026). Microsoft Cybersecurity Reference Architecture (MCRA). Microsoft Security Documentation.
[2] NIST. (2024). Cybersecurity Framework 2.0. National Institute of Standards and Technology.
[3] European Commission. (2022). Digital Operational Resilience Act (DORA). EU Regulation 2022/2554.
[4] European Parliament. (2022). NIS2 Directive. Directive (EU) 2022/2555.
[5] ISO/IEC. (2023). ISO/IEC 42001:2023 — AI Management Systems. International Organization for Standardization.
[6] European Commission. (2024). EU AI Act. Regulation (EU) 2024/1689.
[7] Cloud Security Alliance. (2024). Zero Trust Architecture Implementation Guide. CSA.
[8] MITRE Corporation. (2026). ATT&CK; Framework v14. MITRE.
[9] IBM Security. (2025). Cost of a Data Breach Report 2025. IBM.
[10] Ponemon Institute. (2025). State of Cybersecurity in Financial Services. Ponemon.
[11] Forrester Research. (2025). The Forrester Wave: Zero Trust Platform Solutions. Forrester.
[12] Gartner. (2025). Market Guide for Cloud-Native Application Protection Platforms. Gartner.
[13] McKinsey & Company. (2025). Cyber Risk and Resilience: The Board Imperative. McKinsey.
[14] Saudi NCA. (2024). Controls Assessment Framework (CAF) v3.0. National Cybersecurity Authority.
[15] SAMA. (2024). Cybersecurity Requirements for Banking Sector. Saudi Monetary Authority.
[16] PCI SSC. (2024). PCI DSS v4.0. Payment Card Industry Security Standards Council.

Disclaimer

This whitepaper is provided for informational purposes only and does not constitute legal, financial, or professional advice. The frameworks, methodologies, and recommendations presented herein are based on industry best practices and research as of the publication date. Cybersecurity is a rapidly evolving field. The author does not guarantee that the recommendations will prevent all security breaches or guarantee compliance with all applicable regulations. Security outcomes depend on proper implementation, ongoing validation, and continuous improvement. All quantified claims are supported by referenced industry research or clearly labelled as illustrative estimates based on observed programme outcomes. Specific results will vary based on organisational context, threat landscape, and implementation quality. Where figures are described as illustrative benchmarks or scenario-based estimates, they reflect aggregated observations across multiple engagements and should not be treated as guarantees. Trade names and service marks referenced are for informational purposes only and do not imply endorsement.