



ΠΑΝΕΠΙΣΤΗΜΙΟ
ΑΙΓΑΙΟΥ

ΠΟΛΥΤΕΧΝΙΚΗ ΣΧΟΛΗ
ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ ΠΛΗΡΟΦΟΡΙΑΚΩΝ
ΚΑΙ ΕΠΙΚΟΙΝΩΝΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ



ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ
ΠΟΛΥΤΕΧΝΕΙΟ

ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ
ΜΗΧΑΝΙΚΩΝ ΚΑΙ
ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ

ΔΙΔΡΥΜΑΤΙΚΟ
ΠΡΟΓΡΑΜΜΑ
ΜΕΤΑΠΤΥΧΙΑΚΩΝ
ΣΠΟΥΔΩΝ (Δ.Π.Μ.Σ.)

ΜΒΑ ΣΤΗΝ ΨΗΦΙΑΚΗ
ΚΑΙΝΟΤΟΜΙΑ ΚΑΙ
ΝΕΟΦΥΗ
ΕΠΙΧΕΙΡΗΜΑΤΙΚΟΤΗΤΑ

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ
ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ
ΗΛΕΚΤΡΟΛΟΓΩΝ
ΜΗΧΑΝΙΚΩΝ ΚΑΙ
ΜΗΧΑΝΙΚΩΝ
ΥΠΟΛΟΓΙΣΤΩΝ

ΠΑΝΕΠΙΣΤΗΜΙΟ
ΑΙΓΑΙΟΥ
ΠΟΛΥΤΕΧΝΙΚΗ ΣΧΟΛΗ

ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ
ΠΛΗΡΟΦΟΡΙΑΚΩΝ ΚΑΙ
ΕΠΙΚΟΙΝΩΝΙΑΚΩΝ
ΣΥΣΤΗΜΑΤΩΝ



Quantified OT Supply Chain Assurance
A Continuous C-SCRM Framework
Integrating SBOM Evidence, Provenance
Signals, and FAIR-Inspired Risk
Quantification, with a Counterfeit-
Component Case Study from Civil Aviation

[GAVRIILIDIS NIKOLAOS]

[725002, icsdm725002@icsd.aegean.gr]

Επιβλέπων: Kontos Gerasimos

Η Διπλωματική Εργασία παρουσιάστηκε ενώπιον
του Διδακτικού Προσωπικού του Πανεπιστημίου Αιγαίου & του Εθνικού Μετσόβιου
Πολυτεχνείου
Σε Μερική Εκπλήρωση των απαιτήσεων για την απόκτηση του μεταπτυχιακού διπλώματος
ειδίκευσης «MBA στην Ψηφιακή Καινοτομία και Νεοφυή Επιχειρηματικότητα»

Η ΤΡΙΜΕΛΗΣ ΕΠΙΤΡΟΠΗ ΔΙΔΑΣΚΟΝΤΩΝ
ΕΠΙΚΥΡΩΝΕΙ ΤΗ ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

ΓΑΒΡΙΗΛΙΔΗΣ ΝΙΚΟΛΑΟΣ

Γεράσιμος Κοντός

Βαθμίδα

Τμήμα Μηχανικών Πληροφοριακών
και Επικοινωνιακών Συστημάτων

ΟΝΟΜΑΤΕΠΩΝΥΜΟ, Μέλος

Χαραλαμπίδης Ιωάννης

Καθηγητής

ΟΝΟΜΑΤΕΠΩΝΥΜΟ, Μέλος

Χαράλαμπος Αλεξόπουλος

Επίκουρος Καθηγητής

© 2026

Με επιφύλαξη παντός δικαιώματος.

ΟΝΟΜΑΤΕΠΩΝΥΜΟ

Πανεπιστήμιο Αιγαίου, Εθνικό Μετσόβιο Πολυτεχνείο

Abstract

The problem. Point-in-time supplier certification is inadequate especially for critical infrastructures, operational technology (OT) and civil aviation, where a certificate proves the supplier's state on that day of an audit and says nothing about the months that follow. More specifically, the 2023 AOG Technics affair, in which more than sixty thousand aircraft engine parts reached the global fleet behind forged airworthiness certificates, is that inadequacy realised at a great scale.

Aim and method. The thesis develops a continuous cyber supply chain risk management capability that treats supplier assurance as a perishable asset (i.e., a resource that loses its value) and then it

- prices the cost of losing it,
- and assembles the result into an operating model an organisation can adopt.

What was built. A 4-layer artifacts and one integrating artifact. The 1st layer is a unified control ontology that maps five overlapping rulebooks, NIS2, the EU Cyber Resilience Act, EASA Part-IS, IEC 62443, and NIST SP 800-161r1, together with DO-326A, onto a single traceable control set, paired with an out of band telemetry abstraction that reconciles continuous monitoring with OT isolation.

The 2nd is a 4-stage maturity model and a trust decay metric that turns silence between evidence arrivals into a measurable risk signal.

The 3rd is a FAIR inspired compound frequency and severity model, informed by the FABRICS framework of Slapnicar and Joshi (2026) and extended with a safety critical loss term, expressing risk as a monetary loss distribution a chief financial officer can act on. It takes the scenario based financial framing and the explicit treatment of epistemic uncertainty from FABRICS and does not implement its Bayesian fault tree or its expert elicitation protocol.

Against paper certification in the AOG Technics case, continuous cryptographic provenance prizes the fourth.

The 5th binds the four into one governance function, with phases and gates, named ownership, a single source of truth,, and a maturity path from paper certification to a commercialized capability.

Key results. On the neutral OT scenario, the model returns a median annual loss of zero against an expected annual loss of about 1.3 million dollars, with cyber value at risk of about 6.3 million at the 95th percentile and about 13.6 million at the 99th. Fewer than one percent of simulated years involves a fatality, yet the safety loss term lifts the 99th percentile by roughly four million dollars while barely moving the mean, which is the analytical reason a safety critical adaptation is needed. Continuous cryptographic provenance shifts the whole loss exceedance curve left, cutting the 99th percentile tail by about 7.8 million dollars and returning 33.5 percent a year against its cost. Sensitivity analysis binds both results: the 99th percentile Shifts by a factor of 1.7 across a plausible range for the safety exposure share while the mean barely moves, and the provenance investment reverses below a frequency reduction of about 55 percent or above an annual cost In the range of a million dollars.

Contribution and limits. The contribution isn't a new result in any single field but the integration of four mature fields within one continuous assurance operating model, carried along a single traceable chain from a verified piece of evidence, through a decaying trust score, to a monetary figure and a funded control decision. In practical terms, the thesis shows how an organisation can move from trusting paperwork to operating a continuous assurance function: one that verifies evidence, watches trust decay, prices the risk, and supports a funded control decision. The thesis delivers designs, demonstrations, and defensible methods. It doesn't provide field validated results, calibrated incident or fatality statistics, or an audited return on investment.

Keywords: cyber supply chain risk management; operational technology; software bill of materials; cryptographic provenance; cyber risk quantification; FAIR; cyber value at risk; continuous assurance; counterfeit components; civil aviation; design science research

Περίληψη

Το πρόβλημα. Η πιστοποίηση ενός προμηθευτή σε μία χρονική στιγμή είναι δομικά ανεπαρκής για τις κρίσιμες υποδομές και ποσό μάλλον την πολιτική αεροπορία. Ένα πιστοποιητικό αποδεικνύει την κατάσταση του προμηθευτή την ημέρα του ελέγχου και δεν λέει τίποτα για τους μήνες που ακολουθούν. Η υπόθεση AOG Technics του 2023, όπου περισσότερα από εξήντα χιλιάδες εξαρτήματα αεροπορικών κινητήρων έφτασαν στον παγκόσμιο στόλο με πλαστά πιστοποιητικά πτητικής ικανότητας, αποτελεί την επιβεβαίωση αυτής της ανεπάρκειας σε μεγάλη κλίμακα.

Σκοπός και μεθοδολογία. Η εργασία δημιουργεί μια συνεχή ικανότητα διαχείρισης κινδύνου της εφοδιαστικής αλυσίδα, η οποία αντιμετωπίζει τη διασφάλιση του προμηθευτή ως φθαρτό περιουσιακό στοιχείο, αποτιμά το κόστος της απώλειάς της και τη συγκροτεί σε ένα λειτουργικό μοντέλο που ένας οργανισμός μπορεί να υιοθετήσει.

Τι Δημιουργήθηκε. Το πρώτο artifact αντιστοιχίζει πέντε επικαλυπτόμενα κανονιστικά πλαίσια, τη NIS2, τον EU Cyber Resilience Act, το Part-IS του EASA, το IEC 62443 και το NIST SP 800-161r1, μαζί με το DO-326A.

Το δεύτερο artifact είναι ένα μοντέλο ωριμότητας τεσσάρων σταδίων μαζί με μια μετρική φθοράς της εμπιστοσύνης (trust decay metric). Η μετρική αυτή μετατρέπει την αδράνεια ανάμεσα στα διαδοχικά τεκμήρια σε μετρήσιμο σήμα κινδύνου όπου θα πρέπει να αποτελεί και σημείο διερεύνησης,

Το τρίτο artifact προσαρμόζει τη συνιστώσα προσομοίωσης Monte Carlo του πλαισίου FABRICS και την επεκτείνει με έναν όρο απώλειας ασφάλειας (safety loss term), ώστε η εκτίμηση απώλειας να περιλαμβάνει και τον κίνδυνο ανθρώπινων απωλειών, κάτι που σε βιομηχανικά περιβάλλοντα αποτελεί υψίστη σημασία

Το τέταρτο artifact συγκρίνει τη συνεχή κρυπτογραφική προέλευση (continuous cryptographic provenance), δηλαδή τον διαρκή έλεγχο ψηφιακά υπογεγραμμένων τεκμηρίων, με την έντυπη πιστοποίηση που ελέγχεται μία φορά κατά την παραλαβή.

Το πέμπτο artifact συνδέει τα προηγούμενα τέσσερα σε ένα ενιαίο, λειτουργικό μοντέλο (operating model).

Βασικά αποτελέσματα. Στο τυχαίο σενάριο λειτουργικής τεχνολογίας (OT scenario) η διάμεση ετήσια απώλεια είναι μηδενική, ενώ η αναμενόμενη ετήσια απώλεια ανέρχεται περίπου σε 1,3 εκατομμύρια δολάρια. Η αξία σε κίνδυνο (value at risk) φτάνει περίπου τα 6,3 εκατομμύρια δολάρια στην 95η εκατοστιαία τιμή και τα 13,6 εκατομμύρια δολάρια στην 99η εκατοστιαία τιμή. Λιγότερο από ένα τοις εκατό των προσομοιωμένων ετών περιλαμβάνει θανατηφόρο συμβάν. Παρ' όλα αυτά, ο όρος απώλειας ασφάλειας (safety loss term) ανεβάζει την 99η εκατοστιαία τιμή κατά περίπου τέσσερα εκατομμύρια δολάρια. Με τη συνεχή κρυπτογραφική προέλευση (continuous cryptographic provenance) η ουρά της κατανομής μειώνεται κατά περίπου 7,8 εκατομμύρια δολάρια, και ο έλεγχος αποδίδει 33,5 τοις εκατό ετησίως.

Συμβολή και περιορισμοί. Η συμβολή αυτής της εργασίας δεν είναι κάποιο νέο εύρημα σε ένα καινούργιο επιστημονικό πεδίο. Αποτελεί ένας τρόπος με τον οποίο τέσσερις ώριμοι ήδη τομείς ενώνονται σε ένα ενιαίο λειτουργικό μοντέλο. Το μοντέλο αυτό ακολουθεί μια συνεχή, ελέγξιμη διαδρομή δηλαδή ένα τεκμήριο που έχει επαληθευτεί γίνεται μετρήσιμο σήμα κινδύνου → το σήμα μετατρέπεται σε νομισματικό μέγεθος πιθανής απώλειας → και αυτό με τη σειρά του στηρίζει μια απόφαση ελέγχου που εγκρίνεται και χρηματοδοτείται από τη διοίκηση. Σε κάθε βήμα μπορεί κανείς να ανιχνεύσει προς τα πίσω γιατί άρθηκε η κάθε απόφαση και σε ποιο τεκμήριο στηρίχθηκε.

Αυτό που παραδίδει η εργασία είναι η εφαρμογή της μεθοδολογίας συνεχούς επισκόπησης προμηθευτών σε μια πραγματική περίπτωση έτσι ώστε ο οργανισμός να έχει πλήρη επίγνωσή των ρίσκων και της κατάστασης των προμηθευτών. Δεν παραδίδει αποτελέσματα δοκιμασμένα σε πραγματικό περιβάλλον οργανισμού, ούτε

αριθμούς βαθμονομημένους πάνω σε ιστορικά δεδομένα συμβάντων ή θανάτων. Κάθε χρηματικό μέγεθος που αναφέρεται είναι είτε ενδεικτικό.

Λέξεις κλειδιά: διαχείριση κινδύνου εφοδιαστικής αλυσίδας; λειτουργική τεχνολογία; κατάλογος υλικών λογισμικού; κρυπτογραφική προέλευση; ποσοτικοποίηση κινδύνου ασφάλειας πληροφοριακών συστημάτων; αξία σε κίνδυνο; συνεχής διασφάλιση; πλαστά εξαρτήματα; πολιτική αεροπορία

Table of Contents

Abstract	iii
Περίληψη	v
Table of Contents	vii
List of Figures	x
List of Tables	x
List of Abbreviations	xii
1. Introduction	1
1.1 The Problem	1
1.2 The Core Claim	1
1.3 Research Questions	2
1.4 Scope and Approach	2
1.5 Contribution	2
1.6 An Executive Perspective	3
1.7 Thesis Structure	3
2. Background and Literature Review	5
2.1 Overview and Chapter Structure	5
2.2 The Regulatory Landscape for OT and Aviation Cybersecurity	5
2.3 Cyber Supply Chain Risk Management and Maturity Models	6
2.4 Quantifying Cyber Risk in Monetary Terms	6
2.5 Counterfeit Detection and Provenance	7
2.6 Synthesis: The Gap This Thesis Fills	7
2.7 The AOG Technics Counterfeit Parts Affair	8
2.7.1 What Happened	8
2.7.2 Discovery and Response	8
2.7.3 Legal Outcome and Reported Losses	9
2.7.4 Why This Case Anchors the Thesis	10
2.8 Summary and Bridge	10
3. Research Methodology	11
3.1 Overview and Chapter Structure	11
3.2 Research Questions and Design Objectives	11
3.3 The Design Science Process	11
3.3.1 Process Activities	11
3.3.2 Contribution Boundary: Adapted versus Original	12
3.3.3 Proof of Concept Instances	12
3.4 The Case Study as Illustrative Anchor	13

3.5 Evidence Base and Source Quality Discipline.....	13
3.6 Scope and Boundaries.....	13
3.6.1 Within Scope.....	13
3.6.2 Beyond Scope	13
3.7 Validity, Reflexivity, Ethics and AI Use	14
4. From Regulation to Operation: A Unified Control Ontology and the Out of Band Telemetry Abstraction.....	15
4.1 Overview and Objectives	15
4.2 The Problem: Five Rulebooks, One Plant.....	15
4.3 Artifact Part One: The Unified Control Ontology	15
4.4 Artifact Part Two: The Out of Band Telemetry Abstraction	17
4.5 What This Means in Practice	18
4.6 Summary and Bridge	18
5. The Continuous C-SCRM Maturity Model and the Trust Decay Metric	19
5.1 Overview and Objectives	19
5.2 Foundations: What the Evidence Base Says	19
5.3 Artifact Part One: The Four Stage Maturity Model.....	19
5.4 Artifact Part Two: The Trust Decay Metric.....	20
5.4.1 Definition	20
5.4.2 Behaviour	21
5.4.3 Key Risk Indicators.....	22
5.5 Proof of Concept: Scoring an AOG Style Scenario	22
5.6 What This Means in Practice	23
5.7 Summary and Bridge	23
6. Monetary Risk Quantification: A FAIR Inspired Loss Model Informed by FABRICS	24
6.1 Overview and Objectives	24
6.2 Foundations: What Is Taken from FABRICS and What Is Not	24
6.3 The Model.....	24
6.4 Proof of Concept: A Neutral OT Scenario.....	25
6.5 The Safety Loss Term and Why It Matters.....	26
6.6 Sensitivity: What the Executive Figure Depends On.....	27
6.6.1 Value at Risk and Expected Shortfall	28
6.7 From Figure to Decision	29
6.8 Summary	29
7. Counterfeit Grounding Value at Risk and the Provenance Investment Case	30
7.1 Overview and Objectives	30

7.2 The Counterfeit Grounding Scenario	30
7.3 The Control: Continuous Cryptographic Provenance versus Paper Certification	30
7.4 Method: Re-running the Model Before and After Controls.....	30
7.5 Results: The Figure Falls	31
7.6 The Investment Decision	32
7.7 Sensitivity and Break Even	32
7.8 Evaluation and Limitations	33
7.9 Summary	33
8. The Continuous Assurance Operating Model: Running the Four Layers as One Function	35
8.1 Overview and Objectives	35
8.2 From Four Layers to One Function.....	35
8.3 Standing the Function Up: Phases and Gates.....	35
8.4 Ownership: Who Runs Each Layer.....	36
8.5 Governance: No Silent Risk.....	36
8.6 The Maturity Path From Certification to Commercialisation	36
8.7 The Operating Model on the AOG Pattern	37
8.8 Summary and Bridge	37
9. Conclusion	38
9.1 What the Thesis Built.....	38
9.2 Answering the Research Questions.....	39
9.3 Contribution	39
9.4 Limitations.....	40
9.5 Future Work	39
9.6 Closing	40
References.....	41
Appendix A Neutral OT Cyber Risk Model (Chapter 6).....	44
Appendix B Counterfeit Grounding Value at Risk and Provenance ROI (Chapter 7)	47
Appendix C Operating Model Onboarding Schedule (Chapter 8).....	50

List of Figures

Figure 3.1 The design science research methodology (DSRM) process, after Peffers et al. (2007), run once per artifact, with the build and evaluate feedback loop and the five design iterations mapped to Chapters 4 to 8. The fifth iteration integrates the four layers rather than adding a new layer beside them.	12
Figure 4.1 The out of band telemetry abstraction. Passive taps collect evidence inside the isolated OT zone; a hardware data diode enforces one way flow; the assurance plane consumes the evidence for scoring, quantification, and reporting. No inbound path is created.....	17
Figure 5.1 The trust decay metric on two suppliers over one year: half life 60 days, thresholds at 0.75 and 0.50, verified score 0.88. Evidence that keeps verifying produces a sawtooth in the assured band; evidence that stops produces a silent slide into escalation. All values are illustrative.	21
Figure 6.1 Annual loss exceedance curve for the neutral OT scenario. The horizontal axis is annual loss on a logarithmic scale; the vertical axis is the probability that annual loss is exceeded. All figures are illustrative.	26
Figure 6.2 Effect of the safety loss term on the headline figures. The term barely changes the mean or the 95th percentile, but lifts the 99th percentile materially. All figures are illustrative.	27
Figure 7.1 Annual loss exceedance before and after continuous cryptographic provenance. The whole curve shifts left, and the tail falls sharply. All figures are illustrative.....	31
Figure 7.2 Annual return on investment of the provenance control: gross risk reduction, less the control cost, gives the net benefit. All figures are illustrative.	32

List of Tables

Table 1.1 The four layers, the operating model that integrates them, and where each is developed.	3
Table 2.1 The four fields, the gap each leaves, and the artifact that addresses it.	8
Table 2.2 The AOG Technics affair at a glance. Monetary figures are as stated by the Serious Fraud Office or attributed to prosecutors in contemporary reporting.	9
Table 3.1 Research questions and the artifact that answers each.	11
Table 3.2 Origin of artifact components. FABRICS is the published work of Slapnicar and Joshi. This thesis takes its framing and does not implement its Bayesian fault tree, as Section 6.2 states.	12
Table 3.3 Source tiers and the weight each may carry. Non peer reviewed material and practitioner experience are used for perspective and operating structure only, never to support a core claim.	13
Table 4.1 The ontology pattern: one normalised objective, the frameworks that demand it, and the single control whose evidence answers them all.	16
Table 4.2 Each control objective enriched with the evidence that proves it, how quickly that evidence decays, and what its failure costs. These fields are the inputs the later layers consume.	17
Table 5.1 The four stage continuous C-SCRM maturity model. The stages align with the C2M2 maturity indicator level progression, specialised to the continuity of supplier assurance.....	20
Table 5.2 Key risk indicators generated by the trust decay metric.....	22
Table 6.1 Illustrative model inputs. All figures are illustrative placeholders that demonstrate the method, not calibrated values.....	25

Table 6.2 Headline figures with and without the safety loss term (illustrative).	27
Table 6.3 Sensitivity of the headline figures to the safety exposure share and the value of a statistical life. Same seed in every cell. All figures illustrative.	28
Table 7.1 Headline figures before and after the control (illustrative).	31
Table 7.2 The provenance investment case (illustrative).	32
Table 7.3 Return on investment across control effectiveness, against an unchanged baseline and a fixed annual control cost of \$750,000. Expected annual loss after the control falls from \$0.78M in the low case to \$0.17M in the high case. All figures illustrative.	32
Table C.1 Onboarding schedule for the continuous assurance operating model, by 30 day horizon, with phase gates. All dates are relative; owners are roles, not individuals.	50

List of Abbreviations

C2M2	Cybersecurity Capability Maturity Model (US Department of Energy)
CFM56	The CFM International turbofan engine family at the centre of the case study
CFO	Chief Financial Officer
CRA	Cyber Resilience Act, Regulation (EU) 2024/2847
C-SCRM	Cyber Supply Chain Risk Management
Cyber VaR	Cyber Value at Risk
DO-326A	Airworthiness Security Process Specification (with ED-202A)
DO-356A	Airworthiness Security Methods and Considerations (with ED-203A)
DSR	Design Science Research
DSRM	Design Science Research Methodology
EASA	European Union Aviation Safety Agency
FAA	Federal Aviation Administration (United States)
FABRICS	Fault tree And Bayesian Risk and Impact analysis for Cyber Security, the framework of Slapnicar and Joshi (2026)
FAIR	Factor Analysis of Information Risk
HMI	Human Machine Interface
IEC	International Electrotechnical Commission
KRI	Key Risk Indicator
MIL	Maturity Indicator Level (C2M2)
NIS2	Network and Information Security Directive (EU) 2022/2555
NIST	National Institute of Standards and Technology (United States)
NTIA	National Telecommunications and Information Administration (United States)
OBTA	Out of Band Telemetry Abstraction, the architecture defined in Chapter 4
OT	Operational Technology
Part-IS	EASA information security requirements for civil aviation organisations
PERT	Program Evaluation and Review Technique distribution
PLC	Programmable Logic Controller
ROI	Return on Investment
SBOM	Software Bill of Materials
SFO	Serious Fraud Office (United Kingdom)
VaR	Value at Risk
VSL	Value of a Statistical Life

1. Introduction

Why can't a certificate be considered enough, and what to develop instead.

1.1 The Problem

In 2023, tens of thousands of aircraft engine parts reached the worldwide fleet carrying paperwork that was forged. AOG Technics, sold more than sixty thousand components, most of them for the CFM56 engine that powers the Boeing 737 and the Airbus A320 accompanied by a falsified certificate stating the part was airworthy. The fraud was running for many years. Obviously, It was not stopped by an audit, a regulator, or a control system. It was stopped when one airline happened to ask the engine maker whether a the provided certificate was genuine. The full story is given in Chapter 2.

The above mentioned scandal is not about a single rogue corporation. It is a story about a control that failed silently, and control is one that almost every safety critical industry still relies on: **the point in time certificate**. A supplier is checked once, at onboarding or during delivery, a document is issued, and from that moment the supplier is trusted until something forces a re-examination. The model assumes that trust established at a point in time remains valid across time but AOG Technics is proof that it doesn't.

This thesis argues that the point in time model is weak and structurally inadequate especially for operational technology (OT) and civil aviation, the environments where a compromised or counterfeit component can cause not just financial loss but physical harm and since safety is the most important value in such organizations, the impact can be catastrophic.

The weakness is not that organisations ask for evidence it is that they stop asking too soon and they don't continuously monitor the attack surface of the suppliers.

We have identified 3 forces that make the above mentioned issue even more important and those are the following. 1. opaque and multi-layered supply chains, 2. a shift toward mandatory continuous regulatory oversight, and 3. high-consequence operational environments where the risk of fatalities demands an evidentiary rigor way beyond what an annual certificate can provide.

1.2 The Core Claim

The central claim of this thesis is that certifying a supplier once, on paper, at a point in time is structurally inadequate for safety critical infrastructures, OT environments and civil aviation, and must be replaced by continuous assurance capability that is both academically defensible and deployable in a real organization. That capability blends 3 things the current model keeps apart. A. continuous evidence about what is inside a product and where it came from, B. a measure of supplier trust that decays as that evidence ages, and C. a monetary quantification of the residual risk and of the controls that reduce it. The thesis does not stop at these three pieces. It designs them as the layers of a single operating model that can be adopted and supported by the appropriate tech stack.

1.3 Research Questions

Thesis has been structured with 4 design questions. The first four are each answered by one layer and one chapter. The fifth asks how the four becomes one function and is answered by the integrating operating model in Chapter 8.

1. **RQ1.** How can the fragmented obligations of NIS2, the EU Cyber Resilience Act, EASA Part-IS, IEC 62443, and NIST SP 800-161r1 be mapped to one set of operational controls, traceable in both directions, in a way that allows continuous monitoring without breaking OT isolation?
2. **RQ2.** How can supplier assurance be expressed as a continuous property that decays over time rather than a one-off attestation, and measured through a maturity model and key risk indicators?
3. **RQ3.** How can cyber risk to a safety critical OT environment be expressed in monetary terms under severe data scarcity, in a form a chief financial officer can act on?
4. **RQ4.** How can the economic case for continuous cryptographic provenance¹ over paper certification be quantified, using counterfeit component risk as the worked example?
5. **RQ5.** How can these four layers be assembled into a single assurance function, with the ownership, governance, and maturity path that a safety critical operator would need in order to run it?

1.4 Scope and Approach

The thesis delivers designs, demonstrations, and defensible methods. It does not deliver field validated results, calibrated incident or fatality statistics, or an audited return on investment. Inputs, most importantly the value of a statistical life used in the safety loss term, are treated as explicit and contestable design choices, and every monetary figure is either traced to a source or labelled illustrative. The thesis is also sector agnostic with aviation as the illustrative example. The artifacts are built for OT supply chains in general, and the AOG Technics case demonstrates them rather than bounding them.

1.5 Contribution

The contribution is not a new result in any single field, and it is not four separate artifacts standing side by side. It is the integration of four mature fields into one continuous assurance operating model. More specifically it is a deployable function, with layers that produce evidence, measure capability, and price risk, held together by a governance discipline and a maturity path an organisation can actually climb. Each of the four artifacts is a layer of that model rather than an end in itself, and the model is the unit of contribution. Table 1.1 below lists the four layers and the operating model that integrates them, with the chapter in which each is developed.

¹ Continuous cryptographic provenance refers to the cryptographic assurance of the evidence stream, that is, the signing and verification of SBOMs, provenance attestations, and certificates, not to cryptographic elements embedded in the physical component. The guarantee concerns the authenticity of the evidence, not the physical identity of the component.

Layer	What it delivers	Chapter
Unified control ontology and out of band telemetry abstraction	One traceable control set answering five frameworks, and a one-way evidence pipeline that monitors OT without opening an attack path	Chapter 4
Continuous C-SCRM maturity model and trust decay metric	A capability ladder, and a metric under which supplier trust decays in silence and resets on verified evidence	Chapter 5
FAIR inspired loss model informed by FABRICS, with a safety loss term	A monetary risk figure, a loss exceedance curve and a cyber value at risk, that includes the safety critical tail	Chapter 6
Counterfeit grounding value at risk and provenance ROI	A before and after comparison that prices a control and turns a risk figure into a budget decision	Chapter 7
Continuous assurance operating model	The four layers run as one governed function, with owners, gates, a single source of truth, and a maturity path	Chapter 8

Table 1.1 The four layers, the operating model that integrates them, and where each is developed.

The four layers form a single chain of reasoning that runs evidence, capability, money, and then operation. Chapter 4 produces trustworthy evidence, Chapter 5 turns it into a measured capability, and Chapters 6 and 7 express the risk and the value of acting on it in the language of the balance sheet that top management understand cared for. The closing operating model chapter then does the work that separates a deployable function from a set of good ideas. Especially, it fixes who owns what and assigns responsibility, how the layers hand off to one another, where the single source of truth lives, and how an organisation moves from paper certification to a priced, governed, and ultimately commercialised assurance capability. That integration is what this thesis contributes.

1.6 An Executive Perspective

This thesis is written with business leaders in mind, not just researchers. In safety-critical industries, decisions are not based only on theory. Executive committees need to understand the financial impact of risks, the value of security controls, and how to explain their decisions to regulators and boards. For this reason, each part of the framework supports a business decision. So, the quantification model translates cyber risk into financial terms. The ROI analysis helps turn those numbers into a clear investment decision and to actually understand the value of Third Party Risk Management and the investment to person days.

The work is also designed to be used beyond an academic environment. The four layers can form the foundation of a continuous assurance function and support future consulting services for the OT sector. This is why the thesis uses a business-focused tone throughout and ends with an operating model chapter that explains how all four layers work together, from initial deployment to a measurable and sustainable capability.

1.7 Thesis Structure

The rest of the thesis is organized based on the following:

- **Chapter 2** reviews the four fields the thesis joins, states the gap each leaves, and presents the AOG Technics affair in full as the case that anchors the work.

- **Chapter 3** sets out the design of science methodology, the research questions and the scope boundaries.
- **Chapter 4** develops the unified control ontology and the out of band telemetry abstraction: the evidence layer.
- **Chapter 5** develops the continuous maturity model and the trust decay metric: the capability layer.
- **Chapter 6** develops monetary risk quantification and runs the proof of concept simulation: the economic layer, pricing the risk.
- **Chapter 7** develops the counterfeit grounding model and the provenance investment case: the economic layer, pricing the control.
- **Chapter 8** is the synthesis. It draws the four layers into one continuous assurance operating model, and sets out the phases and gates by which the function is stood up, the roles that own each layer, the single source of truth that keeps the layers consistent, the governance discipline that stops assurance decaying quietly once bought, and the maturity path from point in time certification to a priced and commercialised capability.

Closing the work, Chapter 9 articulates the core contributions, frames the limitations, and charts the path toward operational practice. Written as an unbroken progression, each chapter links directly to the next, reinforcing the overarching argument that assurance must be continuous, quantitatively measured, priced, and executed as a governed discipline.

2. Background and Literature Review

Four fields, one gap, and the real world failure that proves it.

2.1 Overview and Chapter Structure

This chapter performs an analysis to the existing research that forms the foundation of the thesis and identifies the gap it aims to address. The study brings together four areas that are usually examined separately. First of all we have the cybersecurity regulations for critical infrastructure and aviation then we have the cyber supply chain risk management and maturity models, then the cyber risk quantification, and finally the counterfeit part detection through provenance. While each field is well established, none provides a complete way for safety-critical organisations to continuously assess and measure supplier trustworthiness or understand its financial impact or/and value.

Bellow Sections 2.2 to 2.5 discuss each field and highlight its main limitation. Section 2.6 combines these gaps into the central problem addressed by the thesis. Section 2.7 presents the 2023 AOG Technics counterfeit parts case, a real-world example that demonstrates all four gaps in practice. Finally, Section

2.2 The Regulatory Landscape for OT and Aviation Cybersecurity

It is well know that f you're running a critical tech or industrial system in Europe right now in Q4 of 2026, you basically have to deal with a mountain of overlapping laws. For starters, NIS2 makes companies responsible for their own cybersecurity and forces them to report hacks. Then there's the new EU Cyber Resilience Act, which brings in stricter rules for digital products starting in late 2026 and 2027. If you look at aviation specifically, things get even more specific with rules like EASA Part-IS and standards like DO-326A and ED-202A to make sure airborne systems are secure. Underneath all of this, engineers also usually rely on technical guides like IEC 62443 for factories and NIST rules for supply chains.

Recent studies show just how messy and overwhelming this whole setup is for everyone involved. For instance, Siddiqui et al. explain how DO-326A and ED-202A help assess plane security risks, pointing out that old-school safety rules just don't cut it anymore now that avionics are connected to networks and are considered Internet connected things. Meanwhile, Saini et al. (2026) look at industrial control systems and show how IEC 62443 breaks a system down into "zones and conduits" to set different security levels. Finally, Parlov et al. (2025) try to make sense of all these rules by proposing a 5-step model that links NIS2 and the Cyber Resilience Act into one timeline, though they are honest enough to admit that their model still needs to be tested in the real world.

Gap. The frameworks are complementary but fragmented. They use different words for overlapping duties, divide responsibility differently, and each expects its own evidence. Parlov and colleagues name the deeper problem directly. They state that the difficulty is not a shortage of frameworks but the mismatch between what policies require and what operations actually do. No source in the corpus provides a single control set that satisfies all of them at once and produces evidence traceable back to each. That absence is what the artifact in Chapter 4 tries to address.

2.3 Cyber Supply Chain Risk Management and Maturity Models

Cyber supply chain risk management, C-SCRM, is the discipline of governing the risk that flows into an organisation through the hardware, software, and services it buys. Boyson's two studies a decade apart give the field its empirical spine. The 2014 work framed C-SCRM as the fusion of cybersecurity, enterprise risk, and supply chain management, and found in interviews across industry and government that shop floor industrial controls were routinely excluded from IT security audits and that no single executive owned end to end risk. By 2022 the governance picture had shifted: the share of organisations with a C-Suite risk board rose sharply across the survey period, while direct cooperation with supply chain partners barely moved. Organisations learned to govern risk at board level faster than they learned to verify their suppliers.

Maturity models give this progression a structure. The United States Department of Energy C2M2 model (DOE, 2022) defines maturity indicator levels across ten domains, including one dedicated to supply chain and external dependencies, and describes a path from absent, through performed but ad hoc, to documented and managed practices. The corpus adds peer reviewed weight around it: Karabacak et al. (2016) develop a Delphi validated maturity model for national critical infrastructure preparedness, and Rahman et al. (2021) build a supply chain cyber resilience index that links condition states to financial loss. On the evidence that feeds such models, European evidence points the same way. ENISA (2026) surveyed 334 organisations, most of them inside the European Union and most directly affected by the Cyber Resilience Act, and found that while 78 percent had started adopting software bills of materials, a large share still rarely receive one from their own suppliers, with inconsistent data quality and incomplete inventories reported as the leading obstacles. That is the supplier silence problem of this thesis measured at European scale. Nocera et al. (2025) show that software bills of materials are often incomplete or non compliant in practice, so the mere existence of an evidence artifact says little about its quality.

It was identified that, existing maturity models describe capability at a point of assessment. They do not treat supplier assurance as something that decays between assessments, and they do not turn the absence of fresh evidence into a measurable signal. The models score what has been collected, not how current it is. That gap, from periodic assessment to continuous, decaying assurance, is what the maturity model and trust decay metric in Chapter 5 address.

2.4 Quantifying Cyber Risk in Monetary Terms

Expressing cyber risk in money, rather than in heat map colours, is what lets a chief financial officer weigh it against other risks and against the cost of controls. The dominant approach, from the FAIR tradition (Factor Analysis of Information Risk), expresses risk as the product of how often a loss event occurs and how large the loss is, and produces a distribution of possible annual loss rather than a single figure. Slapničar and Joshi (2026) introduce FABRICS, which combines structured expert elicitation, Bayesian fault trees, and Monte Carlo simulation to produce a financial distribution of cyber loss under data scarcity. Erola et al. (2022) demonstrate a cyber value at risk engine on a real case, and Orlando (2021) sets out the compound frequency and severity model that underpins the annual loss calculation.

The evidence that this style of model transfers into OT and industrial contexts is growing. Zhang et al. (2025) apply Monte Carlo cost benefit analysis directly to industrial control system attack paths, Seid et al. (2024) quantify cyber risk for logistics and supply chains, and Ryan et al. (2012) provide the expert judgment elicitation foundation on which quantification under data scarcity depends.

Again it is identified that the quantification literature is strong for enterprise and IT loss, and it is beginning to reach OT, but it does not carry a safety critical loss term that prices potential loss of life, and its published cases are rarely aviation or safety critical. Nor does it connect the resulting figure to the value of a specific control. Chapters 6 and 7 address both. This is why a FAIR inspired loss model informed by FABRICS with a safety loss term, and a before and after comparison that prices a control was built.

2.5 Counterfeit Detection and Provenance

Montecchi et al. (2019) frame provenance as a business problem of knowing an object's origin and journey, and Dwivedi et al. (2020) demonstrate blockchain-based provenance in a pharmaceutical supply chain, an analogical case for regulated goods. For aviation specifically, Ho et al. (2021) build a blockchain system for aircraft spare parts traceability that records the authorized release certificate and the part's history, and Kumar et al. (2025) bind a physical aerospace part to its digital twin through a cryptographic anchor, so a counterfeit that lacks the anchor is detectable.

A parallel strand gives a part an identity that cannot simply be copied onto a document. Wigger et al. (2018, 2020) show that the microscopic surface pattern of a manufactured component can act as a physical unclonable function, an unforgeable fingerprint read by a camera. Sola et al. (2021) review the tagging and fingerprinting strategies available for additively manufactured parts, and on the electronics side Mura et al. (2020), together with Vepraskas and Stirgwolt (2024), address the detection and procurement controls needed to keep counterfeit components out of the supply chain.

However the above mentioned provenance and anti-counterfeit techniques are strong individually, but they are treated as engineering tech solutions in isolation. They are not tied to a continuous assurance capability, they are not expressed in the language of regulatory compliance, and their economic value against paper certification is rarely made. Chapter 7 makes that economic case, and Chapter 4 ties provenance evidence into the compliance and assurance pipeline.

2.6 Synthesis: The Gap This Thesis Fills

Read together, the four fields point at one conclusion from four directions. Regulation demands continuous assurance but arrives fragmented. Maturity models describe capability but treat assurance as static and not continuous. Quantification prices risk but omits the safety critical tail and the value of controls to understand the ROI. Provenance can make a part trustworthy but stays disconnected from compliance and from the balance sheet. The common thread is that supplier trustworthiness is still treated as a fact established once, on paper, at a point in time, when the evidence from every field says it is a perishable, continuous property that should be verified, measured, and priced.

Table 2.1 summarises the four gaps and the artifact that answers each. The contribution of this thesis is not a new result in any single field as it has been already mentioned. It is the joining of the four into one capability that a safety critical operator can deploy and defend, and that a chief financial officer can understand and fund the overall TPRM program.

Field	What it offers	The gap it leaves	Artifact
Regulation	Detailed obligations across NIS2, CRA, Part-IS, IEC 62443, NIST 800-161r1, DO-326A	Fragmented; no single traceable control set; monitoring conflicts with OT isolation	Chapter 4

C-SCRM and maturity	A path from ad hoc to managed supplier assurance	Assurance treated as static; silence is not a signal	Chapter 5
Monetary quantification	Risk expressed as a loss distribution a CFO can act on	No safety critical loss term; control value not priced	Chapters 6 and 7
Counterfeit and provenance	Techniques to make a part verifiably authentic	Disconnected from compliance and from the economic case	Chapters 4 and 7

Table 2.1 The four fields, the gap each leaves, and the artifact that addresses it.

2.7 The AOG Technics Counterfeit Parts Affair

The 2023 AOG Technics affair is the real world event in which all four gaps appear together, and it is the case that anchors the artifact chapters. This section sets out the facts once, in full, so the later chapters can refer back to it. All events listed in the thesis are drawn from the public record, principally statements by the United Kingdom Serious Fraud Office and contemporary reporting; the monetary figures are attributed to their sources, and Section 2.7 closes with a note on how the thesis treats them.

2.7.1 What Happened

AOG Technics was a UK aircraft parts broker run by Jose Alejandro Zamora Yrala. Between 2019 and 2023, the company sold thousands of aircraft engine parts, mainly for CFM56 engines used on Boeing 737 NG and Airbus A320ceo aircraft. These parts were supplied with fake airworthiness certificates (EASA Form 1 and FAA 8130-3), making them appear to be approved and safe for use. As a result, used or unverified parts entered airline inventories as if they were genuine and without detection.

The fraud was operated from a home office. Genuine certificates were altered, fake shipping records were created, and even non-existent employees were invented to make the documents look legitimate. According to prosecutors, the company was created for the initial stage to run in order to committe fraud.

2.7.2 Discovery and Response

This whole mess started blowing up in 2023 when an airline got suspicious about a plane part and checked it with Safran (one of the companies that owns CFM International). Safran realized the paperwork was totally fake and called the authorities. After that, major aviation regulators like the FAA, UK CAA, and EASA rushed out urgent safety alerts and notifications about all parts linked to a company called AOG Technics.

As a result, CFM International, along with GE Aerospace and Safran, had to launch massive checks to track down every engine with fake parts. At first, they found over 100 suspected engines, but later reports showed the number was closer to 180 CFM56 engines carrying forged paperwork.

Equally important, the fallout was a total nightmare for travel. It is apparent that airlines had to ground planes all over the world, pull engines apart, and cancel flights just to make sure everything was safe. Some of the big-name airlines caught up in it across Europe and the US included American, United, Delta, Southwest, Ryanair, TAP, WestJet, and Virgin Australia. Fortunately, even with all the chaos and grounded planes, no actual accidents have been directly tied to the fake parts.

2.7.3 Legal Outcome and Reported Losses

According to valid data the serious Fraud Office (SFO) launched an investigation into AOG Technics in October 2023 after the company was ordered by the court to provide documents it had refused to release. Although the director initially denied the allegations, he pleaded guilty to fraudulent trading on 1 December 2025. He was later sentenced to four years and eight months in prison in February 2026.

Hence, the financial impact was significant. According to the SFO, AOG Technics sold aircraft parts with fake certificates worth around £6.9 million, representing almost 90% of its revenue. The total losses across the aviation industry were estimated at about more than £39 million. Then, American Airlines reportedly suffered the largest losses, around £23 million, due to engine repairs, replacement engine leasing, and aircraft downtime. GE Aerospace, Safran, and Ethiopian Airlines were also among the affected organisations.

Item	Detail	Source basis
Company and director	AOG Technics (UK parts broker); Jose Alejandro Zamora Yrala	SFO and court record
Period	2019 to 2023	SFO
Parts affected	Over 60,000, sold with falsified documentation	SFO
Engine and aircraft	Mainly CFM56; Boeing 737 NG and Airbus A320ceo	SFO and reporting
Forged documents	Authorised Release Certificates (EASA Form 1 / FAA 8130-3)	SFO and reporting
Value of falsified parts	About £6.9M (\$9.3M), around 90% of revenue	SFO
Estimated industry loss	Over £39.3M (about \$53M)	SFO estimate
Discovery	Airline query to Safran; forged certificate identified	SFO and reporting
Engines affected	Over 100 initially; later accounts cite about 180	CFM and reporting
Legal outcome	Guilty plea 1 Dec 2025; sentenced 4 years 8 months on 23 Feb 2026	SFO and court record
Reported operator loss	American Airlines about £23M (\$31M); Ethiopian Airlines among those affected	Prosecutors, via press

Table 2.2 The AOG Technics affair at a glance. Monetary figures are as stated by the Serious Fraud Office or attributed to prosecutors in contemporary reporting.

For that reason, the confirmed facts of the case, the parts, dates, engines, forged certificates, discovery, and legal outcome, are used as established background. Similarly, the monetary figures above are as stated by the Serious Fraud Office or reported from prosecutors, and are cited without any alteration. Where the thesis builds a worked financial example on this case, in Chapter 7, it treats specific per operator inputs as illustrative unless they correspond to these reported figures, and it labels every such figure accordingly.

2.7.4 Why This Case Anchors the Thesis

The AOG Technics affair anchors this thesis because it exposed all four gaps at once firstly the forged certification, then the unchecked supplier trust, then the unpriced industry exposure, and finally the unverifiable part provenance. Each artifact addresses one part of that failure, which is why the case returns throughout the thesis.

2.8 Summary and Bridge

This chapter reviewed the four fields behind the thesis and showed their shared gap which is that supplier trust needs and must be continuously verified, measured, and priced, not checked once on paper.

The next chapters build the solution and more specifically the methodology is presented in Chapter 3, evidence in Chapter 4, trust measurement in Chapter 5, risk and control value in Chapters 6 and 7, the operating model in Chapter 8, and conclusions in Chapter 9.

3. Research Methodology

A practical method for designing and testing the thesis framework.

3.1 Overview and Chapter Structure

This thesis uses Design Science Research (DSR) because the purpose and goal is to build a practical solution to a real business problem. In contrast of testing an existing theory, the study designs five connected artifacts for continuous supplier assurance and evaluates whether they are useful, logical, and supported by evidence.

Additionally, This thesis uses Design Science Research to build and evaluate practical artifacts for continuous supplier assurance.

3.2 Research Questions and Design Objectives

The study asks how continuous supplier assurance can combine regulatory controls, current evidence, risk measurement, and business decisions. The five research questions link directly to the five artifacts developed in Chapters 4 to 8.

Research question	Main output
RQ1: How can different requirements be combined?	Control mapping and safe evidence collection
RQ2: How can supplier trust be measured over time?	Maturity model and trust decay metric
RQ3: How can cyber risk be expressed in financial terms?	FAIR-inspired loss model
RQ4: How can the value of provenance controls be measured?	Counterfeit-risk and ROI model
RQ5: How can all outputs operate as one function?	Continuous assurance operating model

Table 3.1 Research questions and the artifact that answers each.

3.3 The Design Science Process

The thesis applies the six DSRM activities to each artifact as presented below:

3.3.1 Process Activities

1. **Identify the problem:** explain why point-in-time supplier certification is not enough.
2. **Set objectives:** define what each artifact must achieve.
3. **Design:** build each artifact using research, standards, and practice.
4. **Demonstrate:** apply each artifact to a worked example.
5. **Evaluate:** review each artifact against fixed criteria.
6. **Communicate:** present the results in clear academic and business language.




Figure 3.1 The design science research methodology (DSRM) process, after Peffers et al. (2007), run once per artifact, with the build and evaluate feedback loop and the five design iterations mapped to Chapters 4 to 8. The fifth iteration integrates the four layers rather than adding a new layer beside them.

3.3.2 Contribution Boundary: Adapted versus Original

The thesis clearly separates and acknowledges ideas adapted from earlier research from its own contribution. Existing frameworks provide the foundation, while the new contribution is the way they are combined and applied to continuous supplier assurance.

Artifact	Contribution	Main foundation
Control ontology and telemetry model	Original	Regulations and OT standards
Maturity and trust decay model	Original	C2M2 and C-SCRM research
Financial risk model	Adapted and extended	FAIR and FABRICS
Counterfeit-risk and ROI model	Original application	Cyber VaR and provenance research
Operating model	Original integration	The four artifacts and professional practice

Table 3.2 Origin of artifact components. FABRICS is the published work of Slapnicar and Joshi. This thesis takes its framing and does not implement its Bayesian fault tree, as Section 6.2 states.

3.3.3 Proof of Concept Instances

Each artifact is tested through a practical business example. These include examples such as supplier scoring, Monte Carlo risk simulation, comparison of paper and cryptographic provenance, and an end-to-end operating model scenario.

3.4 The Case Study as Illustrative Anchor

The AOG Technics case is used as a practical example because it is public and clearly shows the weakness of paper-based supplier assurance. It helps demonstrate each artifact but is not used as statistical proof.

Case figures are treated as illustrative unless they come from a reliable primary source and where and if possible. The framework is designed for OT supply chains generally, with aviation used only as the main example.

3.5 Evidence Base and Source Quality Discipline

Our study uses sources according to their reliability. Peer-reviewed research and official standards support the main claims. Practitioner material and professional experience are used only for context and practical design.

Source type	Use in the thesis
Peer-reviewed research	Main academic evidence
Official standards and regulations	Authoritative requirements and structures
Practitioner material	Context only
Professional experience	Practical organisation of the operating model only

Table 3.3 Source tiers and the weight each may carry. Non peer reviewed material and practitioner experience are used for perspective and operating structure only, never to support a core claim.

All financial inputs are either linked to a source or clearly marked as assumptions. This helps us avoid presenting illustrative values as proven facts.

3.6 Scope and Boundaries

The scope is strictly limited to designing and demonstrating the framework at a level suitable for a master's thesis.

3.6.1 Within Scope

- design of 4 artifacts and one integrated operating model;
- expert or structured self-review; and
- proof-of-concept examples.

3.6.2 Beyond Scope

- long-term testing in real organisations;
- calibration with confidential operational data; and
- proof of actual financial return.

These limits mean that the thesis presents a defensible design and method, not a fully validated commercial solution.

3.7 Validity, Reflexivity, Ethics and AI Use

Validity: The main concepts and measures are drawn from cited sources, and assumptions are kept separate from evidence. The results should still be read as formative, not generalisable.

Reflexivity: The author's role may influence interpretation. To limit this, the same review criteria are applied to all artifacts, and professional experience is used only to support the practical design.

Ethics: The study relies only on public or published material. It uses no confidential or personal data, and it does not present assumptions as verified facts.

AI use: Generative AI has been used as a support tool during script development and editing. It helped test Python code and improve grammar and clarity, but the author wrote the prompts, reviewed the outputs, tested the code, and approved the final analysis and text.

Summary: DSR fits the study because it supports practical design where data are limited, while keeping the thesis claims tied to the available evidence.

4. From Regulation to Operation: A Unified Control Ontology and the Out of Band Telemetry Abstraction

One control set that answers five regulators, and a way to monitor continuously without opening a door into OT.

4.1 Overview and Objectives

This chapter presents the evidence layer. As such the thesis provide a way to meet overlapping OT and aviation requirements without weakening critical system isolation.

It combines a unified control ontology with OBTA - Out of Band Telemetry Abstraction - so assurance evidence can be mapped, collected, and reused safely.

This evidence then feeds the maturity, risk, and ROI layers that follow.

4.2 The Problem: Five Rulebooks, One Plant

A critical OT operator faces overlapping duties from NIS2, the CRA, EASA Part-IS, IEC 62443, NIST SP 800-161r1, and different aviation security guidance.

The problem is fragmentation. More specifically different terms, owners, and evidence demand create gaps between policy and practice. However, the ontology closes this by turning them into one traceable control set.

4.3 Artifact Part One: The Unified Control Ontology

The ontology is a three-layer structure. At the top are regulatory obligations stated in each framework's own words. In the middle is a normalized set of control objectives, written once, in plain operational language. At the bottom are concrete OT controls and the evidence each produces. It should be noted that every link is bidirectional, from any regulation you can trace down to the controls that satisfy it, and at the same time from any control you can trace up to every obligation it serves.

Building the ontology follows four steps as presented bellow

1. **Decompose.** Break each framework into atomic obligations, single demands that can be satisfied and evidenced on their own.
2. **Normalise.** Group obligations that ask for the same operational outcome under one control objective, keeping a record of each framework's original wording so nothing is lost in translation.
3. **Bind.** Attach to each control objective the concrete OT controls that deliver it, and define the evidence each control produces, for example an SBOM, a provenance record, a monitoring log, or a signed attestation.
4. **Trace.** Publish the mapping in both directions, so an auditor can start from a regulation and an engineer can start from a control, and both arrive at the same place.

Table 4.1 illustrates the structure with four of the control objectives most relevant to this thesis. The full ontology is an operational deliverable; the table shows the pattern.

Control objective (normalised)	Who demands it (examples)	OT control and evidence
Know what is inside every product and component	CRA Annex I (SBOM); NIST SP 800-161r1 SR-3, SR-4; IEC 62443-4-1 SM-9, SM-10; NIS2 Article 21	Require and verify SBOMs from suppliers; evidence is the SBOM itself plus its quality score (Chapter 5)
Prove where every part came from	NIST SP 800-161r1 SR-4 (provenance), SR-11 (component authenticity); IEC 62443-4-1 SM-6; DO-326A; Part-IS IS.I.OR.205	Cryptographic provenance records bound to the part (Chapter 7); evidence is the verification log
Detect and report incidents and anomalies continuously	NIS2 Article 21 and Article 23; Part-IS IS.I.OR.215, IS.I.OR.220; NIST SP 800-161r1 SR-9; IEC 62443 monitoring	Passive OT telemetry via OBTA; evidence is the monitoring record on the assurance plane
Re-verify each supplier over time	NIST SP 800-161r1 SR-6 (supplier assessments and reviews); Part-IS IS.I.OR.205, IS.I.OR.235	Continuous C-SCRM scoring with the trust decay metric (Chapter 5); periodic supplier attestation

Table 4.1 The ontology pattern: one normalised objective, the frameworks that demand it, and the single control whose evidence answers them all.

The crosswalk explains the requirement and who expects it, but it becomes useful only when linked to evidence and business impact. As we mentioned thesis cared about business impact and only academic research that's why the table bellow has been created. Table 4.2 shows the evidence needed, how quickly it becomes outdated, and the cost if the control fails. Chapter 5 uses this information to measure trust over time, while Chapters 6 and 7 estimate the financial risk. This connects each regulation to evidence and, finally, to figure management can use.

Control objective	Evidence that proves it	Freshness: cadence and decay	Failure consequence
Component inventory (SBOM)	A complete, verified SBOM for each release, with all minimum fields present	Per software release; stale the moment a new version ships without one	A hidden or vulnerable component that nobody knows to patch
Provenance and authenticity	A signed provenance record that verifies against the issuing authority, bound to each part	Per delivery; only as fresh as the last verified shipment	An unverified or counterfeit part reaches a critical system, the safety critical tail in Chapter 6
Detection and reporting	Passive OT telemetry and anomaly or incident	Continuous; a gap in the telemetry is itself a signal	An undetected compromise, or a missed signal

	records carried through OBTA		regulatory reporting deadline
Supplier re-verification	Periodic reassessment and attestation, scored by the trust decay metric	Set by supplier criticality, for example a 60 day half life, decaying between arrivals	Assurance that goes stale in silence, the exact pattern behind the AOG Technics affair

Table 4.2 Each control objective enriched with the evidence that proves it, how quickly that evidence decays, and what its failure costs. These fields are the inputs the later layers consume.

4.4 Artifact Part Two: The Out of Band Telemetry Abstraction

The ontology creates a new problem the moment it is put to work. Its third objective in Table 4.1, continuous monitoring, collides with the first rule of OT engineering which is that the plant must be isolated on purpose and supported by different technological stack. IEC 62443 partitions the environment into zones and conduits precisely to keep attack paths out, and decades of practice place the control system behind a demilitarised zone - DMZ, separated from the enterprise network. A naive continuous monitoring deployment, an agent on every controller reporting to a cloud console, would tear that isolation open.

OBTA is the architectural answer. The idea is simple and can be actually stated in one sentence. More specifically assurance evidence flows out of the OT environment through a path on which, by physics rather than by policy, nothing can flow back in. Figure 4.1 shows the pattern.




Figure 4.1 The out of band telemetry abstraction. Passive taps collect evidence inside the isolated OT zone; a hardware data diode enforces one way flow; the assurance plane consumes the evidence for scoring, quantification, and reporting. No inbound path is created.

Three design rules define OBTA.

- **Collection is passive.** Evidence is taken from network taps and existing historians, not from agents installed on controllers. The monitoring adds no software to the devices that run the plant and cannot disturb the process.

- **The boundary is physical.** The evidence crosses the zone boundary through a unidirectional gateway, a data diode, which is enforced in hardware. There is no credential to steal and no rule to misconfigure, because the return path does not exist.
- **Interpretation happens outside.** Everything that analyses, scores, and reports lives on the assurance plane, outside the OT zone. The plant gives evidence; it never takes instructions from the assurance side.

The consequence for compliance is the point of the artifact. The continuous monitoring that NIS2, Part-IS, and IEC 62443 ask for is delivered without adding a single inbound attack path. Isolation and continuous assurance stop being a trade off. The assurance plane then becomes the natural home for everything the rest of this thesis builds. It holds the SBOM and provenance evidence. Also it runs the continuous scoring of Chapter 5, and finally it feeds the risk quantification of Chapters 6 and 7 and the reporting duties of the regulations.

4.5 What This Means in Practice

For an operator, the two parts of the artifact translate into a short, concrete programme as presented bellow in order.

1. **Adopt the ontology.** Map current controls onto the normalised objectives once, and record the gaps. From then on, every audit, from any of the five frameworks, is answered from the same mapping.
2. **Instrument passively.** Place taps at the levels of the OT architecture that already carry the needed signals, typically the supervisory and site operations levels, and route them to a diode.
3. **Stand up the assurance plane.** Land the evidence in one place, attach the supplier SBOM and provenance records to it, and switch scoring and reporting on.

4.6 Summary and Bridge

This chapter turned 5 fragmented rulebooks into one traceable control set and resolved the conflict between continuous monitoring and OT isolation with one-way evidence architecture. Together the two parts produce something the thesis needs from here on. This is a continuous stream of trustworthy assurance evidence to actually feed our TPRM program, flowing to one place, without weakening the plant.

Chapter 5 measures assurance maturity and how supplier trust declines when evidence stops. Chapters 6 and 7 then quantify the remaining risk and the value of controls. In short, the chain moves from evidence to capability and then to money.

5. The Continuous C-SCRM Maturity Model and the Trust Decay Metric

Treating supplier assurance as a perishable asset: a maturity ladder for continuous assurance, and a metric that makes trust expire.

5.1 Overview and Objectives

This chapter introduces the evidence layer of the operating model. The chapter combines different regulatory requirements into a single control framework and uses the OBTA architecture to collect assurance evidence safely from OT environments. The evidence produced by this layer supports the trust model, risk quantification, and ROI analysis in the chapters that follow.

5.2 Foundations: What the Evidence Base Says

The maturity model is built based on the US Department of Energy's C2M2 framework, focusing specifically on how continuous supplier checks are, alongside research by Karabacak et al. and Rahman et al.

The practical need comes from Boyson's studies, which found that while C-suite risk management jumped from 19% to 71% over a decade, direct supply chain cooperation barely budged (29% to 33%). Companies talk about governance but still fail to verify supplier via a structured TPRM program

Furthermore, Nocera et al. (2025) showed that just collecting software bills of materials (SBOMs) is not enough, since most lacked required fields or correct formatting. Therefore, this model scores verified evidence rather than paperwork. Finally, Campbell (2026) inspires the top level of real-time zero-trust attestation, though it carries no heavy academic weight here.

5.3 Artifact Part One: The Four Stage Maturity Model

The model has four stages, deliberately aligned with the C2M2 progression so that any organisations already using C2M2 can adopt it without translation or need to learn and adopt a new framework. Each stage is defined by one question as presented bellow and it concerns what triggers the organisation to re-evaluate trust in a supplier?

Stage	Name	What triggers re-evaluation of supplier trust	Typical evidence posture
Stage 0	Point in time	A calendar: certification at onboarding, perhaps an annual questionnaire or bi annual for critical providers	Paper certificates; no verification between audits; trust is treated as permanent
Stage 1	Reactive	An incident or a headline: trust is revisited only after something goes wrong	Ad hoc evidence requests after events; no baseline; findings are not retained
Stage 2	Scheduled verification	A defined cadence: SBOMs and provenance records are requested and verified on a schedule	Evidence arrives at a cadence and is quality checked; gaps are recorded; scores exist but do not expire

Stage 3	Continuous assurance	The evidence itself: verification runs whenever evidence arrives through the pipeline, and the absence of evidence is itself a signal	OBTA fed evidence stream; trust decay metric live; escalation thresholds wired to procurement and engineering decisions
---------	----------------------	---	---

Table 5.1 The four stage continuous C-SCRM maturity model. The stages align with the C2M2 maturity indicator level progression, specialised to the continuity of supplier assurance.

The decisive move happens between Stage 2 and Stage 3, and it is a change of philosophy, not of tooling. At Stage 2 the organisation asks *is the evidence good when it arrives*. At Stage 3 it also asks *what does it mean that evidence has not arrived* and now silence becomes a signal. That is the capability that would have mattered in the AOG Technics affair. More specifically, no single forged certificate had to be caught, because the pattern of evidence, and its gaps, is itself being watched. Stage 3 is what the assurance plane of Chapter 4 exists to enable.

5.4 Artifact Part Two: The Trust Decay Metric

The trust decay metric is the instrument that makes Stage 3 measurable. It rests on the principle that verification is information about the past, and information ages throughout the time. A certificate verified today says a great deal about today and progressively less about every day after. Evidence does not become false overnight, but it becomes less reliable the longer it is left unrefreshed.

5.4.1 Definition

Every supplier has a trust score. It runs from 0 to 1. Reviewing / auditing the supplier pushes the score up. Time pulls it back down.

The score at any moment is $T(t) = V \times e^{(-\lambda d)}$. V is the score given at the last check. d is the number of days since that check. λ sets how fast the score falls.

The easiest way to talk about λ is the half-life. The half life is the number of days it takes for the score to drop to half its value when no new evidence arrives. λ is the natural logarithm of 2 divided by the half life.

When new evidence arrives and passes, the score jumps back up to the new V . When new evidence arrives and fails, the score is cut straight away. It is not left to drift down slowly.

Three settings make the metric work. All three are choices the organisation makes, and it must be able to explain them to an auditor. They are not buried inside a tool. They are also the point where this layer joins the one before it. Specifically, the evidence that sets V and the decay that sets the half life are the same fields the Chapter 4 ontology records for each control objective in Table 4.2.

The verified score V . This depends on how good the evidence is, not just on whether evidence exists. An SBOM that carries every minimum field and matches the delivered software scores high. An SBOM that is missing fields, or does not comply, scores low. The Nocera study shows this difference matters in practice.

The half life. This is set by how critical the supplier is. A supplier of safety critical parts might get a 60 day half life. A supplier of office consumables might get 365 days. The half life says, in one number, how fast the organisation thinks that supplier's assurance goes stale.

The thresholds. These are bands on the score that trigger action. Above the top band the supplier counts as assured. Between the bands the supplier goes into review. Below the bottom band the supplier is escalated and treated as unverified, with the procurement and engineering consequences agreed in advance.

5.4.2 Behaviour

Figure 5.1 shows the metric on two suppliers over one year, with a 60 day half-life, thresholds at 0.75 and 0.50, and each verification resetting the score to 0.88. The three numbers are stated together because the reset value, the half life, and the thresholds are what make every crossing point in the figure reproducible by hand.

Supplier A delivers evidence that verifies every 30 days as a result, its score follows a sawtooth, decaying between arrivals and resetting on each verification, and it lives in the assured band.

Supplier B is identical until day 90, when its evidence stops. Nothing bad is reported until there is simply silence.

Under point in time certification that silence is invisible because the certificate from day 90 remains on file and the supplier remains, on paper, in good standing. Under the trust decay metric the silence is the event. So the score slides, enters review around day 104, and crosses the escalation threshold near day 139, at which point the organisation is forced to confront a question it would otherwise never have asked to actually identify and mitigate potential risks.




Figure 5.1 The trust decay metric on two suppliers over one year: half life 60 days, thresholds at 0.75 and 0.50, verified score 0.88. Evidence that keeps verifying produces a sawtooth in the assured band; evidence that stops produces a silent slide into escalation. All values are illustrative.

The metric is one expression. The trust score of a supplier is:

$$T(t) = V \times e^{(-\lambda d)}$$

where V is the score set by the last piece of evidence that verified, d is the number of days since that evidence arrived, and λ is the decay rate, obtained from the chosen half life H as $\lambda = \ln(2) / H$. In plain and simple terms, the expression says that the confidence in a supplier falls by a fixed percentage of whatever remains for every day that passes without new verified evidence. Setting the half life to 60 days therefore fixes λ at 0.0116 per day, which means the score loses about 1.15 per cent of its remaining value each day and halves after 60 days of silence. The half life is the only parameter a governance body or executive committee needs to discuss and agree, because it states in one number how perishable the organisation considers that supplier's assurance to be.

Now about exponential decay, a fixed deadline creates an artificial cut-off so the evidence is trusted on day 89 but rejected on day 91. Exponential decay reduces trust gradually as evidence gets older. It is simple to calculate, easy to audit, and can be adjusted using one value the half-life based on the supplier's risk.

5.4.3 Key Risk Indicators

The metric generates a small set of key risk indicators that management can put on one page. Each is computable directly from the assurance plane of Chapter 4.

Indicator	Definition	What it tells management
Mean time to verify provenance	Average days from evidence arrival to completed verification	How fast the organisation converts evidence into assurance; the core Stage 3 speed metric
Portfolio trust index	Criticality weighted average of T(t) across all suppliers	One number for the health of the whole supply base, trending over time
Suppliers below threshold	Count of suppliers currently in the review or escalation bands	Where attention and budget must go this week
Evidence silence	Days since last verified evidence, per critical supplier	The leading indicator: silence precedes every slide in Figure 5.1
Verification failure rate	Share of received evidence that fails verification	Whether the supply base is producing paperwork or proof; the Nocera gap made local

Table 5.2 Key risk indicators generated by the trust decay metric.

5.5 Proof of Concept: Scoring an AOG Style Scenario

The demonstration required by chapter 3 presents the artifact through the failure pattern of the AOG Technics affair, stated here in the model's terms and using only the public shape of the case.

1. **Stage 0 organization, which is the historical reality.** The distributor was certified at onboarding and its paperwork accepted at face value on every delivery. Trust was permanent by construction. Hence, there was no mechanism by which accumulating deliveries could raise or lower confidence, so the fraud ran until a human happened to look closely.
2. **The same events under Stage 3.** Every delivery is checked. If the paper is fake, the supplier is flagged at once. If the part has no proof of origin, the score stays low. If no evidence arrives, the silence also lowers the score and triggers review.
3. **The management view.** On the key risk indicators, the distributor appears as a rising verification failure rate or a lengthening silence, on one page, weeks or months before any incident. The organization's question changes from why nobody noticed, after the fact, to what is our threshold for acting, in advance, which is a governance question with a written answer.

The proof-of-concept claim is deliberately modest, consistent with chapter 3, It does not claim the metric would have caught this specific fraud, whose forensic details exceed the public record. However, It demonstrates that the artifact converts the pattern of that fraud, prolonged supply under unverifiable paperwork, from an invisible condition into a measured, thresholded, and escalated one.

5.6 What This Means in Practice

For an operator, adopting the artifact is a staged programme, and the stages are the maturity model itself.

1. **Locate the organisation on the ladder.** Most organisations that rely on supplier certification sit at Stage 0 or Stage 1, this is the honest self assessment is the starting point.
2. **Reach Stage 2.** Define the evidence cadence per supplier class, verify what arrives, and record quality. This needs the ontology of Chapter 4 but not yet the full pipeline.
3. **Reach Stage 3.** Switch on the trust decay metric over the assurance plane, set half lives and thresholds per criticality class through the governance body, and wire the escalation band to concrete procurement and engineering actions.

The commercial logic completes the thesis chain. Chapter 6 prices the risk of the environment and Chapter 7 prices the value of the provenance control. This chapter supplies the operating discipline that keeps that value real over time. A control that is bought but not fed with verified evidence decays exactly like the red curve in Figure 5.1. The maturity model tells an executive what capability they are buying with the Chapter 7 budget, and the trust decay metric tells them, every day afterwards, whether they still have it.

5.7 Summary and Bridge

This chapter turned the evidence stream of Chapter 4 into a measured capability. The four-stage maturity model positions an organisation by the single question of what triggers re-evaluation of supplier trust, and the trust decay metric makes Stage 3 operational. More specifically, trust becomes a perishable asset that decays in silence, resets on verification, and escalates on thresholds a governance body has set in advance. The AOG pattern, prolonged supply under unverifiable paperwork, becomes visible weeks or months before any incident.

The chain now needs its final links priced. Chapter 6 takes the environment this capability watches over and expresses its cyber risk in money, a loss exceedance curve and a Cyber VaR figure a chief financial officer can act on. Chapter 7 then prices the provenance control itself and shows the figure falling. Evidence, capability, money: the argument of the thesis completes across these three steps.

6. Monetary Risk Quantification: A FAIR Inspired Loss Model Informed by FABRICS

Expressing safety critical cyber risk as a Cyber VaR figure under data scarcity.

6.1 Overview and Objectives

This chapter bring into the table the economic layer of the operating model. It shows how cyber risk in safety-critical OT environments can be converted into financial terms actually helping business leaders understand and manage risk since money is the main value they care. The output is a Cyber VaR figure that estimates potential losses and supports decision-making. A neutral OT example is used here.

6.2 Foundations: What Is Taken from FABRICS and What Is Not

This chapter introduces the economic layer of the operating model. It uses a FAIR-inspired approach, informed by FABRICS, to convert cyber risk in safety-critical OT environments into financial terms. Rather than producing a single estimate, the model generates a range of possible losses, helping decision-makers understand both likely and worst-case outcomes.

The model is designed for environments where data is limited and extends existing approaches by including OT-specific impacts and a dedicated safety loss term. The result is a Cyber VaR figure that helps organisations understand the potential financial impact of cyber incidents and supports investment decisions.

6.3 The Model

The model expresses annual loss as the sum of the losses from each event that occurs in a year. The number of events in a year is modelled as a Poisson process whose annual rate is itself uncertain, so that uncertainty about how often events happen is carried through the simulation rather than fixed at a single number. The loss from a single event is the sum of the severity components listed below.

The safety loss term is being kept separate from the operational categories. It equals the number of fatalities multiplied by the value of a statistical life. The value of a statistical life is taken from an external primary authority, the United States Department of Transportation, at approximately 13.2 million dollars. The United States Environmental Protection Agency figure of roughly 11 million dollars is a reasonable alternative. This figure is an **explicit and contestable design choice**, not an empirical finding. The probability that an event creates a safety exposure, and the number of fatalities given such an exposure, are **elicited or assumed parameters**, not values calibrated against incident data.

Because actually no public incident dataset supports calibration in this domain, every input parameter would in practice be obtained by structured expert elicitation, using the apparatus described in the source framework. The values in Table 6.1 are illustrative only placeholders chosen to demonstrate the method. In a real engagement they would be replaced by values elicited from the operator's own engineers and risk experts.

Parameter	Illustrative distribution	Note
Annual event rate	PERT(0.10, 0.30, 0.80) events per year	Elicited frequency; carries epistemic uncertainty

Outage duration	PERT(4, 24, 120) hours	Unplanned production stoppage
Cost per hour of lost production	PERT(\$20k, \$60k, \$150k)	Process or manufacturing setting
Equipment damage	40% chance, then PERT(\$50k, \$400k, \$3M)	Many events cause none
Incident response and recovery	PERT(\$50k, \$200k, \$800k)	Investigation, restoration, forensics
Regulatory exposure	30% chance, then PERT(\$0, \$150k, \$2M)	For example NIS2 related penalty
Safety exposure per event	5% chance	Elicited, not calibrated
Fatalities given a safety exposure	Poisson(0.7)	Elicited, not calibrated
Value of a statistical life	\$13.2M (US DOT)	Explicit design choice

Table 6.1 Illustrative model inputs. All figures are illustrative placeholders that demonstrate the method, not calibrated values.

6.4 Proof of Concept: A Neutral OT Scenario

The scenario is deliberately sector neutral. A compromised component or firmware in an industrial control system at a process or manufacturing facility leads to an unplanned production outage. In some events there is also equipment damage, and in a small fraction of events there is a safety exposure. The scenario is chosen to keep the quantification method general so it does not depend on any aviation specific feature.

The analysis was carried out using a short Monte Carlo simulation written in Python with the help of Claude models and vibe coding techniques. In plain terms, the model plays out one possible year at a time, many times over, and records the loss each time. The steps were as follows.

1. **Set up.** Take the inputs in Table 6.1. They describe how often an incident might happen and how costly each one could be.
2. **Simulate one year.** Decide, at random and guided by those inputs, how many incidents occur in the year. Most years have none.
3. **Cost each incident.** For every incident in that year, draw its five cost parts: outage cost, equipment damage, incident response, regulatory exposure, and the separate safety cost.
4. **Add them up.** Sum those costs into a single total loss for that year.
5. **Repeat.** Do the previous three steps two hundred thousand times, giving two hundred thousand possible years.
6. **Read the results.** From all those years, take the average loss, the middle value, the worst one year in twenty (the 95th percentile), and the worst one year in a hundred (the 99th percentile). These last two are the Cyber VaR figures.

7. **Draw the curve.** Sort the yearly losses to produce the loss exceedance curve in Figure 6.1.

The same starting point (a fixed random seed) is used each time, so the results are reproducible. It is worth noting that running the script again returns the same figures. The full inputs are illustrative and would be replaced by values elicited from the operator's own experts. Hen




Figure 6.1 Annual loss exceedance curve for the neutral OT scenario. The horizontal axis is annual loss on a logarithmic scale; the vertical axis is the probability that annual loss is exceeded. All figures are illustrative.

We can identify that three results stand out. Hence, in any given year there is roughly a **29 percent chance of at least one event**. Because most years contain no event, the **median annual loss is zero**, while the **expected (mean) annual loss is about 1.3 million dollars**. That gap, between a zero median and a multimillion dollar mean, is the central message for a decision maker. Clearly, budgeting to the average hides the exposure that matters, because the loss lives in infrequent but costly events. The Cyber VaR at the 95th percentile is about 6.3 million dollars, and at the 99th percentile about 13.6 million dollars. In plain terms, there is a five percent chance that annual loss exceeds about 6.3 million dollars, and a one percent chance it exceeds about 13.6 million.

For executives, this is the main point: the quiet year is not proof that the exposure is small; it may simply be one of the many years in which the tail has not yet appeared.

6.5 The Safety Loss Term and Why It Matters

The reason why the safety critical extension is necessary becomes clear and evident when the executive figures are computed with and without the safety loss term. Figure 6.2 and Table 6.2 set out the comparison.




Figure 6.2 Effect of the safety loss term on the headline figures. The term barely changes the mean or the 95th percentile, but lifts the 99th percentile materially. All figures are illustrative.

Measure	Excluding safety term	With safety term
Probability of at least one event in a year	29%	29%
Median annual loss	\$0	\$0
Expected (mean) annual loss	about \$1.1M	about \$1.3M
Cyber VaR (95th percentile)	about \$5.9M	about \$6.3M
Cyber VaR (99th percentile)	about \$9.6M	about \$13.6M

Table 6.2 Headline figures with and without the safety loss term (illustrative).

Less than one percent of simulated years involves a fatality, yet the safety term lifts the 99th percentile Cyber VaR from about 9.6 million dollars to about 13.6 million, an increase of roughly four million, while barely moving the mean or the 95th percentile. The interpretation is direct and can be easily understood by technical and non technical users. In a safety critical OT environment the financial tail is dominated by rare events that carry potential loss of life. A quantification that left the safety term out would systematically understate the very risk that justifies investment in continuous assurance. This is the analytical reason the safety critical extension is needed, and it is the bridge to Chapter 7, where an undetected counterfeit component is exactly this kind of rare, high consequence event.

The probability of safety exposure and the number of fatalities are elicited or assumed. The thesis does not claim a calibrated fatality model. The safety term is a transparent and contestable device for making low probability, high consequence risk visible in monetary terms. It is not a prediction, and it should not be read as one.

6.6 Sensitivity: What the Executive Figure Depends On

A distribution produced from elicited inputs is only as good as the inputs, so the honest question is not what the figure is but how far it moves when the assumptions move. Three inputs carry the safety tail. Firstly, the share of events that create a safety exposure, then the number of fatalities given such an exposure, and finally

the value of a statistical life. The first and the third are the ones an examiner or a board member will challenge, so the model was re-run across both.

Table 6.3 holds everything else fixed and varies only the safety exposure share and the value of a statistical life, using the same random seed in every cell so that the differences are attributable to the parameters rather than to sampling.

Table 6.3 Sensitivity of the headline figures to the safety exposure share and the value of a statistical life. Same seed in every cell. All figures illustrative.

Safety exposure share	Value of a statistical life	Expected annual loss	Cyber VaR (95th)	Cyber VaR (99th)	Expected shortfall (99th)
0.02	\$13.2M (US DOT)	\$1.19M	\$6.05M	\$10.40M	\$15.89M
0.02	\$11.0M (US EPA)	\$1.18M	\$6.05M	\$10.40M	\$14.88M
0.05	\$13.2M (US DOT)	\$1.28M	\$6.28M	\$13.56M	\$21.85M
0.05	\$11.0M (US EPA)	\$1.26M	\$6.28M	\$12.62M	\$19.28M
0.10	\$13.2M (US DOT)	\$1.45M	\$6.76M	\$17.90M	\$28.21M
0.10	\$11.0M (US EPA)	\$1.39M	\$6.76M	\$15.87M	\$24.55M

Three findings follow, and each one is more useful to a decision maker than the single figure it replaces.

The safety exposure share dominates, not the value of a statistical life. Moving the exposure share from 0.02 to 0.10 moves the 99th percentile from about 10.4 million dollars to about 17.9 million, a factor of about 1.7. Switching from the US DOT figure to the lower US EPA figure at a fixed exposure share moves the same percentile by about seven percent. The parameter that honestly deserves the elicitation effort is therefore the exposure share, which is an engineering and operational question about how a compromised component reaches a safety function, not the monetary constant, which is a published policy figure.

The monetary value of life does not change the 95th percentile. We can say an example to make this complex math terms easy so like a rare storm, it appears only in the worst tail years, so a 95th percentile report would miss it.

The average loss looks stable, but the worst-case tail moves sharply. This is why the safety term is shown separately: it reveals the risk that averages hide.

6.6.1 Value at Risk and Expected Shortfall

Value at Risk shows where the bad years start. Expected shortfall shows how bad those bad years usually are.

For this reason, the thesis reports both. Firstly VaR for the risk threshold and expected shortfall for capital, insurance, or contingency planning.

6.7 From Figure to Decision

Cyber VaR turns cyber risk into money, so executives can compare it with other business risks and decide whether a control is worth funding. Chapter 7 applies this logic to continuous provenance.

6.8 Summary

This chapter built a FAIR inspired compound loss model, informed by FABRICS, into a safety critical OT risk quantification that produces a defensible monetary figure under data scarcity, and demonstrated it on a neutral OT scenario. The result is a loss exceedance curve and a Cyber VaR figure that show the exposure in the tail rather than only on average. Consequently, the safety loss term, although it bites in fewer than one percent of years, dominates the executive tail and is the reason the adaptation matters. Afterwards, Chapter 7 applies the same machinery to the aviation counterfeit component case and turns the figure into a return on investment comparison between continuous cryptographic provenance and paper certification.

7. Counterfeit Grounding Value at Risk and the Provenance Investment Case

Turning a risk figure into a budget decision: re-running the model with controls in place.

7.1 Overview and Objectives

Now, this chapter completes the economic layer by showing whether a security control is worth the investment. Using the AOG Technics case, it compares cyber risk before and after implementing controls and calculates the financial benefit.

The process is simple: calculate the risk, apply the control, calculate the risk again, compare the reduction with the control cost, and decide whether the investment is worthwhile.

7.2 The Counterfeit Grounding Scenario

The core issue is a fake or unapproved part slipping into a safety-critical system using forged paperwork. The 2023 AOG Technics scandal is the main real-world example. To be more clear, a fake airworthiness forms were used to push unapproved parts into the CFM56 engine supply chain. With traditional paper checks, this kind of fraud gets caught very late, if ever.

This is a rare but severe risk. Actually, in most years nothing happens, but one counterfeit part in a critical system can create major safety and financial impact.

As in Chapter 6, the monetary figures are illustrative; verified case figures show the scale of the fraud, not the actual clean-up cost.

7.3 The Control: Continuous Cryptographic Provenance versus Paper Certification

Traditional paper certificates prove a part's origin only at one point in time and can be forged, as seen in the AOG Technics case. Continuous provenance provides a stronger alternative by linking a part to a secure digital record and, where possible, to its unique physical characteristics.

It is clear that research has shown that technologies such as blockchain, digital twins, and physical fingerprints can help verify the authenticity of aircraft parts and detect counterfeits more effectively than paper records alone. In this model, the control mainly reduces the likelihood of counterfeit parts entering the supply chain and helps identify suspicious parts earlier, lowering

7.4 Method: Re-running the Model Before and After Controls

The same Monte Carlo machinery from Chapter 6 is used, with one addition. More specifically, a second run that represents the world after the control is in place. The procedure is as follows.

1. **Baseline runs.** Run the model with the inputs that represent paper certification, namely Chapter 6 frequency and severity. This gives the figures before the control.
2. **Represent the control.** Reduce the inputs the control affects. A. lower the annual event rate, because fewer undetected counterfeits reach a critical system, and B.

lower the share of events that reach a safety exposure. How large these reductions are is an elicited assumption about how effective control is.

3. **After run.** Run the model again with the reduced inputs. This gives the figures after the control.
4. **Compare.** Subtract the after figures from the before figures to get the risk reduction, both on the average and in the tail.
5. **Cost it.** Subtract the annual cost of the control from the reduction in expected annual loss, to get the net annual benefit and the return on investment.

The script is provided in Appendix B. As in Chapter 6, a fixed starting point makes the results reproducible.

7.5 Results: The Figure Falls

Figure 7.1 shows the loss exceedance curve before and after the control. Table 7.1 gives the headline figures.




Figure 7.1 Annual loss exceedance before and after continuous cryptographic provenance. The whole curve shifts left, and the tail falls sharply. All figures are illustrative.

Measure	Before (paper certification)	After (continuous provenance)
Chance of an event in a year	about 29%	about 8%
Expected (mean) annual loss	about \$1.3M	about \$0.28M
Cyber VaR (95th percentile)	about \$6.3M	about \$2.4M
Cyber VaR (99th percentile)	about \$13.6M	about \$5.7M

Table 7.1 Headline figures before and after the control (illustrative).

The chance of an event in a year falls from about 29 percent to about 8 percent, and the whole loss exceedance curve shifts left. The headline for a safety critical setting is the tail. The 99th percentile Cyber VaR drops by about 7.8 million dollars, from about 13.6 million to about 5.7 million. That is the figure that matters most, because the tail is where catastrophic, safety driven loss sits.

7.6 The Investment Decision

With the reduction quantified, the chief financial officer can make the decision. Figure 7.2 sets out the annual return on investment, and Table 7.2 summarizes it.




Figure 7.2 Annual return on investment of the provenance control: gross risk reduction, less the control cost, gives the net benefit. All figures are illustrative.

Item	Illustrative annual value
Gross annual risk reduction (expected loss)	about \$1.0M
Annual control cost	about \$750k
Net annual benefit	about \$250k
Return on investment	33.5%

Table 7.2 The provenance investment case (illustrative).

The control costs about \$750k per year and saves about \$1M, leaving a net benefit of about \$250k and a 33.5% return. In simple terms, it costs less than the risk it removes.

The strongest point is the tail. We can evidently understand that the control cuts the worst-case 99th percentile loss by about \$7.8M, which is the safety-critical outcome the board needs to see.

7.7 Sensitivity and Break Even

This result lays on three simple assumptions: how much the control reduces counterfeit events, how much it reduces safety exposure, and how much it costs.

The model was then tested under weaker and stronger control scenarios to see when the investment still makes sense.

Table 7.3 Return on investment across control effectiveness, against an unchanged baseline and a fixed annual control cost of \$750,000. Expected annual loss after the control falls from \$0.78M in the low case to \$0.17M in the high case. All figures illustrative.

Scenario and safety exposure share after	Frequency cut	Risk reduction	Net benefit	Return
Low effectiveness, share unchanged at 0.050	40%	\$0.51M	minus \$0.24M	minus 32%
Modest effectiveness, share 0.035	60%	\$0.78M	\$0.03M	5%
Central case of Section 7.5, share 0.020	about 76%	\$1.00M	\$0.25M	33.5%
High effectiveness, share 0.010	85%	\$1.11M	\$0.36M	48%

The table shows the simple rule: the control is worth funding only if it reduces risk enough.

The break even point is identified at a \$750k yearly cost, the control needs about a 55–60% risk reduction to make financial sense.

Cost limit -> Under the main case, the control saves about \$1M per year, so paying more than that would not be justified.

These are not proven results however they are decision targets that can be tested in a pilot or written into supplier requirements.

The model covers recurring cost only so a full business case should also include implementation, onboarding, integration, and replacement costs.

As it has been mentioned, all figures remain illustrative, including the fatality probability and the value of a statistical life.

7.8 Evaluation and Limitations

The artifact is assessed against the five criteria from Section 3.3.

- **Grounded.** It uses provenance research and the Chapter 6 risk model.
- **Coherent.** The before and after cases use the same model, so the comparison is fair.
- **Demonstrable.** The simulation shows the risk falling and turns that reduction into ROI.
- **Useful.** It gives management the risk, the reduction, and the cost in one decision view.
- **Clear on uncertainty.** Effectiveness, cost, and all figures are illustrative assumptions.

The main limits are simple: the cost and effectiveness are assumed, the scenario is single, and the AOG financial figures are hypothetical.

The control is also a chain, not one tool. Its weakest point is linking the physical part to the digital record; if that link fails, the whole control becomes weaker.

7.9 Summary

To sum up, chapter 6 produced the risk figure. Then, This chapter closed the loop by showing control reducing it and returning more than it costs, using continuous cryptographic provenance against the aviation counterfeit problem. It is clear that the before and after re-run turns a risk number into a budget decision. the loss

exceedance curve shifts left, the 99th percentile tail falls by about 7.8 million dollars, and the control returns 33.5 percent a year and is net positive every year. Together, Chapters 6 and 7 give an executive both the size of the risk and the value of acting on it, which is the core argument of the thesis expressed in money.

8. The Continuous Assurance Operating Model: Running the Four Layers as One Function

How the evidence, capability, and money layers become a governed, priced function an operator can stand up, run, and grow.

8.1 Overview and Objectives

This chapter brings the four layers, as presented above, together as one operating model that an organisation can easily own, fund, and improve over time. The point is simple and the explanation follows. Evidence, metrics, and risk figures only create value when someone is responsible for keeping them current and acting on them.

Supplier assurance remains the main use case, but the same model can also support internal control assurance where evidence, ownership, and decisions need to stay connected.

The chapter explains how the layers work together, how governance prevents assurance from going stale, and how an organisation can move from periodic checks to continuous, measured assurance.

8.2 From Four Layers to One Function

The model works as a continuous loop. Chapter 4 provides verified evidence, Chapter 5 turns it into supplier trust and maturity scores, and Chapters 6 and 7 translate the risk and control value into financial terms that top management can understand.

All evidence flows into one source of truth. This avoids conflicting versions of the same information and lets each decision be traced back to the evidence that supported it.

At each handoff, the model creates an audit point: evidence is verified, assumptions are agreed, risk is priced, and controls are funded only when the case is clear.

8.3 Standing the Function Up: Phases and Gates

The function starts through a short set of phases, each ending with a clear gate before the next phase begins. The phases are the following:

Setup. Assign owners, choose the single source of truth, load the control ontology, and prepare the templates for scoring, pricing, and reporting.

Strategic intake. Agree the current maturity level, key dependencies, and material risk scenarios.

Goal alignment. Price the main scenarios, agree risk reduction targets, and assign clear owners.

Long-term planning. Turn the targets into a roadmap of controls, budgets, and milestones.

After setup, the loop runs in steady state: evidence is verified, trust scores move, risk figures update, and control decisions are recorded.

8.4 Ownership: Who Runs Each Layer

One of the most important things, is clear ownership is essential for the operating model to work. The assurance lead owns the overall function, while OT teams support the evidence layer, finance supports the economic layer, and business owners make the final risk and investment decisions.

The assurance function measures and reports risk but does not accept it. When risk exceeds agreed limits, the responsible risk owner decides whether to accept, transfer, or reduce it. This separation keeps the assurance process independent and accountable.

8.5 Governance: No Silent Risk

The operating model follows one simple rule. The rule is that no risk should remain invisible. A risk becomes dangerous when it is not measured, owned, or acted on. For this reason, missing information or overdue decisions are treated as warning signs, not as silence.

This principle applies to both suppliers and the assurance function itself. Supplier trust falls when fresh evidence stops arriving. In the same way, the function begins to weaken when risk estimates are outdated, decisions are not recorded, maturity assessments are not refreshed, or the system of record no longer reflects reality.

A small set of indicators helps management see whether the function is genuinely operating. These include whether material risks have current financial estimates, whether risks above appetite have an owner's decision, whether control decisions are recorded, and whether maturity assessments remain current. Put simply, continuous assurance must also continuously check itself.

8.6 The Maturity Path From Certification to Commercialisation

Chapter 5 measures how well an organisation monitors its suppliers. The operating model adds a second maturity path, focused on how the assurance function itself develops. As a result the two should be assessed separately, because an organisation may be strong in supplier monitoring but still lack a mature internal function.

The function develops through five stages. Firstly, It begins as an advisory service based on periodic document reviews. Next, It then builds a governance foundation, collects verified evidence, and introduces continuous supplier monitoring. Then the next stage adds financial risk quantification, allowing management to compare exposure with the cost of controls. Once this is established, the function can manage risk and investment across multiple scenarios and business units. At the final stage, the capability is stable enough to be offered as an external service.

We have identified that the order matters. Reliable financial estimates require verified evidence, portfolio decisions require consistent risk figures, and commercialization requires a capability that has already proved itself internally. Put simply, each stage creates the foundation for the next.

8.7 The Operating Model on the AOG Pattern

The AOG case shows why the four layers must work together. Under paper certification, forged documents were accepted for years because no one rechecked the evidence, again no supplier score changed, and no financial exposure was visible, and no decision was forced.

In the operating model, the same problem becomes visible earlier. Certificates and provenance are checked, weak or missing evidence lowers the supplier score, the financial risk is already priced, and any risk above appetite goes to a named owner. The model does not prove it would have caught AOG, but it shows how long-lasting silence becomes a warning signal instead of being ignored.

8.8 Summary and Bridge

This chapter showed how the four layers become one operating model that actually all organizations need to either embed appropriate technology or improve their day to day actions. Evidence is checked, trust is measured, risk is priced, and decisions are assigned to owners. In simple terms, the model makes sure supplier risk does not stay quiet.

Chapter 9 closes the thesis by summarising the contribution, answering the research questions, and setting out the remaining limits and next steps.

9. Conclusion

What was built, what it answers, what it does not, and where it goes next.

9.1 What the Thesis Built

The AOG Technics case shows what that silence costs. A distributor supplied aircraft parts with forged paperwork for years. The paperwork looked correct, so it was accepted on every delivery. The fraud was not caught by a control. It was caught only by chance. In a plant or an aircraft, that kind of gap can hurt people and not only budgets.

The thesis answers this with a capability that never stops working. It has four layers and one operating model that runs them as a single function that every organization can actually adopt.

The evidence layer takes five rulebooks, NIS2, the EU Cyber Resilience Act, EASA Part-IS, IEC 62443, and NIST SP 800-161r1, and maps them onto one set of controls. An auditor can start from a rule and find the control that answers it. An engineer can start from a control and find every rule it satisfies. The layer also solves a practical problem. Plants are kept isolated on purpose, so the evidence is collected inside the plant and pushed out through a one-way device. As a result evidence leaves and at the same time nothing gets in.

The capability layer treats supplier trust as something that goes off. A verified piece of evidence gives the supplier a score. From that day the score falls on its own. On the other hand, fresh evidence lifts it again. But if nothing arrives, the score keeps sliding until it crosses a threshold and someone has to act. Silence stops being invisible and becomes a number on a page. A four-stage ladder then tells an organisation where it stands today and what the next step is.

The economic layer puts a price on what is left. The model runs two hundred thousand possible years and reports what those years look like in money. In the illustrative run, the average year costs about 1.3 million dollars, the worst year in twenty costs about 6.3 million, and the worst year in a hundred costs about 13.6 million. Safety loss is kept as its own line, so the possibility of people being harmed appears in the figure instead of being hidden inside an average.

The fourth layer prices the control rather than the risk. It runs the same model twice, once with paper certification and once with continuous cryptographic provenance. The control costs 750,000 dollars a year. It lowers the average year by about 1.0 million dollars, lowers the worst year in a hundred by about 7.8 million, and leaves a net benefit of about 251,000 dollars a year, a return of 33.5 percent. Hence, that is the moment the work stops being a risk report and becomes a budget request.

Regarding the operating model, it makes the four layers into a function someone can actually run. It names an owner for each layer, sets out the phases and sign off gates for starting the function up, keeps one and only one system of record, and applies one rule throughout: no risk stays silent. It also separates the people who measure risk from the people who accept it.

In one sentence, the thesis moves the assurance from a certificate in a filing cabinet to a working function that verifies, scores, prices, and decides, over and over.

9.2 Answering the Research Questions

The first question asked how five overlapping rulebooks could become one set of controls that can be traced in both directions, without breaking the isolation that keeps a plant or an OT environment safe. The control ontology does the mapping, and the out of band telemetry abstraction moves the evidence out through a one-way path.

The second question asked how supplier assurance could be measured as something continuous instead of a single attestation. The maturity model and the trust decay metric do this by giving trust a score that falls when fresh evidence stops arriving (giving an alert), with thresholds that force a decision.

The third question asked how risk to a safety critical plant could be stated in money when there is very little data. The loss model does this by combining how often something happens with how much it costs, and by keeping the safety costs separate so the worst case is visible to a finance director.

The fourth question asked whether continuous provenance is worth more than paper certification. The before and after running answers it with a figure a board can approve or reject the mitigation plan and understand the investment of the process and its importance .

The fifth question asked how the four layers become one function. Chapter 8 answers it with owners, phases, gates, a single source of truth, the no silent risk rule, and a five stage path from an advisory service to a capability that can be sold.

9.3 Contribution

It needs to be clear that the contribution is not a new finding in any single field. Most of the parts already existed and are used here as published. FABRICS is the work of Slapnicar and Joshi and is used for its framing only. The FAIR family gives the loss structure. C2M2 gives the shape of the maturity ladder. The standards give the obligations.

Two ideas carry the weight and are the thesis's own. The first is that assurance goes off, and that this can be measured. The second is that one unbroken chain can run from a single verified piece of evidence, through a falling trust score, to a monetary figure, to a funded decision, in a way a regulator and a board can both follow. Joining four mature fields into one function that is priced, governed, and possible to operate is what the thesis offers.

9.5 Future Work

The main task is to run the whole model inside a real organisation, with real elicited numbers, and to see whether exposure actually falls over time.

On the evidence side, the mapping should be taken down to clause level and the one-way path built in a live plant.

On the money side, two fixes matter most. Adding a Bayesian fault tree would connect the monetary figure to the specific control that failed, which turns a number into a repair list. Dropping the independence assumption would give a more honest, worst case. Additionally, real data collection should start with the safety exposure share, because that input moves the tail more than any other.

Finally, the same function can be pointed at internal controls rather than only at suppliers, which is the step that turns this design into a service.

9.6 Closing

A certificate speaks for the day it was signed. It cannot speak for months after. The AOG affair shows clearly and validates what that silence is worth.

This thesis argues that assurance must instead be continuous, measured, priced, and owned. It has built that function layer by layer, shown the layers working together on a real case, and set out the conditions under which paying it makes sense.

Point in time certification treats trust as a fact settled once. This work treats trust as it really behaves. More specifically it is something that goes off, secondly has to be renewed with evidence, thirdly can be watched as it fades, fourth can be priced when it fades too far, and finally it belongs to someone whose job is to act before silence turns into an incident.

References

- Boyson, S. (2014). Cyber supply chain risk management: Revolutionizing the strategic control of critical IT systems. *Technovation*, 34, 342-353.
- Boyson, S., Corsi, T. M., and Paraskevas, J. P. (2022). Defending digital supply chains: Evidence from a decade-long research program. *Technovation*, 118, 102380.
- Campbell, R. (2026). AI supply chain security: MBOM-PQC provenance, PQC attestation, and a maturity model for quantum-resistant assurance. Preprints.org (preprint, not peer reviewed). doi:10.20944/preprints202603.1963.v2.
- Dwivedi, S. K., Amin, R., and Vollala, S. (2020). Blockchain based secured information sharing protocol in supply chain management system with key distribution mechanism. *Journal of Information Security and Applications*, 54, 102554.
- Erola, A., Agraftiotis, I., Nurse, J. R. C., Axon, L., Goldsmith, M., and Creese, S. (2022). A system to calculate Cyber Value-at-Risk. *Computers and Security*, 113, 102545.
- EUROCAE and RTCA. (2014). *Airworthiness Security Process Specification (DO-326A / ED-202A) and Airworthiness Security Methods and Considerations (DO-356A / ED-203A)*. RTCA, Washington, DC, and EUROCAE, Paris.
- European Commission. (2022). Directive (EU) 2022/2555 (NIS2) concerning measures for a high common level of cybersecurity across the Union. *Official Journal of the European Union*.
- European Commission. (2023). Commission Implementing Regulation (EU) 2023/203 and Delegated Regulation (EU) 2022/1645 (Part-IS) on information security requirements for civil aviation organisations. *Official Journal of the European Union*.
- ENISA. (2026). *SBOM Adoption State of Play 2026*. European Union Agency for Cybersecurity.
- European Commission. (2024). Regulation (EU) 2024/2847 (Cyber Resilience Act) on horizontal cybersecurity requirements for products with digital elements. *Official Journal of the European Union*.
- Gregor, S., and Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337-355.
- Hevner, A. R., March, S. T., Park, J., and Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105.
- Ho, G. T. S., Tang, Y. M., Tsang, K. Y., Tang, V., and Chau, K. Y. (2021). A blockchain-based system to enhance aircraft parts traceability and trackability for inventory management. *Expert Systems with Applications*, 179, 115101.
- IEC. (2018). *IEC 62443-4-1: Security for industrial automation and control systems, Part 4-1: Secure product development lifecycle requirements*. International Electrotechnical Commission, Geneva.
- Karabacak, B., Ozkan Yildirim, S., and Baykal, N. (2016). A vulnerability-driven cyber security maturity model for measuring national critical infrastructure protection preparedness. *International Journal of Critical Infrastructure Protection*, 15, 47-59.

A Continuous C-SCRM Framework Integrating SBOM Evidence, Provenance Signals, and FAIR-Inspired Risk Quantification, with a Counterfeit-Component Case Study from Civil Aviation

- Kumar, M., Epiphaniou, G., and Maple, C. (2025). Securing additive manufacturing with blockchain-based cryptographic anchoring and dual-lock integrity auditing. *Computers in Industry*, 173, 104395.
- Montecchi, M., Plangger, K., and Etter, M. (2019). It's real, trust me! Establishing supply chain provenance using blockchain. *Business Horizons*, 62(3), 283-293.
- Mura, G., Murru, R., and Martines, G. (2020). Analysis of counterfeit electronics. *Microelectronics Reliability*, 114, 113793.
- NIST. (2022). *Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations (NIST SP 800-161r1)*. National Institute of Standards and Technology, Gaithersburg, MD.
- Nocera, S., Romano, S., Di Penta, M., Francese, R., and Scanniello, G. (2025). On the adoption of software bill of materials in open-source software projects. *The Journal of Systems and Software*, 230, 112540.
- Orlando, A. (2021). Cyber risk quantification: Investigating the role of cyber value at risk. *Risks*, 9(10), 184.
- Parlov, N., Akrap, G., and Esterhajer, D. (2025). Supply chain security and AI risk governance model for critical infrastructure. *Applied Cybersecurity and Internet Governance*, 4(1). doi:10.60097/ACIG/211823.
- Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77.
- Rahman, S., Hossain, N. U. I., Govindan, K., Nur, F., and Bappy, M. (2021). Assessing cyber resilience of additive manufacturing supply chain leveraging data fusion technique: A model to generate cyber resilience index of a supply chain. *CIRP Journal of Manufacturing Science and Technology*, 35.
- Ryan, J. J. C. H., Mazzuchi, T. A., Ryan, D. J., Lopez de la Cruz, J., and Cooke, R. (2012). Quantifying information security risks using expert judgment elicitation. *Computers and Operations Research*, 39(4), 774-784.
- Saini, A., Krishan, K., and Gaur, M. S. (2026). Analyzing ICS security: A survey of design principles, risks, threats, and mitigation methods. *Computers and Electrical Engineering*, 132, 110967.
- SecurityScorecard (Yampolskiy, A., and Routh, J.). Why the FAIR model can be so unfair: CISOs need better risk assessment tools. SecurityScorecard (industry commentary, not peer reviewed).
- Seid, E., Satheesh, S., Popov, O., and Blix, F. (2024). FAIR: Cyber security risk quantification in the logistics sector. *Procedia Computer Science*, 237, 783-792.
- Serious Fraud Office (United Kingdom). (2023-2026). Investigation and prosecution concerning AOG Technics and falsified aircraft parts documentation. Serious Fraud Office, London.
- Siddiqui, F., Ahlbrecht, A., Khan, R., Yengce Tasdemir, S., Hui, H., Sonigara, B., Sezer, S., McLaughlin, K., Zaeske, W., and Durak, U. . Cybersecurity engineering: Bridging the security gaps in avionics architectures and DO-326A/ED-202A. Conference paper produced under the XANDAR project, EU Horizon 2020 grant 957210. Queen's University Belfast and German Aerospace Center (DLR).
- Slapnicar, S., and Joshi, C. (2026). The FABRICS framework: A Bayesian approach to financial quantification of cyber risk. *Computers and Security*, 167, 104900.

A Continuous C-SCRM Framework Integrating SBOM Evidence, Provenance Signals, and FAIR-Inspired Risk Quantification, with a Counterfeit-Component Case Study from Civil Aviation

Sola, A., Sai, Y., Trinchi, A., Chu, C., Shen, S., and Chen, S. (2021). How can we provide additively manufactured parts with a fingerprint? A review of tagging strategies in additive manufacturing. *Materials*, 14.

U.S. Department of Energy. (2022). *Cybersecurity Capability Maturity Model (C2M2) Version 2.1*. Office of Cybersecurity, Energy Security, and Emergency Response, Washington, DC.

U.S. Department of Transportation. (2021). *Departmental guidance on the treatment of the value of a statistical life in economic analysis*. U.S. Department of Transportation, Washington, DC.

Vepraskas, M. J., and Stirgwort, B. W. (2024). A system dynamics approach to building counterfeit avoidance programs. *Computers and Security*, 139, 103711.

Wigger, B., Meissner, T., Foerste, A., Jetter, V., and Zimmermann, A. (2018). Using unique surface patterns of injection moulded plastic components as an image based physical unclonable function for secure component identification. *Scientific Reports*, 8.

Wigger, B., Koinzer, I., Meissner, T., Barth, M., and Zimmermann, A. (2020). Robust and fast part traceability in a production chain exploiting inherent individual surface patterns. *Robotics and Computer-Integrated Manufacturing*, 63, 101925.

Zhang, J., Zio, E., Ma, C., Liu, K., and Wang, W. (2025). A probabilistic cost-benefit analysis approach for cyberattack path evaluation. *Reliability Engineering and System Safety*, 263, 111255.

Appendix A Neutral OT Cyber Risk Model (Chapter 6)

Python source for the Chapter 6 Monte Carlo model. Its baseline run is reused unchanged as the before case in Appendix B.

```
# Appendix A. Neutral OT cyber risk model (Chapter 6).
# Compound frequency and severity Monte Carlo with a separate safety loss term.
# The baseline it produces (seed 20260607) is reused unchanged as the "before"
# case in Appendix B, so the Chapter 7 before figure equals the Chapter 6 headline.
#
# Verified output with this seed:
# probability of at least one event 0.2894
# expected annual loss 1,282,670
# median annual loss 0
# Cyber VaR 95 6,275,735
# Cyber VaR 99 13,555,251
# excluding the safety loss term 1,124,666 / 5,928,771 / 9,600,647

import json
import numpy as np
import matplotlib
matplotlib.use("Agg")
import matplotlib.pyplot as plt
from matplotlib.ticker import FuncFormatter

rng = np.random.default_rng(20260607)
N = 200_000 # simulated years
VSL = 13_200_000 # US DOT value of a statistical life, an explicit design choice

def pert(a, b, c, size):
    """Three point PERT distribution: minimum, most likely, maximum."""
    a, b, c = float(a), float(b), float(c)
    if c <= a:
        return np.full(size, a)
    alpha = 1 + 4 * (b - a) / (c - a)
    beta = 1 + 4 * (c - b) / (c - a)
    return a + rng.beta(alpha, beta, size) * (c - a)

# Frequency: the annual event rate is itself uncertain, so epistemic
# uncertainty is carried through the simulation rather than fixed.
lam = pert(0.10, 0.30, 0.80, N)
k = rng.poisson(lam)
M = int(k.sum())
trial = np.repeat(np.arange(N), k)

# Severity: five components drawn once per event.
downtime = pert(4, 24, 120, M) * pert(20_000, 60_000, 150_000, M)
equipment = (rng.random(M) < 0.40) * pert(50_000, 400_000, 3_000_000, M)
response = pert(50_000, 200_000, 800_000, M)
regulatory = (rng.random(M) < 0.30) * pert(0, 150_000, 2_000_000, M)
safety_exposure = rng.random(M) < 0.05
fatalities = safety_exposure * rng.poisson(0.7, M)
safety_loss = fatalities * VSL

event_total = downtime + equipment + response + regulatory + safety_loss
event_no_safety = downtime + equipment + response + regulatory

annual_total = np.zeros(N)
annual_no_safety = np.zeros(N)
np.add.at(annual_total, trial, event_total)
np.add.at(annual_no_safety, trial, event_no_safety)

out = {
    "p_at_least_one_event": float(np.mean(k > 0)),
    "mean_total": float(annual_total.mean()),
    "median_total": float(np.median(annual_total)),
```

```

"var95_total": float(np.percentile(annual_total, 95)),
"var99_total": float(np.percentile(annual_total, 99)),
"mean_no_safety": float(annual_no_safety.mean()),
"var95_no_safety": float(np.percentile(annual_no_safety, 95)),
"var99_no_safety": float(np.percentile(annual_no_safety, 99)),
}
print(json.dumps(out, indent=2))

def money(x, _=None):
    if abs(x) >= 1_000_000:
        return f"${x / 1_000_000:.1f}M"
    if abs(x) >= 1_000:
        return f"${x / 1_000:.0f}k"
    return f"${x:.0f}"

def exceedance(a):
    s = np.sort(a)[::-1]
    return s, np.arange(1, len(s) + 1) / len(s)

# Figure 6.1: annual loss exceedance curve.
st, pt = exceedance(annual_total)
sn, pn = exceedance(annual_no_safety)
fig, ax = plt.subplots(figsize=(9.6, 4.9), dpi=220)
ax.plot(st, pt, color="#14314E", lw=2.4, label="Total annual loss, with safety loss term")
ax.plot(sn, pn, color="#2F6F5E", lw=2.1, ls=(0, (6, 3)), label="Excluding the safety loss term")
ax.set_xscale("log")
ax.set_xlim(10_000, max(out["var99_total"] * 1.4, 5_000_000))
ax.set_ylim(0, 0.35)
ax.set_xlabel("Annual loss (USD, log scale)")
ax.set_ylabel("Probability annual loss is exceeded")
ax.xaxis.set_major_formatter(FuncFormatter(money))
ax.grid(True, which="both", ls=":", color="#D5DDE4", alpha=0.8)
for v, lab, col in [(out["var95_total"], "Cyber VaR 95", "#9A6A2F"),
                    (out["var99_total"], "Cyber VaR 99", "#9E2B25")]:
    ax.axvline(v, color=col, ls=":", lw=1.5)
    ax.annotate(f"{lab}\n{money(v)}", xy=(v, 0.305), fontsize=9.2, color=col,
               ha="center", va="top",
               bbox=dict(boxstyle="round,pad=0.28", fc="white", ec=col, lw=1))
ax.legend(fontsize=9.4, loc="upper right", framealpha=0.96)
fig.savefig("fig_61.png", dpi=220, bbox_inches="tight", facecolor="white")
plt.close(fig)

# Figure 6.2: the effect of the safety loss term on the executive figures.
labels = ["Expected annual loss", "Cyber VaR (95th)", "Cyber VaR (99th)"]
with_safety = [out["mean_total"], out["var95_total"], out["var99_total"]]
without_safety = [out["mean_no_safety"], out["var95_no_safety"], out["var99_no_safety"]]
x = np.arange(3)
w = 0.36
fig, ax = plt.subplots(figsize=(9.2, 4.5), dpi=220)
b1 = ax.bar(x - w / 2, without_safety, w, label="Excluding the safety loss term",
            color="#E4EDE9", edgecolor="#2F6F5E", linewidth=1.4)
b2 = ax.bar(x + w / 2, with_safety, w, label="With the safety loss term",
            color="#E8EEF4", edgecolor="#14314E", linewidth=1.4)
ax.set_xticks(x)
ax.set_xticklabels(labels)
ax.yaxis.set_major_formatter(FuncFormatter(money))
ax.set_ylabel("Annual loss (USD)")
ax.grid(True, axis="y", ls=":", color="#D5DDE4", alpha=0.8)
for bars in (b1, b2):
    for r in bars:
        ax.annotate(money(r.get_height()),
                    xy=(r.get_x() + r.get_width() / 2, r.get_height()),
                    xytext=(0, 3), textcoords="offset points",
                    ha="center", fontsize=9.2, color="#14314E")
ax.legend(fontsize=9.4, loc="upper left", framealpha=0.96)

```

```
fig.savefig("fig_62.png", dpi=220, bbox_inches="tight", facecolor="white")  
plt.close(fig)  
  
print("figures written: fig_61.png, fig_62.png")
```

Appendix B Counterfeit Grounding Value at Risk and Provenance ROI (Chapter 7)

Python source for the Chapter 7 before and after model. The baseline uses the same seed as Chapter 6, so the before figure equals the Chapter 6 headline.

```
# Appendix B. Counterfeit grounding value at risk and the provenance investment case.
#
# The "before" run is the SAME baseline as Chapter 6, using an identical seed and
# draw order rather than an independent re draw. That is what makes the Chapter 7
# before figure equal the Chapter 6 headline of about 13.6M at the 99th percentile.
#
# The control is a recurring annual cost set against a recurring reduction in
# expected loss, so the case is reported as a net annual benefit and an annual
# return on investment. There is no upfront investment in this model, so no
# payback period is claimed. If an operator also faces a one time implementation
# cost, set IMPLEMENTATION_COST below and a payback period is then computed
# against the net annual benefit.
#
# Verified output with these seeds:
# before p(event) 0.2894 mean 1,282,670 VaR95 6,275,735 VaR99 13,555,251
# after p(event) 0.0792 mean 281,396 VaR95 2,421,512 VaR99 5,714,404
# reduction in expected annual loss 1,001,275
# reduction in the 99th percentile 7,840,848
# net annual benefit 251,275
# annual return on investment 33.5 percent

import json
import numpy as np
import matplotlib
matplotlib.use("Agg")
import matplotlib.pyplot as plt
from matplotlib.ticker import FuncFormatter

N = 200_000
VSL = 13_200_000
CONTROL_COST = 750_000 # illustrative annual cost of continuous provenance
IMPLEMENTATION_COST = 0 # optional one time cost; 0 means not modelled
BASELINE_SEED = 20260607 # identical to the Chapter 6 baseline
POST_SEED = 202

def simulate(rate_lo, rate_mode, rate_hi, safety_p, seed):
    """One run of the Chapter 6 model with a given frequency and safety exposure."""
    g = np.random.default_rng(seed)

    def pert(a, b, c, size):
        a, b, c = float(a), float(b), float(c)
        if c <= a:
            return np.full(size, a)
        alpha = 1 + 4 * (b - a) / (c - a)
        beta = 1 + 4 * (c - b) / (c - a)
        return a + g.beta(alpha, beta, size) * (c - a)

    lam = pert(rate_lo, rate_mode, rate_hi, N)
    k = g.poisson(lam)
    M = int(k.sum())
    trial = np.repeat(np.arange(N), k)

    downtime = pert(4, 24, 120, M) * pert(20_000, 60_000, 150_000, M)
    equipment = (g.random(M) < 0.40) * pert(50_000, 400_000, 3_000_000, M)
    response = pert(50_000, 200_000, 800_000, M)
    regulatory = (g.random(M) < 0.30) * pert(0, 150_000, 2_000_000, M)
    safety = (g.random(M) < safety_p) * g.poisson(0.7, M) * VSL

    event = downtime + equipment + response + regulatory + safety
    annual = np.zeros(N)
```

```

np.add.at(annual, trial, event)
return annual

# Before: paper certification, a counterfeit component reaching a critical
# system. This is the Chapter 6 baseline itself.
base = simulate(0.10, 0.30, 0.80, 0.05, BASELINE_SEED)

# After: continuous cryptographic provenance cuts the frequency sharply and
# lowers the safety exposure share, because an undetected counterfeit is far
# less likely to reach a critical system.
post = simulate(0.02, 0.07, 0.20, 0.02, POST_SEED)

def stats(a):
    return {"mean": float(a.mean()),
            "var95": float(np.percentile(a, 95)),
            "var99": float(np.percentile(a, 99)),
            "p_event": float(np.mean(a > 0))}

b, p = stats(base), stats(post)
reduction_mean = b["mean"] - p["mean"]
reduction_var99 = b["var99"] - p["var99"]
net_annual_benefit = reduction_mean - CONTROL_COST
annual_roi_pct = net_annual_benefit / CONTROL_COST * 100.0
payback_years = (IMPLEMENTATION_COST / net_annual_benefit
                  if IMPLEMENTATION_COST > 0 and net_annual_benefit > 0 else None)

print(json.dumps({"baseline": b, "post": p, "control_cost": CONTROL_COST,
                  "reduction_mean": reduction_mean,
                  "reduction_var99": reduction_var99,
                  "net_annual_benefit": net_annual_benefit,
                  "annual_roi_pct": annual_roi_pct,
                  "payback_years": payback_years}, indent=2))

def money(x, _=None):
    if abs(x) >= 1_000_000:
        return f"${x / 1_000_000:.1f}M"
    if abs(x) >= 1_000:
        return f"${x / 1_000:.0f}k"
    return f"${x:.0f}"

def exceedance(a):
    s = np.sort(a)[::-1]
    return s, np.arange(1, len(s) + 1) / len(s)

# Figure 7.1: loss exceedance before and after the control.
sb, qb = exceedance(base)
sp, qp = exceedance(post)
fig, ax = plt.subplots(figsize=(9.6, 4.9), dpi=220)
ax.plot(sb, qb, color="#9E2B25", lw=2.5, label="Before: paper certification")
ax.plot(sp, qp, color="#2F6F5E", lw=2.5, label="After: continuous cryptographic provenance")
ax.set_xscale("log")
ax.set_xlim(10_000, max(b["var99"] * 1.4, 5_000_000))
ax.set_ylim(0, 0.35)
ax.set_xlabel("Annual loss (USD, log scale)")
ax.set_ylabel("Probability annual loss is exceeded")
ax.xaxis.set_major_formatter(FuncFormatter(money))
ax.grid(True, which="both", ls=":", color="#D5DDE4", alpha=0.8)
ax.axvline(b["var99"], color="#9E2B25", ls=":", lw=1.4)
ax.axvline(p["var99"], color="#2F6F5E", ls=":", lw=1.4)
ax.annotate("VaR 99n" + money(b["var99"]), xy=(b["var99"], 0.305), fontsize=9.2,
            color="#9E2B25", ha="center", va="top",
            bbox=dict(boxstyle="round,pad=0.28", fc="white", ec="#9E2B25", lw=1))
ax.annotate("VaR 99n" + money(p["var99"]), xy=(p["var99"], 0.20), fontsize=9.2,

```

```

        color="#2F6F5E", ha="center", va="top",
        bbox=dict(boxstyle="round,pad=0.28", fc="white", ec="#2F6F5E", lw=1))
ax.legend(fontsize=9.4, loc="upper right", framealpha=0.96)
fig.savefig("fig_71.png", dpi=220, bbox_inches="tight", facecolor="white")
plt.close(fig)

# Figure 7.2: the investment case as gross reduction, control cost, net benefit.
fig, ax = plt.subplots(figsize=(8.6, 4.5), dpi=220)
ax.bar(0, reduction_mean, 0.58, color="#E8EEF4", edgecolor="#14314E", linewidth=1.4)
ax.bar(1, -CONTROL_COST, 0.58, bottom=reduction_mean, color="#F6E7E6",
        edgecolor="#9E2B25", linewidth=1.4)
ax.bar(2, net_annual_benefit, 0.58, color="#E4EDE9", edgecolor="#2F6F5E", linewidth=1.4)
ax.plot([0.29, 0.71], [reduction_mean, reduction_mean], color="#6B7A88", lw=1, ls=":")
ax.plot([1.29, 1.71], [net_annual_benefit, net_annual_benefit], color="#6B7A88", lw=1, ls=":")
ax.set_xticks(range(3))
ax.set_xticklabels(["Gross annual\nrisk reduction", "Annual control\ncost",
                    "Net annual\nbenefit"])
ax.yaxis.set_major_formatter(FuncFormatter(money))
ax.set_ylim(0, reduction_mean * 1.22)
ax.set_ylabel("Annual value (USD)")
ax.grid(True, axis="y", ls=":", color="#D5DDE4", alpha=0.8)
ax.text(2, reduction_mean * 1.12,
        f"annual return on investment {annual_roi_pct:.1f}%",
        ha="center", fontsize=9.6, color="#2F6F5E", style="italic")
fig.savefig("fig_72.png", dpi=220, bbox_inches="tight", facecolor="white")
plt.close(fig)

print("figures written: fig_71.png, fig_72.png")

```

Appendix C Operating Model Onboarding Schedule (Chapter 8)

A worked schedule for standing up the continuous assurance operating model of Chapter 8 in a single business unit over the first quarter. It is anonymized and illustrative. More specifically, the horizons, owners, and gates show the pattern by which the function is deployed, and not a calendar bound to any one organization. The four layers of the thesis map onto the activities as marked. Each phase ends in a gate that must be signed off before the next phase begins. Finally, gate rows are the sign off points at which the operating model can be audited.

Table C.1 Onboarding schedule for the continuous assurance operating model, by 30 day horizon, with phase gates. All dates are relative; owners are roles, not individuals.

Horizon	Phase / gate	Activity	Output	Owner
Pre start	Setup	Assign a role to each layer; choose and stand up the single source of truth; load the control ontology (evidence layer) as the control backbone; prepare scoring, pricing and reporting templates.	Operating basis: owners named, one system of record, ontology loaded.	Assurance lead
Pre start	SETUP GATE	Confirm there is one owner per layer and one and only one system of record before any assurance work begins.	Setup sign off.	Assurance lead + sponsor
0-30	Strategic intake	Capture business objectives and the supply chain and operational dependencies that threaten them; run an honest maturity self assessment (capability layer).	Baseline: current maturity position and material risk scenarios.	Assurance lead + business
0-30	PHASE 1 GATE	Agree the as-is maturity baseline and the material scenarios the function will manage.	Phase 1 sign off.	Assurance lead + sponsor
31-60	Goal alignment	Price each material scenario with the economic layer to a Cyber VaR range; attach a risk reduction objective with a residual exposure movement target to each; assign a named owner per scenario.	Priced scenario portfolio; risk reduction objectives with owners.	Assurance lead + finance
31-60	Goal alignment	Run the business and assurance handshake: business and the function agree the scenarios, the figures, and the reduction targets.	Handshake sign off; owned reduction targets.	Assurance lead + business
31-60	PHASE 2 GATE	Confirm scenarios, prices, and reduction targets are agreed and owned before planning.	Phase 2 sign off.	Assurance lead + sponsor

A Continuous C-SCRM Framework Integrating SBOM Evidence, Provenance Signals, and FAIR-Inspired Risk Quantification, with a Counterfeit-Component Case Study from Civil Aviation

Horizon	Phase / gate	Activity	Output	Owner
61-90	Long-term planning	Turn agreed targets into a multi horizon roadmap of controls; justify budgeted controls with the provenance investment case (economic layer); ensure two way traceability objectives to gaps to scenarios to targets.	Sponsor ready multi horizon roadmap; funded control decisions.	Assurance lead + sponsor
61-90	PHASE 3 GATE	Confirm a traceable, sponsor-ready plan; hand the function into steady state.	Phase 3 sign off; steady state start.	Assurance lead + sponsor
Steady state	Operate (loop)	Evidence arrives and is verified; trust scores move; figures update; controls are proposed, funded, and re-measured. No silent risk indicators reported on one page with supplier indicators; degradation of the function's own artefacts escalates to the sponsor.	Running assurance function; one page governance and supplier indicators.	Assurance lead

The schedule is the function's birth. Specifically, the steady state loop is its life. After the final gate the phases are not repeated unless the function is re-based or extended to a new business unit.