MRIS v2.0 — Mythos-Resistant Information Security

MRIS

Effectiveness assessment of existing ISO 27001 Annex A controls under a Gen-AI-accelerated offensive

Framework-spanning · Evidence-based

Version 2.0 | August 2026

Author: Richard Peddi (produced with AI assistance)

AUDIENCE

CISOs, ISMS owners and security architects who need to test their existing control implementations against the new reality of agentic and Gen-AI-accelerated attacks

License and disclaimer

Disclaimer

This handbook is a technical and organizational reference. It does not replace an individual risk analysis, legal advice or an audit-specific examination. The assessment of the effectiveness of individual controls in their typical implementation refers to the state of agentic AI threats publicly documented at the time of writing. The assessment may shift as new evidence emerges.

Note on how this document was produced

This document was produced with the support of AI tools. The selection and evaluation of sources, the technical classifications and the responsibility for the content rest with the author.

Recommended citation

Peddi, Richard (2026): MRIS – Mythos-Resistant Information Security, Version 2.0. mris.info

About the author

Richard Peddi is a certified IT specialist for systems integration (Fachinformatiker) and a Certified IT Business Manager (Operative Professional), and holds an M.Sc. in Information Security Management. Among other qualifications he is an ISO/IEC 27001 Lead Auditor. He has worked as a security consultant in the critical infrastructure (KRITIS) sector and as a CISO before taking up his current role as Group CISO. MRIS is produced independently of his professional employment.

Table of contents

Foreword

In November 2025, Anthropic documented GTG-1002, the first cyberattack in which Claude Code autonomously executed 80 to 90 percent of the tactical attack work across roughly 30 targets. In April 2026, Claude Mythos Preview and the Glasswing defensive program confirmed the pattern: AI models find vulnerabilities and chain them into working exploits at a pace that structurally overwhelms classic response cycles.

This raises the central question of this handbook: which of the controls anchored in an ISMS today still hold up against that reality — and which do not?

Methodologically, MRIS builds on ISO/IEC 27001:2022 and the risk methodology of ISO/IEC 27005:2022. It assesses all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape, compares them primarily with NIST, BSI C5:2026, DORA, the CRA, the NIS2 Directive and the NIS2 Implementing Regulation, and delivers a catalogue of thirteen Mythos Hardening Controls (MHC) with maturity levels and audit-ready thresholds. Further standards and regulatory frameworks drawn on — including BSI IT-Grundschutz, BSI TR-02102 and TR-03183, ISO/IEC 42001, MITRE ATT&CK and D3FEND, OWASP, the CIS Controls, SLSA, FIDO2/WebAuthn and the GDPR — are set out in full in Chapters 3.3 and 3.4.

MRIS is not a full ISMS framework and replaces neither ISO/IEC 27001 nor ISO/IEC 27005. It adds an effectiveness and reality layer to an existing ISMS (see Chapter 1.4 on its position in the ISMS stack).

In April 2026, a working group from the Cloud Security Alliance, the SANS Institute, [un]prompted and the OWASP Gen AI Security Project published the strategy paper "The AI Vulnerability Storm: Building a Mythos-ready Security Program". MRIS and Mythos-Ready are complementary: Mythos-Ready defines the risk landscape at the level of industry consensus, while MRIS translates it into a control-by-control assessment and an audit-ready hardening catalogue.

Since May 2026 this capability has no longer been a single-vendor observation: with OpenAI Daybreak and Microsoft MDASH, agentic vulnerability discovery is a multi-vendor reality, and with Claude Fable 5 the Mythos class is generally available. Chapter 2.5 places the developments since first publication in context.

This handbook is produced as an independent work outside my professional employment. It is published at mris.info; comments and criticism are expressly welcome.

My thanks go to Daniel Heldt for the critical technical review of version 1.8 and the improvements that grew out of it.

Richard Peddi

August 2026

Management summary

Central finding

The established methods of information security have not become wrong — but one assumption underlying them no longer holds: that an attack is slow, laborious and expensive for the attacker. That braking effect disappears as soon as attackers deploy artificial intelligence. A substantial part of today's protection level therefore rests on a foundation that no longer carries.

The threat landscape has shifted

Four changes matter at leadership level:

- **Speed:** The time between a vulnerability becoming known and its exploitation is shrinking from weeks to hours. Response windows that used to be sufficient have become too short.
- **Automation:** Attackers have software perform the bulk of the work autonomously — faster than a team can react.
- **Concealment:** Attacks are broken down into many small steps, each unremarkable on its own, that individually trigger no alarm.
- **Availability:** Attack capabilities once available only to state actors are now broadly accessible. The circle of attackers to be taken seriously is growing.

This calls for an adjustment of security management

The answer is not a new security strategy but a targeted recalibration of the existing information security management system (ISMS). Where protection has so far relied on attacks being too laborious for the adversary, it must be moved onto more resilient foundations. This document sets out, in a structured and prioritized way, which existing measures still hold, which need to be sharpened, and where targeted additions are needed. The effort builds on what already exists and is manageable — but it requires the deliberate decision to take it on now.

1 Introduction

1.1 Why this handbook exists

The occasion for this handbook is a concrete and publicly confirmed change in the reality of attacks. Between August 2025 and April 2026, Anthropic documented in several threat intelligence reports how Claude models were used in real campaigns: as autonomous orchestrators of a cyber espionage campaign against roughly 30 global targets, for exploit generation, and for the automated creation of phishing content and malware. With the publication of the Glasswing program and the first assessment of Claude Mythos Preview in April 2026, the vendor stated publicly that models can now find and exploit vulnerabilities in code at a pace that structurally overwhelms classic response processes.

For information security this represents a break that has already occurred. The attacking side operates in time windows that were considered unrealistic three years ago: patch reversing to a working exploit in hours instead of days, intrusion operations with 80 to 90 percent autonomous execution, and the decomposition of complex attacks into micro-steps that each appear legitimate.

This raises the question that every ISMS ownership function must answer: which of the controls anchored in an ISMS today still hold up against this reality? And where does the mere existence of a control suggest a protective effect that the Mythos attacker long ago stopped respecting?

This handbook answers that question through a systematic assessment of all 93 controls of ISO/IEC 27002:2022, supplemented by the perspectives of the five primary comparison frameworks: NIST, the BSI C5:2026 audit catalogue, DORA, the CRA, and the NIS2 Directive together with Implementing Regulation (EU) 2024/2690; the full framework stack including secondary references is set out in Chapter 3.3. The added value lies in the convergence analysis: wherever a framework states a requirement, or states it more explicitly, that ISO 27002 does not contain in that form or only conditionally implies, and where that requirement holds up under Mythos conditions, there is a hardening proposal for ISO-centric ISMS.

What began in April 2026 as an Anthropic-centric event has been a multi-vendor market since May 2026 (Chapter 2.5). The core theses of this handbook remain unaffected; their empirical basis has broadened.

Terminology: "MRIS" (Mythos-Resistant Information Security) denotes the independent framework presented here; this handbook is its reference edition. Wherever "MRIS" appears below, the framework is meant, not this single document.

1.2 The central thesis: methodology intact, inputs destabilized

ISO/IEC 27005:2022 defines risk as the "effect of uncertainty on objectives" (3.1.3); the *level of risk* (3.1.15) results from the combination of consequence and likelihood. For estimating the likelihood of deliberate risk sources, the standard (Clause 7, cf. Annex A.2.3) points in particular to motivation (the cost-benefit calculation of the attack), the capabilities and resources of the actor, and the exploitability of the vulnerability. This methodology is not wrong — what has changed are its inputs.

Note on an addition by the author: To the risk variables that ISO/IEC 27005:2022 draws on for deliberate sources — likelihood, consequence and attacker capability (motivation, capabilities, resources) — this handbook introduces a further dimension: the *time between threat event and organizational response*. The standard does not name this time factor explicitly as a risk parameter; it serves below as an additional assessment criterion, because Mythos attacks structurally push response time below the threshold a human process can handle.

Mythos-class AI substantially lowers the capacity barrier between opportunistic and advanced actors for many attack classes, compresses the time between vulnerability disclosure and exploit to hours, fragments attacks into micro-events that individually raise no alarm, and systematically devalues controls whose protective effect rests only on limited attacker patience. Physical and hardware-level attacks — fault injection on chips, or access to isolated systems — remain largely unaffected. The central thesis is therefore: the ISO/IEC 27005 methodology remains applicable, but the likelihood and effectiveness parameters fed into it must be reassessed — each control examined as to whether it derives its protective effect from a hard barrier or merely from friction, and whether that friction still suffices under Mythos conditions.

1.3 The central claim

Wherever NIST, C5, DORA, the CRA, the NIS2 Directive or the NIS2 Implementing Regulation states a requirement, or states it more explicitly, that ISO 27002 does not contain in that form or only conditionally implies, and that requirement holds up under Mythos conditions, it constitutes a concrete hardening proposal for any ISO-centric ISMS. The CISO task shifts from depth of compliance to an effectiveness-oriented search for controls.

1.4 Position in the ISMS stack and limits of this handbook

MRIS is not a management system standard and does not claim to represent a complete ISMS. This self-positioning is explicitly part of the methodology.

MRIS as an effectiveness and reality layer on an ISMS foundation. MRIS presupposes an established ISMS — typically to ISO/IEC 27001:2022, BSI IT-Grundschutz, BSI C5 2026 or an equivalent framework — and on that basis answers a single question: do the controls implemented today still work against a Gen-AI-accelerated offensive? Without an ISMS foundation, MRIS is not applicable.

What MRIS delivers:

- An effectiveness assessment of all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape (Part II).
- A gap analysis against the primary frameworks (BSI C5:2026, NIST, DORA, the CRA, NIS2 with its Implementing Regulation) (Chapter 8).
- An audit-ready catalogue of thirteen Mythos Hardening Controls with maturity levels, tooling and thresholds (Chapter 9).
- Operationalization in four annexes: assessment matrix (A), standalone MHC worksheet (C), RACI model (E), KPI definitions (G). The assessment matrix and the MHC catalogue are additionally available in a machine-readable edition (YAML, CC BY-SA 4.0).

What MRIS deliberately does not deliver:

- No complete ISMS to ISO/IEC 27001 Clauses 4–10 (context, leadership, planning, support, operation, evaluation, improvement). These requirements are presupposed, not replaced.
- No risk management process of its own. Through the assessment categories (Chapter 3) and the risk assessment bridge (Annex F), MRIS supplies inputs to the risk process to ISO/IEC 27005 but does not replace it.
- No complete PDCA management system. Through the maturity levels (Chapter 9.5), versioning and the continuous audit recommendations (MHC-10), MRIS supplies building blocks for Plan-Do-Check-Act but does not replace a management system in its own right.
- No compliance mapping tool. The framework comparisons for individual controls are effectiveness cross-references, not a complete requirement-to-standard mapping with an audit trail. For certification and regulatory evidence, specialized GRC tools should be used.
- No organization-specific policy hierarchy. Policies, procedures and work instructions must be developed within the organization; MRIS supplies substantive anchors for them, but no templates.

Correct use. MRIS is intended as a layer on an existing ISMS: the ISMS supplies governance, the risk management process, the policy hierarchy and the PDCA cycle. MRIS supplies the effectiveness assessment of the individual controls against Mythos and concrete, audit-ready hardening recommendations. Keeping the two layers separate is methodologically decisive — MRIS is not meant to duplicate ISMS frameworks but to supplement them precisely where their technology-agnostic design reaches its limits in the face of Gen-AI-accelerated threats.

2 Findings: the four shifts in the reality of attacks

The following four findings are not hypotheses. They rest on the publicly documented threat intelligence reports published by Anthropic between August 2025 and April 2026, and on the defensive recommendations of the Anthropic security engineering team within the Glasswing program.

2.1 Collapse of the patch gap window

Classic vulnerability management programs assume a time window between the publication of a patch and the availability of a working exploit. In practice this window was usually measured in days to weeks. It permitted staggered patch cycles, manual approval processes and selective prioritization by criticality.

Anthropic describes patch reversing to a working exploit as precisely the kind of mechanical analysis in which current models excel. The published defensive recommendation consistently calls for internet-exposed systems to be patched within 24 hours (as a guide value) of an exploit becoming available. For every control that relies on an orderly patch cycle with manual approval, the effectiveness assumption is broken.

2.2 Compression of the time axis

ISO 27005, like many incident response playbooks, presupposes that a humanly manageable amount of time lies between detection, triage and decision. That time no longer exists.

In the documented GTG-1002 case, Claude Code executed 80 to 90 percent of the tactical attack steps autonomously, at request rates unattainable for human operators. Controls with a human in the loop in the immediate response chain do not become ineffective because of poor quality, but because their response time is no longer of the same order of magnitude as the speed of the attack.

2.3 Fragmentation of the threat event

The risk model of ISO 27005 presupposes a discrete, identifiable threat event against which controls take effect. That precondition no longer holds under agentic attack patterns. For GTG-1002, Anthropic explicitly describes how the attacker decomposed the overall attack into sub-tasks — vulnerability scanning, credential validation, lateral movement, data extraction — each of which appeared as a legitimate technical request.

The risk materializes only in the aggregate. Controls based on evaluating discrete events — signature-based detection without behavioural context, rate limits with naive thresholds, policy-based alerting on individual

actions — lose effectiveness even though they function technically. They see what they are meant to see. They do not see that the sum amounts to an attack.

2.4 Decoupling of capability and actor

Classic likelihood assessment rests substantially on an estimate of attacker capability: a nation-state actor was considered unlikely against a mid-sized German software vendor, not because the interest was absent but because the technical capacity for an individualized campaign was scarce. Mythos-class AI has largely dissolved this coupling between actor type and available capacity.

The threat intelligence reports show that actors with little technical competence of their own are now able to carry out operations that would previously have been attributed only to professional APT groups. The likelihood axis in every risk matrix shifts upward — not because of new threat actors, but because the capacity barrier facing existing actors has fallen.

2.5 Developments since first publication: from a single-vendor event to a multi-vendor market

Between May and July 2026, an Anthropic-centric observation became a market. Four developments are relevant to the assessments in this handbook:

OpenAI Daybreak (11 May 2026). OpenAI combines the GPT-5.5 model family with "Codex Security" as an agentic harness into a system for threat modelling, vulnerability discovery, exploit validation in isolated environments and patch generation; access is tiered (including "GPT-5.5-Cyber" for authorized red teaming).

Microsoft MDASH (12 May 2026). The "Multi-Model Agentic Scanning Harness" orchestrates more than 100 specialized agents across an ensemble of frontier and distilled models. First volume evidence: 16 previously unknown vulnerabilities in the Windows networking and authentication stack, among them four critical RCEs (including CVE-2026-33824, CVSS 9.8), patched in the May patch day; 88.45% on the public CyberGym benchmark against 83.1% for Mythos Preview.

General availability of the Mythos class (9 June / 1 July 2026). With Claude Fable 5, a Mythos-class model is generally available — with cyber safeguard classifiers; the variant without those safeguards (Claude Mythos 5) remains reserved for Glasswing partners. What matters for the assessment: the capability class is diffusing broadly, and the defensive lead conferred by an exclusive preview is progressively disappearing. This sharpens the exploit economics derived in 2.1 and 2.2.

Disclosure waves (from July 2026). Project Glasswing reported more than 10,000 high/critical findings for the first program phase; by 22 May 2026, 1,596 findings had been disclosed to maintainers of 281 open source projects, 97 of them patched. The announced public 90-day report (July 2026) and expiring embargo periods are producing CVE waves in H2 2026. In parallel, the NVD reports structural overload (submissions up 263% against 2020; since April 2026 enrichment has been prioritized to the highest risks), and HackerOne paused its Internet Bug Bounty program in March 2026. The volume shift forecast in MHC-02 (SBOM query capability) and MHC-09 (AI-supported testing) has thus materialized.

Consequence for this handbook. None of the 93 individual assessments and none of the thirteen MHC rationales is rendered wrong by these developments — they are confirmed and sharpened. From this version onward, "Mythos" is explicitly treated as a class term (see glossary): agentic frontier models with autonomous vulnerability discovery and exploit chaining, irrespective of vendor. Sources: list of references (accessed 19 July 2026).

2.6 Consequence for ISO/IEC 27005

The methodology of ISO 27005 — risk identification, analysis, evaluation, treatment — remains intact as a process. Its inputs, however, are affected: the likelihood axis shifts upward as the capacity barrier disappears. The time parameters between patch and exploit collapse. The aggregation logic of individually assessed events no longer holds. And the effectiveness assumption behind many existing controls — particularly those resting on attacker friction rather than on hard impossibility — has been empirically refuted.

Note: In Clause 7.2.1, ISO/IEC 27005:2022 distinguishes two approaches to risk identification — the event-based and the asset-based approach. The event-based approach is particularly affected under Mythos conditions,

because it constructs scenarios from individual threat sources that, in a micro-step attack, become visible only in the aggregate.

3 Assessment methodology

3.1 The four-category grid

Each control is assessed along four categories. The terminology follows the language in which Anthropic itself describes the limit of effectiveness. In the Glasswing defensive post (April 2026), the security engineering team states that mitigations whose value lies in generating friction lose considerable effectiveness against an adversary with unlimited patience.

3.1.1 Robust

Controls whose protective effect derives from a hard barrier: cryptographic impossibility, physical separation, the architectural non-existence of an attack path. These controls hold because they structurally exclude attacks or make them fail at an objective limit.

3.1.2 Partially degraded

Controls that continue to provide protection but whose originally assumed strength falls significantly under Mythos conditions. They remain worthwhile but require supplementation — through additional controls, changed parameters or architectural flanking.

3.1.3 Friction-only

Controls whose classic or commonly encountered implementation form no longer provides sufficient standalone protective effect under Mythos conditions, because the underlying protective mechanism either (a) relies on limited attacker capacity (classic friction) or (b) presupposes human reaction time within an automatically executed kill chain. The classification targets the prevailing implementation form, not the control text: an automated, tool-supported implementation of the same control can indeed provide protective effect — that implementation is then precisely the hardening MRIS describes in the corresponding replacement recommendation. In both cases, the asymmetry between a model-driven attacking side and a defending side that is not automated to the same degree causes the control to lose its core effect. These controls must either be replaced or supplemented by barriers based not on effort or reaction time but on structural impossibility. The classification refers to the operational mechanism of a control. Organizational process shells — ownership, SLA definition, escalation paths — can remain effective even for friction controls; they then do not count as standalone protective effect but are carried forward as a precondition of the flanking MHC (example: A.8.8, whose accountability and escalation structure forms the basis for the automation under MHC-09).

3.1.4 Not affected

Controls that are neutral with respect to Mythos-specific threats — typically organizational, documentary, governance-oriented or physically bound controls whose effectiveness is systematically changed neither by friction arguments nor by agentic acceleration. They remain important but receive no separate Mythos hardening recommendation.

3.1.5 Note on the context dependency of the classification

The four-category classification applies expressly with respect to the Mythos threat landscape in the form described in Chapter 2. A control may fall into different categories against different attack vectors. Example: A.5.17 (authentication information) is robust against credential stuffing once phishing-resistant MFA is implemented — the origin binding of WebAuthn is a cryptographic hard barrier. Against session hijacking after a successful login, the same control is effectiveness-neutral, because authentication has already been completed. Against agentic push notification storms (MFA fatigue), classic push MFA generates only friction. The assessment in Chapters 4 to 7 takes the dominant Mythos effectiveness mode of a control as the basis for classification. Before adopting the results into the Statement of Applicability, auditors and ISMS owners check whether divergent classifications are justified in their specific threat situation.

3.2 Assessment criteria

What is assessed is not the abstract control text of ISO/IEC 27001 Annex A, but the effectiveness of the control in its typical implementation according to the guidance of ISO/IEC 27002; "degraded" therefore denotes a weakened effectiveness of the measure, not a deficiency of the control itself. For readability, the short form "control" is used throughout.

Structural protective capability and implementation maturity. MRIS distinguishes between the structural protective capability of a control and the organization-specific implementation maturity. The assessment describes the structural effectiveness a control can achieve under Mythos conditions given an implementation that reflects the state of the art. The classification "Robust" therefore expressly does not mean that every existing implementation of that control is robust — it means that the control can create a hard barrier when implemented properly. Where the classification presupposes a particular implementation form (for example append-only logging for A.8.15 or declarative configuration management for A.8.9), this is named in the Mythos finding. Actual effectiveness must be examined organization-specifically against the measures implemented; MRIS does not replace that examination.

Each control is assessed along four criteria derived from the four findings in Chapter 2:

- Attacker patience: does the protective effect rest on an attack being too laborious for an adversary? If so, the control is degraded under Mythos conditions.
- Time compression: does the control presuppose human reaction time in a chain in which the attacker now operates autonomously? If so, effectiveness is impaired by the latency differential.
- Capability decoupling: does the control's likelihood assumption depend on a historical correlation between actor and capability that Mythos has dissolved? If so, the likelihood must be reassessed.
- Aggregation resistance: can the control still take effect when the attack is decomposed into many micro-steps that each appear legitimate? If not, the control is blind to fragmented attack patterns. Operationalization: aggregation resistance is deemed present if the control either (a) makes integrated use of correlation mechanisms across time, identities, resources or action chains — such as SIEM correlation rules with time windows, UEBA baselines, graph-based identity and data flow analysis, or kill chain tracking to MITRE ATT&CK — or (b) is structurally indivisible because its protective effect applies fully to each individual transaction (cryptographic operations, origin-bound authentication, hardware-attested identities). A control that only checks thresholds on individual actions, without correlation across multiple actions, does not count as aggregation-resistant.

A control is deemed robust if it shows no structural weakness against any of the four criteria. It is deemed partially degraded if one or two criteria reveal weaknesses that can be compensated by supplementation. It is deemed friction-only if its core effect disappears under Mythos (see 3.1.3). It is deemed not affected if none of the four criteria applies directly.

3.3 Framework stack and prioritization

The core assessment is made against ISO/IEC 27002:2022 as the most internationally established generic control catalogue. For each control, a cross-comparison examines whether other frameworks relevant in regulatory or audit terms state a requirement that covers the same protective purpose while showing higher Mythos resilience.

Primary frameworks:

- BSI C5:2026 (Cloud Computing Compliance Criteria Catalogue). The BSI audit catalogue for cloud services (updated March 2026). Not legally binding as such, but made binding in practice through contractual anchoring (for example with public sector cloud customers) or through sector regulation (BaFin BAIT, the German KRITIS regulation). With OPS-32/33 (confidential computing), CRY-01.01AC (post-quantum), DEV-13 (SBOM) and container sub-criteria, it contains explicit Mythos-relevant requirements.
- NIST Cybersecurity Framework 2.0 and NIST SP 800-53 Rev. 5. These supply the technical depth on detection, response and adversarial testing.
- DORA (Digital Operational Resilience Act, Regulation (EU) 2022/2554). Legally binding for the financial sector. Relevant for the resilience aspect and the TLPT requirements (Art. 26–27).

- The CRA (EU Cyber Resilience Act). Legally binding for manufacturers of products with digital elements placed on the EU market. Vulnerability handling obligations (Art. 13, 14), SBOM as part of the essential requirements (Annex I Part II).
- NIS2 Directive (EU) 2022/2555 with Implementing Regulation (EU) 2024/2690. The NIS2 Directive defines the measure categories (Art. 21(2)(a)–(j)) and the notification deadlines for significant incidents (Art. 23(4): 24 h early warning, 72 h full notification, one month final report). Under its Article 1, Implementing Regulation (EU) 2024/2690 applies exclusively to certain digital providers: DNS service providers, TLD name registries, cloud computing service providers, data centre service providers, content delivery network providers, managed service providers, managed security service providers, online marketplaces, online search engines, social networking services platforms and trust service providers. For other NIS2 addressees (critical infrastructure, manufacturing, health), national transposition governs.

Note on citation: In the framework comparisons for individual controls, the Implementing Regulation is cited with its specific annex number (example: "NIS2 IR No. 3.2.2"). References to the Directive itself are cited as "NIS2 Directive Art. …".

Secondary frameworks (supplementary layer, Part V):

- BSI IT-Grundschutz Compendium (German ISMS deep dive).
- MITRE ATT&CK and D3FEND (attack and defence vocabulary).
- ISO/IEC 42001 (AI management system).
- CIS Controls v8 (prioritized implementation aid).
- OWASP ASVS and SAMM (secure development).

3.4 Source basis

The sources of this handbook fall into three categories.

Empirical basis for the threat side: Anthropic threat intelligence reports from August 2025 to April 2026, in particular the GTG-1002 report (November 2025) and the publications on Claude Mythos Preview and the Glasswing defensive program (April 2026).

Normative basis: ISO/IEC 27001:2022, ISO/IEC 27002:2022, ISO/IEC 27005:2022, the relevant NIST publications, BSI C5:2026, DORA (EU) 2022/2554 with its RTS, the CRA, and NIS2 Directive (EU) 2022/2555 with Implementing Regulation (EU) 2024/2690. The BSI IT-Grundschutz Compendium serves as a deep-dive reference in Part V.

Industry consensus and supplementary practice references: The strategy paper "*The AI Vulnerability Storm: Building a Mythos-ready Security Program*" (Cloud Security Alliance, SANS Institute, [un]prompted, OWASP Gen AI Security Project, first published 14 April 2026; cited here in version v0.95) is treated in MRIS as the central industry consensus reference for the Mythos threat landscape. The risk assessment defined there (thirteen prioritized risks at Critical, High and Medium levels) has fed directly into MRIS specifics, in particular into MHC-13 (AI agent governance, addressing Mythos-Ready Risk 3 "Unmanaged AI Agent Attack Surface"), into the extension of A.5.9 by a shadow AI inventory (Risk 6), into Chapter 10.4 on a permanent VulnOps function and innovation acceleration governance (Risks 9 and 11), and into Chapter 10.5 on the shift in the standard of care (Risk 12). Mythos-Ready is neither a standard nor an audit catalogue but a consensus paper by senior practitioners; its status in the MRIS source hierarchy lies between a threat intelligence report and a normative basis. Additional industry references used: OWASP Top 10 for LLM Applications 2026 and the OWASP Agentic Security Initiative (ASI01–ASI10), MITRE ATLAS, NIST AI RMF 1.0, ISO/IEC 42001 Annex A.

3.5 Limits

The assessment is a snapshot, does not replace an organization-specific risk analysis, and makes no claim to completeness of hardening measures (see also Chapter 1.4).

Threat scope. MRIS expressly addresses only Gen-AI-accelerated and agentic attack patterns. Other threat classes — malicious insiders, low-technology attacks such as physical manipulation or social engineering without an AI component, unintentional human error in the supply chain, natural events — are covered by the ISMS foundation that MRIS presupposes (see Chapter 1.4). A classification as "Friction-only" or "Partially degraded" refers to effectiveness against Mythos attackers and not to the effectiveness of the control overall. A control may be considerably degraded against a Mythos attacker while continuing to be fully effective against insider threats or unsophisticated external actors.

Part II – Main analysis: the 93 controls of ISO/IEC 27002:2022

The following four chapters assess all 93 controls against the Mythos threat picture established in Chapter 2. The ordering follows the four categories from Chapter 3, not the numerical ISO sequence. Readers looking up the classification of an individual control will find the complete matrix in Annex A.

Every individual control assessment follows the same format: category classification, Mythos finding with a rationale based on the four criteria from Chapter 3.2, framework cross-comparison against the primary comparison frameworks C5:2026, NIST, DORA, the CRA and NIS2 (Directive plus Implementing Regulation), supplemented for individual controls by further relevant references, and — for degraded and friction-based controls — a concrete hardening or replacement recommendation.

Overall distribution of the 93 controls:

Category	Count	Share
Robust	29	31%
Partially degraded	37	40%
Friction-only	4	4%
Not affected	23	25%

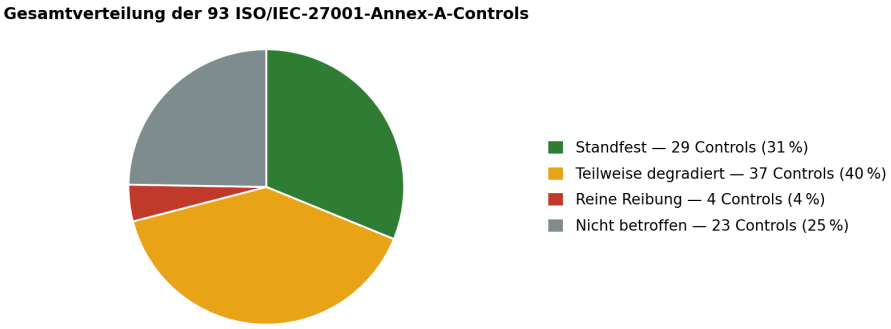

The distribution is expressly not a catastrophic finding. It says that ISO/IEC 27002:2022 remains sound at its core: 29 controls (31%) are robust, 23 (25%) are not affected. Around 44% of the controls (41 of 93 — 37 partially degraded, 4 friction-only) require supplementary hardening or realignment under Mythos conditions in order to continue fulfilling their original protective purpose.

4 Robust controls – the sound foundation

4.1 Selection logic

A control is classified as robust if its protective effect derives from a hard barrier that persists under Mythos conditions: cryptographic impossibility, hardware-bound identity, the architectural non-existence of an attack path, or a structural reduction of the blast radius. Robust controls do not rest on limited attacker capacity but on objective limits.

Twenty-nine controls are assessed as robust in this chapter. The presentation follows the ISO numbering; the thematic grouping is summarized in Chapter 4.3.

4.2 Controls in detail

A.5.9 Inventory of information and other associated assets

Category	ROBUST
Mythos finding	A complete, current asset inventory is the precondition for any targeted defence. In the Glasswing post, Anthropic states that systems the organization does not know about cannot be defended either. The control is

Category	ROBUST
	Mythos-robust because the existence of an inventory is devalued neither by attacker speed nor by micro-step fragmentation.
Framework comparison	C5:2026 AM-02, AM-03, AM-04. NIST CSF 2.0 ID.AM-1 to ID.AM-5. DORA Art. 8. NIS2 IR No. 12.
Hardening recommendation	Automate inventory updating (CMDB integration, cloud asset inventory tools such as AWS Config, Azure Resource Graph). Manual inventories are structurally out of date at Mythos iteration speed. Add shadow IT scans and external perimeter inventorying (certificate transparency logs, ASM tools). Maintain shadow AI as a separate inventory category: unauthorized AI coding assistants and browser plug-ins (for example Cursor, Continue.dev, Codeium extensions), MCP servers with data access, VS Code and JetBrains extensions with AI functionality, browser-based AI sidebar tools with tab read access. Capture them via endpoint software inventory (for example through EDR telemetry), browser extension audits (for example via MDM or browser enterprise policies) and network egress monitoring to known AI provider endpoints (api.openai.com, api.anthropic.com, generativelanguage.googleapis.com).

A.5.12 Classification of information

Category	ROBUST
Mythos finding	Classification creates the basis for prioritized hardening. An attacker executing 80–90 percent of the tactical work autonomously ultimately encounters the same data set — whether that set is classified as highly confidential and protected accordingly determines the actual impact. Classification as an act is not time-critical.
Framework comparison	C5:2026 AM-09. NIST SP 800-60. DORA Art. 8 (implicit). CRA Annex I. NIS2 IR No. 12.1.
Hardening recommendation	Extend the classification scheme to Mythos-relevant aspects: exfiltration sensitivity (what is particularly critical under AI-assisted mass extraction) and integrity sensitivity (what is particularly critical under model manipulation or agent misdirection). Use automated classification through data discovery tools.

A.5.16 Identity management

Category	ROBUST
Mythos finding	Unique, cryptographically anchored identities are a hard barrier. Under Mythos, the centre of gravity shifts from human user identities to workload and agent identities. The control covers both categories, provided the organization consistently uses workload identity (SPIFFE/SPIRE, AWS IAM roles, Azure managed identities).
Framework comparison	NIST SP 800-63. NIST NCCoE Agent Identity (2026, announced). C5:2026 IAM-01, IAM-02. DORA Art. 9. NIS2 IR No. 11.

A.5.27 Learning from information security incidents

Category	ROBUST
Mythos finding	Post-incident analysis remains fully effective under Mythos — its importance even increases, because Mythos attacks establish new TTPs that must be collected systematically and fed back into detection logic. Deriving

Category	ROBUST
	detections from lessons learned (linked to A.8.16 monitoring) must be carried out in hardened form.
Framework comparison	NIST SP 800-61 Revision 3. C5:2026 SIM-06. DORA Art. 13 (root cause analysis). NIS2 IR No. 3.6.

A.5.28 Collection of evidence

Category	ROBUST
Mythos finding	Forensic evidence collection is robust as long as the logging basis (A.8.15) and time synchronization (A.8.17) are intact. Mythos attacks leave traces. The challenge lies not in collection but in correlation (A.8.16).
Framework comparison	NIST SP 800-86. C5:2026 SIM-04. DORA Art. 17. ISO/IEC 27037 (chain of custody). NIS2 IR No. 3.5.4.

A.5.30 ICT readiness for business continuity

Category	ROBUST
Mythos finding	Architecture-driven resilience — redundant systems, failover, defined RTO/RPO — is structural, not friction-based. In the Glasswing post, Anthropic recommends tabletop exercises for three to five parallel incidents (as a guide value), because Mythos attacks can run on several vectors at once; the underlying ICT readiness capability remains robust if it has been dimensioned accordingly.
Framework comparison	DORA Art. 11–12 (the strictest regime, including TLPT Art. 26–27). C5:2026 BCM-01 to BCM-04. NIS2 IR No. 4. NIST CSF RC.RP. ISO 22301.
Hardening recommendation	Re-dimension RTO/RPO targets for Mythos-relevant scenarios: parallel ransomware attacks on primary and backup sites, agentic modification of DR configurations, supply chain compromise of the DR service provider. Run tabletops for multiple simultaneous incidents.

A.6.5 Responsibilities after termination or change of employment

Category	ROBUST
Mythos finding	The control is robust provided it is coupled with automated deprovisioning. The classic attack vector — the former employee with still-active credentials — remains equally relevant under Mythos and is if anything sharpened by an attacker's ability to find and chain such credentials automatically.
Framework comparison	C5:2026 HR-05. NIST SP 800-53 PS-4. DORA Art. 9. NIS2 IR No. 10, No. 11.2.

A.7.1 Physical security perimeters

Category	ROBUST
Mythos finding	Physical barriers are Mythos-orthogonal: an agentic attacker does not overcome a concrete wall with AI. The protective effect is unchanged. Relevant for data centres, server rooms and workplaces with access to highly sensitive data.
Framework comparison	C5:2026 PS-03. NIST SP 800-53 PE-3. DORA Art. 9(4). NIS2 IR No. 13.

A.7.2 Physical entry

Category	ROBUST
Mythos finding	Hardware-bound entry systems (badges, biometrics) are a hard barrier in physical space.
Framework comparison	C5:2026 PS-04. NIST PE-2/PE-3. DORA Art. 9(4). NIS2 IR No. 13.2.

A.7.3 Securing offices, rooms and facilities

Category	ROBUST
Mythos finding	Physical hardening of work areas remains fully effective.
Framework comparison	C5:2026 PS-01, PS-08. NIST PE-5. NIS2 IR No. 13.

A.7.4 Physical security monitoring

Category	ROBUST
Mythos finding	CCTV, intrusion detection and motion detectors are Mythos-orthogonal.
Framework comparison	C5:2026 PS-07. NIST PE-6. NIS2 IR No. 13.

A.7.6 Working in secure areas

Category	ROBUST
Mythos finding	Zone-based physical security remains effective. Mythos-neutral.
Framework comparison	C5:2026 PS-08. NIST PE-19.

A.7.10 Storage media

Category	ROBUST
Mythos finding	Media lifecycle management (secure storage, transport, deletion) is robust. Cryptographic shredding (destruction of the keys) is Mythos-robust.
Framework comparison	C5:2026 AM-12. NIST SP 800-88. NIS2 IR No. 12.

A.7.14 Secure disposal or re-use of equipment

Category	ROBUST
Mythos finding	Irreversible data destruction before disposal or reassignment is a cryptographically hard barrier (DoD-compliant wipe, cryptographic shredding).
Framework comparison	C5:2026 AM-07. NIST SP 800-88. BSI TR-03183-1. CRA Annex I.

A.8.2 Privileged access rights

Category	ROBUST
Mythos finding	Least privilege for admin accounts is one of the hardest barriers against Mythos attacks. An agentic attacker collecting credentials automatically runs into a wall if those credentials carry no far-reaching privileges. Robust where implemented with hardware attestation and just-in-time elevation (PAM); password protection alone would be degraded.
Framework comparison	NIST SP 800-53 AC-6, IA-5. NIST SP 800-207. C5:2026 IAM-06 (just-in-time assignment). DORA Art. 9, 15. NIS2 IR No. 11.2.

A.8.9 Configuration management

Category	ROBUST
Mythos finding	Versioned, automatically monitored configuration baselines (infrastructure as code, GitOps) are robust. Drift detection identifies unauthorized changes in real time — a hard structural barrier against agentic modification.
Framework comparison	NIST SP 800-53 CM-2 to CM-8. C5:2026 OPS-26, DEV-03. CIS Controls v8.4. SLSA Level 3+. NIS2 IR No. 6.

A.8.10 Information deletion

Category	ROBUST
Mythos finding	Cryptographic deletion is a hard barrier. Mythos attackers cannot subsequently exfiltrate what has already been irreversibly deleted. Essential against supply chain attacks on backups and archives.
Framework comparison	NIST SP 800-88. C5:2026 PI-03, CRY-14. GDPR Art. 17. NIS2 IR No. 12.

A.8.11 Data masking

Category	ROBUST
Mythos finding	Data masking reduces the blast radius structurally: in the event of successful exfiltration, the data that leaves is devalued to the extent that attribution to natural persons is not possible without additional information. A distinction must be drawn here: anonymized data offers no possibility of re-identification, whereas pseudonymized data remains re-identifiable with the mapping table and must therefore be protected like personal data. The protective effect is structural in the case of anonymization; in the case of pseudonymization it is tied to keeping the mapping separate.
Framework comparison	ISO/IEC 20889. NIST SP 800-188. GDPR Art. 32(1)(a). C5:2026 OPS-30, OPS-31. NIS2 IR No. 6.

A.8.13 Information backup

Category	ROBUST
Mythos finding	Immutable, offline-capable backups (immutable, air-gapped) are among the most important hard barriers against modern ransomware and Mythos-driven data integrity attacks. An agentic attacker with infrastructure access can manipulate backups reachable online; air-gapped backups remain protected.
Framework comparison	C5:2026 OPS-06 to OPS-09 (OPS-09: physically separate sites). NIST SP 800-34. DORA Art. 12(2). NIS2 IR No. 4.
Hardening recommendation	Extend the backup strategy to the 3-2-1-1-0 principle: 3 copies, 2 media types, 1 offsite, 1 offline/immutable, 0 errors in the restore test. Restore tests at least quarterly. Separate IAM paths for the backup infrastructure.

A.8.14 Redundancy of information processing facilities

Category	ROBUST
Mythos finding	Architectural redundancy is a structural property. Mythos DDoS, automated exploit campaigns and parallel ransomware attacks target availability; redundancy as a foundational principle (no single points of failure) remains effective.
Framework comparison	DORA Art. 11. NIST SP 800-53 CP-7. C5:2026 PS-02, BCM-03, BCM-04. NIS2 IR No. 4.1.

A.8.15 Logging

Category	ROBUST
Mythos finding	Central, append-only-protected logs are the non-negotiable precondition for detection, forensics and learning. Robust as long as log integrity is cryptographically secured (WORM, Merkle chains, HSM integration). Without logs, every downstream detection is blind. A further effectiveness dimension alongside integrity: logs may contain sensitive authentication artefacts (for example WebAuthn assertions in cleartext, CVE-2026-34348); control read and write rights through separate identity paths (A.5.33) and mask sensitive fields before forwarding. Robust only where the confidentiality of the logs is preserved alongside their integrity.
Framework comparison	NIST SP 800-92. C5:2026 OPS-10 to OPS-17 (policy, SIEM, retention, accountability). DORA Art. 10(3). NIS2 IR No. 3.2 (minimum content 3.2.3 a–l).

A.8.17 Clock synchronization

Category	ROBUST
Mythos finding	NTP-based time synchronization is Mythos-relevant: without precisely synchronized timestamps across all systems, correlating fragmented attack actions is impossible. Mythos attacks appear in micro-steps across distributed systems; a clock skew of a few seconds destroys the ability to reconstruct events.
Framework comparison	NIST SP 800-53 AU-8. C5:2026 OPS-16. NIS2 IR No. 3.2.6 (cross-system time synchronization, a precondition for correlatable logs).

A.8.18 Use of privileged utility programs

Category	ROBUST
Mythos finding	Restricting and monitoring legitimate tools that can override system controls (debuggers, admin utilities, configuration management tooling) is a hard barrier. Mythos attackers use these tools by preference (living off the land); restricting them remains fully effective.
Framework comparison	NIST SP 800-53 AC-3, SC-18. C5:2026 IAM-06. Application allow-listing (CIS 2.5, 2.6).

A.8.19 Installation of software on operational systems

Category	ROBUST
Mythos finding	Application allow-listing with signature verification is a hard barrier against unauthorized software installation — including against AI-generated malware variants, provided these are not signed.
Framework comparison	NIST SP 800-167. C5:2026 OPS-26, DEV-10. CIS Controls v8.2.

A.8.24 Use of cryptography

Category	ROBUST
Mythos finding	Strong, correctly implemented cryptography is the hardest barrier in information security. An agentic attacker cannot read AES-256-encrypted data without the key. The vulnerability lies in key management. Post-quantum readiness must be considered on the 2030+ horizon.
Framework comparison	NIST SP 800-175, FIPS 140-3. NIST PQC standards. BSI TR-02102. C5:2026 CRY-01 to CRY-19 (CRY-01.01AC: post-quantum strategy). DORA Art. 9(4)(e). CRA Annex I Part II. NIS2 IR No. 9.
Hardening recommendation	Cryptographic inventory (crypto agility audit): which algorithms, key lengths and implementations are used where? Define a post-quantum migration path for asymmetric methods. Hardware security modules (HSM) for keys with long-term relevance.

A.8.27 Secure system architecture and engineering principles

Category	ROBUST
Mythos finding	Defence in depth, least privilege, fail-secure and secure by default are architectural principles, not point controls. They are the foundation of any Mythos resistance, because they structurally assume that individual controls can fail. In the Glasswing post, Anthropic names "design for breach" explicitly as a core idea.
Framework comparison	NIST SP 800-160. NIST SP 800-207. C5:2026 DEV-01, DEV-05. CISA Secure by Design. OWASP SAMM. NIS2 IR No. 6.

A.8.29 Security testing in development and acceptance

Category	ROBUST
Mythos finding	Automated security testing in the CI/CD pipeline (SAST, DAST, SCA, AI-supported vulnerability scanning) is robust — and, according to the Anthropic Glasswing post, the single most important measure against an AI-accelerated offensive. The goal: scan your own code with the same model class attackers use, before they do.
Framework comparison	OWASP ASVS. NIST SP 800-218 SSDF. C5:2026 DEV-07, OPS-22, OPS-25 (OPS-25.01AS: daily scans). SLSA Level 3+. CRA Annex I Part II. NIS2 IR No. 6, No. 7.
Hardening recommendation	Activate pre-merge blockers for high-confidence findings. Establish AI vulnerability scanning as a stage of its own (isolated agent deployment, verification step, integration into the triage process). Vulnerability reports must disclose the use of AI.

A.8.31 Separation of development, test and production environments

Category	ROBUST
Mythos finding	Strict environment separation is a hard barrier against cross-environment propagation of compromises. Mythos-relevant in particular for supply chain attacks.
Framework comparison	NIST SP 800-53 CM-4, SC-2. C5:2026 DEV-11, DEV-12. SLSA build levels. NIS2 IR No. 6.

A.8.33 Test information

Category	ROBUST
Mythos finding	Not using production PII in test environments is a hard reduction of the blast radius. Test environments have historically been less well protected than production environments; with agentic exfiltration capability, that discrepancy becomes a critical risk where production data resides there.
Framework comparison	GDPR Art. 5(1)(c). NIST SP 800-53 SA-15(9). C5:2026 DEV-11. NIS2 IR No. 6.

4.3 Summary

The 29 robust controls can be grouped into five thematic clusters that together form the structural foundation of a Mythos-resilient security architecture:

- Cryptographic barriers: A.8.24, A.8.10, A.8.11, A.7.14.
- Identity and inventory foundation: A.5.9, A.5.12, A.5.16, A.8.2.
- Architectural resilience: A.5.30, A.8.14, A.8.27, A.8.31, A.8.13.
- Forensic traceability: A.8.15, A.8.17, A.5.28, A.5.27.
- Integrity gates in development and operations: A.8.9, A.8.18, A.8.19, A.8.29, A.8.33. Anthropic highlights A.8.29 as the single most effective measure.

Mythos-robust controls anchor security either in cryptography, in architecture, or in automated integrity checking. Controls that anchor security in human patience, manual response or episodic review are the subject of Chapters 5 and 6.

Prioritization note. The structural controls summarized in this chapter — zero trust architecture, strong cryptography, least privilege and secrets management, micro-segmentation, immutable backups — form the most robust layer of Mythos defence. They cannot be replaced by new AI-specific controls. A resource-efficient Mythos hardening programme prioritizes the consistent implementation of these structural controls ahead of introducing new detection or automation capabilities (Chapter 9). The risks classified as CRITICAL in Mythos-Ready (CSA, SANS, OWASP, April 2026) are already addressed to a considerable extent by consistently implementing structurally robust controls; the additional MHC close the gaps that structural controls alone do not cover.

5 Partially degraded controls – effectiveness with reservations

5.1 Overview and degradation patterns

A control is deemed partially degraded if its basic protective effect is preserved under Mythos conditions but the strength originally assumed falls significantly. Such controls need not be replaced; they need supplementation through additional mechanisms, changed parameters or architectural flanking. The degradation can be observed in four recurring patterns derived directly from the four assessment criteria in Chapter 3.2.

First — time compression: Controls that take effect within a response time appropriate for human operators fail structurally when the attacker has already completed the next attack step within that time. Affected are

incident response playbooks, patch cycles with manual approvals, quarterly access recertifications and annual audits.

Second — aggregation blindness: Controls that check individual events for policy conformity do not take hold when the attack is decomposed into micro-steps that each appear legitimate. Affected are signature-based monitoring, content-based DLP, role-based access controls without behavioural context, and classic segregation of duties.

Third — credential compromisability: Controls relying on knowledge- or possession-based factors without phishing resistance are considerably less effective against high-quality AI-generated phishing and credential stuffing.

Fourth — supply chain vulnerability: Controls relying on contractual assurances, annual questionnaires or point-in-time supplier audits do not hold against the speed at which agentically supported supply chain compromises propagate.

The following 37 controls are presented in ISO numbering. Typical supplementation patterns are consolidated in Chapter 5.3.

5.2 Controls in detail

A.5.3 Segregation of duties

Category	PARTIALLY DEGRADED
Mythos finding	Segregation of duties remains structurally effective: a single account cannot exercise conflicting roles under Mythos either, provided the separation is cleanly implemented. What is degraded is not the separation itself but its supervision: if an attacker takes over several accounts in different roles, agentic automation executes the sub-steps in such a distributed and time-stretched way that classic anomaly review no longer recognizes them as connected misuse. The criterion violated is the aggregation resistance of the control, not the separation principle. In addition, the fragmented approach produces individual actions separated in time and context, each of which looks policy-conformant. Without behaviour-based correlation, segregation of duties remains formally observed but no longer protects against actual misuse.
Framework comparison	C5:2026 OIS-04 (Segregation of Duties). NIS2 IR No. 11.2.2(a) names segregation of duties as a principle of access assignment alongside need-to-know and need-to-use. NIST SP 800-53 AC-5. DORA Art. 9 implicitly.
Hardening recommendation	Deploy behaviour-based detection (UEBA) as a supplement to identify compromised accounts through atypical activity patterns. Context-sensitive authorization with dynamic anomaly checking. Cross-role activity within short time windows as an alerting criterion.

A.5.5 Contact with authorities

Category	PARTIALLY DEGRADED
Mythos finding	The core of the control — established, maintained contacts with authorities — remains effective unchanged under Mythos; time-critical incident handling itself sits in A.5.24 to A.5.26. What is degraded is solely the practical usability of the notification chain under deadline pressure: NIS2 Directive Art. 23(4) requires an early warning within 24 hours and a full notification within 72 hours of becoming aware of a significant incident. Manual escalation chains via standing meetings and hand-written incident reports collide with these deadlines as soon as Mythos attacks escalate on several levels in parallel.
Framework comparison	C5:2026 OIS-06 (Contact with Relevant Government Agencies and Interest Groups). NIS2 Directive Art. 23(4) sets the notification deadlines; the NIS2 IR specifies in No. 3.3 the internal reporting mechanism (a simple mechanism for employees, suppliers and customers) and in No. 3.3.2 the obligation to inform those groups about it. DORA Art. 19 requires major incident reports to the competent supervisory authorities.

Category	PARTIALLY DEGRADED
Hardening recommendation	Automated notification chain with prepared templates for different incident types. Dedicated reporting role with deputy arrangements and 24/7 availability. Integration of the reporting system with the SIEM for automatic notification triggers.

A.5.7 Threat intelligence

Category	PARTIALLY DEGRADED
Mythos finding	Threat intelligence is indispensable under Mythos. The control itself prescribes no update rate; what is degraded is the prevailing implementation form: TI feeds with weekly or monthly update rates and manual evaluation lag structurally behind the iteration speed of AI-supported attacker techniques. Threat intelligence that is integrated automatically and evaluated in near real time satisfies the same control and is not degraded. The TTPs of a campaign can change within the time a classic feed needs for delivery.
Framework comparison	C5:2026 OIS-05 (Threat Intelligence). NIS2 IR No. 2 (risk management policy) requires an up-to-date threat assessment as the basis for risk treatment. NIST SP 800-150 Guide to Cyber Threat Information Sharing.
Hardening recommendation	Use TI feeds with automated SIEM integration and real-time updating (MISP, STIX/TAXII). Join sector-specific sharing communities (ISACs, CERT networks). Use AI-supported TI correlation platforms (for example Recorded Future, Mandiant Advantage, GreyNoise) for automated IOC enrichment and prioritization. Audit threshold: new high-severity IOCs must be available in the SIEM within 60 minutes of publication.

A.5.14 Information transfer

Category	PARTIALLY DEGRADED
Mythos finding	Securing data transfers through transport encryption is Mythos-robust. The policy layer of the control — which data may be transferred where — is degraded by aggregation blindness: classic DLP systems do not recognize legitimate-looking, fragmented small transfers as part of an aggregated exfiltration.
Framework comparison	C5:2026 COS-08 (Policies for Data Transmission) and CRY-04 (Protection of Data for Transmission). NIS2 IR No. 8 (cyber hygiene) and No. 9 (cryptography). NIST SP 800-53 SC-8.
Hardening recommendation	Add volume- and pattern-based egress anomaly detection. Data loss prevention with behavioural context instead of pure content inspection. Egress traffic analysis with ML-supported baseline detection.

A.5.15 Access control

Category	PARTIALLY DEGRADED
Mythos finding	The access control policy remains the foundation of all authorization. Under Mythos the control is degraded in two dimensions: aggregation blindness in role-based models that do not recognize fragmented access patterns, and time compression in static access decisions that do not react to behavioural change. A formally correctly configured RBAC does not protect against a compromised account acting within its role.
Framework comparison	C5:2026 IAM-01 (Policy for Identities and Access Rights) and IAM-07 (Access to Cloud Service Customer Data). NIS2 IR No. 11 (access control) — in particular No. 11.1 (policy) and No. 11.2 (management) — requires risk-based assignment and regular review. DORA Art. 9. NIST SP 800-53 AC family.

Category	PARTIALLY DEGRADED
Hardening recommendation	Risk-adaptive access with dynamic evaluation drawing on identity signals, device posture and behavioural context. Continuous recertification instead of quarterly review. Transition to zero trust principles with explicit verification of every transaction.

A.5.17 Authentication information

Category	PARTIALLY DEGRADED
Mythos finding	The control covers all forms of authentication information. The classification as partially degraded is deliberately an average: phishing-resistant MFA (FIDO2, passkeys, hardware tokens to WebAuthn) is Mythos-robust. Passwords, SMS-based MFA and time-based OTPs are considerably weakened against AI-quality phishing and credential stuffing. The control takes full effect only in its phishing-resistant implementation.
Framework comparison	C5:2026 IAM-08 (Authentication Mechanisms) and IAM-09 (Confidentiality of Authentication Information). NIS2 IR No. 11.6 (authentication) and No. 11.7 (multi-factor authentication); for privileged accounts additionally No. 11.3.2(a) (strong identification and authentication). NIST SP 800-63B authentication assurance levels; AAL2 and AAL3 recommended for Mythos environments.
Hardening recommendation	Phishing-resistant MFA (FIDO2, passkeys) mandatory for all privileged accounts and externally reachable services. Exclude SMS-based MFA for new deployments. Hardware tokens for administrative access.

A.5.18 Access rights

Category	PARTIALLY DEGRADED
Mythos finding	Provisioning and deprovisioning function in principle. Classic recertification on quarterly cycles is degraded by time compression: an agentic attacker uses the time between a role change and recertification to exploit accumulated entitlements (permission creep). The latency between detecting an anomaly and withdrawing entitlements is too long for Mythos speed.
Framework comparison	C5:2026 IAM-04 (Withdrawal or Adjustment of Access Rights) and IAM-05 (Regular Review of Access Rights). NIS2 IR No. 11.2 requires timely adjustment on role change. NIST SP 800-53 AC-2.
Hardening recommendation	Continuous recertification instead of quarterly review. Automated withdrawal of entitlements triggered by HR events. Just-in-time entitlement for privileged access with automatic revocation on expiry.

A.5.19 Information security in supplier relationships

Category	PARTIALLY DEGRADED
Mythos finding	The control remains necessary and is considerably widened by the Mythos threat picture. Supply chain attacks are easier to find, more precisely executable and more broadly scalable with AI support. Classic supplier questionnaires at onboarding do not hold against the continuously changing attack surface of the supply chain. Time compression takes effect here on a large scale.
Framework comparison	C5:2026 SSO-01 (Policies and Procedures for Controlling and Monitoring Service Organisations) and SSO-02 (Risk Assessment). NIS2 IR No. 5 (supply chain security) anchors supplier security as a pillar in its own right. DORA Art. 28–30 as the strictest regime for critical ICT third-party providers. CRA Annex I requires SBOM transparency along the supply chain.

Category	PARTIALLY DEGRADED
Hardening recommendation	Anchor an SBOM obligation contractually. Automated vendor security monitoring with continuous perimeter scans of suppliers. A tiering approach with differentiated requirements by criticality. Incident notification SLAs of no more than 24 hours.

A.5.20 Addressing information security within supplier agreements

Category	PARTIALLY DEGRADED
Mythos finding	The contractual control remains necessary; its effectiveness depends on contractual depth. ISO 27002 formulates minimum requirements that are not sufficient under Mythos. Without binding provisions on SBOM, vulnerability disclosure deadlines, right to audit and incident notification, the control remains a documentation shell.
Framework comparison	DORA Art. 28–30 contains the currently strictest binding contractual requirements, including exit strategies and subcontractor provisions. C5:2026 SSO-01 and SSO-06 (Contract Termination Strategy for Service Organisations). NIS2 IR No. 5 requires contractual anchoring of security requirements on third parties. CRA Annex I requires contractual SBOM obligations.
Hardening recommendation	Draft contractual clauses to DORA standards even outside the financial sector: right to audit, incident SLAs (24 h notification, 72 h detailed report), subcontracting approval requirements, exit clauses with data return and proof of deletion. Reality check: with hyperscalers (AWS, Azure, GCP) an individual right to audit is not enforceable in practice. Acceptable substitutes from an audit perspective: ISO 27001 certificate plus SOC 2 Type II, a BSI C5 attestation (Type 2), TISAX, pooled audits through industry consortia (CSA STAR, ENISA EUCS), and contractual sub-processor lists with change notification.

A.5.21 Managing information security in the ICT supply chain

Category	PARTIALLY DEGRADED
Mythos finding	The control is doubly degraded, by time compression and by aggregation blindness. Classic audits and questionnaires cover point-in-time states, not the continuous change of ICT supply chains. Automated supply chain attacks — malicious dependencies, compromised build pipelines, smuggled-in Docker layers — structurally escape these mechanisms.
Framework comparison	C5:2026 DEV-13 (Transparency about Software Components) requires SBOM provision; DEV-14 (Secure Use of Third Party Hardware and Software). SSO-05 (Monitoring of Compliance with Requirements). NIS2 IR No. 5. CRA Annex I Part II makes the SBOM an obligation. SLSA framework for build provenance.
Hardening recommendation	Integrate SBOM generation into the CI pipeline (SPDX, CycloneDX). SLSA Level 3+ for critical build paths. Continuous dependency scans with automated vulnerability correlation (EUVD, CVE, OSV). Verify package provenance attestation.

A.5.22 Monitoring, review and change management of supplier services

Category	PARTIALLY DEGRADED
Mythos finding	Periodic reviews are structurally too slow against the iteration speed of Mythos attacks on supply chains. The control remains sensible as a process but takes effect only if it is operated continuously and in automated form.
Framework comparison	C5:2026 SSO-05 (Monitoring of Compliance with Requirements) and SSO-07 (Ensuring Transparency within Service Organisations). NIS2 IR No. 5. DORA Art. 28 requires continuous monitoring for critical third-party providers.
Hardening recommendation	External attack surface management and security rating services (for example BitSight, SecurityScorecard, Black Kite, RiskRecon) for continuous vendor monitoring. Integration

Category	PARTIALLY DEGRADED
	of threat intelligence on supplier companies. Deviations above defined score thresholds automatically raise tickets in the SIEM or GRC system. A scorecard approach with quantified security indicators (at minimum: patch hygiene, TLS configuration, leaked credentials in public breach databases, compromised hosts on botnet lists).

A.5.23 Information security for use of cloud services

Category	PARTIALLY DEGRADED
Mythos finding	Cloud security is partially degraded because the shared responsibility model creates dependencies on provider-side detection. If the cloud provider is itself compromised, or an agentic attacker successfully disguises API abuse on legacy endpoints, customer-side controls must be able to intervene. Individual provider controls such as hardware-attested identities are robust; the overall control depends on the quality of the provider audits.
Framework comparison	C5:2026 is the defining framework: general conditions controls GC-01 to GC-06 for transparency and 17 domains as the audit object. NIS2 IR No. 5 (supply chain security). DORA Art. 28. ENISA EUCS as the forthcoming European certification.
Hardening recommendation	Require provider assurance to C5:2026 or ISO 27017. On the customer side, deploy cloud security posture management (CSPM) and cloud workload protection (CWPP). Document a multi-cloud exit strategy. Evaluate confidential computing for highly sensitive workloads.

A.5.24 Information security incident management planning and preparation

Category	PARTIALLY DEGRADED
Mythos finding	The existence of an IR plan remains critical. Many classic playbooks are designed for response times on the order of hours and are degraded by time compression. Where an agentic attacker completes the tactical work within minutes, playbooks must be designed for second- and minute-scale windows. In addition, many plans underestimate the possibility of parallel incidents.
Framework comparison	C5:2026 SIM-01 (Policy for Security Incident Management) and SIM-02 (Security Incident Response Plans). NIS2 IR No. 3.1 (incident handling policy). DORA Art. 17. NIST SP 800-61 Revision 3.
Hardening recommendation	Revise playbooks for minute-scale windows with predefined automated containment actions. SOAR integration. Tabletop exercises for three to five parallel incidents. Predefined communication templates for customers, authorities and the public.

A.5.26 Response to information security incidents

Category	PARTIALLY DEGRADED
Mythos finding	IR execution itself is partially degraded because the human decision chain between detection and response dominates response latency. Where an attacker operates 80 to 90 percent autonomously and triggers actions in seconds, a human SOC analyst is often left with too little time for an informed decision. The control becomes robust only if defined escalation stages are executed fully automatically.
Framework comparison	C5:2026 SIM-03 (Processing of Security Incidents) and SIM-04 (Documentation and Reporting). NIS2 IR No. 3.5 (incident response) names three response phases in 3.5.2 (containment, eradication and, where necessary, recovery); post-incident review and lessons learned are regulated separately in No. 3.6. DORA Art. 17–18. NIST SP 800-61 Revision 3.

Category	PARTIALLY DEGRADED
Hardening recommendation	SOAR platform for automated containment actions (account lockout, network isolation, key rotation). Mythos-specific playbooks: mass data exfiltration, parallel multi-vector attack, supply chain compromise. A dedicated threat hunter function.

A.5.29 Information security during disruption

Category	PARTIALLY DEGRADED
Mythos finding	The control addresses maintaining security measures during contingency operations. Under Mythos it is doubly degraded: first by parallel attack scenarios that classic BCM plans do not assume; second by the possibility that the attacker has provoked the BCM event and targets the emergency infrastructure itself. DR systems are often less well protected than production systems.
Framework comparison	C5:2026 BCM-01 to BCM-04. CRY-16 (Operational Continuity for Key Management). NIS2 IR No. 4 (business continuity and crisis management). DORA Art. 11–12.
Hardening recommendation	Build parallel incident scenarios into tabletops. DR infrastructure at a security level equivalent to production. Separate IAM paths for DR administration. Regular DR failover tests including security controls.

A.5.33 Protection of records

Category	PARTIALLY DEGRADED
Mythos finding	Retention as such is Mythos-orthogonal; integrity assurance is threatened by agentic manipulation. An attacker with infrastructure access can subsequently alter or deliberately falsify records if they are not cryptographically integrity-protected.
Framework comparison	C5:2026 OPS-12 (Logging and Monitoring – Access, Retention and Deletion) and OPS-14 (Retention of the Logging Data). COM-03 (Internal Audits of the ISMS). NIS2 IR No. 3.2.5 (retention and protection of logs). DORA Art. 10.
Hardening recommendation	Append-only storage with cryptographic integrity assurance (WORM, object lock). Hash chains for tamper detection. Separate IAM paths for read and write access to archives. Separate custody with external third parties for critical records.

A.5.34 Privacy and protection of PII

Category	PARTIALLY DEGRADED
Mythos finding	The policy control remains necessary; its effectiveness depends on the technical controls A.8.11 (masking) and A.8.24 (cryptography). Under Mythos, GDPR-compliant data minimization becomes the most important lever, because it structurally limits the blast radius in the event of successful exfiltration. Data not collected cannot be exfiltrated.
Framework comparison	GDPR Art. 5 and Art. 32. C5:2026 PI-03 (Secure Deletion of Data) and OPS-30/31 (Separation of Datasets). ISO/IEC 27701 as the privacy extension of the ISMS. NIS2 IR No. 8 and No. 9 implicitly.
Hardening recommendation	Enforce data minimization technically (schema validation, ORM rules, API response filtering). Purpose limitation as a mandatory attribute check on data access. Pseudonymization as the default in analytics and ML pipelines. Regular privacy impact assessments using Mythos threat scenarios.

A.5.35 Independent review of information security

Category	PARTIALLY DEGRADED
Mythos finding	Annual or biennial audits are structurally too slow against the iteration speed of Mythos threats. The control remains sensible as a governance mechanism but no longer delivers a timely effectiveness assessment. Between two audits, the attack surface, the TTPs and control effectiveness may each have changed significantly several times.
Framework comparison	C5:2026 COM-03 (Internal Audits of the ISMS) and COM-04 (Information on Information Security Performance). NIS2 IR No. 2.3 (independent review of network and information security). DORA Art. 6(5) considerably stricter for the financial sector.
Hardening recommendation	Continuous audit mechanisms through automated control monitoring. Metric-based effectiveness measurement with real-time dashboards. Supplementary penetration tests and red teaming between formal audits. Purple team exercises for ongoing validation.

A.6.3 Information security awareness, education and training

Category	PARTIALLY DEGRADED
Mythos finding	Classic awareness training does not prepare staff for the quality of Mythos-generated social engineering. AI-generated phishing messages are barely distinguishable from legitimate communication in grammar, context and style. Deepfake-based vishing and impersonation attacks bypass the detection patterns taught in classic training. The control remains necessary but must be sharpened in content.
Framework comparison	C5:2026 HR-03 (Security Training and Awareness Programme) and DEV-04 (Safety Training Regarding Continuous Software Delivery). NIS2 IR No. 8 (cyber hygiene and cybersecurity training). DORA Art. 13(6).
Hardening recommendation	Mythos-specific training modules: AI phishing recognition, deepfake vishing, impersonation scenarios. Regular AI-generated phishing simulations instead of static templates. Stronger emphasis on process verification rather than sender recognition (out-of-band verification for payment instructions and credential resets).

A.6.7 Remote working

Category	PARTIALLY DEGRADED
Mythos finding	Remote working policies remain necessary but are degraded by the widened attack surface of private networks, unmanaged devices and weaker physical protection. Mythos attackers can exploit these weaknesses in automated form. The control becomes robust only with hardware-attested access control and strict device compliance checking.
Framework comparison	C5:2026 HR-07 (Remote Working – Policy) and HR-08 (Remote Working – Implementation). NIS2 IR No. 11 (access control) requires differentiated protection depending on the access context. NIST SP 800-46.
Hardening recommendation	Hardware-attested endpoints for privileged remote access. Zero trust network access (ZTNA) instead of classic VPN. Conditional access with device posture checking. Encrypted, managed workstations with strict application allow-listing.

A.6.8 Information security event reporting

Category	PARTIALLY DEGRADED
Mythos finding	User reporting is necessary as a channel but loses detection significance under Mythos. Agentic attacks often run entirely below the threshold of user perception: API abuse, credential stuffing against service accounts, fragmented data exfiltration. The user sees nothing they could report.

Category	PARTIALLY DEGRADED
Framework comparison	C5:2026 SIM-05 (Duty of the Personnel to Report Security Incidents to a Central Body). NIS2 IR No. 3.3 (event reporting) requires a simple mechanism for employees, suppliers and customers; No. 3.3.2 obliges the organization to inform those groups about the mechanism.
Hardening recommendation	Focus on automated detection rather than user reporting. Simple one-click reporting in the mail client for suspicious messages. A feedback channel for the user, so that reports are seen to be taken seriously and reporting behaviour is encouraged. Mythos-specific training emphasis: train staff explicitly to report unusual authentication indicators — unexpected MFA push notifications without a login attempt of their own, unfamiliar active sessions in the account overview, unexplained account lockouts, sudden push notification storms (MFA fatigue attacks). These indicators are often the only visible sign of an agentic credential stuffing campaign.

A.7.9 Security of assets off-premises

Category	PARTIALLY DEGRADED
Mythos finding	The control primarily addresses physical theft protection for mobile devices; in that dimension it remains effective. Under Mythos, the greater risk is remote compromise of mobile devices, which A.7.9 covers only in part (otherwise A.8.1). The physical layer alone is not sufficient.
Framework comparison	C5:2026 AM-12 (Removable Media and Endpoint Devices). NIS2 IR No. 12 (asset management) and No. 13 (physical security). NIST SP 800-124.
Hardening recommendation	Full disk encryption with HSM-backed key management. Remote wipe capability via MDM. Tracking capability (find-my-device). Endpoint-level DLP with volume anomaly detection: egress thresholds per device class (alert at more than 2σ deviation from the 30-day baseline of outbound data volume, automatic block at more than 3σ). Geofencing for highly sensitive device classes.

A.8.1 User endpoint devices

Category	PARTIALLY DEGRADED
Mythos finding	Classic signature-based endpoint protection is heavily degraded against AI-generated malware variants: the ability to generate functionally equivalent but signature-resistant mutated malware in automated form devalues the signature paradigm. EDR with behavioural analysis remains robust. The control as a whole is partially degraded because implementation depth decides effectiveness.
Framework comparison	C5:2026 OPS-04/05 (Protection Against Malware) and OPS-26 (System Hardening). AM-12 (Removable Media and Endpoint Devices). NIS2 IR No. 11 (access control) and No. 8 (cyber hygiene). NIST SP 800-53 SI-3 and SI-4.
Hardening recommendation	EDR with behavioural analysis and AI-supported detection. Application allow-listing instead of block-listing. Hardware-based identity (TPM, secure enclave) as a precondition for access. Automated network isolation on high-confidence suspicion (clear definition: EDR confidence score of 80% or higher, or correlation of three or more ATT&CK techniques within a 60-minute window). Zero trust endpoint architecture.

A.8.3 Information access restriction

Category	PARTIALLY DEGRADED
Mythos finding	Role-based data access control takes effect against individual accesses that conform to policy, but is aggregation-blind: a compromised account acting within its role and retrieving data in micro-steps raises no alarm. The degradation particularly affects read

Category	PARTIALLY DEGRADED
	access to large data holdings — the classic symptom of mass exfiltration through legitimately authenticated accounts.
Framework comparison	C5:2026 IAM-07 (Access to Cloud Service Customer Data) and IAM-05 (Regular Review of Access Rights). NIS2 IR No. 11 (access control). DORA Art. 9. NIST SP 800-53 AC-3 and AC-4.
Hardening recommendation	Attribute-based access control (ABAC) with behavioural context and volume limits. Data access governance with ML anomaly detection for read patterns. Bloat detection for unusual data volumes within permitted roles. Automatic step-up of queries based on access history.

A.8.4 Access to source code

Category	PARTIALLY DEGRADED
Mythos finding	Source code access is formally secured where IAM protection is strong. Under Mythos, mass exfiltration through compromised developer accounts is a significant risk: wholesale cloning of entire repositories through agentically automated Git operations leaves no anomalous individual events, only patterns visible in the aggregate.
Framework comparison	C5:2026 DEV-11 (Protection of Development and Test Environments) and IAM-06 (Privileged Access Rights). NIS2 IR No. 6 (security measures in development) requires appropriate access protection for development artefacts.
Hardening recommendation	Repository anomaly detection for clone velocity and unusual access patterns. SSO with mandatory FIDO2 MFA for code repositories. Dedicated developer workstations with device attestation. Code signing with hardware keys. Audit trails for all clone and fork operations.

A.8.5 Secure authentication

Category	PARTIALLY DEGRADED
Mythos finding	Like A.5.17, this control addresses the technical implementation of authentication. Phishing-resistant MFA to the WebAuthn standard is Mythos-robust. SMS-based MFA, time-based OTPs and push notifications without number matching are degraded against AI-quality phishing.
Framework comparison	C5:2026 IAM-08 (Authentication Mechanisms) and PSS-05 (Authentication Mechanisms) with product-side requirements on cloud services. NIS2 IR No. 11.6 and No. 11.7. NIST SP 800-63B; AAL2 or AAL3 for Mythos-relevant systems. FIDO Alliance standards.
Hardening recommendation	Phishing-resistant MFA (FIDO2, passkeys, hardware tokens to WebAuthn) mandatory for privileged and externally reachable accounts. Certificate-based authentication for service-to-service communication. Continuous authentication with behavioural biometrics for sensitive sessions.

A.8.7 Protection against malware

Category	PARTIALLY DEGRADED
Mythos finding	Signature-based malware protection is heavily degraded under Mythos, because AI assistance has dramatically reduced the cost of generating functionally equivalent malware variants. Behaviour-based detection and sandbox analysis are robust. The classification of the overall control depends on the implementation depth chosen.
Framework comparison	C5:2026 OPS-04 (Protection Against Malware – Policies and Procedures) and OPS-05 (Implementation). NIS2 IR No. 8 (cyber hygiene). NIST SP 800-53 SI-3. CIS Control 10.

Category	PARTIALLY DEGRADED
Hardening recommendation	EDR with behaviour-based detection and an ML engine (for example CrowdStrike Falcon, SentinelOne Singularity, Microsoft Defender for Endpoint with Defender XDR, Palo Alto Cortex XDR). Sandboxing for unknown files and processes with ML-based detonation analysis (for example Joe Sandbox, ANY.RUN, Hatching Triage). Application allow-listing as complementary protection. Deep inspection at network level for command-and-control traffic.

A.8.12 Data leakage prevention

Category	PARTIALLY DEGRADED
Mythos finding	Content-based DLP takes effect against classic exfiltration patterns but is aggregation-blind: fragmented exfiltration in small, unremarkable transfers over longer periods goes undetected. Behaviour-based DLP with volume and pattern anomaly analysis is robust.
Framework comparison	C5:2026 COS-08 (Policies for Data Transmission) and PI-01 (Safety of Input and Output Interfaces). NIS2 IR No. 8 (cyber hygiene). NIST SP 800-53 AC-4 and SI-4.
Hardening recommendation	DLP with behavioural context and volume anomaly detection. UEBA integration. Egress traffic analysis with an ML-supported baseline. Combination of content-based, context-based and behaviour-based DLP. Automated quarantine at a high risk score.

A.8.16 Monitoring activities

Category	PARTIALLY DEGRADED
Mythos finding	Signature-based monitoring is degraded by aggregation blindness against micro-step attacks. Behaviour-based monitoring with AI support is robust. The classification reflects the fact that many ISMS implement the control in classic form and therefore run it ineffectively against Mythos TTPs.
Framework comparison	C5:2026 OPS-13 (Security Information and Event Management) requires SIEM-based correlation. OPS-18 to OPS-25 for vulnerability management integration. NIS2 IR No. 3.2 (monitoring and logging) with explicit requirements for automated, continuous monitoring. DORA Art. 10.
Hardening recommendation	UEBA with an ML-supported baseline. Kill-chain-oriented detection instead of an isolated alert focus. Correlation of fragmented individual events over longer periods. Integration of MITRE ATT&CK as a detection framework with measurable coverage (validated via DeTT&CT or Atomic Red Team): target value at least 60% coverage of the ATT&CK techniques relevant to the documented organization-specific threat profile. Detection priorities against Mythos TTPs: living-off-the-land binaries (PowerShell with Base64 encoding, schtasks persistence, an Office application spawning cmd or powershell, unusual curl or wget calls from service accounts) and beacon-less C2 (DNS tunnelling, HTTPS C2 with long sleep intervals, cloud-API-based C2 channels such as abused SaaS integrations). A dedicated threat hunting function with a documented methodology (PEAK framework or TaHiTI), minimum capacity 0.5 FTE up to 500 employees and 1 or more FTE above that; at least twelve hypothesis-based hunts per year with ATT&CK mapping. Practicality note: for organizations without a dedicated SOC, managed detection and response (MDR) services (for example CrowdStrike Falcon Complete, Arctic Wolf, SentinelOne Vigilance, Sophos MDR) are a valid alternative with a contractually anchored MTTC SLA.

A.8.20 Networks security

Category	PARTIALLY DEGRADED
Mythos finding	Classic zone architectures with perimeter firewalls and pivot routes between trusted internal zones offer a Mythos attacker lateral movement space that can be scouted in automated form. A zero trust architecture with workload identity is robust.
Framework comparison	C5:2026 COS-01 (Technical Safeguards) to COS-07 (Documentation of the Network Topology). NIS2 IR No. 6.7 (network security) requires a documented network architecture, remote access control and deactivation of services not required; No. 11 (access control) and No. 8 (cyber hygiene) supplement this. NIST SP 800-207 Zero Trust Architecture.
Hardening recommendation	A maturity path to zero trust network architecture in three stages, instead of a big-bang migration. Stage 1 (immediate, bridging concept for classic data centres): internal segment firewalls without a trust pivot between internal zones, identity-aware proxy (for example Cloudflare Access, Tailscale, Zscaler ZPA) for privileged paths. Stage 2 (medium term): SASE for remote access, service mesh with mTLS for the most critical service-to-service paths. Stage 3 (long term): identity-based micro-segmentation at workload level across the estate.

A.8.21 Security of network services

Category	PARTIALLY DEGRADED
Mythos finding	Perimeter-oriented protection of network services is degraded. Identity-based protection through mTLS, SPIFFE identities and service mesh is robust. The control depends on the architectural discipline chosen — classic VPN and jump host models are partially degraded.
Framework comparison	C5:2026 COS-02 (Security Requirements for Connections in the Cloud Service Provider Network) and COS-05 (Networks for Administration). NIS2 IR No. 6.7 (network security) and No. 8. NIST SP 800-207 and SP 800-53 SC family.
Hardening recommendation	mTLS for all service-to-service connections. SPIFFE/SPIRE for workload identity. Service mesh (Istio, Linkerd) for central policy enforcement. API gateway with token-based authentication instead of IP allow-listing.

A.8.22 Segregation of networks

Category	PARTIALLY DEGRADED
Mythos finding	VLAN-based segmentation is too coarse against the lateral movement of an agentic adversary: within a zone, classic network separation mechanisms are often ineffective. Micro-segmentation at identity rather than network level is robust.
Framework comparison	C5:2026 COS-05 (Networks for Administration) and COS-06 (Separation of Data Traffic in Jointly Used Network Environments). NIS2 IR No. 6.7 (network security) and No. 11. NIST SP 800-53 AC-4 and SC-7.
Hardening recommendation	Micro-segmentation at workload level (for example Calico, Cilium). Identity-aware firewalls. East-west traffic inspection with ML-based anomaly detection. Separate admin networks with dedicated infrastructure.

A.8.25 Secure development life cycle

Category	PARTIALLY DEGRADED
Mythos finding	The classic secure SDLC concept with an occasional security review before releases is partially degraded. The speed at which AI-supported attack vectors develop requires security checking in every CI iteration.

Category	PARTIALLY DEGRADED
Framework comparison	C5:2026 DEV-01 to DEV-15 cover the entire SDLC, including DEV-04 (Safety Training), DEV-05 (Design Documentation for Security Features) and DEV-06 (Risk Assessment). NIS2 IR No. 6 (security measures in acquisition, development and maintenance). NIST SP 800-218 SSDF. OWASP SAMM.
Hardening recommendation	AI-supported pre-commit scanning with the current frontier model generation. Concrete tool classes: SAST (Semgrep, SonarQube, Checkmarx), DAST (OWASP ZAP, Burp Suite Enterprise), SCA (Snyk, Dependency-Track, OWASP Dependency-Check), next-generation AI-supported code scanning (GitHub Copilot Autofix, Snyk Code DeepCode AI, Semgrep Pro AI-Assist, Endor Labs Code, Socket.dev). Secure-by-default frameworks and libraries. Threat modelling as a mandatory phase in design reviews with a documented methodology (STRIDE, PASTA, LINDDUN for privacy threats).

A.8.26 Application security requirements

Category	PARTIALLY DEGRADED
Mythos finding	Formalized security requirements remain necessary but take effect only through automated verification. Without a link to security testing in acceptance (A.8.29), the control remains documentation without effect.
Framework comparison	C5:2026 DEV-05 (Design Documentation for Security Features) and DEV-07 (Testing Changes). PSS-01 to PSS-12 for cloud product security requirements. NIS2 IR No. 6. OWASP ASVS as an established requirements catalogue.
Hardening recommendation	Formalize security requirements as executable tests (security as code). Automated compliance checks in CI. Threat modelling as a requirements source using STRIDE or PASTA. OWASP ASVS as the baseline requirements catalogue.

A.8.28 Secure coding

Category	PARTIALLY DEGRADED
Mythos finding	Training and guidelines alone are degraded against Mythos; automated SAST with AI-supported verification is robust. Manual code reviews without tool support are too slow and systematically miss the defect classes that Mythos attackers find in automated form.
Framework comparison	C5:2026 DEV-04 (Safety Training Regarding Continuous Software Delivery). NIS2 IR No. 6 (security measures in development). NIST SP 800-218 SSDF. OWASP Secure Coding Practices. CERT Secure Coding Standards.
Hardening recommendation	SAST in the IDE (shift left) with developer feedback before commit. AI code review as a mandatory stage before merge. Secure coding guidelines enforced via pre-commit hooks. Dependency pinning and supply chain scanning.

A.8.30 Outsourced development

Category	PARTIALLY DEGRADED
Mythos finding	Contractual control alone is degraded under Mythos. External development teams can themselves become an attack vector: through compromise, social engineering or the unintentional introduction of insecure dependencies. Automated artefact verification and SBOM obligations harden the control.
Framework comparison	C5:2026 DEV-02 (Outsourcing of the Development). NIS2 IR No. 5 (supply chain security) and No. 6. CRA Annex I Part II. SLSA for build attestation.
Hardening recommendation	SBOM obligation for all delivered artefacts. SLSA Level 3+ for build provenance. Contractually established code audit rights. Automatic artefact verification in the

Category	PARTIALLY DEGRADED
	organization's own pipeline. Security standards identical to those for internal development.

A.8.32 Change management

Category	PARTIALLY DEGRADED
Mythos finding	Classic change management processes with approval cycles running to several weeks have themselves become a risk against Mythos time compression. Security patches cannot wait weeks when exploits become public within hours. Anthropic explicitly recommends preparing emergency change procedures.
Framework comparison	C5:2026 DEV-03 (Policies for Changes to System Components), DEV-06 (Risk Assessment, Categorisation and Prioritisation of Changes) and DEV-15 (Exceptions to the Change Management Process). NIS2 IR No. 6. ITIL as the classic reference.
Hardening recommendation	Define emergency change procedures in advance and make them auditable. Automated change validation with CI-supported security regression tests. GitOps as the standard change path with automated rollback capability. Separate SLAs for security-critical patches (a maximum of 24 to 48 hours for KEV listings).

5.3 Summary

The 37 partially degraded controls fall into five recurring supplementation patterns which, in combination, enable the transition from partial degradation to Mythos robustness:

- Behaviour-based detection instead of signatures: UEBA, ML-supported anomaly detection, kill-chain correlation. Affects A.8.7, A.8.12, A.8.16 and, supplementarily, A.5.3, A.5.14, A.5.15, A.8.3.
- Phishing-resistant and hardware-bound authentication: FIDO2, passkeys, hardware tokens, mTLS with SPIFFE. Affects A.5.17, A.8.5, A.8.21 and, supplementarily, A.6.7, A.8.1, A.8.4.
- Identity-based network architecture (zero trust): micro-segmentation, identity-aware firewalls, service mesh. Affects A.8.20, A.8.21, A.8.22 and, supplementarily, A.5.15, A.8.3, A.6.7.
- Automated, continuous security testing: AI-supported pre-commit scanning, SAST/DAST in CI, SBOM pipelines. Affects A.8.25, A.8.26, A.8.28, A.8.30, A.8.32 and, supplementarily, A.5.19, A.5.21.
- Supply chain transparency and continuous vendor monitoring: SBOM, SLSA, automated vendor security monitoring. Affects A.5.19, A.5.20, A.5.21, A.5.22, A.5.23.

These five supplementation patterns are not redundant but complementary. They are referenced repeatedly in the hardening recommendations for the individual controls and are brought together in Chapter 9 into the synthesis of the Mythos Hardening Controls.

6 Friction-only – controls whose classic implementation form must be replaced or fundamentally hardened

6.1 Overview and risk profile

A control is deemed friction-only if its core effect against Mythos attackers structurally disappears — either because it rests on limited attacker capacity or because it presupposes human reaction time within an automatically executed kill chain (see the definition in Chapter 3.1.3). The core statement of the Anthropic security engineering team in the Glasswing defensive post (April 2026) — that mitigations whose value lies in generating friction lose considerable effectiveness against an adversary with unlimited patience — applies directly to the following four controls.

What the four controls share is that their protective effect derives not from structural impossibility or cryptographic hardness but either from the assumption of limited attacker capacity or from the assumption of

humanly manageable reaction times. Both assumptions have been empirically refuted under Mythos conditions. The controls must be replaced or fundamentally redesigned in their underlying logic.

Important for interpretation: Classification as friction-only does not mean that these controls are obsolete. They remain useful against certain attacker types and in certain contexts. Under the specific Mythos threat landscape, however, the following applies: the classic or friction-based implementation form must not be treated as the sole effective mitigation for Mythos-relevant risks. An automated or otherwise hardened implementation of the same control can indeed deliver effective risk reduction; its effectiveness must be demonstrated organization-specifically (see Chapters 3.1.3 and Annex F).

6.2 Controls in detail

A.5.25 Assessment and decision on information security events

Category	FRICTION-ONLY
Mythos finding	The control addresses the human triage level: a SOC analyst or security engineer who reviews incoming alerts, classifies them and decides on escalation. Under Mythos, attacker speed systematically outpaces human triage time. In the documented GTG-1002 case, Claude Code executed 80 to 90 percent of the tactical attack steps at request rates physically unattainable by humans. A triage chain that takes decisions in minutes to hours faces an attacker who has already completed several attack stages in the same time. The control remains formally intact but no longer creates an effective response barrier — human reaction time within the automatically executed kill chain is the structural weakness under Chapter 3.1.3.
Framework comparison	C5:2026 SIM-03 (Processing of Security Incidents) and OPS-13 (Security Information and Event Management). NIS2 IR No. 3.4 (assessment and classification of events) requires documented triage criteria and a correlation and analysis procedure for logs. NIST SP 800-61 Revision 3. DORA Art. 17.
Replacement recommendation	Replace human first-line triage with AI-supported tier 1 automation. A SOAR platform with predefined playbooks for autonomous containment actions on unambiguous signals (known IOCs, volume anomalies above thresholds, credential stuffing patterns). Concentrate human escalation on ambiguous cases and high-impact scenarios. Reposition the SOC as a threat hunting and incident commander function rather than an alert review function. Measurable SLA: mean time to containment below 10 minutes for high-confidence alerts. Practicality note: SOAR platforms (Splunk SOAR, Palo Alto XSOAR, Microsoft Sentinel with automation rules) require dedicated engineering for playbook maintenance. For organizations without a 24/7 SOC, managed detection and response (for example CrowdStrike Falcon Complete, Arctic Wolf, SentinelOne Vigilance, Sophos MDR) with a contractually anchored MTTC SLA is the more workable alternative.

A.5.36 Compliance with policies, rules and standards for information security

Category	FRICTION-ONLY
Mythos finding	The control checks activities and system states against documented policies. The fundamental weakness: the check rests on the assumption that deviating attack actions are recognizable as such. An agentic attacker who decomposes an attack into micro-steps that each look policy-conformant structurally bypasses this check. The aggregation blindness of classic compliance checks — quarterly reviews, policy comparisons against configuration baselines — makes them friction-only: they create effort for the attacker but no protection where that effort can be automated.
Framework comparison	C5:2026 COM-03 (Internal Audits of the Information Security Management System) and COM-01 (Identification of Applicable Legal, Regulatory, Self-imposed or Contractual Requirements). NIS2 IR No. 2.2 (monitoring of compliance) requires effective reporting and regular conformity checking. DORA Art. 6(5) and Art. 24.

Category	FRICTION-ONLY
Replacement recommendation	Continuous control monitoring in addition to periodic audits. Automated, real-time-capable compliance checks against baselines. Policy as code (OPA/Rego, Sentinel) with CI integration and runtime enforcement. Integration with UEBA and SIEM to detect activities that are policy-conformant but behaviourally anomalous. Compliance dashboards with quantified deviation indicators instead of PDF audit reports. Linkage to automated incident tickets on deviation.

A.8.8 Management of technical vulnerabilities

Category	FRICTION-ONLY
Mythos finding	Classic vulnerability management with monthly or quarterly patch cycles, change advisory board dates and three-stage test phases is friction-only under Mythos against the collapse of the patch gap window described in Chapter 2.1. Where only hours lie between the publication of a patch and a working exploit, a process with approval cycles running to several weeks is not a protective measure but a structural weakness. The existence of an orderly, documented vulnerability management process produces compliance evidence but not Mythos resistance. Clarification: the friction classification concerns the cyclical patch mechanism. The organizational process shell of the control — vulnerability ownership, SLA accountability, escalation paths — remains effective and is a precondition for the automation under MHC-09 (cf. Chapter 3.1.3).
Framework comparison	C5:2026 offers the most granular comparison standard: OPS-18 (Managing Vulnerabilities – Policies), OPS-25 (Vulnerability Scans) with the additional sharpening level OPS-25.01AS (daily scan frequency), OPS-27 (Patch Management Policies) and OPS-28 (Patch Management Implementation). NIS2 IR No. 6 (security measures in acquisition, development and maintenance). CRA Art. 13/14 requires vulnerability handling with documented response times for software manufacturers. NIST SP 800-40.
Replacement recommendation	EPSS-based prioritization (Exploit Prediction Scoring System) instead of pure CVSS weighting. A 24-hour SLA for KEV listings (CISA Known Exploited Vulnerabilities). Automated patch pipelines with canary deployments, automatic rollback and security regression tests. A separate hotfix channel outside regular change management for KEV entries. AI-supported pre-ship vulnerability scanning (see A.8.29) as a shift-left supplement. Emergency change procedures approved in advance and auditable (see A.8.32).

A.8.23 Web filtering

Category	FRICTION-ONLY
Mythos finding	URL- and reputation-based web filters fail structurally under Mythos. Mythos-generated phishing infrastructure rotates faster than reputation feeds can be updated: domain generation algorithms, short-lived lookalike domains with automatically generated, legitimate-looking landing pages, payloads hosted on compromised legacy infrastructure. The friction that web filters generate — that the attacker must register more phishing domains — is not significant effort for an automated adversary.
Framework comparison	C5:2026 COS-04 (Cross-Network Access). NIS2 IR No. 8 (cyber hygiene). NIST SP 800-53 SC-7 (Boundary Protection). CIS Control 9 (Email and Web Browser Protections).
Replacement recommendation	DNS security as a protective DNS layer with AI-supported detection (Cloudflare Gateway, Cisco Umbrella, Quad9) instead of pure URL block lists. Remote browser isolation (RBI) for all external links, particularly from mail. DNS over HTTPS exclusively to controlled, trusted resolvers. Structurally: phishing-resistant MFA (A.5.17, A.8.5) as a hard barrier that takes effect even where the user reaches a phishing page.

6.3 Summary

The four friction controls share a logic: their protective effect rests on the assumption of disproportionate attacker effort (A.8.8, A.8.23, A.5.36) or of humanly manageable reaction time (A.5.25). Mythos-class AI removes both, because the marginal cost of a further attack iteration is close to zero. Consequence for the Statement of Applicability: these controls should not be recorded as the sole mitigation for Mythos-relevant risks; the risk assessment must document which supplements from the Robust and Partially degraded categories take over the actual protective effect.

7 Not affected controls – the neutral baseline

7.1 Controls without Mythos relevance

Twenty-three controls of ISO/IEC 27002:2022 are neutral with respect to the Mythos threat landscape — predominantly organizational, documentary, governance-oriented or physically bound controls. They remain important for building an ISMS and for compliance, but receive no separate Mythos hardening recommendation, because their classification stands unchanged.

7.2 Overview of the 23 not affected controls

ID	Title	Short rationale
A.5.1	Policies for information security	Policy artefact. Effectiveness arises from implementation in downstream controls.
A.5.2	Information security roles and responsibilities	Governance artefact. Role assignment as such is not affected by Mythos.
A.5.4	Management responsibilities	Leadership accountability and management commitment. Not Mythos-sensitive.
A.5.6	Contact with special interest groups	Community exchange with professional bodies. Mythos-neutral.
A.5.8	Information security in project management	Procedural control for integrating security into project life cycles.
A.5.10	Acceptable use of information and other associated assets	Usage policy. Effectiveness arises from technical enforcement in other controls.
A.5.11	Return of assets	Offboarding artefact. Mythos-neutral.
A.5.13	Labelling of information	Operational artefact of the classification from A.5.12.
A.5.31	Legal, statutory, regulatory and contractual requirements	Compliance inventory.
A.5.32	Intellectual property rights	License and IP administration.
A.5.37	Documented operating procedures	Operational artefact.
A.6.1	Screening	Pre-employment check. Background checks address insider risk, not AI-accelerated external attacks.
A.6.2	Terms and conditions of employment	Contractual artefact.
A.6.4	Disciplinary process	HR process.

ID	Title	Short rationale
A.6.6	Confidentiality or non-disclosure agreements	Legal artefact.
A.7.5	Protecting against physical and environmental threats	Environmental protection against fire, water and power failure.
A.7.7	Clear desk and clear screen	Addresses physical presence attacks (shoulder surfing, unattended documents); Mythos attacks are by definition remote and agentic.
A.7.8	Equipment siting and protection	Physical siting question.
A.7.11	Supporting utilities	Uninterruptible power supply, air conditioning.
A.7.12	Cabling security	Physical layer of the infrastructure.
A.7.13	Equipment maintenance	Operational maintenance with physical access.
A.8.6	Capacity management	Capacity planning against overload; defence against targeted exhaustion attacks is covered by A.8.14 (redundancy) and A.8.20/A.8.21 (network security). Note: automated exhaustion patterns — API abuse, autoscaling and cost exhaustion — are relevant under Mythos conditions but concern the controls named, not capacity planning as such.
A.8.34	Protection of information systems during audit testing	Operational audit artefact for avoiding audit side effects.

Part III – Convergence analysis: gaps and synthesis

The individual assessment of all 93 controls in Part II shows which controls hold, which must be supplemented and which must be replaced. Part III systematically identifies those requirements that other regulatorily relevant frameworks state explicitly, that ISO 27002:2022 does not contain in that form or covers only implicitly through the risk perspective, and that provide resilient protective effect under Mythos conditions. This is not about gaps in the ISO methodology itself but about effectiveness aspects that other frameworks name more concretely. The result is a prioritized supplementary catalogue of concrete controls.

The convergence analysis proceeds in two steps: Chapter 8 groups the gaps thematically into seven clusters and explains for each why the requirement concerned is Mythos-robust. Chapter 9 consolidates the results into a numbered catalogue of Mythos Hardening Controls.

8 Gap analysis – what other frameworks require explicitly that ISO 27002 covers only implicitly

8.1 Methodology of the gap analysis

The gap analysis proceeds in three stages. First, all requirements of the five primary comparison frameworks — C5:2026, NIST, DORA, the CRA and NIS2 with its Implementing Regulation — are compared against the ISO 27002:2022 catalogue. Second, those requirements are extracted that ISO 27002 does not contain, or contains only in very abstract form. Third, each extracted requirement is examined for whether it provides hard protective effect under Mythos conditions — that is, whether it satisfies one of the criteria in Chapter 3.2: cryptographic impossibility, structural reduction of the blast radius, aggregation resistance, or a hard time barrier.

Requirements that merely restate controls already present in ISO 27002 are not counted as gaps. Nor are requirements that, while more detailed in another framework, are themselves degraded under Mythos conditions. The following analysis is confined to substantive differences from the ISO text that are simultaneously Mythos-robust.

The gaps group into seven thematic clusters, ordered by operational proximity to the Mythos core threat.

8.2 Cluster 1: post-quantum strategy and crypto agility

ISO 27002 reference: A.8.24 (use of cryptography) requires a documented, risk-based approach to cryptographic mechanisms but remains at an abstract level. No explicit requirement for post-quantum preparation, a cryptographic inventory or crypto agility.

Requirement in other frameworks: C5:2026 CRY-01.01AC requires, in binding form, a documented post-quantum strategy with four elements: a cryptographic inventory with priority levels, ongoing observation of the state of the art, use of hybrid cryptography models, and defined triggers, resources, transition plans and success criteria for the PQC migration. CRY-01.03AC requires at least annual review. CRY-02 (Cryptographic Change Management) anchors crypto agility operationally.

Mythos robustness: The gap is Mythos-robust for two reasons. First, the store-now-decrypt-later threat model is concrete and effective today. Exfiltrated data can be stored for years until quantum computing is operationalized. Second, crypto agility is a structural precondition for reacting to suddenly devalued algorithms without entering weeks of re-engineering cycles. A qualification applies: the core of this gap — the quantum threat — does not result from Mythos but exists independently of it; through the risk perspective, ISO already implicitly requires appropriate methods. A third aspect is Mythos-specific: PQC methods are comparatively new and less battle-tested than RSA or ECDH, so AI-supported cryptanalysis could weaken immature schemes faster. That further raises the requirement for crypto agility — for example several trust anchors with different algorithms and the ability to deactivate individual ones at short notice.

Key management facets. An analysis by the Cloud Security Alliance on PQC key management (Cloud Key Management Working Group, draft, August 2026) concretizes four aspects that a pure algorithm migration overlooks. First, the trust-now-forge-later variant of the harvesting attack: not only encrypted data but also signatures valid today — certificates, code signatures, long-lived signed documents — can be harvested and later retroactively forged or repudiated using a cryptographically relevant quantum computer; long-lived signatures therefore belong in the inventory and the migration scope. Second, the downgrade risk of hybrid protocols: an attacker can force a hybrid negotiation (for example hybrid TLS) back onto the classical component; what is required is enforced quantum-safe negotiation, rejection of purely classical sessions, and logged downgrade detection. Third, the migration order within the key hierarchy: the root must be migrated to PQC first (master or key encryption key, root CA), then the derived keys — PQC only at lower levels with a classical root leaves the trust anchor vulnerable, particularly for archived data with long protection requirements. Fourth, quantum-safe timestamps: a classical timestamping authority (TSA) weakens the non-repudiation chain irrespective of the signature algorithm; signature and TSA certificates must be migrated to PQC, with periodic timestamp renewal for long-term archives.

8.3 Cluster 2: supply chain transparency and SBOM obligation

ISO 27002 reference: A.5.19 to A.5.23 address the supply chain at policy and contract level. No requirement for technical transparency about the software components in use, no SBOM obligation and no binding build provenance requirement.

Requirement in other frameworks: C5:2026 DEV-13 (Transparency about Software Components). Under EU CRA Annex I Part II, producing an SBOM is one of the essential requirements for vulnerability handling; satisfying the applicable requirements is in turn a precondition of conformity, on the basis of which CE marking takes place. This is binding only for manufacturers of products with digital elements placed on the EU market; for other organizations the SBOM remains a risk-based decision. DORA Art. 28 requires detailed transparency for ICT third-party providers. NIS2 IR No. 5 (supply chain security) requires documented dependencies. The SLSA framework defines build provenance levels.

Mythos robustness: An SBOM enables structural detection that does not rest on attacker friction. The attack becomes recognizable as soon as the compromised component is published in EUVD, CVE or OSV. The reduction in detection latency is structural, not temporal.

8.4 Cluster 3: containers, confidential computing and multi-tenancy

With its March 2026 revision, C5:2026 introduced three structurally new control domains that ISO 27002 does not contain and that are explicitly tailored to modern cloud workload architectures.

Container security

ISO 27002 reference: No container-specific controls. General environment and configuration controls (A.8.9, A.8.19, A.8.31) are applicable but do not address the specific container attack surfaces.

Requirement: C5:2026 OPS-34 (Container Management – Policies and Procedures) and OPS-35 (Implementation), together with PSS-11 (Images for Virtual Machines and Containers), require documented image sources, signed images, runtime protection and network segmentation at container level. NIST SP 800-190 as the reference.

Mythos robustness: Container signing creates a cryptographic and structural barrier against compromised base images from public registries.

Confidential computing

ISO 27002 reference: No requirements for trusted execution environments or remote attestation.

Requirement: C5:2026 OPS-32 (Confidential Computing – Policies and Procedures) and OPS-33 (Remote Attestation). NIST trusted computing concepts as the reference.

Mythos robustness: Confidential computing is currently the hardest barrier against attackers with infrastructure access, because it protects the confidentiality of data in processing — not only in transit and at rest.

Multi-tenancy isolation

ISO 27002 reference: No specific multi-tenancy controls. Generic separation controls (A.8.22, A.8.31) are not sufficient.

Requirement: C5:2026 OPS-30 (Separation of Datasets – Policies and Procedures) and OPS-31 (Implementation) explicitly require the separation of customer data in multi-tenant environments with demonstrable implementation. PSS-10 (Software Defined Networking) supplements the network side.

Mythos robustness: Structural limitation of the blast radius: an attacker who gains a foothold in one tenant is demonstrably prevented from accessing other tenants.

8.5 Cluster 4: binding notification deadlines and root cause analysis

ISO 27002 reference: A.5.5 (contact with authorities) and A.5.25 (assessment of events) address notification channels without concrete deadlines. A.5.27 (learning from incidents) does not require root cause analysis.

Requirement in other frameworks: NIS2 Directive Art. 23(4) requires an early warning within 24 hours, a full notification within 72 hours for significant incidents, and a final report after one month. DORA Art. 19 requires major incident reports to supervisory authorities with detailed content requirements. DORA Art. 17 makes root cause analysis a mandatory phase in the incident management life cycle. The Cyber Resilience Act (CRA) additionally obliges those placing products with digital elements on the market to report actively exploited vulnerabilities (including a 24-hour early warning to ENISA or the competent CSIRT). NIS2 IR No. 3.6 requires post-incident reviews including root cause determination.

Mythos robustness: Time-based notification deadlines are Mythos-robust because they mirror, on the regulatory side, the time compression that benefits Mythos attackers: an organization that must report within 24 hours needs corresponding automation on the detection and triage side. Root cause analysis creates the empirical basis for detection improvement. The deadlines do not, however, apply equally to every organization — what matters is the role (for example placing products on the market under the CRA) and the classification under NIS2; for organizations not covered, the gap is correspondingly smaller. Through the compliance obligations of an ISMS (consideration of the expectations of interested parties, including legal requirements), the applicable deadlines feed into the security objectives in any case.

8.6 Cluster 5: continuous checking in addition to periodic audits

ISO 27002 reference: A.5.35 (independent review) describes periodic audits. A.5.36 (compliance with policies) and A.8.8 (vulnerability management) are classified as friction-only in Chapter 6.

Requirement in other frameworks: C5:2026 OPS-25.01AS requires daily vulnerability scans as an additional sharpening. COM-03 and COM-04 require continuous ISMS effectiveness checking. NIS2 IR No. 2.2 (monitoring of compliance) requires an effective reporting system. DORA Art. 24–26 requires a digital operational resilience testing programme; Art. 26–27 requires threat-led penetration testing (TLPT) for critical functions.

Mythos robustness: Continuous checking offsets time compression on the defensive side. Detection latency falls from months and quarters to hours and days. TLPT with Mythos-specific scenarios is the best-known regulatory regime to address the empirical level explicitly: it tests whether the controls work, not whether they are documented.

8.7 Cluster 6: phishing-resistant identity and zero trust architecture

ISO 27002 reference: A.5.17 and A.8.5 require authentication appropriate to the protection requirement, without naming concrete methods or assurance levels. A.5.16 requires identity management without an explicit workload identity requirement.

Requirement in other frameworks: NIST SP 800-63B defines three authentication assurance levels; AAL2 and AAL3 are relevant for Mythos environments. FIDO Alliance standards (WebAuthn, passkeys). NIST SP 800-207 Zero Trust Architecture defines workload identity as a structural principle. NIST NCCoE work on agent identity. SPIFFE/SPIRE as the de facto standard.

Mythos robustness: Phishing-resistant MFA is the direct answer to AI-quality phishing: even if a user lands on a perfectly crafted phishing page, the attacker cannot obtain a usable authentication credential, because FIDO2 performs origin binding. Workload identity is structural: it replaces shared secrets with hardware-anchored, short-lived identities.

8.8 Cluster 7: automation obligation and resilience testing

ISO 27002 reference: No express automation obligation. A.5.30 (ICT readiness for business continuity) and A.8.16 (monitoring) mention no concrete test scenarios for parallel incidents.

Requirement in other frameworks: NIS2 IR No. 3.2.2 requires automated monitoring to the extent feasible. NIS2 IR No. 3.5.5 requires testing of the response procedures. DORA Art. 24–26 requires a comprehensive digital operational resilience testing programme. DORA Art. 26–27 requires TLPT with scenarios that reflect real threats. In the Glasswing defensive post, Anthropic explicitly recommends tabletop exercises for three to five parallel incidents.

Note on the strength of the NIS2 IR wording: The Implementing Regulation frames the automation obligation as a "should" provision ("to the extent feasible", "subject to operational capacity") and not as an absolute obligation. For Mythos-exposed organizations, feasibility should be construed restrictively, since manual response times fail structurally against agentic attackers.

Mythos robustness: The automation obligation is the only realistic answer to the asymmetry between an agentic attacker and a human defensive chain. Parallel incident testing directly addresses the fragmentation strategy.

8.9 Summary: the gap delta of ISO 27002

Together, the seven clusters describe the systematic delta between ISO/IEC 27002:2022 and the regulatory requirements of the past three years. This delta was not designed into ISO 27002 as a deficiency — it reflects the framework-by-framework specialization of the standards landscape.

Under Mythos conditions, this delta becomes an operationally relevant problem. An organization relying on the pure ISO 27002 baseline does not have the concrete requirements that C5:2026, NIST, DORA, the CRA and NIS2 with its Implementing Regulation impose in response to the current threat landscape. Closing the delta is not a compliance exercise but the core task of operational Mythos hardening.

The seven clusters can be condensed into two movements: a structural modernization of the security architecture (cluster 1 cryptography, 3 cloud infrastructure, 6 identity) and a temporal shortening of the defensive loops (cluster 2 supply chain, 4 notification obligations, 5 continuous checking, 7 automation). Chapter 9 brings these movements together into the prioritized catalogue of Mythos Hardening Controls.

9 Synthesis – catalogue of Mythos Hardening Controls

9.1 Purpose and use of the catalogue

This chapter consolidates the findings from Chapter 8 into a compact numbered catalogue of thirteen Mythos Hardening Controls (MHC).

Note: use the MRIS IMPLEMENTATION GUIDE as the implementation guideline for the thirteen Mythos Hardening Controls.

Each MHC is worded so that it can be adopted directly into the Statement of Applicability of an existing ISO 27001-certified ISMS.

The thirteen MHC form a prioritized supplementary catalogue for the Mythos threat landscape. They are applied in addition to the existing ISO controls but address protective objectives that ISO does not explicitly contain or describes only in abstract terms. Each MHC references an ISO 27002 anchor.

The numbering MHC-01 to MHC-13 follows operational proximity to the Mythos core threat, not an implementation sequence. Each MHC stands on its own and can be introduced independently; the overall effect arises from combined application.

The catalogue makes no claim to completeness. It is confined to the currently most robust and operationally clearly describable additions. The version notes document the development of this catalogue.

Prioritization relative to existing controls. The thirteen MHC supplement; they do not replace. Before investing in new Mythos Hardening Controls comes the consistent implementation of the existing structural controls from Chapter 4 (in particular zero trust architecture, strong cryptography, least privilege, secrets management, micro-segmentation, immutable backups). Under Mythos, the structural controls remain the most robust defensive layer. An organization that has not implemented these controls consistently gains less Mythos resilience from introducing new MHC than from hardening what already exists. The assessment language in Chapters 5 and 6 is therefore to be read as "selectively degraded" — not as "obsolete". Mythos weakens classic controls in specific effectiveness dimensions but they retain their full protective effect in others. The correct response is hardening, not replacement.

9.2 Thirteen Mythos Hardening Controls

MHC-01 Post-quantum strategy and cryptographic inventory

Positioning: a cross-cutting resilience control with heightened relevance under AI-accelerated crypto discovery and exploit economics. The quantum threat itself exists independently of Mythos; AI accelerates the discovery of weak cryptography, migration pressure and exploitation economics.

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: flanks A.8.24 (use of cryptography).
Framework comparison	Framework basis: C5:2026 CRY-01.01AC to CRY-01.03AC. BSI TR-02102. FIPS 203 (ML-KEM), FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA) as finalized PQC standards (13 August 2024); HQC as a backup KEM (selected 11 March 2025, standardization ongoing). NIST SP 1800-38 as the NCCoE migration guide. European Commission roadmap for post-quantum migration.
Hardening recommendation	Core idea: a documented PQC strategy with four elements — a cryptographic inventory of all algorithms and keys, ongoing observation of the state of the art, hybrid cryptography models, and defined trigger events and transition paths (root of the key hierarchy first, downgrade-resistant hybrid negotiation). Mythos effect: addresses store-now-decrypt-later and creates structural crypto agility; the protective effect rests on cryptographic impossibility, not on attacker friction. → Implementation, maturity path and evidence: Implementation Guide, MHC-01.

MHC-02 SBOM and build provenance

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily flanks A.5.21 (ICT supply chain management) and A.8.30 (outsourced development); for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 DEV-13. CRA Annex I Part II as a legally binding SBOM obligation for manufacturers of products with digital elements (vulnerability notification obligations from September 2026, main obligations from 11 December 2027). DORA Art. 28. NIS2 IR No. 5. SLSA framework. BSI TR-03183-2 as the SBOM quality standard. Market

Category	MYTHOS HARDENING CONTROL
	position 2026: the disclosure waves now beginning confirm the volume demand for SBOM query capability (Chapter 2.5).
Hardening recommendation	Core idea: automated SBOM generation in the CI pipeline (SPDX/CycloneDX, CRA-conformant) for all self-developed and distributed artefacts, a contractual SBOM obligation for critical suppliers, automated vulnerability correlation (CVE/EUVD/OSV) with separate exploitability communication via VEX/CSAF, and SLSA Level 3+ for build provenance. Mythos effect: creates structural detection capability for supply chain attacks; detection latency falls to the publication time in vulnerability databases. → Implementation, maturity path and evidence: Implementation Guide, MHC-02.

MHC-03 Phishing-resistant multi-factor authentication

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily flanks A.5.17 (authentication information) and A.8.5 (secure authentication); for the complete mapping see Annex A.
Framework comparison	Framework basis: NIST SP 800-63B Revision 4 (final 31 July 2025), AAL2/AAL3 — AAL3 requires device-bound, non-exportable keys; synchronizable passkeys reach AAL2 only. Terminological clarification: a device-bound passkey satisfies the multi-factor criterion only together with local user verification (PIN or biometrics); without it, only a possession factor is present. What matters for protection against Mythos phishing is origin binding, not the number of factors. C5:2026 IAM-08. FIDO Alliance standards (WebAuthn, passkeys, CTAP2). NIS2 IR No. 11.6 and No. 11.7; for privileged accounts additionally No. 11.3.2(a).
Hardening recommendation	Core idea: origin-bound methods (WebAuthn/FIDO2) for all privileged and externally reachable access; for the highest protection requirement only attested, device-bound authenticators (attestation enforcement), with synchronizable and software passkeys (passkey vaults) excluded; SMS MFA and push without number matching excluded for new deployments. Mythos effect: origin binding renders authentication credentials worthless on phishing pages; a structural answer to AI-quality phishing. Effectiveness limit: phishing-resistant is not phishing-proof — endpoint compromise and session theft after login remain to be addressed through MHC-05 and A.8.7. → Implementation, maturity path and evidence: Implementation Guide, MHC-03.

MHC-04 Workload identity and zero trust network architecture

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily flanks A.5.16 (identity management), A.8.20 (networks security) and A.8.21 (security of network services); for the complete mapping see Annex A.
Framework comparison	Framework basis: NIST SP 800-207 Zero Trust Architecture. NIST NCCoE work on agent identity (2026, announced). SPIFFE/SPIRE. Service mesh architectures (Istio, Linkerd).
Hardening recommendation	Core idea: workload identity — every machine, service and container receives a unique, cryptographically verifiable identity with short-lived, automatically rotating credentials instead of shared access data — (SPIFFE/SPIRE, mTLS) in three maturity stages rather than a big bang: privileged workloads, production mainstream, full coverage. AI agents are a special case with their own time-limited, capability-scoped identities per tool call instead of personal developer credentials, with a complete audit trail. Mythos effect: structurally reduces classic credential-based lateral movement paths and limits the reachable blast radius; shared secrets are replaced by hardware-anchored, short-lived credentials. A compromised legitimate workload can still act within its own entitlements — MHC-04 must therefore be combined with detection (MHC-05) and entitlement

Category	MYTHOS HARDENING CONTROL
	minimization. → Implementation, maturity path and evidence: Implementation Guide, MHC-04.

MHC-05 Behaviour-based detection and kill-chain correlation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily flanks A.8.16 (monitoring activities) and A.8.12 (data leakage prevention); for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 OPS-13 (SIEM). NIS2 IR No. 3.2. DORA Art. 10. MITRE ATT&CK as the operational detection framework (continuously updated, content version v19.2). NIST SP 800-94 only as dated background (2007 edition; the draft Revision 1 was withdrawn and no current successor exists).
Hardening recommendation	Core idea: UEBA with an ML baseline and ATT&CK-based detection rules (measurable coverage), kill-chain correlation across events and time windows; priorities against Mythos TTPs (living off the land, beacon-less C2, WebAuthn and passkey anomalies); threat hunting as a function in its own right; adversarial robustness of the organization's own detection AI. Mythos effect: addresses the aggregation blindness of classic detection; fragmented micro-step attacks become recognizable through pattern correlation. Realism note: detection is not prevention — low-and-slow attacks, agents disguised within legitimate baselines and adversarial ML evasion remain residual risk; always to be assessed in combination with the structural controls from Chapter 4. → Implementation (including the MDR alternative), maturity path and evidence: Implementation Guide, MHC-05.

MHC-06 Container security and confidential computing

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: A.5.23 (cloud services, flanked) and A.8.31 (separation of environments).
Framework comparison	Framework basis: C5:2026 OPS-34/35 (Container Management), OPS-32/33 (Confidential Computing) and PSS-11. NIST SP 800-190 (2017 edition, still the authoritative container baseline). Confidential Computing Consortium.
Hardening recommendation	Core idea: for containers, signed base images from controlled registries (Sigstore/cosign), admission controller policies (OPA Gatekeeper, Kyverno) and image scans before deployment; for confidential computing, TEE and secure enclaves (Intel TDX, AMD SEV-SNP, ARM CCA) with remote attestation for elevated protection requirements. Mythos effect: container signing creates a cryptographic hard barrier against manipulated images; confidential computing is currently the hardest barrier against attackers with access to provider infrastructure. → Implementation, maturity path and evidence: Implementation Guide, MHC-06.

MHC-07 Multi-tenancy isolation with demonstrable separation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: A.5.23 (cloud services, flanked) and A.8.22 (segregation of networks); extended by data- and data-centre-specific aspects of tenant separation.
Framework comparison	Framework basis: C5:2026 OPS-30 and OPS-31. PSS-10 (Software Defined Networking). CSA Cloud Controls Matrix. ISO/IEC 27017 (information security for cloud services).
Hardening recommendation	Core idea: documented separation concepts for data (tenant-specific keys, logical or physical separation), networks (VPCs, namespaces, SDN rules) and compute (tenant-specific scheduling policies), with demonstrable separation through regular, ideally

Category	MYTHOS HARDENING CONTROL
	automated tests. Mythos effect: structural limitation of the blast radius — an attacker who gains a foothold in one tenant is demonstrably prevented from accessing others. → Implementation, maturity path and evidence: Implementation Guide, MHC-07.

MHC-08 Immutable backups and recovery validation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily deepens A.8.13 (information backup), which is robust in principle but does not explicitly state concrete requirements for immutability; for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 OPS-06 to OPS-09 with separate treatment of policies, monitoring, testing and storage. DORA Art. 12. NIS2 IR No. 4.1. Industry practice 3-2-1-1-0 model.
Hardening recommendation	Core idea: the 3-2-1-1-0 model (three copies, two media types, at least one offsite and one immutable or offline, zero errors in the restore test), separate IAM paths for the backup infrastructure, and quarterly documented restore tests. Mythos effect: immutable backups are one of the most effective hard barriers against ransomware and agentic data integrity attacks. → Implementation, maturity path and evidence: Implementation Guide, MHC-08.

MHC-09 AI-supported security testing in the pipeline

Scope note: the control presupposes automated security testing (SAST, DAST, SCA) as its basis and describes the Mythos-specific addition on top — next-generation AI-supported scanning methods and assurance for AI-generated code. The remediation latency metric (Annex G.4) measures not the test method itself but its effect in operation.

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: primarily flanks A.8.29 and A.8.25; for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 DEV-07, OPS-22 and OPS-25 with the sharpening OPS-25.01AS (daily scans). NIST SP 800-218 SSDF, supplemented by NIST SP 800-218A (Secure Software Development for Generative AI, 26 July 2024); SSDF v1.2 (SP 800-218r1) in draft (17 December 2025). OWASP ASVS. CRA Annex I Part II. Multi-system position: comparable agentic testing capability is available from several vendors (including Daybreak and MDASH; Chapter 2.5); the choice of tooling should be made vendor-independently.
Hardening recommendation	Core idea: SAST/DAST/SCA plus next-generation AI-supported code scanning in the pipeline, a pre-merge block for high and critical findings at 80% confidence or above (KEV blocking immediately), daily scans, and Mythos-specific penetration tests. Address the secondary threat as well: AI coding assistants generate insecure patterns, hence pre-commit AI code audit and treatment of AI-generated code as a distinct risk factor in the secure SDLC standard and in the risk assessment (risk register or application risk classification), not in the Statement of Applicability — that document records the applicability of controls, not risks. Mythos effect: shift-left removes the attacker's window between exploit availability and the production patch; a structural answer to the collapse of the patch gap window. → Implementation, maturity path and evidence: Implementation Guide, MHC-09.

MHC-10 Continuous control monitoring and policy as code

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: largely replaces the effectiveness of A.5.35 and A.5.36 under Mythos conditions; for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 COM-03 and COM-04. NIS2 IR No. 2.2. DORA Art. 6(5). OWASP SAMM. NIST SP 800-137/137A (Information Security Continuous Monitoring). NIST OSCAL (machine-readable control catalogues and evidence). Industry practice Open Policy Agent (OPA/Rego), HashiCorp Sentinel.
Hardening recommendation	Core idea: policy as code (OPA, Sentinel) as executable rules, integrated into CI (pre-deploy) and runtime enforcement; continuous compliance checks against baselines, real-time dashboards with quantified deviation indicators, and automatic incident tickets on deviation. Mythos effect: shortens detection latency for compliance deviations from months or quarters to seconds or minutes. → Implementation, maturity path and evidence: Implementation Guide, MHC-10.

MHC-11 SOAR-based tier 1 automation and parallel response playbooks

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: largely replaces the effectiveness of A.5.25 and supplements A.5.24 and A.5.26; for the complete mapping see Annex A.
Framework comparison	Framework basis: C5:2026 SIM-02 and SIM-03. NIS2 IR No. 3.5 (incident response, three phases under 3.5.2). DORA Art. 17. NIST SP 800-61 Revision 3 (incident response as a CSF 2.0 profile, 3 April 2025; supersedes the 2012 edition). Anthropic recommendation from the Glasswing defensive post on parallel incident scenarios.
Hardening recommendation	Core idea: a SOAR platform with predefined playbooks for fully automated containment on unambiguous signals and human escalation for ambiguous cases; Mythos-specific playbooks (mass exfiltration, multi-vector, supply chain, MFA fatigue), tabletop exercises for three to five parallel incidents; hardening of the automation layer itself (signed triggers, an audit trail with execution identity, rate limits, human in the loop where the blast radius is high, playbook protection equivalent to source code). Mythos effect: closes the time lag between agentic attacker speed and the human response chain; MTTC from hours to minutes. Effectiveness limit: the SOAR logic can itself become a target — without the hardening of the automation layer described, MHC-11 is not effective. → Implementation (including the MDR alternative), maturity path and evidence: Implementation Guide, MHC-11.

MHC-12 Threat-led penetration testing with Mythos scenarios

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: supplements A.5.35 and A.8.29.
Framework comparison	Framework basis: DORA Art. 26 and 27 as currently the only binding regulatory TLPT regime, specified by the TLPT RTS (Delegated Regulation (EU) 2025/1190, applicable since 8 July 2025). The ECB's TIBER-EU framework (aligned with DORA since 11 February 2025, with mandatory purple teaming; in Germany TIBER-DE via the Bundesbank and BaFin). C5:2026 OPS-22 (Penetration Tests). NIS2 IR No. 3.5.5 with a requirement for regular testing of the response procedures.
Hardening recommendation	Core idea: threat-led penetration testing with external red teamers based on real scenarios, extended in Mythos-specific form (AI spear phishing with deepfake voice, fragmented attack chains, multi-vector, supply chain, cloud API abuse); at least annually for critical functions, with purple teaming between cycles and ATT&CK coverage; the scenario pool should be built multi-systemically (at least two independent agentic systems). Mythos effect: TLPT tests effectiveness against real Mythos TTPs rather than

Category	MYTHOS HARDENING CONTROL
	against documentation; it reduces false-positive compliance assurances. → Implementation (including lower-cost formats for organizations outside DORA), maturity path and evidence: Implementation Guide, MHC-12.

MHC-13 AI agent governance and harness security

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 anchor: extends A.5.9 (inventory, shadow AI capture), A.5.16 (identity management), A.8.27 (secure system architecture) and MHC-04 by the specific governance and security requirements of agentic AI systems operated by the organization.
Framework comparison	Framework basis: OWASP Top 10 for LLM Applications 2026 (in particular LLM01 prompt injection, LLM03 excessive agency). OWASP Agentic Security Initiative (ASI01 to ASI10). MITRE ATLAS (AML.T0051 LLM Prompt Injection, AML.T0053 AI Agent Tool Invocation (formerly LLM Plugin Compromise; ATLAS status August 2026), AML.T0086 Exfiltration via AI Agent Tool Invocation, AML.T0110 AI Agent Tool Poisoning). Anthropic, OpenAI and Google responsible use guidelines. NIST AI RMF 1.0. ISO/IEC 42001 Annex A.
Hardening recommendation	Core idea: governance of AI agents across five dimensions — harness audit (prompts, tools and retrieval treated like production code, versioned, signed, with pre-production penetration tests), blast radius limits (capability scoping per tool call, rate limits, circuit breakers, human approval for irreversible actions), human override (an owner with a kill switch, an audit trail with human correlation), a supply chain inventory of agentic components (MCP servers, IDE extensions, agentic skills) and pre-production checks (scope, threat model, rollback, monitoring). Mythos effect: closes the gap classified as CRITICAL in Mythos-Ready, "unmanaged AI agent attack surface" — simultaneously a defensive risk and a supply chain risk. Maturity realism: conceptually mature, but operational implementation readiness is currently at Initial level in most organizations (a 12 to 24 month transition); prioritize inventory, audit trail and capability scoping first. → Implementation, maturity path and evidence: Implementation Guide, MHC-13.

9.3 Derivation from the gap clusters

Twelve of the thirteen MHC can be traced primarily to the seven gap clusters from Chapter 8; MHC-08 arises primarily from the control degradation analysis and is additionally supported by external framework requirements:

- Cluster 1 (cryptography and post-quantum): MHC-01.
- Cluster 2 (supply chain transparency): MHC-02; MHC-13 adds agentic supply chain components (MCP servers, IDE extensions, agentic skills).
- Cluster 3 (containers, confidential computing, multi-tenancy): MHC-06 and MHC-07.
- Cluster 4 (notification obligations and root cause): MHC-11 (response side); notification deadline compliance is not listed as a separate MHC, because it must be implemented in any case as a direct legal obligation under NIS2 and DORA.
- Cluster 5 (continuous checking): MHC-10 and MHC-12.
- Cluster 6 (phishing-resistant identity and zero trust): MHC-03, MHC-04 and MHC-13 (identity life cycle of AI agents).
- Cluster 7 (automation and resilience testing): MHC-05, MHC-09 and MHC-11.

MHC-08 (immutable backups) results not primarily from the external framework gap analysis but from the in-depth assessment of A.8.13, A.5.29, A.5.30 and the adjacent resilience controls in Chapters 4 and 5: ISO requires backup and recoverability but neither immutability nor separate access paths for the backup infrastructure — and it is precisely these two properties that determine effectiveness under Mythos conditions. MHC-05 is the

operational basis of the automation obligation (cluster 7) and at the same time the answer to the aggregation blindness of classic detection.

9.4 Compact overview of the MHC catalogue

The following overview summarizes the thirteen MHC and assigns each its primary framework basis and its ISO 27002 anchor. The anchors are consistent with the assessment matrix in Annex A and the individual assessments in Chapters 4 to 6.

MHC	Title	Framework basis	ISO anchor
MHC-01	Post-quantum strategy and cryptographic inventory	C5:2026 CRY-01.01AC	A.8.24
MHC-02	SBOM and build provenance	C5:2026 DEV-13, CRA Annex I, SLSA	A.5.19, A.5.20, A.5.21, A.5.22, A.8.4, A.8.30
MHC-03	Phishing-resistant MFA	NIST SP 800-63B Rev. 4 (AAL2/3), FIDO2	A.5.17, A.6.7, A.8.1, A.8.5, A.8.23
MHC-04	Workload identity and zero trust	NIST SP 800-207, SPIFFE	A.5.15, A.5.16, A.6.7, A.8.2, A.8.20, A.8.21, A.8.22
MHC-05	Behaviour-based detection	C5:2026 OPS-13, MITRE ATT&CK	A.5.3, A.5.7, A.5.14, A.5.15, A.8.1, A.8.3, A.8.4, A.8.7, A.8.12, A.8.15, A.8.16
MHC-06	Containers and confidential computing	C5:2026 OPS-32 to OPS-35	A.5.23, A.8.31
MHC-07	Multi-tenancy isolation	C5:2026 OPS-30/31, PSS-10	A.8.22, A.5.23
MHC-08	Immutable backups	C5:2026 OPS-06 to OPS-09	A.8.13, A.5.29, A.5.30, A.5.33, A.8.14
MHC-09	AI-supported security testing	C5 OPS-25.01AS, CRA, SSDF	A.8.8, A.8.25, A.8.26, A.8.28, A.8.29, A.8.32
MHC-10	Continuous control monitoring	C5 COM-03/04, NIS2 IR 2.2	A.5.18, A.5.35, A.5.36, A.8.9
MHC-11	SOAR and tier 1 automation	C5 SIM-02/03, NIS2 IR 3.5	A.5.5, A.5.24, A.5.25, A.5.26
MHC-12	Threat-led penetration testing	DORA Art. 26/27, TIBER-EU	A.5.35, A.8.29
MHC-13	AI agent governance and harness security	OWASP LLM Top 10 2026, ASI01–ASI10, MITRE ATLAS, ISO 42001 Annex A	A.5.9, A.5.16, A.8.27, supplements MHC-04

The catalogue is the subject of the recommendations for action in Chapter 10. These notes are not a roadmap, because the concrete introduction sequence depends on ISMS maturity phase, the existing control landscape and prioritized risk scenarios.

9.5 Maturity levels per MHC

The thirteen MHC are not binary yes/no decisions. Three maturity levels are defined for each MHC, to be maintained as a target-versus-actual comparison in the Statement of Applicability. An organization may reach different levels for different MHC; the levels are cumulative (level 2 presupposes the requirements of level 1).

Assessment grid: Initial = implementation path begun, ad hoc or documented; Defined = standardized and introduced; Managed = automated, measured and continuously improved.

Binding reading of the levels. The Initial level documents the implementation path begun and does not necessarily represent full fulfilment of the respective MHC control statement. Only from Defined onwards is an

MHC deemed implemented in principle; Managed describes the advanced or continuously effective implementation. An organization therefore cannot present the implementation of an MHC as achieved by reference to the Initial level — in that case the Statement of Applicability must record the implementation status as "partial" with a target level and target date. The control statements of the MHC modules in the Implementation Guide describe the minimum requirements of the Defined level; requirements going beyond that are expressly marked there as advanced hardening (Managed level).

MHC	Initial	Defined	Managed
MHC-01	Cryptographic inventory documented	PQC strategy adopted, hybrid cryptography in pilot	Hybrid cryptography in production, annual review automated
MHC-02	SBOM generation in individual pipelines	SBOM for all own artefacts, vulnerability correlation	SBOM from critical suppliers, SLSA Level 3+, automated provenance verification
MHC-03	FIDO2/passkeys for privileged accounts	FIDO2 mandatory for all privileged access and all externally reachable services	Passwordless authentication organization-wide, SMS MFA switched off
MHC-04	SPIFFE/SPIRE for privileged workloads, mTLS on the 3–5 most critical paths	Workload identity for all production microservices and for production AI agents and agentic workflows falling under MHC-13, ZTNA for privileged access	Workload identity across the estate including dev/test, AI agents with capability-scoped identities
MHC-05	MITRE ATT&CK as the detection framework, coverage not systematically recorded	Coverage of the ATT&CK techniques relevant to the organization's own threat profile documented and prioritized, threat hunting documented	Coverage validated by testing (purple teaming, BAS), adversarially robust ML detection, continuous updating
MHC-06	Container images signed, vulnerability scan before deployment	Runtime protection with admission controllers, documented TEE use cases	Confidential computing in production for the workloads classified as requiring particular protection, remote attestation automated
MHC-07	Logical tenant separation documented	Tenant-specific keys, network and compute isolation	Demonstrable separation through automated tests, regular audits
MHC-08	Backups in place, restore tests annual	3-2-1-1-0 model implemented (at least one immutable or network-separated copy), separate backup administration access, quarterly documented restore tests	A dedicated administration and identity plane for the backup infrastructure, automated integrity and recovery validation, adversarial deletion tests, continuous immutability monitoring
MHC-09	SAST/DAST/SCA in CI	At least one AI-supported test method demonstrably in production for suitable applications, plus a pre-merge block on high findings	Daily vulnerability scans, AI code audit for vibe-coded code, automated KEV hotfix channel
MHC-10	Compliance dashboards maintained manually	Policy as code in CI, real-time compliance checks against baselines, alerting and tracking of deviations	Runtime enforcement, automated incident tickets on deviation

MHC	Initial	Defined	Managed
MHC-11	Response playbooks documented, manual execution	SOAR with partially automated playbooks, MTTC below 60 minutes	Fully automated tier 1 containment, MTTC below 10 minutes, semi-annual multi-incident tabletops; alternatively an MDR service with a contractual SLA
MHC-12	Annual penetration tests with standard scope	Threat-led tests with Mythos scenarios, purple team between cycles	Continuous adversary emulation, BAS platform in production, closure rate of 90% or above within the SLA window
MHC-13	Production AI agents and agentic workflows inventoried (including SaaS, low-code and no-code agents), manual inventory of MCP servers and extensions	Capability-scoped identities and audit trail in production, a named owner and a tested kill switch per agent, a documented threat model per agent use case, an allow list for extensions and MCP servers	A full harness code review process, automated pre-production checks, mean time to revoke below 5 minutes, regular prompt injection penetration tests

The levels are intended as an audit aid: an external auditor sees at a glance where the organization stands per MHC and what improvement path is planned. The target state per MHC should derive from the organization's own risk assessment — not every organization needs level Managed everywhere.

Part IV – Recommendations for action and outlook

Parts I to III set out the Mythos threat landscape, the assessment of all 93 ISO controls, the gaps relative to other frameworks and the catalogue of thirteen Mythos Hardening Controls. Part IV formulates the recommendations derived from them and closes with a reflection on the limits of this work and an outlook.

The recommendations are deliberately not formulated as a work plan or project structure. Every ISMS has its own maturity phase and its own priorities. A concrete introduction path must be developed within the organization on the basis of its own risk assessment, resource position and stakeholder structure.

10 Recommendations for adapting the ISMS

10.1 Scope and limits of the recommendations

The recommendations formulated in this chapter are a prioritized selection, not an exhaustive catalogue. They cover five subject areas: immediate reassessment of existing controls, structural hardening, maturing of the ISMS process, board communication and metrics work. Other subject areas (sector specifics, special regulatory questions, concrete supplier relationships) are deliberately left aside because they cannot be generalized.

Note on non-completeness: The recommendations deliberately contain no time indications. The realistic implementation duration depends on ISMS maturity, team strength, budget position and tooling landscape. The order within the recommendations is not an implementation sequence but a thematic structure.

10.2 Immediate reassessment of existing controls

It follows from the analysis in Part II that the four controls from Chapter 6 (friction-only) and the 37 controls from Chapter 5 (partially degraded) must be reassessed in the risk assessment of a Mythos-exposed ISMS. The reassessment is carried out using the four assessment criteria from Chapter 3.2 and has concrete consequences for the Statement of Applicability:

- For the four controls A.5.25, A.5.36, A.8.8 and A.8.23, it is examined whether they are recorded in the SoA in their existing form as the sole mitigation for Mythos-relevant risks. Where that is the case, the entry is supplemented or corrected with the respective replacement recommendation from Chapter 6.

- For the 37 partially degraded controls, it is examined which of the five supplementation patterns from Chapter 5.3 has already been implemented in each case and where gaps remain. The result is a gap list per control as input to the structural hardening.
- For the 29 robust controls from Chapter 4, consistent implementation is examined. A robust control formally recorded in the SoA but not implemented throughout offers no protection under Mythos conditions. Example: an asset inventory without automated updating quickly becomes out of date, which is why Mythos-specific anomaly detection then has blind spots.

The reassessment does not produce a new ISMS document but an update of the existing SoA with two additions: first, notes on Mythos resilience per control; second, mappings to the thirteen MHC from Chapter 9. The update is a subject of the next management review.

10.3 Structural hardening: adopting the MHC into the SoA

The thirteen MHC concretize Mythos-relevant technical and organizational hardening patterns that ISO/IEC 27002:2022 does not prescribe at this technical depth or with this degree of obligation. Several MHC are expressly hardened forms of existing ISO controls — for example MHC-08 for A.8.13, MHC-03 for A.5.17 and A.8.5, or MHC-10 for A.5.35, A.5.36 and A.8.9; others address requirements that ISO covers only implicitly through the risk perspective. Adoption into the SoA follows a simple pattern:

- Each MHC is added as an additional SoA entry, with a reference to the flanking ISO controls from the compact MHC table (Chapter 9.4).
- The implementation status is documented honestly: already implemented, partially implemented, planned, or not addressed. The category "planned" should carry an explicit justification of why immediate implementation is not possible.
- For each MHC it is examined whether it is applicable at all in the organizational context. MHC-07 (multi-tenancy isolation), for example, is not relevant for organizations without a multi-tenant architecture. Such non-applicability is likewise recorded explicitly.

Adoption creates no compliance obligation, because the MHC have no normative basis in the ISO 27001 sense. It documents the deliberate engagement with the Mythos threat landscape and gives internal and external auditors a traceable basis. For organizations subject to regulation that already requires an MHC content, adoption coincides with an existing obligation. That applies only within the respective scope, however: MHC-12 corresponds to the TLPT requirements for those DORA-regulated financial entities identified for TLPT under Art. 26(8) — not for all DORA addressees. MHC-11 operationalizes the NIS2 and Implementing Regulation requirements on detection, response and recovery; a SOAR system or automated tier 1 containment is not prescribed there across the board.

10.4 Maturing the ISMS process

Several observations from the Mythos threat landscape have direct implications for how the ISMS is operated:

- The risk assessment under ISO 27005 remains the basis but must be updated in shorter cycles. What was classified as a low threat a year ago may be acute today. An annual formal risk assessment with supplementary in-year review of specific risks is appropriate.
- The management review under ISO 27001 Clause 9.3 should take up the Mythos threat landscape as a recurring agenda item. The version notes of this handbook and relevant threat intelligence reports can serve as the basis.
- The internal audit function should be supplemented by continuous audit elements (substantively identical with MHC-10). The results of automated monitoring feed into the reduced but still annual formal audits; they do not replace them.
- The ISMS KPI landscape is supplemented by Mythos-relevant indicators (see Chapter 10.6).
- Establish innovation acceleration governance: a cross-functional body drawn from security, legal, engineering, data protection and, where relevant, procurement, with the explicit mandate to evaluate and put into production defensive AI tools and new security controls on an accelerated basis. Without this body, every Mythos hardening measure runs into the usual approval friction (procurement cycles, legal review, data protection clearance) which, under Mythos conditions, becomes a structural weakness — the NIS2 IR implicitly recognizes this in No. 1.1 as a governance obligation. Recommended cadence:

fortnightly, with a defined time-to-decision target of no more than 30 days for tools holding a valid ISO 27001, SOC 2 or C5 attestation.

- Build a permanent VulnOps function analogous to DevOps, instead of running vulnerability management as a part-time task within the SOC. Responsibility: continuous discovery of zero-days across the entire software estate (own code, third-party software, cloud configuration), automated remediation pipelines, and EPSS- and KEV-based prioritization. Minimum capacity from roughly 1,000 employees onwards, or from the point at which the patch load reaches 100 or more critical findings per month. This function addresses the risk of continuous vulnerability management maturity classified as CRITICAL in Mythos-Ready (CSA, SANS, OWASP).

10.5 Board communication and risk dialogue

The challenge is to present the threat without alarmism, but also without downplaying it. The following points have proved useful:

- The four findings from Chapter 2 lend themselves as a structure for a board presentation: collapse of the patch gap window, time axis compression, fragmentation, capability decoupling. All four are empirically documented.
- The distinction between robust, partially degraded and friction-based controls conveys that security is not failing across the board but that targeted adjustment is required.
- The MHC from Chapter 9 are suitable as a basis for discussing investment decisions. They supply concrete reference points supported by framework references.
- The regulatory embedding should be made transparent. Mythos hardening is, to a considerable extent, a legal obligation rather than merely best practice — depending, however, on sector and addressee group. For CRA-obligated software manufacturers, DORA-regulated financial institutions and NIS2 IR addressees (certain digital providers, see Chapter 3.3), it applies directly; for other NIS2 entities, the respective national transposition governs.
- Address the possible evolution of the standard of care — as a strategic assessment by the author, not as a legal obligation in force. The EU AI Act creates no duty to deploy AI-supported threat hunting, code scanning or vulnerability discovery tools. With the broad availability of defensive AI tools, however, it is to be expected that the yardstick of what counts as an appropriate security measure will shift. Governing bodies may increasingly need to demonstrate to supervisors, insurers and courts which AI tools they deploy for defensive purposes (code audit, threat hunting, vulnerability discovery) — and to justify not using tools that are available. Recommended as a recurring board briefing element: "Which defensive AI tools have we introduced in the reporting period? Which are we evaluating? Which are we deliberately not using — and on what grounds?" This addresses the risk of regulatory and liability exposure classified explicitly as HIGH in Mythos-Ready (CSA, SANS, OWASP).

10.6 Mythos-relevant metrics

The following metrics are a selection, not an exhaustive KPI set. They are oriented to the thirteen MHC and the four assessment criteria from Chapter 3.2.

Metric	Description	Reference
Mean time to containment (MTTC)	Average time from detection to automated or manual containment of an incident. Target for high-confidence alerts below 10 minutes.	MHC-11, MHC-05
Phishing-resistant account coverage	Share of all privileged and externally reachable in-scope accounts with phishing-resistant MFA (account coverage, not logins). Target for privileged accounts: 100% or a documented, time-limited and risk-approved exception; externally reachable typically at least 95%.	MHC-03
SBOM coverage	Share of the organization's own software artefacts with an automatically generated SBOM, and share of critical suppliers with a contractually assured SBOM obligation.	MHC-02

Metric	Description	Reference
Remediation latency for KEV listings	Time between the availability of a robust remediation or effective compensating option and its verified implementation. Patch, workaround, compensating control, isolation or replacement of the affected asset count as effective measures. Target below 24 hours (mean). Definition in Annex G.4.	MHC-09
Pipeline gate coverage and escape rate	Share of production pipelines with security checking and a pre-merge block; supplementary, the share of vulnerabilities that reach production despite pipeline checking, and the share of suitable applications with an AI-supported test method in production.	MHC-09
Restore test success rate	Share of successful quarterly restore tests from immutable backups. Target 100%.	MHC-08
Continuous monitoring coverage	Share of the SoA requirements classified as automatable whose implementation is checked by automated continuous monitoring in addition to periodic audits. The reference base is automation eligibility, not the total number of controls (definition in Annex G.6a/G.6b).	MHC-10
Cryptographic inventory coverage	Share of the cryptographic methods in use that are recorded in the central inventory and carry a priority classification for PQC migration.	MHC-01
TLPT findings closure rate	Share of findings from threat-led penetration tests closed within the defined remediation period — contractual, internal or regulatory.	MHC-12
Container policy conformity	Share of production container workloads that pass signature verification and admission control policies at deployment. Target: 100% in critical or fully in-scope production namespaces, and at least 95% technical coverage organization-wide; remaining exceptions are documented, risk-assessed and approved.	MHC-06
Agent inventory coverage	Share of production AI agents and agentic workflows within the organization's area of responsibility — including SaaS, low-code and no-code agents — that are fully inventoried. Target 100%. Corresponds to the Initial maturity level.	MHC-13
Agent governance coverage	Share of the inventoried production AI agents and agentic workflows for which at least a named owner, a capability-scoped identity, documented permission boundaries and a tamper-resistant execution log are implemented and documented. Target 100%. Corresponds to the Defined maturity level.	MHC-13

The ten primary metrics do not cover all thirteen MHC (supplementary, the automation eligibility reference metric is collected for MHC-10 and governance coverage for MHC-13). For the structural controls MHC-04 (workload identity) and MHC-07 (multi-tenancy isolation), this is a deliberate design decision: implementation evidence is provided through implementation stage gates in project controlling rather than through continuous metrics.

Gate criteria (guide values, per maturity level): MHC-04 — Initial: SPIFFE/SPIRE identities for 80% or more of privileged service workloads; Defined: mTLS mandatory for east-west traffic — 100% of in-scope connections or documented, time-limited and risk-approved exceptions; 95% or above counts expressly only as a transition threshold and not as full fulfilment of Defined; Managed: attestation-supported access policies in production. MHC-07 — one passed, documented tenant separation test per level (Initial: annually, internally; Defined: per major release; Managed: additionally by independent third parties). Gate fulfilment is recorded in project controlling as a yes/no attestation.

10.7 Summary of the recommendations

The five recommendation fields — immediate reassessment, structural hardening, ISMS maturing, board communication and metrics work — can be addressed independently of one another; organizations with limited

resources can prioritize individual fields. What binds them together is the withdrawal of the assumption that attacker effort is a reliable protective quantity: protective effect must instead be anchored in hard barriers — cryptographic, architectural, aggregation-resistant or automation-based.

11 Reflection and outlook

11.1 Limits of this work

The limits of MRIS apply as described in Chapter 1.4 (position in the ISMS stack) and Chapter 3.5 (threat scope): a snapshot rather than a permanent state, no substitute for an organization-specific risk analysis, and no claim to completeness of the MHC catalogue. The analysis relies exclusively on publicly available primary sources; internal incident data from individual organizations are not part of the assessment basis.

11.2 What this handbook does not deliver

MRIS is not a tool for certification or compliance auditing, does not replace vendor, product selection or architecture decisions, is not a legal document and is not a threat intelligence source; it operates at the level of categories and principles.

11.3 Expected developments

Rising agent capability: The pace of progress implies that the share of tactical attack work models can take over autonomously will continue to grow. Classifications as partially degraded may shift towards friction-only.

Build-out of defensive AI: The same capabilities are increasingly available on the defensive side (MHC-05, MHC-09, MHC-11). This will partly close the asymmetry again but brings new challenges: adversarial robustness of the defensive AI, governance of AI-supported security functions, and new training requirements.

Regulatory consolidation: The recent developments — C5:2026, the NIS2 IR, the CRA, DORA — are part of a recognizable consolidation of cybersecurity regulation in Europe. This consolidation will continue, particularly in the areas of AI system governance (AI Act), product security (CRA) and sector-specific regulation.

Maintenance and review commitment. MRIS is reviewed quarterly. Unscheduled reviews are triggered by: (a) general availability or material capability jumps in the Mythos model class; (b) new releases of MITRE ATT&CK or ATLAS; (c) publication of the announced BSI cross-reference table for C5:2026; (d) new regulatory requirements with control relevance (for example NIS2 IR supplements, CRA implementing acts).

Trigger status July 2026: Trigger (a) has been activated by the multi-vendor development and the general availability of the Mythos class and was worked through in version 1.8 (Chapter 2.5). Triggers (b) to (d) were not activated as at 19 July 2026, or remain under observation (ATT&CK v19 is the reference state; the BSI cross-reference table for C5:2026 has been announced but not published).

11.4 Threat horizon

Four lines of development are mapped for coming versions but not yet assessed as control-effective:

- Agent-to-agent protocols (A2A): delegation chains between agents create new trust and attack surfaces (task injection via intermediary agents, transitive entitlements); prospectively affects MHC-04 and MHC-13.
- The MCP ecosystem as a supply chain: tool servers, manifests and registries are becoming the software supply chain of agents; the provenance and signature logic from MHC-02 must be carried over to tool servers.
- Agentic attacks on OT/ICS: transferring the exploit economics to cyber-physical systems has safety consequences and requires an assessment axis of its own beyond Annex A.
- Convergence of PQC and agentic cryptanalysis: automated analysis of protocol and implementation flaws shortens the period in which "harvest now, decrypt later" is economically viable; this increases the pressure on MHC-01.

These points are candidates for a version 2.0; their assessment follows the review cadence from Chapter 11.3.

11.5 Closing remark

MRIS supplies no new theoretical framework, no methodology of its own and no compliance statement, but a working aid: it sets existing frameworks against the current threat landscape, makes gaps visible, and names concrete reference points for closing them. Classic controls do not thereby become obsolete — they are weakened in specific effectiveness dimensions and are to be hardened in targeted fashion, not replaced; the MHC supplement them, they do not take their place.

Part V – Supplementary layer: further frameworks

Part V supplements the primary frameworks from Chapter 3.3 with five further ones that are useful for in-depth Mythos hardening work. They are cited in Parts II and III where they apply; the following table summarizes what each delivers and when it is relevant. Readers concentrating on the primary frameworks can skip Part V.

Framework	What it delivers	When relevant
BSI IT-Grundschutz Compendium	Technically concrete ISMS modules (basic, standard and high requirements), broader than the abstract ISO 27002	Public sector, critical infrastructure, IT-Grundschutz certification; as a second primary reference alongside C5:2026
MITRE ATT&CK and D3FEND	Operational vocabulary for detection (ATT&CK) and defence (D3FEND); the kill chain structure represents fragmented attacks coherently	Concretization of SIEM and detection (C5 OPS-13, NIS2 IR 3.2) and DORA TLPT scenarios; the basis of MHC-05 and MHC-12
ISO/IEC 42001	Management system for AI with Annex A controls (risk management, data quality, life cycle, transparency, robustness)	Governance of defensive AI components and assurance for the organization's own models; particularly for SaaS providers with a high AI share
CIS Controls v8	18 prioritized control categories with implementation groups (IG1–IG3) and concrete safeguards	Pragmatic prioritization aid for organizations at an early maturity stage (IG1 as the entry point)
OWASP ASVS and SAMM	ASVS: an operational verification catalogue for secure software (L1–L3); SAMM: a maturity model for the secure SDLC organization	A basis for SAST/DAST that integrates directly into CI (MHC-09); a benchmark for software manufacturers with CRA obligations

None of these frameworks replaces the primary basis from Chapter 3.3; the selection depends on the sector, size and strategic orientation of the organization.

C5:2026 to IT-Grundschutz correspondence (orientation aid). For mapping the C5 references named in Parts II and III onto typical IT-Grundschutz modules:

C5:2026 domain	IT-Grundschutz modules (typical)
C5 OIS (Organisation of Information Security)	ISMS layer, ORP
C5 HR (Personnel)	ORP.2
C5 AM (Asset Management)	CON.10, OPS.1.1.3
C5 PS (Physical Security)	INF.1 to INF.7
C5 OPS (Operations)	OPS.1.1.1 to OPS.1.2.5, DER
C5 IAM (Identity and Access Management)	ORP.4
C5 CRY (Cryptography and Key Management)	CON.1, CON.6
C5 COS (Communication Security)	NET.1, NET.3
C5 DEV (Development)	CON.8
C5 SSO (Service Providers and Suppliers)	OPS.2, CON.14
C5 SIM (Security Incident Management)	DER.2.1 to DER.2.3

C5:2026 domain	IT-Grundschutz modules (typical)
C5 BCM (Business Continuity Management)	DER.4
C5 COM (Compliance)	ISMS layer
C5 PSS (Product Safety and Security)	APP modules depending on the system

The list is an orientation aid, not a binding one-to-one mapping.

Annexes – working materials

The annexes bundle the content developed in Parts I to V into a form directly usable for practical ISMS work. Annex A supplies the complete assessment matrix of all 93 controls with category classification and MHC mapping. Annex B supplements this matrix with a cross-reference between the primary frameworks. Annex C provides the MHC catalogue as a standalone worksheet. Annex D lists indicator and monitoring sources. Annex E supplies the RACI model for MHC implementation. Annex F describes the risk assessment bridge under ISO/IEC 27005. Annex G standardizes the KPI definitions from Chapter 10.6 for audit-ready measurement that is reproducible across organizations.

Annex A – Assessment matrix of all 93 controls

The following table lists all 93 controls of ISO/IEC 27002:2022 in ISO numbering. For each control it states: ID, title (abbreviated), Mythos category (S = Robust, T = Partially degraded, R = Friction-only, N = Not affected) and — where applicable — the flanking Mythos Hardening Controls from Chapter 9.

The MHC mappings in this matrix are consistent with the compact overview in Chapter 9.4.

ID	Control title	Cat.	Flanking MHC
A.5.1	Policies for information security	N	-
A.5.2	Information security roles and responsibilities	N	-
A.5.3	Segregation of duties	T	MHC-05
A.5.4	Management responsibilities	N	-
A.5.5	Contact with authorities	T	MHC-11
A.5.6	Contact with special interest groups	N	-
A.5.7	Threat intelligence	T	MHC-05
A.5.8	Information security in project management	N	-
A.5.9	Inventory of information and other associated assets	S	MHC-13
A.5.10	Acceptable use of information and other associated assets	N	-
A.5.11	Return of assets	N	-
A.5.12	Classification of information	S	-
A.5.13	Labelling of information	N	-
A.5.14	Information transfer	T	MHC-05
A.5.15	Access control	T	MHC-04, MHC-05
A.5.16	Identity management	S	MHC-04, MHC-13
A.5.17	Authentication information	T	MHC-03
A.5.18	Access rights	T	MHC-10
A.5.19	Information security in supplier relationships	T	MHC-02
A.5.20	Addressing information security within supplier agreements	T	MHC-02

ID	Control title	Cat.	Flanking MHC
A.5.21	Managing information security in the ICT supply chain	T	MHC-02
A.5.22	Monitoring, review and change management of supplier services	T	MHC-02
A.5.23	Information security for use of cloud services	T	MHC-06, MHC-07
A.5.24	Information security incident management planning and preparation	T	MHC-11
A.5.25	Assessment and decision on information security events	R	MHC-11
A.5.26	Response to information security incidents	T	MHC-11
A.5.27	Learning from information security incidents	S	-
A.5.28	Collection of evidence	S	-
A.5.29	Information security during disruption	T	MHC-08
A.5.30	ICT readiness for business continuity	S	MHC-08
A.5.31	Legal, statutory, regulatory and contractual requirements	N	-
A.5.32	Intellectual property rights	N	-
A.5.33	Protection of records	T	MHC-08
A.5.34	Privacy and protection of PII	T	-
A.5.35	Independent review of information security	T	MHC-10, MHC-12
A.5.36	Compliance with policies, rules and standards for information security	R	MHC-10
A.5.37	Documented operating procedures	N	-
A.6.1	Screening	N	-
A.6.2	Terms and conditions of employment	N	-
A.6.3	Information security awareness, education and training	T	-
A.6.4	Disciplinary process	N	-
A.6.5	Responsibilities after termination or change of employment	S	-
A.6.6	Confidentiality or non-disclosure agreements	N	-
A.6.7	Remote working	T	MHC-03, MHC-04
A.6.8	Information security event reporting	T	-
A.7.1	Physical security perimeters	S	-
A.7.2	Physical entry	S	-
A.7.3	Securing offices, rooms and facilities	S	-
A.7.4	Physical security monitoring	S	-
A.7.5	Protecting against physical and environmental threats	N	-
A.7.6	Working in secure areas	S	-
A.7.7	Clear desk and clear screen	N	-
A.7.8	Equipment siting and protection	N	-
A.7.9	Security of assets off-premises	T	-
A.7.10	Storage media	S	-
A.7.11	Supporting utilities	N	-
A.7.12	Cabling security	N	-

ID	Control title	Cat.	Flanking MHC
A.7.13	Equipment maintenance	N	-
A.7.14	Secure disposal or re-use of equipment	S	-
A.8.1	User endpoint devices	T	MHC-03, MHC-05
A.8.2	Privileged access rights	S	MHC-04
A.8.3	Information access restriction	T	MHC-05
A.8.4	Access to source code	T	MHC-02, MHC-05
A.8.5	Secure authentication	T	MHC-03
A.8.6	Capacity management	N	-
A.8.7	Protection against malware	T	MHC-05
A.8.8	Management of technical vulnerabilities	R	MHC-09
A.8.9	Configuration management	S	MHC-10
A.8.10	Information deletion	S	-
A.8.11	Data masking	S	-
A.8.12	Data leakage prevention	T	MHC-05
A.8.13	Information backup	S	MHC-08
A.8.14	Redundancy of information processing facilities	S	MHC-08
A.8.15	Logging	S	MHC-05
A.8.16	Monitoring activities	T	MHC-05
A.8.17	Clock synchronization	S	-
A.8.18	Use of privileged utility programs	S	-
A.8.19	Installation of software on operational systems	S	-
A.8.20	Networks security	T	MHC-04
A.8.21	Security of network services	T	MHC-04
A.8.22	Segregation of networks	T	MHC-04, MHC-07
A.8.23	Web filtering	R	MHC-03
A.8.24	Use of cryptography	S	MHC-01
A.8.25	Secure development life cycle	T	MHC-09
A.8.26	Application security requirements	T	MHC-09
A.8.27	Secure system architecture and engineering principles	S	MHC-13
A.8.28	Secure coding	T	MHC-09
A.8.29	Security testing in development and acceptance	S	MHC-09, MHC-12
A.8.30	Outsourced development	T	MHC-02
A.8.31	Separation of development, test and production environments	S	MHC-06
A.8.32	Change management	T	MHC-09
A.8.33	Test information	S	-
A.8.34	Protection of information systems during audit testing	N	-

Annex B – Framework mapping for core topics

The following table summarizes, for twelve central subject areas of information security, how the primary frameworks address each. It is an orientation aid, not a complete mapping.

Subject area	ISO 27002	C5:2026	NIST CSF	DORA	CRA	NIS2 IR
Asset management	A.5.9, A.5.12	AM-02, AM-03, AM-04, AM-09	ID.AM	Art. 8	Annex I	No. 12
Identity and access	A.5.15, A.5.16, A.8.2	IAM-01, IAM-06, IAM-08	PR.AA	Art. 9	Annex I	No. 11
Cryptography	A.8.24	CRY-01 to CRY-19	PR.DS	Art. 9(4)(e)	Annex I II	No. 9
Logging and monitoring	A.8.15, A.8.16	OPS-10 to OPS-17	DE.CM	Art. 10	-	No. 3.2
Incident response	A.5.24, A.5.26	SIM-01 to SIM-06	RS.	Art. 17, 19	-	No. 3.1 to 3.6
Business continuity	A.5.29, A.5.30, A.8.13	BCM-01 to BCM-04, OPS-06 to OPS-09	RC.	Art. 11, 12	-	No. 4
Supply chain	A.5.19 to A.5.23	SSO-01 to SSO-08, DEV-13	GV.SC	Art. 28-30	Annex I II	No. 5
Development	A.8.25 to A.8.29	DEV-01 to DEV-15	PR.IR	Art. 8(3)	Annex I II	No. 6
Physical security	A.7.1 to A.7.14	PS-01 to PS-08	PR.IR	Art. 9(4)	-	No. 13
Risk management	A.5.1 to A.5.8	OIS-07, OIS-08, OIS-09	GV.RM, ID.RA	Art. 6, 7	-	No. 2
Compliance	A.5.31 to A.5.36	COM-01 to COM-04	GV.OC	Art. 6(5)	Art. 13, 14	No. 2.2, 2.3
Cloud use	A.5.23	GC-01 to GC-06, OPS-30 to OPS-35	GV.SC-07	Art. 28	-	No. 5

Empty cells (-) mean that the framework states no specific requirement of its own for the subject area concerned. The NIST CSF column names the respective categories from CSF 2.0; concrete implementation is via SP 800-53.

Note on the scope of the NIS2 IR: The numbers cited in the "NIS2 IR" column are taken from Implementing Regulation (EU) 2024/2690, which under its Article 1 applies exclusively to certain digital providers (see Chapter 3.3). For other NIS2 addressees, the notification deadlines from NIS2 Directive Art. 23(4) in conjunction with the respective national transposition govern.

Annex C – MHC catalogue as a standalone worksheet

The following overview provides the Mythos Hardening Control catalogue from Chapter 9 as a standalone worksheet. Annex C can be extracted from the overall document and used as an independent worksheet for the Statement of Applicability. The "Status" column is deliberately left blank and can be completed organization-specifically (for example already implemented, partially implemented, planned, not applicable).

MHC	Title	Framework basis	ISO anchor	Status
MHC-01	Post-quantum strategy and cryptographic inventory	C5:2026 CRY-01.01AC	A.8.24	
MHC-02	SBOM and build provenance	CRA, C5:2026 DEV-13, SLSA	A.5.19, A.5.20, A.5.21, A.5.22, A.8.4, A.8.30	
MHC-03	Phishing-resistant MFA	NIST 800-63B, FIDO2	A.5.17, A.6.7, A.8.1, A.8.5, A.8.23	
MHC-04	Workload identity and zero trust	NIST 800-207, SPIFFE	A.5.15, A.5.16, A.6.7, A.8.2, A.8.20, A.8.21, A.8.22	
MHC-05	Behaviour-based detection	MITRE ATT&CK, C5 OPS-13	A.5.3, A.5.7, A.5.14, A.5.15, A.8.1, A.8.3, A.8.4, A.8.7, A.8.12, A.8.15, A.8.16	
MHC-06	Containers and confidential computing	C5:2026 OPS-32 to OPS-35	A.5.23, A.8.31	
MHC-07	Multi-tenancy isolation	C5:2026 OPS-30/31	A.5.23, A.8.22	
MHC-08	Immutable backups	C5:2026 OPS-06 to OPS-09	A.5.29, A.5.30, A.5.33, A.8.13, A.8.14	
MHC-09	AI-supported security testing	C5 OPS-25.01AS, CRA, SSDF	A.8.8, A.8.25, A.8.26, A.8.28, A.8.29, A.8.32	
MHC-10	Continuous control monitoring	C5 COM-03/04, NIS2 IR 2.2	A.5.18, A.5.35, A.5.36, A.8.9	
MHC-11	SOAR and tier 1 automation	C5 SIM-02/03, NIS2 IR 3.5	A.5.5, A.5.24, A.5.25, A.5.26	
MHC-12	Threat-led penetration testing	DORA Art. 26/27, TIBER-EU	A.5.35, A.8.29	
MHC-13	AI agent governance and harness security	OWASP LLM Top 10 2026, ASI01–ASI10, MITRE ATLAS, ISO 42001 Annex A	A.5.9, A.5.16, A.8.27, supplements MHC-04	

Annex D – Indicators and monitoring sources

The following table lists public indicator sources useful for the ongoing observation of the Mythos threat landscape and for detecting Mythos-relevant events. The list is a selection, not a complete catalogue.

Source	Description	URL / access
CVE – Common Vulnerabilities and Exposures	MITRE's central vulnerability library, published via the NVD.	cve.org / nvd.nist.gov
EUVD – European Union Vulnerability Database	ENISA-operated vulnerability database for the EU.	euvd.enisa.europa.eu
KEV – Known Exploited Vulnerabilities Catalog	CISA list of actively exploited vulnerabilities, for patch prioritization.	cisa.gov/known-exploited-vulnerabilities-catalog

Source	Description	URL / access
EPSS – Exploit Prediction Scoring System	Probability that a CVE will be exploited within 30 days.	first.org/epss
OSV – Open Source Vulnerabilities	Vulnerability feed for open source dependencies.	osv.dev
MITRE ATT&CK	Taxonomy of adversary tactics and techniques.	attack.mitre.org
CISA advisories	US agency notifications on acute threats.	cisa.gov/news-events/cybersecurity-advisories
BSI warnings and situation reports	German security warnings and the annual situation report.	bsi.bund.de
CERT-EU advisories	Warnings for EU institutions with broad relevance.	cert.europa.eu
Anthropic threat intelligence	Publications on AI-supported attack campaigns.	anthropic.com/news
GitHub security advisories	Central point for open source vulnerability reports.	github.com/advisories
FIRST.org	Forum of Incident Response and Security Teams.	first.org

Annex E – RACI model for MHC implementation

The following table assigns to each of the thirteen Mythos Hardening Controls the typical responsibilities in a mid-sized company. The assignment is to be understood as a proposal and must be adjusted organization-specifically — particularly in organizations without a dedicated AI governance lead or without a board reporting line for cybersecurity.

The RACI below is an example assignment and not a normative governance requirement. It shows an allocation common in mid-sized companies. Depending on the form of organization, accountability for individual MHC may rest with other roles — for example with the CTO for workload architecture (MHC-04), with the Head of Engineering for secure development (MHC-09), or with a product security function for AI agent governance (MHC-13).

Assessment grid: R = Responsible (performs the work), A = Accountable (answerable, one person per row), C = Consulted (asked beforehand), I = Informed (informed afterwards).

Role columns: CISO = Chief Information Security Officer, ISMS = ISMS manager or officer, AI-Gov = AI governance lead, IT = IT operations and platform engineering, Dev = development, SOC = security operations centre and IR function, Legal = legal and compliance function, Board = executive management or supervisory body.

MHC	Title (short)	CISO	ISMS	AI-Gov	IT	Dev	SOC	Legal	Board
MHC-01	Post-quantum strategy / crypto inventory	A	R	C	R	C	I	C	I
MHC-02	SBOM and build provenance	A	C	I	C	R	I	C	I
MHC-03	Phishing-resistant MFA	A	C	I	R	C	I	I	I
MHC-04	Workload identity / zero trust	A	C	C	R	R	C	I	I
MHC-05	Behaviour-based detection	A	I	C	C	I	R	I	I
MHC-06	Containers / confidential computing	A	I	C	R	R	C	I	I
MHC-07	Multi-tenancy isolation	A	C	I	R	R	C	C	I
MHC-08	Immutable backups	A	C	I	R	I	C	I	I
MHC-09	AI-supported security testing	A	I	C	C	R	C	I	I

MHC	Title (short)	CISO	ISMS	AI-Gov	IT	Dev	SOC	Legal	Board
MHC-10	Continuous control monitoring	A	R	C	R	C	C	I	I
MHC-11	SOAR / tier 1 automation	A	C	C	C	I	R	C	I
MHC-12	Threat-led penetration testing	A	C	C	C	C	R	C	I
MHC-13	AI agent governance	A	C	R	C	R	C	C	C

Notes on application

- Exactly one A per row: accountability cannot be delegated and is not split. In the template, the CISO is recorded as accountable throughout; in larger organizations, accountability for individual MHC can be delegated to domain owners (for example the CTO for MHC-04, the Head of Engineering for MHC-09).
- Several R entries are possible where execution spans several domains. Example MHC-04: IT sets up SPIFFE/SPIRE while Dev migrates the service identities — both are responsible.
- MHC-13 is the only row in which the board is recorded as consulted. Rationale: the governance of AI agents entering production concerns risk acceptance questions above the CISO mandate threshold (standard of care, see Chapter 10.5).
- External service providers (MDR, MSSP, TLPT provider, GRC consultancy) do not appear as a separate column. They are integrated for control purposes through the respective responsible role, not through a RACI column of their own.

Annex F – Risk assessment bridge under ISO/IEC 27005

This annex describes how an MRIS assessment feeds into the risk management process under ISO/IEC 27005:2022. It closes the gap made explicit in Chapter 1.4 — "no risk management process of its own" — through a concrete interface definition between the MRIS effectiveness layer and the ISMS risk process.

F.1 Input: MRIS assessment result per control

Applying MRIS to the Statement of Applicability yields four data points per control:

- Mythos category (Robust / Partially degraded / Friction-only / Not affected) — from Chapters 4 to 7.
- Weaknesses identified along the four assessment criteria from Chapter 3.2 (attacker patience, time compression, capability decoupling, aggregation resistance).
- Flanking Mythos Hardening Controls from Annex A.
- Current maturity level per flanking MHC under Chapter 9.5 (Initial / Defined / Managed).

F.2 Processing steps in the risk process

Step 1: Risk reassessment (ISO/IEC 27005:2022 Clause 7.3)

For every control with the Mythos category "Partially degraded" or "Friction-only", the associated risk position in the risk register is reassessed. Specifically:

The methodological starting point is the previously assumed effectiveness of the controls concerned (control effectiveness), not directly the inherent likelihood of occurrence. The residual likelihood, or residual risk, is then determined afresh using the organization-specific risk methodology.

MRIS finding	Effect on the risk assessment
Robust	No automatic adjustment.
Partially degraded	Reassess the assumed control effectiveness; redetermine residual risk using the organization's own methodology.
Friction-only	Do not treat as the sole effective mitigation; identify supplementary measures or accept the residual risk explicitly.
Not affected	No Mythos-specific adjustment.

The consequence axis remains unchanged unless the damage scenario itself (for example mass exfiltration through agentic tools) suggests a higher consequence level.

No blanket level increase ("+1", "+2") is deliberately prescribed: organizations work with different models (3×3, 4×4, 5×5, 10×10, quantitative under FAIR), in which a fixed number of steps would carry no transferable meaning. Double counting must also be avoided where several degraded controls address the same risk — in that case the effectiveness of the mitigation bundle is assessed jointly, not control by control.

Step 2: Treatment options (ISO/IEC 27005:2022 Clause 8.1)

For each reassessed risk, the four ISO 27005 treatment options are examined:

- Modification: adopting the flanking MHC into the SoA (see Chapter 10.3) — the standard answer for most Mythos risks.
- Retention: acceptance of elevated residual risks by the risk owner body where treatment costs are disproportionate. The rationale must be documented in writing.
- Avoidance: decommissioning the function or service concerned. Rarely feasible in Mythos contexts, but it can be relevant for legacy systems without a patch path.
- Sharing: insurance solutions (cyber risk insurance), outsourcing to specialized providers with their own Mythos hardening.

Step 3: Residual risk assessment and acceptance (ISO/IEC 27005:2022 Clause 8.6)

After the treatment decision, the residual risk is assessed and formally accepted by the risk owner. Where residual risks above the acceptance threshold are accepted, a time-limited acceptance decision with a re-review date (typically six months) must be taken.

F.3 Example application

Example: risk R-042 "mass exfiltration from code repositories through a compromised developer account". MRIS assessment of the relevant controls:

Control	MRIS finding	Consequence
A.5.17 Authentication information	Partially degraded (aggregation blindness, credential compromisability). Flanked by MHC-03.	Assumed effectiveness against credential phishing significantly reduced; without MHC-03 (phishing-resistant MFA, FIDO2) at least at maturity level Defined, the control works only to a limited extent.
A.8.4 Access to source code	Partially degraded (aggregation blindness). Flanked by MHC-02 and MHC-05.	Assumed effectiveness against fragmented mass access reduced; it is restored only in combination with repository anomaly detection (MHC-05) and an SBOM obligation (MHC-02) at maturity level Managed.
A.8.16 Monitoring activities	Partially degraded (aggregation blindness). Flanked by MHC-05.	Assumed detection effectiveness reduced; fragmented mass clones are recognized only with behaviour-based detection at maturity level Managed.

Resulting treatment decision: after reassessing control effectiveness, the residual risk lies above the acceptance threshold. Adoption of MHC-02, MHC-03 and MHC-05 into the SoA with target maturity level Managed within twelve months. Until that is achieved, a temporary residual risk is accepted by the risk owner with a re-review after six months.

F.4 Output into the ISMS

The risk assessment bridge produces four artefacts for the ISMS:

- Updated entries in the risk register with an MRIS reference per risk position.

- An updated Statement of Applicability with additional MHC entries and a reference to the respective Mythos category of the flanked ISO controls.
- A treatment plan with maturity target values per MHC and defined milestones.
- Residual risk acceptance decisions with re-review dates.

These four artefacts close the bridge between the MRIS assessment layer and the ISO 27005 risk process. They replace neither the risk register nor the ISMS process but supply the risk process with the inputs necessary under Mythos conditions.

Annex G – KPI definitions and measurement standardization

This annex standardizes the ten primary Mythos metrics from Chapter 10.6 — for MHC-10 the automation eligibility reference metric (G.6a) and for MHC-13 the supplementary governance metric (G.10b) are additionally collected — so that they can be measured reproducibly across organizations and over time. Without this standardization, two organizations can operate the same KPI with an identical target value and still produce results that are not comparable.

KPI	Definition, measurement points and data source	Target value	Frequency · responsibility	Reference
G.1 Mean time to containment (MTTC)	Time from the first high-confidence alert (SIEM/EDR, confidence of 80% or above, or a high-severity flag) to the confirmed first containment action; mean and median. Source: SOAR audit log (primary), SIEM, IR ticket system.	Mean below 10 min; median below 5 min	Monthly · SOC lead / CISO	MHC-11, MHC-05
G.2 Phishing-resistant account coverage	Number of in-scope accounts secured with phishing-resistant methods divided by the number of all in-scope accounts, as a percentage; privileged and externally reachable accounts with AAL2+ phishing-resistant MFA (FIDO2/WebAuthn/passkey/CTAP2) are counted. Source: IdP (Entra ID, Okta, Ping) plus the CMDB of privileged accounts.	Privileged 100% or a documented, time-limited and risk-approved exception; externally reachable 95% or above; new SMS enrolments 0%	Monthly · IAM lead / ISMS	MHC-03
G.3 SBOM coverage	(a) The organization's own release artefacts with an automatic SBOM in CI; (b) critical suppliers (tier 1/2) with a contractual SBOM obligation. Source: CI reports (Snyk, Dependency-Track); contract database.	(a) 95% or above (12 M); (b) 80% or above (18 M)	Quarterly · Head of Eng. / CISO	MHC-02
G.4 Remediation latency for KEV listings	Time from the availability of a robust remediation to the demonstrated remediation or effective compensation on all affected hosts; mean and 95th percentile. Patch, workaround, compensating control, disconnection from the network or replacement count as remediation. Source: KEV catalogue and vendor information (start), patch and vulnerability management (stop).	Mean below 24 h; P95 below 72 h (exposed)	Per entry / monthly · VulnOps lead / CISO	MHC-09, A.8.8
G.5 Restore test success rate	Share of successful restores from immutable backups within RTO with error-free data. Source: backup audit log plus IT operations	100%; every failure is a control failure and is handled	Quarterly · IT operations / ISMS	MHC-08, A.8.13

KPI	Definition, measurement points and data source	Target value	Frequency · responsibility	Reference
	test records. Minimum frequency quarterly per critical system.	as a finding with root cause analysis and a deadline, with risk-based escalation		
G.6a Automation eligibility	Number of requirements classified as automatically verifiable divided by the number of assessed SoA requirements, as a percentage. The classification is documented and reviewed annually. Source: GRC platform.	No target value; the reference base for G.6b	Annually · ISMS manager	MHC-10
G.6b Continuous monitoring coverage	Number of requirements actively monitored in automated form divided by the number of requirements classified as automatable, as a percentage (checking at least daily, with alerting on deviation, in addition to periodic audits). Source: OPA/Sentinel repo, GRC platform.	60% or above (12 M); 80% or above (24 M)	Quarterly · ISMS manager	MHC-10
G.7 Cryptographic inventory coverage	Share of production cryptographic methods in the crypto inventory carrying a PQC priority level; minimum attributes algorithm, key length, library and version, deployment context. Source: CMDB crypto table, crypto discovery scanner.	80% or above (12 M); 95% or above (24 M)	Half-yearly · architecture lead / CISO	MHC-01
G.8 TLPT findings closure rate	Share of TLPT findings confirmed closed by re-test within the SLA window; standard SLA critical 30 / high 60 / medium 90 / low 180 days; the standard SLA periods are MRIS guide values and apply so far as no stricter internal, contractual or regulatory deadline takes precedence. Source: TLPT report, findings tracker, re-test.	90% or above in aggregate	Quarterly · CISO / risk owner	MHC-12
G.9 Container policy conformity	Share of production containers with a valid image signature, a passed scan and no open-ended privilege exception. Source: admission controller (Kyverno, OPA Gatekeeper), Sigstore/cosign, registry scans.	Critical namespaces 100%; overall 95% or above; unreviewed exceptions 0	Monthly · platform lead / ISMS	MHC-06
G.10a Agent inventory coverage	Number of fully inventoried production AI agents and agentic workflows within the organization's area of responsibility (self-developed, SaaS, low-code and no-code agents) divided by their total number, as a percentage; shadow AI findings (A.5.9) count towards the denominator. Source: agent inventory, discovery scans, CMDB.	100% (agents not inventoried are deactivated)	Monthly · AI governance lead / ISMS	MHC-13
G.10b Agent governance coverage	Share of the inventoried agents and workflows with a named owner, a capability-scoped identity, documented permission boundaries and a tamper-	100% (agents not inventoried are deactivated)	Monthly · AI governance lead / ISMS	MHC-13

KPI	Definition, measurement points and data source	Target value	Frequency · responsibility	Reference
	resistant log (A.8.15). Source: agent inventory, IdP and workload identities, execution logs.			

Standardization notes. The KPIs are reproducible across organizations provided the data sources named are available and the scope definitions are applied consistently. For comparability over time, KPI definitions are versioned on material change and calculation changes are marked as a break in the time series. For industry benchmarks or ISAC sharing, the definition version is transmitted with the value and sector-specific adjustments (for example higher thresholds for DORA-regulated institutions) are declared.

Glossary

The following glossary explains the central terms of this handbook in compact form. Terms are sorted alphabetically.

Term	Explanation
Aggregation resistance	The property of a control of taking effect even where an attack is decomposed into many individually legitimate-looking micro-steps. One of the four assessment criteria (Chapter 3.2).
AI – artificial intelligence	Rendered "KI" in German usage. This work uses "AI" throughout, to align with the MHC designations and the frameworks cited.
Attacker patience	Assessment criterion: does the protective effect rest on an attack being too laborious for an adversary? Structurally weakened under Mythos.
Capability decoupling	Assessment criterion: does the likelihood assumption depend on a historical correlation between actor and capability that Mythos has dissolved?
Crypto agility	The ability of a system to change cryptographic mechanisms at short notice, for example on compromise of an algorithm or during PQC migration.
FIDO2 / WebAuthn	FIDO Alliance authentication standards for phishing-resistant MFA. They anchor authentication cryptographically to the origin domain.
Friction-only	Category for controls whose classic or commonly encountered implementation form no longer provides sufficient standalone protective effect against Mythos attackers, because it rests on limited attacker capacity or on humanly manageable reaction times. Four controls in Chapter 6.
MHC – Mythos Hardening Control	The category, introduced in this handbook, of a supplementary control that closes a Mythos-relevant effectiveness gap which ISO 27002:2022 covers only implicitly. Thirteen MHC in Chapter 9.
Mythos	In this handbook a class term: agentic frontier models with autonomous vulnerability discovery and exploit chaining, irrespective of vendor (including Claude Mythos/Fable 5, OpenAI Daybreak with GPT-5.5-Cyber, Microsoft MDASH). Originally the designation of the reference model Claude Mythos Preview (Chapter 2.5).
Not affected	Category for controls that are neutral with respect to Mythos-specific threats. 23 controls in Chapter 7.
Partially degraded	Category for controls whose protective effect is preserved but whose originally assumed strength falls under Mythos. 37 controls in Chapter 5.
PQC – post-quantum cryptography	Cryptographic methods resistant to attackers with quantum computing capability. NIST standardization (ML-KEM, ML-DSA, SLH-DSA).
Robust	Category for controls whose protective effect derives from a hard barrier that persists under Mythos. 29 controls in Chapter 4.
SBOM – software bill of materials	A machine-readable list of all software components of an artefact. Formats: SPDX, CycloneDX. Required under CRA Annex I Part II.
SLSA – supply-chain levels for software artifacts	A framework for build provenance with four maturity levels. Level 3+ recommended for critical build paths.
SOAR	Security orchestration, automation and response. The platform category for automated incident response workflows.
SPIFFE / SPIRE	The standard and reference implementation for workload identities. The basis for zero trust architectures in cloud-native environments.
Store-now-decrypt-later	A threat model in which encrypted data exfiltrated today is retained in order to decrypt it once quantum computing becomes available.

Term	Explanation
TEE – trusted execution environment	A hardware-supported, isolated execution area for confidential data processing. The basis for confidential computing.
Time compression	Assessment criterion: does the control presuppose human reaction time in a chain in which the attacker operates autonomously?
TLPT – threat-led penetration testing	Penetration testing based on real threat scenarios, anchored in DORA Art. 26/27 and the TIBER-EU framework.
UEBA – user and entity behavior analytics	A detection category based on behavioural baselines and ML-supported anomaly detection.
Zero trust	The architectural principle that no actor, device or network is trusted implicitly. Every transaction is verified explicitly. NIST SP 800-207.

List of references

Standards and regulatory sources

- ISO/IEC 27001:2022 – Information security management systems – Requirements.
- ISO/IEC 27002:2022 – Information security controls.
- ISO/IEC 27005:2022 – Information security risk management.
- ISO/IEC 42001:2023 – Artificial intelligence management system.
- ISO/IEC 20889 – Privacy enhancing data de-identification terminology.
- ISO/IEC 27017 – Code of practice for information security controls for cloud services.
- ISO/IEC 27037 – Guidelines for identification, collection, acquisition and preservation of digital evidence.
- ISO/IEC 27701 – Privacy information management system.
- ISO 22301 – Business continuity management systems.
- BSI Cloud Computing Compliance Criteria Catalogue (C5:2026), March 2026.
- BSI IT-Grundschutz Compendium (current edition).
- BSI TR-02102 Cryptographic Mechanisms: Recommendations and Key Lengths.
- BSI TR-03183 Cyber Resilience Requirements for Manufacturers and Products, Part 1 (General Requirements) and Part 2 (SBOM). bsi.bund.de (accessed 19 July 2026).
- Directive (EU) 2022/2555 (NIS2 Directive) on measures for a high common level of cybersecurity.
- Commission Implementing Regulation (EU) 2024/2690 of 17 October 2024 (NIS2 IR).
- Regulation (EU) 2022/2554 (DORA) on digital operational resilience for the financial sector.
- Regulation on horizontal cybersecurity requirements (EU Cyber Resilience Act, CRA).
- DSGVO – Verordnung (EU) 2016/679.
- AI Act – Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence.

NIST publications

- NIST Cybersecurity Framework 2.0 (CSF 2.0), Februar 2024.
- NIST SP 800-34 Rev. 1 – Contingency Planning Guide for Federal Information Systems.
- NIST SP 800-40 Rev. 4 – Guide to Enterprise Patch Management Planning.
- NIST SP 800-46 Rev. 2 – Guide to Enterprise Telework, Remote Access, and BYOD Security.
- NIST SP 800-53 Rev. 5 – Security and Privacy Controls for Information Systems and Organizations.
- NIST SP 800-60 – Guide for Mapping Types of Information and Information Systems to Security Categories.

- NIST SP 800-61 Revision 3 – Incident Response Recommendations and Considerations for Cybersecurity Risk Management.
- NIST SP 800-63B-4 (Juli 2025) – Digital Identity Guidelines: Authentication and Authenticator Management.
- NIST SP 800-86 – Guide to Integrating Forensic Techniques into Incident Response.
- NIST SP 800-88 Rev. 1 – Guidelines for Media Sanitization.
- NIST SP 800-92 – Guide to Computer Security Log Management.
- NIST SP 800-94 – Guide to Intrusion Detection and Prevention Systems (2007; draft Revision 1 withdrawn).
- NIST SP 800-124 Rev. 2 – Guidelines for Managing the Security of Mobile Devices.
- NIST SP 800-150 – Guide to Cyber Threat Information Sharing.
- NIST SP 800-160 Vol. 1 – Systems Security Engineering.
- NIST SP 800-167 – Guide to Application Whitelisting.
- NIST SP 800-175B – Guideline for Using Cryptographic Standards.
- NIST SP 800-188 – De-Identifying Government Datasets.
- NIST SP 800-190 – Application Container Security Guide.
- NIST SP 800-207 – Zero Trust Architecture.
- NIST SP 800-218 – Secure Software Development Framework (SSDF) v1.1.
- NIST SP 1800-38 (Draft) – Migration to Post-Quantum Cryptography.
- FIPS 140-3 – Security Requirements for Cryptographic Modules.
- NIST NCCoE Project – Agent Identity Framework (2026; announced, primary source not publicly verified at the editorial deadline).

Threat intelligence and industry publications

- Anthropic – Detecting and countering misuse of AI: August 2025 Threat Intelligence Report.
- Anthropic – Disrupting the first reported AI-orchestrated cyber espionage campaign (GTG-1002), November 2025.
- Anthropic – Preparing your security program for AI-accelerated offense (Glasswing-Post), April 2026. claude.com/blog/preparing-your-security-program-for-ai-accelerated-offense (accessed 19 July 2026).
- Anthropic – Claude Mythos Preview, April 2026. anthropic.com/news (accessed 19 July 2026).
- OpenAI – Daybreak: AI-supported vulnerability discovery and patch validation (GPT-5.5, Codex Security), 11 May 2026. Secondary analyses: thehackernews.com, futurumgroup.com (accessed 19 July 2026).
- Microsoft – MDASH (Multi-Model Agentic Scanning Harness), 12 May 2026: 16 Windows vulnerabilities in the May patch day, CyberGym 88.45%. Secondary analyses: csoonline.com, thehackernews.com (accessed 19 July 2026).
- Anthropic – Claude Fable 5 and Claude Mythos 5, 9 June 2026. anthropic.com/news/claude-fable-5-mythos-5 (accessed 19 July 2026).
- Anthropic – Redeploying Claude Fable 5, 1 July 2026. anthropic.com/news/redeploying-fable-5 (accessed 19 July 2026).
- Anthropic – Project Glasswing: An initial update, 22 May 2026. anthropic.com/research/glasswing-initial-update (accessed 19 July 2026).
- Cloud Security Alliance – Project Glasswing: AI Discovery Outpaces Open Source Patching Capacity, June 2026. cloudsecurityalliance.org (accessed 19 July 2026).
- Cloud Security Alliance, Cloud Key Management Working Group – Post-Quantum Cryptography Key Management. Draft, 13 August 2026. cloudsecurityalliance.org (accessed 24 August 2026).
- Cloud Security Alliance, SANS Institute, [un]prompted, OWASP Gen AI Security Project (eds.) – The „AI Vulnerability Storm": Building a „Mythos-ready" Security Program. Expedited Strategy Briefing. Public release 14 April 2026; cited edition v0.95 of 18 April 2026. cloudsecurityalliance.org (accessed 19 July 2026).

- OWASP Gen AI Security Project – Top 10 for LLM Applications 2026 (LLM01–LLM10) and Agentic Security Initiative (ASI01–ASI10).
- MITRE ATLAS – Adversarial Threat Landscape for Artificial-Intelligence Systems, atlas.mitre.org (accessed 19 July 2026).
- NIST AI Risk Management Framework 1.0, AI RMF 1.0.
- MITRE ATT&CK Enterprise Matrix – attack.mitre.org (accessed 19 July 2026).
- MITRE D3FEND – d3fend.mitre.org (accessed 19 July 2026).
- OWASP Application Security Verification Standard (ASVS) – owasp.org/www-project-application-security-verification-standard (accessed 19 July 2026).
- OWASP Software Assurance Maturity Model (SAMM) – owaspsamm.org (accessed 19 July 2026).
- Center for Internet Security (CIS) – CIS Critical Security Controls v8.
- CSA Cloud Controls Matrix v4.
- SLSA – Supply-chain Levels for Software Artifacts, slsa.dev.
- ENISA – European Cybersecurity Certification Scheme for Cloud Services (EUCS).
- ENISA – EU Vulnerability Database (EUVD), euvd.enisa.europa.eu (accessed 19 July 2026).
- FIDO Alliance – WebAuthn Specifications.
- TIBER-EU – European framework for Threat Intelligence-based Ethical Red Teaming, European Central Bank.

Online sources last checked: 19 July 2026 (PQC key management: 24 August 2026; MITRE ATT&CK and ATLAS: 27 August 2026).

Version notes

This handbook is maintained as a living document. The following table serves as the change log for subsequent versions. The present edition is version 2.0.

Version	Date	Changes
1.0	April 2026	First publication. Assessment of all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape. Framework cross-comparison with BSI C5:2026, NIST, DORA, the CRA and NIS2 with its Implementing Regulation. Catalogue of twelve Mythos Hardening Controls (MHC-01 to MHC-12).
1.1	April 2026	Revised edition for audit robustness. Correction of incorrect NIS2 IR references (No. 1.2.5 does not exist, corrected to 11.2.2(a); deadlines not in IR 3.3 but in NIS2 Directive Art. 23(4); misinterpretation of 3.3.2; number of phases in 3.5 corrected from four to three). Addition of No. 6.7 (network security) for A.8.20 to A.8.22. Clarification of the scope of the NIS2 IR. Reclassification of A.7.7 from "Friction-only" to "Not affected" (new count 29/37/4/23). Sharpening of Chapter 3.1.3 to resolve the contradiction between definition and application. Marking of the time factor addition as an author's interpretation of ISO/IEC 27005. Positioning of BSI C5 as an audit catalogue. Consistency of Chapter 9.4 against Annex A. Addition of Annexes A to G. Harmonization of spelling and typography.
1.2	April 2026	Concretization of the hardening recommendations for audit robustness. Tool classes named explicitly (EDR, SAST/DAST/SCA, MDR, SOAR, ZTNA, EASM). Thresholds and quantity figures added (MTTC below 10 min, ATT&CK coverage 60% or above, egress anomaly 2σ/3σ, EDR confidence 80% or above, at least twelve threat hunts per year). Three-stage maturity paths (Initial/Defined/Managed) for all twelve MHC in Chapter 9.5. AI agents addressed as an insider risk in MHC-04 (capability-scoped identities). Vibe-coded vulnerable code addressed in MHC-09 as a secondary threat. Living off the land and beacon-less C2 as detection priorities in MHC-05. Practicality notes added (MDR alternative, hyperscaler reality, BAS platforms). Methodology references concretized (PEAK framework, TaHiTI, DeTT&CT, STRIDE). Mythos-specific user reporting indicators in A.6.8.

Version	Date	Changes
1.3	April 2026	Coverage extension relative to the CSA/SANS/OWASP report "Building a Mythos-ready Security Program" (April 2026). New MHC-13 (AI agent governance and harness security) closes the risk of an unmanaged AI agent attack surface classified as CRITICAL in Mythos-Ready: harness audit, blast radius limits, human override, supply chain inventory for MCP servers, extensions and agentic skills, pre-production checks. A.5.9 extended by a shadow AI inventory (browser plug-ins, IDE extensions, MCP servers). Chapter 10.4 extended by innovation acceleration governance (a cross-functional body with a 30-day decision target) and a permanent VulnOps function. Chapter 10.5 supplemented by the standard of care shift through the EU AI Act and a board briefing element. Maturity table in Chapter 9.5 extended by MHC-13; assessment matrix (Annex A) and MHC standalone worksheet (Annex C) updated accordingly. The Mythos-Ready strategy paper positioned in the foreword and Chapter 3.4 as the central industry consensus reference (status between a threat intelligence report and a normative basis).
1.4	April 2026	Closure of the structural gaps identified by the audit assessment ("completeness = medium"). New Chapter 1.4 "Position in the ISMS stack and limits of this handbook" positions MRIS explicitly as an effectiveness and reality layer on an ISMS foundation; it makes clear what MRIS delivers (effectiveness assessment, gap analysis, MHC catalogue, operationalization) and what it deliberately does not (a complete ISMS, a risk management process of its own, a PDCA system, a compliance mapping tool, an organization-specific policy hierarchy). New Annex E "RACI model for MHC implementation" with eight roles for all thirteen MHC. New Annex F "Risk assessment bridge under ISO/IEC 27005" as an interface definition between MRIS and the ISMS risk process. New Annex G "KPI definitions and measurement standardization" with full standardization of the Mythos metrics (data source, clock start and stop, calculation formula, reporting frequency, measurement responsibility).
1.5	April 2026	Sharpening of the classification language and the effectiveness statements after a critical audit assessment. Classic controls are clearly described as "selectively degraded", not as "obsolete" — corrective paragraphs in Chapters 9.1, 4.3 and 11.4 emphasize the priority of structural controls ahead of new MHC. New Chapter 3.1.5 on the context dependency of the classification (the same control may fall into different categories against different attack vectors, example A.5.17). Aggregation resistance operationalized in Chapter 3.2 (SIEM correlation, UEBA, graph analysis, kill chain tracking, or structural indivisibility). Threat scope in Chapter 3.5 limited explicitly to Gen-AI-accelerated attacks. MHC-05 supplemented by a realism note (detection is not prevention; low-and-slow, disguised agents and adversarial ML evasion as residual risk). MHC-11 extended by hardening of the automation layer (signed triggers, audit trail, rate limits, human in the loop for actions with a high blast radius). MHC-13 supplemented by maturity realism.
1.6	June 2026	Reference update without substantive reassessment of the controls. Chapter 9.2: cryptographic standards switched to the finalized FIPS 203/204/205 and to HQC as selected for standardization (MHC-01); NIST SP 800-94 positioned as the 2007 edition with a withdrawn revision draft (MHC-05); MITRE ATLAS IDs corrected in MHC-13 (AML.T0047 → AML.T0053 for LLM Plugin Compromise, addition of the agentic techniques AML.T0086/T0110) and OWASP LLM numbering cleaned up (LLM08 "Excessive Permissions" removed; covered by LLM06 Excessive Agency). Updated references: NIST SP 800-63B Rev. 4 (MHC-03), CycloneDX/SPDX standardization, BSI TR-03183-2, CRA timeline and VEX/CSAF (MHC-02), Sigstore, admission controller and TEE examples (MHC-06), ISO/IEC 27017 (MHC-07), NIST SP 800-218A (MHC-09), NIST SP 800-137/137A and OSCAL (MHC-10), NIST SP 800-61 Rev. 3 (MHC-11), DORA TLPT RTS (EU) 2025/1190 and TIBER-EU/TIBER-DE (MHC-12). Reference status: June 2026.
1.7	July 2026	Consistency release after an external dual assessment (auditor and CISO): mapping in Chapter 9.4 harmonized with Annexes A and C (Annex A authoritative; A.8.4, A.8.27, A.5.9); annex references to E and G corrected; MHC count consistently "thirteen"; two new metrics (container policy conformity, agent inventory coverage) in Chapter 10.6 and

Version	Date	Changes
		Annex G; process shell clarification for friction controls (Chapters 3.1.3, A.8.8); review triggers in Chapter 11.3; guide value marking for unsubstantiated figures; license changed to CC BY-SA 4.0.
1.8	July 2026	Multi-vendor update (response to review trigger a): new Chapter 2.5 (Daybreak, MDASH, general availability of the Mythos class, disclosure waves); "Mythos" as a class term; MHC-02, MHC-09 and MHC-12 sharpened multi-systemically; new Chapter 11.4 threat horizon; gate criteria for MHC-04 and MHC-07 in Chapter 10.6; reference clean-up (Mythos-Ready dating, TR-03183, access date per web source); machine-readable edition (YAML) referenced.
1.9	August 2026	External evidence update (pass-the-passkey attack class, SpecterOps, August 2026; a classic implementation and endpoint weakness, no control reclassification). MHC-03 supplemented by attestation enforcement and an effectiveness limit; A.8.15 supplemented by the confidentiality dimension of logs (CVE-2026-34348); MHC-05 supplemented by detection signals for WebAuthn and passkey abuse. MHC-13 references pinned to OWASP Top 10 for LLM Applications 2026 and numbering corrected (Excessive Agency LLM06 → LLM03; Chapters 9.2, 9.4, Annex C, foreword, list of references).
1.10	August 2026	Condensation for manageability (extent significantly reduced) without loss of substance at the core. Chapter 9.2: hardening recommendations for the 13 MHC reduced to core idea and Mythos effect, with operational implementation delegated to the Implementation Guide by reference (MHC ID); Mythos finding and framework comparison unchanged. Part V (supplementary frameworks) and Annex G (KPIs) condensed from prose into tables. Duplications on limits and theses consolidated (Chapters 3.5, 11.1/11.2), closing remark 11.5 stripped back, summaries 6.3 and 10.7 and Chapter 7 tightened, citation density of the robust framework comparisons reduced (framework mapping retained, glosses and explanations removed). Core unchanged: Chapters 5 and 6, Annex A, MHC findings.
1.11	August 2026	Implementation of the external technical review (Daniel Heldt) of version 1.8. Correction of the risk concept in Chapter 1.2 against ISO/IEC 27005:2022 (risk = "effect of uncertainty on objectives", 3.1.3; likelihood × consequence = level of risk, 3.1.15; motivation, capability and exploitability located more precisely in Clause 7 and Annex A.2.3). Central claim and gap analysis worded more softly (explicitly required versus implicitly covered by ISO); the capacity barrier statement moderated (physical and hardware attacks excluded). PQC gap positioned honestly (quantum-driven, not Mythos-accelerated; crypto agility with several trust anchors). Continuous monitoring "in addition to" rather than "instead of" periodic audits. The CRA added as a source of notification obligations and their applicability. Framing of "effectiveness of the measure, not the abstract control" (Chapter 3.2); terminological clarification of the MRIS framework versus this handbook. Pie chart of the control distribution. Thanks to Daniel Heldt.
1.12	August 2026	PQC key management deepened (source: CSA Cloud Key Management Working Group, draft 08/2026; own wording with attribution). Chapter 8.2 supplemented by four key-management-specific facets: trust now, forge later (harvesting of signed artefacts), downgrade resistance of hybrid negotiation, migration order within the key hierarchy (root and KEK first), quantum-safe timestamps. MHC-01 core idea refined accordingly; CSA source added to the list of references. Operational implementation in the Implementation Guide, MHC-01.
1.13	August 2026	Clarification of the ISO linkage in Chapter 9.2: where the finding cell names only the primary anchor controls, this is now marked as a selection ("primarily flanks") with a reference to the complete mapping in Annex A. Linguistic corrections from the technical review (Heldt): role reference opened up to the ISMS ownership function, AI systems no longer described as authors of phishing and malware, scope of the CRA SBOM obligation narrowed to manufacturers of products with digital elements, multi-factor criterion for

Version	Date	Changes
		device-bound passkeys clarified, workload identity explained at first mention, Chapter 2.5 titled for low maintenance.
2.0	August 2026	Quality release after an external quality, CISO and audit review. Methodology: separation of structural protective capability and organization-specific implementation maturity (Chapter 3.2); the "Friction-only" category defined in implementation terms (Chapter 3.1.3); binding reading of the maturity levels — Initial is not evidence of implementation (Chapter 9.5); risk assessment bridge without a blanket level increase, replaced by reassessment of control effectiveness using the organization's own methodology (Annex F). Corrections: share of controls requiring hardening corrected to 44% (41 of 93); NIS2 IR references for authentication switched to No. 11.6/11.7 and 11.3.2(a); DORA TLPT limited to the entities identified under Art. 26(8); the SoA no longer named as the place where risks are recorded (MHC-09); CRA SBOM and AI Act standard of care refined in regulatory terms. Further sharpening: MHC-01 positioned as a cross-cutting resilience control, MHC-04 identity/credential wording and lateral movement effect clarified, MHC-08 derived from the control degradation analysis, MHC-09 delimited against the testing baseline, RACI marked as an example assignment, statement on ISO 27002 gap coverage moderated, rationales for A.5.3, A.5.5, A.5.7, A.8.6 and A.8.11 sharpened, metrics G.2, G.4, G.5 and ATT&CK coverage revised, spelling of "AI" harmonized.

Occasions for updating future versions will include in particular: publication of new BSI C5 editions, ISO 27002 revisions, changes in NIS2 implementing acts, new DORA RTS, updates to the CRA implementing acts, significantly new findings on the Mythos threat landscape, and lessons learned from applying this handbook in practice.