MYTHOS-RESISTANT INFORMATION SECURITY

MRIS

Effectiveness Assessment of Existing ISO 27001 Annex A Controls under Gen-AI-Accelerated Offence

Cross-Framework · Evidence-Based

Version 1.6 | June 2026
Author: Richard Peddi

TARGET AUDIENCE
CISOs, ISMS officers and security architects who need to assess existing controls against the new reality of agentic and Gen-AI-accelerated attacks

Licence and Disclaimer

Disclaimer

This guide constitutes a technical and organisational reference. It does not replace an individual risk analysis, legal advice or an audit-specific review. The assessment of individual controls presented here refers to the publicly documented state of agentic AI threats at the time of writing. The assessment may shift as new evidence emerges.

Recommended Citation

Peddi, Richard (2026): MRIS – Mythos-Resistant Information Security, Version 1.6.

Table of Contents

Foreword

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

This raises the core question of this guide: which of the controls anchored in an ISMS today still withstand this reality – and which do not?

MRIS assesses all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape, reconciles them with NIST, BSI C5:2026, DORA, 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.

MRIS is not a fully-fledged ISMS framework and replaces neither ISO/IEC 27001 nor ISO/IEC 27005. It supplements the existing ISMS with an effectiveness and reality layer (see Ch. 1.4 on positioning in the ISMS stack).

In April 2026, a working group comprising the Cloud Security Alliance, 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 industry consensus level, MRIS translates it into a control-by-control assessment and an audit-ready hardening catalogue.

This guide was written in the course of my work as Group CISO at Proalpha. Suggestions and criticism are expressly welcome.

Richard Peddi
April 2026

Management Summary

Central Finding

The ISO/IEC 27005 methodology remains valid. The effectiveness and likelihood parameters with which existing ISMS calculate have been structurally destabilised by Gen-AI-accelerated attacks. This guide makes this destabilisation visible control by control and closes it through a prioritised supplementary catalogue.

Four Shifts in the Threat Landscape (Ch. 2)

- Collapse of the patch-gap window: the time between patch release and a working exploit drops to hours.
- Compression of the response timeline: agents autonomously execute 80 to 90 per cent of the tactical attack work (documented in GTG-1002, November 2025).
- Fragmentation of the threat event into micro-steps that individually appear policy-compliant.
- Decoupling of capability and actor: individual perpetrators possess capabilities previously attributed only to nation-state actors.

Assessment of the 93 ISO 27001 Annex A Controls

Category	Number	Meaning
Robust	29	Foundation: cryptographic, architectural and aggregation-resistant hard barriers.
Partially degraded	37	Still effective, but require supplementation through five Mythos hardening patterns (Chapter 5.3).
Friction-only	4	Must be redesigned or replaced (A.5.25, A.5.36, A.8.8, A.8.23).
Not affected	23	Mythos-neutral governance and infrastructure base (including A.7.7).

Seven Systematic Gaps Relative to ISO 27002:2022 (Ch. 8)

- Post-quantum strategy and crypto-agility (C5:2026 CRY-01.01AC).
- Supply chain transparency with mandatory SBOM (CRA Annex I, C5 DEV-13, SLSA).
- Container security, confidential computing, multi-tenancy isolation (C5:2026 OPS-30 to OPS-35).
- Time-based reporting deadlines and mandatory root cause analysis (NIS2 Directive Art. 23(4), DORA Art. 17/19).
- Continuous assurance instead of periodic audits (C5 OPS-25.01AS, DORA Art. 26/27 TLPT).
- Phishing-resistant MFA and workload identity (NIST SP 800-63B, FIDO2, SPIFFE).

- Mandatory automation and parallel incident testing (NIS2 Implementing Regulation No. 3.2.2, DORA Art. 24 to 27).

Thirteen Mythos Hardening Controls (MHC) as a Supplementary Catalogue (Ch. 9)

- MHC-01: Post-quantum strategy. MHC-02: SBOM and build provenance. MHC-03: Phishing-resistant MFA. MHC-04: Workload identity and Zero Trust.
- MHC-05: Behaviour-based detection. MHC-06: Containers and confidential computing. MHC-07: Multi-tenancy isolation. MHC-08: Immutable backups.
- MHC-09: AI-supported security testing. MHC-10: Continuous control monitoring. MHC-11: SOAR and tier-1 automation. MHC-12: Threat-led penetration testing.
- MHC-13: AI agent governance and harness security (response to Mythos-Ready, CSA/SANS/OWASP, April 2026).

Five Action Areas for ISMS Adaptation (Ch. 10)

- Immediate reassessment of existing controls along the four assessment criteria.
- Structural hardening: adoption of the thirteen MHC into the Statement of Applicability.
- ISMS maturation: shorter risk assessment cycles, continuous audit, management review with a Mythos agenda.
- Board communication along the four findings from Chapter 2.
- Mythos-relevant metrics: MTTC, phishing-resistant MFA share, SBOM coverage, patch latency for KEV listings, restore test success rate, continuous monitoring coverage, cryptographic inventory coverage, TLPT findings closure.

Methodology and Limitations

Primary assessment axis: ISO/IEC 27002:2022, BSI C5:2026 (audit catalogue), NIST CSF 2.0 and SP 800-53, DORA, CRA, NIS2 Directive (EU) 2022/2555 with Implementing Regulation (EU) 2024/2690. Supplementary layer: BSI IT-Grundschutz, MITRE ATT&CK/D3FEND, ISO/IEC 42001, CIS v8, OWASP ASVS/SAMM. This guide does not replace an organisation-specific risk analysis and makes no claim to completeness of the MHC catalogue. As at: April 2026.

1 Introduction

1.1 Why This Guide Exists

The occasion for this guide is a concrete and publicly confirmed change in attack reality. 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, as exploit generators, as phishing and malware authors. With the publication of the Glasswing programme and the first assessment of Claude Mythos Preview in April 2026, the vendor publicly stated that models can now find and exploit vulnerabilities in code at a pace that structurally overwhelms classic response processes.

For information security this means a rupture that has already occurred. The attacker 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 per cent autonomous execution, decomposition of complex attacks into micro-steps that individually appear legitimate.

This raises the question every CISO must answer: which of the controls anchored in an ISMS today still withstand this reality? And where does the mere existence of a control feign a protective effect that the Mythos attacker has long ceased to respect?

This guide delivers the answer as a systematic assessment of all 93 controls of ISO/IEC 27002:2022, supplemented by the perspectives of NIST, the BSI C5:2026 audit catalogue, DORA, CRA, the NIS2 Directive and the NIS2 Implementing Regulation (EU) 2024/2690. The added value lies in the convergence analysis: where one framework imposes a requirement that another does not know and that holds up under Mythos conditions, there lies a hardening proposal for ISO-centred ISMS.

1.2 The Central Thesis: Methodology Intact, Inputs Destabilised

ISO/IEC 27005:2022 defines risk as a function of likelihood and impact. For intentional threat sources, according to ISO/IEC 27005:2022 Annex A.2, the analysis particularly considers the motivation and capabilities of the actor as well as the exploitability of the vulnerability. This methodology is not wrong. What has changed are its inputs.

Note on the author's addition: This guide additionally introduces the parameter *time between threat event and organisational response* as a fourth dimension. ISO/IEC 27005:2022 does not explicitly name this time factor as a risk parameter; in the following it serves as an assessment criterion because Mythos attacks structurally push response time below the humanly manageable threshold.

Mythos-class AI largely removes the capacity barrier between opportunistic offenders and nation-state actors. It compresses the timeline between vulnerability disclosure and available exploit to hours. It fragments attack actions into micro-events that individually do not raise alarms. And it systematically devalues those controls whose protective effect rests not on impossibility but on the limited patience of the attacker.

The central thesis is: the ISO/IEC 27005 methodology remains applicable, but the likelihood and effectiveness parameters that have so far fed into an ISMS must be reassessed. Each control must be examined as to whether it derives its protective effect from a hard barrier or merely from friction – and whether this friction still suffices under Mythos conditions.

1.3 The Central Claim

Where NIST, C5, DORA, CRA, the NIS2 Directive or the NIS2 Implementing Regulation imposes a requirement that ISO 27002 does not explicitly know, and this requirement holds up under Mythos conditions, it constitutes a concrete hardening proposal for every ISO-centred ISMS. The CISO task shifts from compliance depth to effectiveness-oriented control discovery.

1.4 Position in the ISMS Stack and Limitations of This Guide

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 based on ISO/IEC 27001:2022, BSI IT-Grundschutz, BSI C5 2026 or an equivalent framework – and on this basis answers a single question: do the controls implemented today still work against Gen-AI-accelerated offence? Without an ISMS foundation, MRIS is not applicable.

What MRIS provides:

- Effectiveness assessment of all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape (Part II).
- Gap analysis against primary frameworks (BSI C5:2026, NIST, DORA, CRA, NIS2 with Implementing Regulation) (Ch. 8).
- Audit-ready catalogue of thirteen Mythos Hardening Controls with maturity levels, tools and thresholds (Ch. 9).
- Operationalisation in four annexes: assessment matrix (A), standalone MHC worksheet (C), RACI model (H), KPI definitions (J).

What MRIS deliberately does not provide:

- No complete ISMS under 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. Via the assessment categories (Ch. 3) and the risk assessment bridge (Annex F), MRIS delivers inputs for the risk process under ISO/IEC 27005, but does not replace it.
- No complete PDCA management system. Via the maturity levels (Ch. 9.5), versioning and the continuous audit recommendations (MHC-10), MRIS delivers building blocks for Plan-Do-Check-Act, but does not replace a standalone management system.
- No compliance mapping tool. The framework comparisons of individual controls are effectiveness cross-references, not a complete requirement-to-standard mapping with audit trail. For certification and regulatory evidence, specialised GRC tools should be used.
- No organisation-specific policy hierarchy. Policies, procedures and work instructions must be developed within the organisation; MRIS provides content anchors for this, but no templates.

Correct use. MRIS is intended as a layer on an existing ISMS: the ISMS provides governance, risk management process, policy hierarchy and PDCA cycle. MRIS provides the effectiveness assessment of the individual controls against Mythos and concrete, audit-ready hardening recommendations. The separation of both layers is methodologically decisive – MRIS is not meant to duplicate ISMS frameworks, but to supplement them precisely where these

frameworks, owing to their technology-agnostic design, reach their limits against Gen-AI-accelerated threats.

2 Findings: The Four Shifts in Attack Reality

The following four findings are not hypotheses. They rest on the publicly documented threat intelligence reports by Anthropic from the period August 2025 to April 2026, as well as on the defensive recommendations of the Anthropic security engineering team within the Glasswing programme.

2.1 Collapse of the Patch-Gap Window

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

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

2.2 Compression of the Timeline

ISO 27005, like many incident response playbooks, presupposes humanly manageable time between detection, triage and decision. This time is no longer available.

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

2.3 Fragmentation of the Threat Event

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

The risk materialises only in aggregation. Controls based on the evaluation of 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 supposed to see. They do not see that the sum constitutes an attack.

2.4 Decoupling of Capability and Actor

Classic likelihood assessment relies substantially on the estimate of attacker capability: a nation-state actor was considered unlikely against a mid-sized German software manufacturer, not for lack of interest, but because the technical capacity for an individualised 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 execute operations that would previously have been attributed only to professional APT groups. The likelihood axis in every risk matrix shifts upwards – not because of new threat actors, but because the capacity barrier of existing actors has fallen.

2.5 Consequence for ISO/IEC 27005

The methodology of ISO 27005 – risk identification, analysis, evaluation, treatment – remains intact as a process. The inputs into this methodology, however, are affected: the likelihood axis shifts upwards through the removal of the capacity barrier. The time parameters between patch and exploit collapse. The aggregation logic of individually assessed events no longer applies. And the effectiveness assumption of many existing controls – especially those based on attacker friction rather than hard impossibility – has been empirically refuted.

Note: ISO/IEC 27005:2022 distinguishes in Section 7.2.1 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 aggregation.

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 effectiveness boundary. 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 stems from a hard barrier: cryptographic impossibility, physical separation, architectural non-existence of an attack path. These controls hold because they structurally exclude attacks or make them fail at an objective boundary.

3.1.2 Partially Degraded

Controls that continue to provide protection, but whose originally assumed strength drops significantly under Mythos conditions. They remain useful, but need supplementation – through additional controls, changed parameters or architectural flanking.

3.1.3 Friction-Only

Controls whose effectiveness against Mythos attackers is structurally eliminated because the original protective mechanism either (a) rests on limited attacker capacity (classic friction) or (b) presupposes human response time within an automatically traversed kill chain. In both cases, the asymmetry between a model-driven attacker side and a defence side not correspondingly automated causes the control to lose its core effect. These controls must either be replaced or supplemented by barriers based not on effort or response time, but on structural impossibility.

3.1.4 Not Affected

Controls that are neutral towards Mythos-specific threats – typically organisational, documentation-oriented, governance-oriented or physically bound controls whose effectiveness is systematically altered 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 to the Mythos threat landscape as described in Ch. 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 assessments in Chapters 4 to 7 each take a control's dominant Mythos effectiveness mode as the basis for classification. Auditors and ISMS officers check before SoA adoption whether deviating classifications are justified in their specific threat situation.

3.2 Assessment Criteria

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 response time in a chain in which the attacker now operates autonomously? If so, effectiveness is impaired by the latency difference.
- Capability decoupling: does the control's likelihood assumption depend on a historical actor-capability correlation that Mythos has removed? 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 individually appear legitimate? If not, the control is blind to fragmented attack patterns. Operationalisation: aggregation resistance is considered given 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 analyses or kill chain tracking according to MITRE ATT&CK – or (b) is structurally indivisible because its protective effect applies fully per individual transaction (cryptographic operations, origin-bound authentication, hardware-attested identities). A control that only checks thresholds of individual actions, without correlation across multiple actions, is not considered aggregation-resistant.

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

3.3 Framework Stack and Prioritisation

The core assessment is based on ISO/IEC 27002:2022 as the internationally most established generic control catalogue. For each control, a cross-comparison checks whether other frameworks relevant for regulation or audit impose a requirement that covers the same protection purpose while exhibiting higher Mythos robustness.

Primary frameworks:

- BSI C5:2026 (Cloud Computing Compliance Criteria Catalogue). BSI audit catalogue for cloud services (updated March 2026). Not legally binding per se, but becomes practically binding through contractual anchoring (e.g. for cloud customers in the public sector) or through sector regulation (BaFin BAIT, 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. Provide the technical depth for detection, response and adversarial testing.
- DORA (Digital Operational Resilience Act, Regulation (EU) 2022/2554). Legally binding for the financial sector. For the resilience aspect and the TLPT requirements (Art. 26–27).
- CRA (EU Cyber Resilience Act). Legally binding for software manufacturers with CE marking. Vulnerability handling obligations (Art. 13, 14), SBOM (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 reporting deadlines for significant incidents (Art. 23(4): 24 h early warning, 72 h full notification, 1 month final report). According to its Article 1, the Implementing Regulation (EU) 2024/2690 applies exclusively to certain digital providers: DNS providers, TLD name registries, cloud computing providers, data centre services, CDN operators, managed service providers, managed security service providers, online marketplaces, online search engines, social networking platforms and trust service providers. For other NIS2 addressees (critical infrastructure, manufacturing, health), the national transposition is decisive.

Note on citation: In the framework comparisons of individual controls, the Implementing Regulation is cited with the 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/defence vocabulary).
- ISO/IEC 42001 (AI management system).
- CIS Controls v8 (prioritised implementation aid).
- OWASP ASVS and SAMM (secure development).

3.4 Source Base

The sources of this guide fall into three categories.

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

Normative basis: ISO/IEC 27001:2022, ISO/IEC 27002:2022, ISO/IEC 27005:2022, relevant NIST publications, BSI C5:2026, DORA (EU) 2022/2554 with RTS, CRA, as well as the 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, Version 0.95, 18 April 2026) is treated in MRIS as the central industry consensus reference for the Mythos threat landscape. The risk assessment defined there (thirteen prioritised risks at the levels Critical, High, Medium) has directly informed MRIS concretisations, in particular MHC-13 (AI agent governance, addressing Mythos-Ready Risk 3 "Unmanaged AI Agent Attack Surface"), the extension of A.5.9 to include a shadow AI inventory (Risk 6), Ch. 10.4 on the permanent VulnOps function and innovation acceleration governance (Risks 9 and 11), and Ch. 10.5 on the standard-of-care shift (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 threat intelligence report and normative basis. Additionally used industry references: OWASP LLM Top 10 and OWASP Agentic Security Initiative (ASI01–ASI10), MITRE ATLAS, NIST AI RMF 1.0, ISO/IEC 42001 Annex A.

3.5 Limitations

The assessment presented reflects a snapshot in time. The Mythos threat landscape continues to evolve, and individual classifications may shift. The assessment does not replace an organisation-specific risk analysis. It disregards controls that must be implemented anyway for regulatory reasons but require no separate assessment under Mythos aspects. It makes no claim to completeness of possible hardening measures.

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

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. Anyone wishing to look up the classification of an individual control will find the complete matrix in Annex A.

Each individual control assessment follows the same format: category classification, Mythos finding with justification based on the four criteria from Chapter 3.2, framework cross-comparison against C5:2026, NIST, DORA, CRA and NIS2 (Directive plus Implementing Regulation), and – for degraded and friction-based controls – a concrete hardening or replacement recommendation.

Overall distribution of the 93 controls:

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

The distribution is expressly not a finding of catastrophe. It states that ISO/IEC 27002:2022 remains fundamentally resilient, but that around two thirds of the controls require realignment towards the Mythos threat in order to continue fulfilling their original protection purpose.

4 Robust Controls – the Resilient Foundation

4.1 Selection Logic

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

29 controls are assessed as robust in this chapter. The presentation follows the ISO numbering; the thematic grouping is summarised 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, up-to-date asset inventory is the prerequisite of any targeted defence. Anthropic states in the Glasswing post that systems the organisation does not know about cannot be defended either. The control is Mythos-robust because the existence of an inventory is devalued neither by attacker speed nor by micro-step fragmentation.
Framework comparison	C5:2026 concretises via AM-02 (Asset Inventory), AM-03 (Hardware Asset Inventory) and AM-04 (Software Asset Inventory). NIST CSF 2.0 ID.AM-1 to ID.AM-5. DORA Art. 8 requires an ICT asset inventory as a core obligation. NIS2 IR No. 12 (asset management) mandates a documented inventory and regular updating.
Hardening recommendation	Automate inventory updating (CMDB integration, cloud asset inventory tools such as AWS Config, Azure Resource Graph). Manual inventories are structurally outdated at Mythos iteration speed. Additionally deploy shadow IT scans and external perimeter inventorying (Certificate Transparency logs, ASM tools). Maintain shadow AI as a dedicated inventory category: unauthorised AI coding assistants and browser plug-ins (e.g. 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 via endpoint software inventory (e.g. via EDR telemetry), browser extension audits (e.g. via MDM or browser enterprise policies) and network egress monitoring towards 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 prioritised hardening. An attacker who executes 80–90 per cent of the tactical work autonomously ultimately encounters the same data set – whether it 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 (Asset Classification and Labelling). NIST SP 800-60. DORA Art. 8 implicitly. CRA Annex I requires security-appropriate documentation of critical components. NIS2 IR No. 12.1 mandates classification according to protection needs.
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). Deploy automated classification via data discovery tools.

A.5.16 Identity management

Category	ROBUST
Mythos finding	Unique, cryptographically anchored identities are a hard barrier. Under Mythos, the emphasis shifts from human user identities to workload and agent identities. The control covers both categories, provided the organisation consistently deploys workload identity (SPIFFE/SPIRE, AWS IAM Roles, Azure Managed Identities).
Framework comparison	NIST SP 800-63 Digital Identity Guidelines. NIST NCCoE publications on agent identity (2026). C5:2026 IAM-01 (Policy for Identities and Access Rights) and IAM-02 (Granting and Change). DORA Art. 9. NIS2 IR No. 11 (access control) concretises operational requirements for identities and their management.

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 systematically collected and fed back into detection logic. The derivation of detections from lessons learned (linked to A.8.16 monitoring) must be executed in hardened form.
Framework comparison	NIST SP 800-61 Computer Security Incident Handling Guide. C5:2026 SIM-06 (Evaluation and Learning Process). DORA Art. 13 requires root cause analysis as a mandatory component of the incident management lifecycle. NIS2 IR No. 3.6 (post-incident reviews) requires feeding the findings back into the improvement of policies and procedures.

A.5.28 Collection of evidence

Category	ROBUST
Mythos finding	Forensic evidence collection is robust as long as the logging foundations (A.8.15) and time synchronisation (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 Guide to Integrating Forensic Techniques. C5:2026 SIM-04 (Documentation and Reporting of Security Incidents). DORA Art. 17 and ISO/IEC 27037 provide chain-of-custody requirements. NIS2 IR No. 3.5.4 mandates logging of the response measures taken during security incidents.

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, because Mythos attacks can run on several vectors simultaneously; the underlying ICT readiness capability remains robust if dimensioned accordingly.
Framework comparison	DORA Art. 11–12 is the strictest regime: ICT Business Continuity Policy, Digital Operational Resilience Testing Programme and TLPT (Art. 26–27) go well beyond ISO. C5:2026 BCM-01 to BCM-04 require a complete BCM regime including regular testing. NIS2 IR No. 4 (business continuity and crisis management) concretises contingency plans, backup strategy and recovery procedures. 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. Tabletop exercises for multiple 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 aggravated by an attacker's ability to find and chain such credentials automatically.
Framework comparison	C5:2026 HR-05 (Responsibilities in the Event of Termination or Change of Employment). NIST SP 800-53 PS-4. DORA Art. 9. NIS2 IR No. 10 (personnel security) and No. 11.2 (management of access rights) link offboarding to the immediate deactivation of all access rights.

A.7.1 Physical security perimeters

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

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 (Physical Site Access Control). NIST PE-2, PE-3. DORA Art. 9(4). NIS2 IR No. 13.2 concretises access roles and revocation of authorisations.

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 (Physical Security and Environmental Control Requirements) and PS-08 (Workplace Security). NIST PE-5. NIS2 IR No. 13.

A.7.4 Physical security monitoring

Category	ROBUST
Mythos finding	CCTV, intrusion detection and motion sensors are Mythos-orthogonal.
Framework comparison	C5:2026 PS-07 (Surveillance of Operational and Environmental Parameters). NIST PE-6. NIS2 IR No. 13 requires detection measures against unauthorised entry.

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 (Workplace Security Requirements). 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 (Removable Media and Endpoint Devices). NIST SP 800-88 Guidelines for Media Sanitization. NIS2 IR No. 12 covers the lifecycle of storage media.

A.7.14 Secure disposal or re-use of equipment

Category	ROBUST
Mythos finding	Irrevocable data destruction before disposal or reassignment is a cryptographically hard barrier (DoD-compliant wipe, cryptographic shredding).
Framework comparison	C5:2026 AM-07 (Decommissioning of Hardware). NIST SP 800-88. BSI TR-03183-H. CRA Annex I requires provisions for secure decommissioning.

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 harvesting credentials automatically runs into a wall if those credentials carry no far-

	reaching privileges. Robust when hardware-attested and implemented with just-in-time elevation (PAM); pure password protection would be degraded.
Framework comparison	NIST SP 800-53 AC-6, IA-5. NIST SP 800-207 Zero Trust Architecture. C5:2026 IAM-06 (Privileged Access Rights) requires a separate policy, just-in-time granting and dedicated monitoring. DORA Art. 9 and 15. NIS2 IR No. 11.2 (access rights) concretises separation of roles, approval and regular review.

A.8.9 Configuration management

Category	ROBUST
Mythos finding	Versioned, automatically monitored configuration baselines (Infrastructure as Code, GitOps) are robust. Drift detection identifies unauthorised changes in real time – a hard structural barrier against agentic modifications.
Framework comparison	NIST SP 800-53 CM-2 to CM-8. C5:2026 OPS-26 (System Hardening) and DEV-03 (Policies for Changes to System Components). CIS Controls v8.4. SLSA Level 3+ for build provenance. NIS2 IR No. 6 (security in acquisition, development and maintenance) mandates documented configuration and hardening specifications.

A.8.10 Information deletion

Category	ROBUST
Mythos finding	Cryptographic deletion is a hard barrier. Mythos attackers cannot retroactively exfiltrate what was previously irreversibly deleted. Essential against supply chain attacks on backups and archives.
Framework comparison	NIST SP 800-88. C5:2026 PI-03 (Secure Deletion of Data) and CRY-14 (Secure Deactivation of Cryptographic Keys). GDPR Art. 17. NIS2 IR No. 12 implies secure data destruction at the end of the asset lifecycle.

A.8.11 Data masking

Category	ROBUST
Mythos finding	Pseudonymisation and anonymisation structurally reduce the blast radius: even after successful exfiltration, the leaked data set is of reduced value without a re-identification capability.
Framework comparison	ISO/IEC 20889. NIST SP 800-188. GDPR Art. 32(1)(a). C5:2026 OPS-30 and OPS-31 (Separation of Datasets). NIS2 IR No. 6 mandates data-minimising use in development and test environments.

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 (Data Backup and Recovery) require policies, monitoring, regular testing and storage; OPS-09 explicitly requires storage at physically separate sites. NIST SP 800-34. DORA Art. 12(2). NIS2 IR No. 4 concretises backup strategy, tests and redundancy requirements.
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 restore tests. Restore tests at least quarterly. Separate IAM paths for 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 as the strictest regime. NIST SP 800-53 CP-7. C5:2026 PS-02 (Redundancy Model) requires a documented redundancy concept; BCM-03 and BCM-04 ensure implementation maturity. NIS2 IR No. 4.1 mandates recovery capacities.

A.8.15 Logging

Category	ROBUST
Mythos finding	Central, append-only-secured logs are the non-negotiable prerequisite for detection, forensics and learning. Robust as long as log integrity is cryptographically secured (WORM, Merkle chains, HSM integration). Without logs, all downstream detection is blind.
Framework comparison	NIST SP 800-92 Guide to Computer Security Log Management. C5:2026 OPS-10 to OPS-17 form the most comprehensive logging stack in the framework comparison: from the policy (OPS-10) through SIEM integration (OPS-13), retention (OPS-14) and accountability (OPS-15) to the availability of the monitoring software (OPS-17). DORA Art. 10(3). NIS2 IR No. 3.2 (monitoring and logging) with detailed minimum logging contents (3.2.3 a–l).

A.8.17 Clock synchronisation

Category	ROBUST
Mythos finding	NTP-based time synchronisation is Mythos-relevant: without precisely synchronised timestamps across all systems, correlating fragmented attack actions is impossible. Mythos attacks occur in micro-steps across distributed systems; clock skew of a few seconds destroys the ability to reconstruct.
Framework comparison	NIST SP 800-53 AU-8. C5:2026 OPS-16 (Logging and Monitoring – Configuration). NIS2 IR No. 3.2.6 mandates – where feasible – system-wide time synchronisation as a prerequisite for correlatable logs. Kerberos and other time-dependent protocols require it structurally.

A.8.18 Use of privileged utility programs

Category	ROBUST
Mythos finding	Restriction and monitoring of legitimate tools capable of overriding system controls (debuggers, admin utilities, configuration management tools) is a hard barrier. Mythos attackers preferentially use these tools (living off the land); restricting them remains fully effective.
Framework comparison	NIST SP 800-53 AC-3, SC-18. C5:2026 IAM-06 (Privileged Access Rights) addresses the use of privileged tools. Application allowlisting (CIS 2.5, 2.6).

A.8.19 Installation of software on operational systems

Category	ROBUST
Mythos finding	Application allowlisting with signature verification is a hard barrier against unauthorised software installation – including against AI-generated malware variants, provided they are unsigned.
Framework comparison	NIST SP 800-167 Guide to Application Whitelisting. C5:2026 OPS-26 (System Hardening) and DEV-10 (Approvals for Provision in the Production Environment). 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 for the

	2030+ horizon.
Framework comparison	NIST SP 800-175, FIPS 140-3. NIST PQC standards. BSI TR-02102. C5:2026 concretises cryptography in CRY-01 to CRY-19; CRY-01.01AC in particular requires a post-quantum cryptography strategy, a cryptographic inventory and hybrid cryptography models. DORA Art. 9(4)(e). CRA Annex I Part II. NIS2 IR No. 9 mandates policies for algorithms, key management and key lengths.
Hardening recommendation	Cryptographic inventory (crypto-agility audit): which algorithms, key lengths and implementations are in use where? Define a post-quantum migration path for asymmetric schemes. 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 explicitly names "design for breach" as a core idea.
Framework comparison	NIST SP 800-160 Systems Security Engineering. NIST SP 800-207 Zero Trust Architecture. C5:2026 DEV-01 (Policies for the Development / Procurement of System Components) and DEV-05 (Design Documentation for Security Features). CISA Secure by Design. OWASP SAMM. NIS2 IR No. 6 requires documented security architecture principles for acquisition, development and maintenance.

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 AI-accelerated offence. Objective: scan your own code with the same model class the attackers use, before they do.
Framework comparison	OWASP ASVS. NIST SP 800-218 SSDF. C5:2026 DEV-07 (Testing Changes), OPS-22 (Penetration Tests) and OPS-25 (Vulnerability Scans); the Additional level OPS-25.01AS tightens the scan frequency to daily. SLSA Level 3+. CRA Annex I Part II. NIS2 IR No. 6 and No. 7 (effectiveness assessment) mandate documented test processes before go-live.
Hardening recommendation	Activate pre-merge blockers for high-confidence findings. AI vulnerability scanning as a dedicated stage (isolated agent deployment, verification step, integration into the triage process). Vulnerability reports must include disclosure of AI use.

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 especially for supply chain attacks.
Framework comparison	NIST SP 800-53 CM-4, SC-2. C5:2026 DEV-11 (Protection of Development and Test Environments) and DEV-12 (Separation of Environments). SLSA Build Levels. NIS2 IR No. 6 mandates secure separation of development, test and production systems.

A.8.33 Test information

Category	ROBUST
Mythos finding	Non-use of 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, this discrepancy becomes a critical risk if production data resides there.
Framework comparison	GDPR Art. 5(1)(c) data minimisation. NIST SP 800-53 SA-15(9). C5:2026 DEV-11. NIS2 IR No. 6 implies protecting development and test data at the protection level of the respective production data.

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 most effective single measure.

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

Prioritisation note. The structural controls summarised 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 prioritises the consistent implementation of these structural controls before introducing new detection or automation capabilities (Chapter 9). The risks rated CRITICAL in Mythos-Ready (CSA, SANS, OWASP, April 2026) are already addressed to a considerable extent by the

consistent implementation of 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 considered partially degraded if its fundamental protective effect persists under Mythos conditions, but the originally assumed strength drops significantly. Such controls need not be replaced; they require 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 effect when the attack is decomposed into micro-steps that individually 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 occasional supplier audits do not keep pace with 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 works against misuse by individuals. Under Mythos, the criterion of aggregation resistance is violated: a compromised account that, through agentic automation, serves several roles in a fragmented yet seemingly legitimate manner produces temporally and contextually separated individual actions that each look policy-compliant. Without behaviour-based correlation, segregation of duties remains formally observed but no longer protects against the actual misuse.
Framework comparison	C5:2026 OIS-04 (Segregation of Duties). NIS2 IR No. 11.2.2 point (a) names segregation of duties as a principle of access granting alongside need-to-know and need-to-use. NIST SP 800-53 AC-5. DORA Art. 9 implicitly.
Hardening recommendation	Additionally deploy behaviour-based detection (UEBA) to identify compromised accounts by atypical activity patterns. Context-

	sensitive authorisation with dynamic anomaly checking. Cross-role activity within short time windows as an alert criterion.

A.5.5 Contact with authorities

Category	PARTIALLY DEGRADED
Mythos finding	Reporting channels to authorities are intact as a control, but the criterion of time compression applies: 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 recurring meetings and handwritten incident reports collide with these deadlines as soon as Mythos attacks escalate in parallel on several levels.
Framework comparison	C5:2026 OIS-06 (Contact with Relevant Government Agencies and Interest Groups). NIS2 Directive Art. 23(4) sets the reporting deadlines; the NIS2 IR concretises 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 these groups about the reporting mechanism. DORA Art. 19 requires major incident reports to the competent supervisory authorities.
Hardening recommendation	Automated reporting chain with prepared templates for different incident types. Dedicated reporting role with deputy arrangement and 24/7 availability. Integration of the reporting system with the SIEM for automatic reporting triggers.

A.5.7 Threat intelligence

Category	PARTIALLY DEGRADED
Mythos finding	Threat intelligence is indispensable under Mythos, but in its classic form degraded by time compression: TI feeds with weekly or monthly update rates structurally lag behind the iteration speed of AI-supported attacker techniques. 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 framework) 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	Deploy TI feeds with automated SIEM integration and real-time updating (MISP, STIX/TAXII). Join sector-specific sharing communities (ISAC, CERT association). AI-supported TI correlation platforms (e.g. Recorded Future, Mandiant Advantage, GreyNoise) for automated IOC enrichment and prioritisation. 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 level of the control – which data may be transferred where – is degraded by aggregation blindness: classic DLP systems do not recognise 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	Additionally deploy volume- and pattern-based egress anomaly detection. Data loss prevention with behavioural context instead of pure content inspection. Egress traffic analyses with ML-supported baseline detection.

A.5.15 Access control

Category	PARTIALLY DEGRADED
Mythos finding	The access control policy remains the foundation of all authorisation. Under Mythos, the control is degraded in two dimensions: aggregation blindness in role-based models that fail to recognise fragmented access patterns, and time compression in static access decisions that do not react to behavioural changes. 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.
Hardening recommendation	Risk-adaptive access with dynamic evaluation of identity signals, device state and behavioural context. Continuous recertification instead of quarterly reviews. 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 averaged: phishing-resistant MFA (FIDO2, passkeys, hardware tokens under WebAuthn) is Mythos-robust. Passwords, SMS MFA and time-based OTPs are considerably weakened against AI-quality phishing and credential stuffing. The control only works fully in its phishing-resistant implementation.

Framework comparison	C5:2026 IAM-08 (Authentication Mechanisms) and IAM-09 (Confidentiality of Authentication Information). NIS2 IR No. 11.2 requires strong authentication. NIST SP 800-63B Authentication Assurance Levels; AAL2 and AAL3 recommended for Mythos environments.
Hardening recommendation	Phishing-resistant MFA (FIDO2, passkeys) as mandatory for all privileged accounts and externally reachable services. Exclude SMS MFA for new deployments. Hardware tokens for administrative access.

A.5.18 Access rights

Category	PARTIALLY DEGRADED
Mythos finding	Provisioning and deprovisioning fundamentally take effect. Classic recertification in quarterly cycles is degraded by time compression: an agentic attacker uses the time between role change and recertification to exploit accumulated authorisations (permission creep). The latency between detecting an anomaly and revoking authorisations 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 prompt adjustment upon role change. NIST SP 800-53 AC-2.
Hardening recommendation	Continuous recertification instead of quarterly reviews. Automated revocation of authorisations on HR event triggers. Just-in-time authorisation for privileged access with automatic withdrawal on expiry.

A.5.19 Information security in supplier relationships

Category	PARTIALLY DEGRADED
Mythos finding	The control remains necessary but is greatly expanded by the Mythos threat picture. Supply chain attacks are faster to discover, more precise to execute and more broadly scalable with AI support. Classic supplier questionnaires at onboarding do not address the continuously changing attack surfaces of the supply chain. Time compression acts massively here.
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 standalone pillar. DORA Art. 28–30 as the strictest regime for critical ICT third-party providers. CRA Annex I requires SBOM transparency along the supply chain.
Hardening recommendation	Anchor mandatory SBOM contractually. Automated vendor security monitoring with continuous perimeter scans of suppliers. Tiering approach with differentiated requirements by criticality. Incident

	notification SLAs of at most 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 do not suffice 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 the security requirements for third parties. CRA Annex I requires contractual SBOM obligations.
Hardening recommendation	Design contract clauses to DORA standards even outside the financial sector: right to audit, incident SLAs (24 h notification, 72 h detail report), subcontracting approval requirement, exit clauses with data return and deletion evidence. Reality note: with hyperscalers (AWS, Azure, GCP), an individual right to audit is de facto unenforceable. Acceptable substitutes from an audit perspective: ISO 27001 certificate plus SOC 2 Type II, BSI C5 attestation (Type 2 report), TISAX, pooled audits by industry consortia (CSA STAR, ENISA EUCS), as well as 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 aggregation blindness. Classic audits and questionnaires capture point-in-time states, not the continuous change of ICT supply chains. Automated supply chain attacks – malicious dependencies, compromised build pipelines, injected Docker layers – structurally evade 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 SBOM mandatory. 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 useful as a process, but only takes effect when 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 (e.g. BitSight, SecurityScorecard, Black Kite, RiskRecon) for continuous vendor monitoring. Integration of threat intelligence on supplier companies. Deviations above defined score thresholds automatically trigger tickets in the SIEM or GRC system. Scorecard approach with quantified security indicators (at minimum: patch hygiene, TLS configuration, leaked credentials in public breach databases, compromised hosts in 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 itself is compromised, or an agentic attacker successfully disguises API misuse 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 audit subject matter. NIS2 IR No. 5 (supply chain security). DORA Art. 28. ENISA EUCS as the forthcoming European certification.
Hardening recommendation	Demand provider assurance under 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. Assess 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. When an agentic attacker completes the tactical work within minutes, playbooks must be

	designed for second- and minute-scale time 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 800-61r2.
Hardening recommendation	Playbook revision for minute-scale time 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 per cent autonomously and triggers actions in seconds, a human SOC analyst is often left with too little time for an informed decision. The control only becomes robust when defined escalation levels 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); the subsequent review and lessons learned are regulated separately in No. 3.6. DORA Art. 17–18. NIST SP 800-61r2.
Hardening recommendation	SOAR platform for automated containment actions (account lock, network isolation, key rotation). Mythos-specific playbooks: mass data exfiltration, parallel multi-vector attack, supply chain compromise. Dedicated threat hunter function.

A.5.29 Information security during disruption

Category	PARTIALLY DEGRADED
Mythos finding	The control addresses maintaining security measures in contingency operations. Under Mythos it is doubly degraded: first, parallel attack scenarios that classic BCM plans do not assume. Second, the possibility that the attacker has deliberately provoked the BCM case and targets the contingency 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 tabletop exercises. 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	Mere retention is Mythos-orthogonal; integrity assurance is threatened by agentic manipulation capabilities. An attacker with infrastructure access can retroactively alter or selectively 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 protection (WORM, Object Lock). Hash chains for manipulation detection. Separate IAM paths for read and write access to archives. Separate retention 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 minimisation becomes the most important lever because it structurally limits the blast radius of a successful exfiltration. Data never 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 minimisation technically (schema validation, ORM rules, API response filtering). Purpose limitation as a mandatory attribute check in data access. Pseudonymisation as the default in analytics and ML pipelines. Regular privacy impact assessments with 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 useful as a governance mechanism, but no longer provides a timely effectiveness assessment. Between two audits, the attack surface, TTPs and control effectiveness may 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 via 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 employees for Mythos-generated social engineering quality. AI-generated phishing messages are grammatically, contextually and stylistically barely distinguishable from legitimate communication. Deepfake-based vishing and impersonation attacks bypass the recognition patterns of 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 instead of sender recognition (out-of-band verification for financial instructions, credential resets).

A.6.7 Remote working

Category	PARTIALLY DEGRADED
Mythos finding	Remote working policies remain necessary, but are degraded by the extended attack surface of private networks, unmanaged devices and weaker physical security. Mythos attackers can exploit these weaknesses in automated fashion. The control only becomes robust 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 allowlisting.

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 user perception threshold: API misuse, credential stuffing on service accounts, fragmented data exfiltration. The user sees nothing they could report.
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 mandates informing these groups about the mechanism.
Hardening recommendation	Focus on automated detection instead of user reporting. Simple one-click reporting in the mail client for suspicious messages. Feedback channel for the user so that reports are taken seriously and reporting behaviour is encouraged. Mythos-specific training focuses: explicitly train employees to report unusual authentication indicators – unexpected MFA push notifications without an own login attempt, unknown active sessions in the account overview screen, unexplainable 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 this dimension it remains effective. Under Mythos, remote compromise of mobile devices is the greater risk, which is only partially covered by A.7.9 (otherwise A.8.1). The physical level 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-secured key management. Remote wipe capability via MDM. Tracking capability (find my device). DLP at endpoint level 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 strongly degraded against AI-generated malware variants: the ability to automatically

	produce functionally equivalent but signature-resistant mutated malware devalues the signature paradigm. EDR with behavioural analysis remains robust. The control as a whole is partially degraded because implementation depth determines 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 allowlisting instead of blocklisting. Hardware-based identity (TPM, Secure Enclave) as a prerequisite 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 conforming 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 concerns read access to large data sets – 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 tiering of queries based on access history.

A.8.4 Access to source code

Category	PARTIALLY DEGRADED
Mythos finding	Source code access is formally secured given strong IAM protection. Under Mythos, mass exfiltration through compromised developer accounts is a significant risk: wholesale cloning of entire repositories via agentically automated Git operations leaves no anomalous individual events, only patterns recognisable in aggregate.
Framework comparison	C5:2026 DEV-11 (Protection of Development and Test Environments) and IAM-06 (Privileged Access Rights). NIS2 IR No. 6 (security 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 of 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 under 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 for cloud services. NIS2 IR No. 11.2. NIST SP 800-63B; AAL2 or AAL3 for Mythos-relevant systems. FIDO Alliance standards.
Hardening recommendation	Phishing-resistant MFA (FIDO2, passkeys, hardware tokens under WebAuthn) as 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 strongly degraded under Mythos because AI assistance has dramatically reduced the cost of producing functionally equivalent malware variants. Behaviour-based detection and sandbox analysis are robust. The classification of the overall control depends on the chosen implementation depth.
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.
Hardening recommendation	EDR with behaviour-based detection and ML engine (e.g. 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 (e.g. Joe Sandbox, ANY.RUN, Hatching Triage). Application allowlisting 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, inconspicuous transfers over extended periods remains 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 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 that many ISMS implement the control classically and thereby execute 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 ML-supported baseline. Kill-chain-oriented detection instead of an isolated alert focus. Correlation of fragmented individual events over extended periods. Integration of MITRE ATT&CK as a detection framework with measurable coverage (validation via DeTT&CT or Atomic Red Team): target value at least 60 % coverage of the top 10 techniques of one's own industry. Detection focuses against Mythos TTPs: living-off-the-land binaries (PowerShell with Base64 encoding, schtasks persistence, Office app spawning cmd or powershell, unusual curl/wget calls from service accounts), beacon-less C2 (DNS tunnelling, HTTPS C2 with long sleep intervals, cloud-API-based C2 channels such as misused SaaS integrations). Dedicated threat hunting function with documented methodology (PEAK framework or TaHiTI), minimum capacity 0.5 FTE up to 500 employees, 1 or more FTE above; at least twelve hypothesis-based hunts per year with ATT&CK mapping. Practicability note: for organisations without a dedicated SOC, managed detection & response (MDR) services (e.g. 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 spaces that can be reconnoitred automatically. 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 documented network architecture, remote access control and deactivation of unneeded services; No. 11 (access control) and No. 8 (cyber hygiene) supplement this. NIST SP 800-207 Zero Trust Architecture.
Hardening recommendation	Maturity path to a 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 trust pivot between internal zones, identity-aware proxy (e.g. 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 board.

A.8.21 Security of network services

Category	PARTIALLY DEGRADED
Mythos finding	Perimeter-oriented protection of network services is degraded. Identity-based protection via mTLS, SPIFFE identities and service mesh is robust. The control depends on the chosen architectural discipline – 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 allowlisting.

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 not effective. Micro-segmentation at identity level instead of 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 (e.g. 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 occasional security review before releases is partially degraded. The speed of AI-supported attack vector development requires security testing in every CI iteration.
Framework comparison	C5:2026 DEV-01 to DEV-15 cover the entire SDLC, incl. DEV-04 (Safety Training), DEV-05 (Design Documentation for Security Features) and DEV-06 (Risk Assessment). NIS2 IR No. 6 (security 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	Formalised security requirements remain necessary, but only take effect through automated verification. Without linkage 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	Formalise security requirements as executable tests (security as code). Automated compliance checks in CI. Threat modelling as a requirements source with STRIDE or PASTA. OWASP ASVS as the base 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 automatically.
Framework comparison	C5:2026 DEV-04 (Safety Training Regarding Continuous Software Delivery). NIS2 IR No. 6 (security 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 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	Mandatory SBOM for all delivered artefacts. SLSA Level 3+ for build provenance. Contractually fix code audit rights. Automatic artefact verification in one'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 multi-week approval cycles 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 (at most 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. Concerns A.8.7, A.8.12, A.8.16 and additionally 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. Concerns A.5.17, A.8.5, A.8.21 and additionally A.6.7, A.8.1, A.8.4.
- Identity-based network architecture (Zero Trust): micro-segmentation, identity-aware firewalls, service mesh. Concerns A.8.20, A.8.21, A.8.22 and additionally A.5.15, A.8.3, A.6.7.
- Automated, continuous security testing: AI-supported pre-commit scanning, SAST/DAST in CI, SBOM pipelines. Concerns A.8.25, A.8.26, A.8.28, A.8.30, A.8.32 and additionally A.5.19, A.5.21.
- Supply chain transparency and continuous vendor monitoring: SBOM, SLSA, automated vendor security monitoring. Concerns 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 recurringly in the hardening recommendations of the individual controls and consolidated in Chapter 9 into the synthesis of the Mythos Hardening Controls.

6 Friction-Only – Controls That Must Be Replaced or Fundamentally Redesigned

6.1 Overview and Risk Profile

A control is considered friction-only if its core effect against Mythos attackers is structurally eliminated – either because it rests on limited attacker capacity or because it presupposes human response time in an automatically traversed kill chain (see definition in Ch. 3.1.3). The core statement of the Anthropic security engineering team formulated 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 unites the four controls is that their protective effect stems neither from structural impossibility nor from cryptographic hardness, but either from the assumption of limited attacker capacity or from the assumption of humanly manageable response times. Both assumptions have been empirically refuted under Mythos conditions. The controls must be replaced in their fundamental logic or fundamentally redesigned.

Important for context: The 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, they no longer offer resilient protection and must not be entered in any risk assessment as an effective mitigation measure for Mythos-relevant risks.

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 per cent of the tactical attack steps at request rates physically unattainable for humans. A triage chain that makes 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 provides an effective response barrier – human response time in the automatically traversed kill chain is the structural weakness per Ch. 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) mandates documented triage criteria and the correlation/analysis procedure for logs. NIST SP 800-61r2. DORA Art. 17.
Replacement recommendation	Replace human first-line triage with AI-supported tier-1 automation. 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-damage scenarios. Reposition the SOC as a threat hunting and incident commander function, not as alert review. Measurable SLA:

	mean time to containment under 10 minutes for high-confidence alerts. Practicability note: SOAR platforms (Splunk SOAR, Palo Alto XSOAR, Microsoft Sentinel with Automation Rules) require dedicated engineering for playbook maintenance. For organisations without a 24/7 SOC, managed detection & response (e.g. CrowdStrike Falcon Complete, Arctic Wolf, SentinelOne Vigilance, Sophos MDR) with a contractually anchored MTTC SLA is the more viable 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 recognisable as such. An agentic attacker who decomposes an attack into micro-steps that each look policy-compliant structurally bypasses this check. The aggregation blindness of classic compliance checks – quarterly reviews, policy comparisons against configuration baselines – makes them pure friction: they create effort for the attacker, but no protection when 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 (compliance monitoring) requires effective reporting and regular conformity checking. DORA Art. 6(5) and Art. 24.
Replacement recommendation	Continuous control monitoring instead of 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 policy-compliant but behaviourally anomalous activities. Compliance dashboards with quantified deviation indicators instead of PDF audit reports. Linkage with an automated incident ticket upon 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 meetings and triple-staggered test phases is, under Mythos, pure friction against the collapse of the patch-gap window described in Chapter 2.1. When only hours lie between the release of a patch and a working exploit, a process with multi-week approval cycles is not a protective measure but a structural weakness. The existence of an orderly, documented vulnerability management process produces compliance evidence, but no Mythos resistance.
Framework comparison	C5:2026 offers the most granular benchmark: 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 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 prioritisation (Exploit Prediction Scoring System) instead of pure CVSS weighting. 24-hour SLA for KEV listings (CISA Known Exploited Vulnerabilities). Automated patch pipelines with canary deployments, automatic rollback and security regression tests. 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 – the attacker having to register more phishing domains – is no 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 blocklists. Remote browser isolation (RBI) for all external links, especially 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 works even when the user reaches a phishing page.

6.3 Summary: Why Friction Fails under Mythos

The four controls of this chapter share a common logic: their protective effect arises either from the assumption that an attack would involve disproportionate effort for the adversary (A.8.8, A.8.23, A.5.36) or from the assumption that the response chain runs within humanly manageable time (A.5.25). Both assumptions are undermined by Mythos-class AI: the effort is borne by the model, whose marginal cost of an additional iteration is close to zero, and the attacker's response chain is faster than any human triage.

It is remarkable that three of the four friction controls – A.5.25, A.8.8 and A.8.23 – were each considered central pillars of operational cyber defence over the past two decades: SOC triage,

vulnerability management and web filtering. The recommendation is not to dismantle this infrastructure, but to supplement it with structural Mythos hardening as described in the respective replacement recommendations.

For the Statement of Applicability of an ISO 27001-certified ISMS, the classification has a clear consequence: these four controls should not be recorded as the sole mitigation for Mythos-relevant risks. The risk assessment must document which supplements from the categories Robust and Partially degraded assume the actual protective effect.

7 Not Affected Controls – Neutral Base

7.1 Controls without Mythos Relevance

23 controls of ISO/IEC 27002:2022 are neutral towards the Mythos-specific threat landscape. Their effectiveness is systematically altered neither by attacker speed nor by fragmentation, neither by capability decoupling nor by the collapse of the patch-gap window. These are predominantly organisational, documentation-oriented, governance or physically bound controls.

These controls remain important for the overall ISMS structure and for regulatory compliance. They receive no separate Mythos hardening recommendation because their classification remains unchanged against the Gen-AI-accelerated threat landscape.

7.2 Overview of the 23 Not Affected Controls

ID	Title	Brief 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 responsibility 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	Identification of legal, statutory, regulatory and contractual requirements	Compliance inventory.
A.5.32	Intellectual property rights	Licence and IP management.

ID	Title	Brief rationale
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.
A.6.6	Confidentiality or non-disclosure agreements	Legal artefact.
A.7.5	Protecting against physical and environmental threats	Environmental protection against fire, water, 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. DoS defence is covered by A.8.14 redundancy.
A.8.34	Protection of information systems during audit testing	Operational audit artefact to avoid audit side effects.

7.3 Summary: Why They Still Matter

The classification as not affected expressly does not mean that these controls can be neglected. They form the governance and infrastructure base on which all Mythos-relevant controls operate. Without documented policies (A.5.1), without clear role assignment (A.5.2), without legal compliance (A.5.31) and without physically protected infrastructure (A.7.5, A.7.11, A.7.12), every technical Mythos hardening loses its foundation.

Rather, the assessment makes clear: these controls make their unchanged and important contribution, but they are not the place where CISOs should prioritise additional investment in response to the Mythos threat. The focus belongs on the controls of the categories Robust (consistent implementation) and Partially degraded (concrete supplementation).

Part III – Convergence Analysis: Gaps and Synthesis

The individual assessment of all 93 controls in Part II shows which controls hold up, which must be supplemented and which must be replaced. Part III systematically identifies which requirements other regulatorily relevant frameworks impose that ISO 27002:2022 does not explicitly know and that deliver resilient protection under Mythos conditions. The result is a prioritised supplementary catalogue of concrete controls.

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

8 Gap Analysis – What Other Frameworks Require That ISO 27002 Does Not Cover

8.1 Gap Analysis Methodology

The gap analysis proceeds in three stages. First, all requirements of the five primary comparison frameworks – C5:2026, NIST, DORA, CRA and NIS2 with Implementing Regulation – are reconciled against the ISO 27002:2022 catalogue. Second, those requirements are extracted that ISO 27002 does not know or knows only very abstractly. Third, each extracted requirement is checked for whether it delivers a hard protective effect under Mythos conditions – i.e. fulfils one of the criteria from Chapter 3.2: cryptographic impossibility, structural reduction of the blast radius, aggregation resistance or a hard time barrier.

Requirements that are merely a different formulation of controls already present in ISO 27002 are not counted as a gap. Nor are requirements that are more detailed in another framework but are themselves degraded under Mythos conditions. The following analysis is limited to substantive differences from the ISO text combined with Mythos robustness.

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 handling of 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 bindingly requires a documented post-quantum cryptography 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) operationally anchors crypto-agility.

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 operationalised. Second: crypto-agility is a structural prerequisite for reacting to suddenly devalued algorithms without falling into weeks-long re-engineering cycles.

8.3 Cluster 2: Supply Chain Transparency and Mandatory SBOM

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 deployed software components, no SBOM obligation and no binding build provenance requirement.

Requirement in other frameworks: C5:2026 DEV-13 (Transparency about Software Components). EU CRA Annex I Part II makes mandatory SBOM a statutory component of CE marking. DORA Art. 28 requires detailed transparency for ICT third-party providers. NIS2 IR No. 5 (supply chain security) mandates documented dependencies. The SLSA framework defines build provenance levels.

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

8.4 Cluster 3: Containers, Confidential Computing and Multi-Tenancy

With the March 2026 revision, C5:2026 introduced three structurally new control domains that ISO 27002 does not know 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), as well as 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 a reference.

Mythos robustness: Container signing creates a cryptographic-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 concepts on trusted computing as a 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 Reporting Deadlines and Root Cause Analysis

ISO 27002 reference: A.5.5 (contact with authorities) and A.5.25 (assessment of events) address reporting channels without concrete deadlines. A.5.27 (learning from incidents) imposes no root cause analysis obligation.

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 lifecycle. NIS2 IR No. 3.6 mandates post-incident reviews including cause determination.

Mythos robustness: Time-based reporting deadlines are Mythos-robust because they mirror, on the regulatory side, the time compression that benefits Mythos attackers: anyone who must report within 24 hours needs corresponding automation on the detection and triage side. Root cause analysis creates the empirical basis for detection improvement.

8.6 Cluster 5: Continuous Assurance instead of 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 a sharpening. COM-03 and COM-04 require continuous ISMS effectiveness checking. NIS2 IR No. 2.2 (compliance monitoring) requires an effective reporting system. DORA Art. 24–26 requires a Digital Operational Resilience Testing Programme; Art. 26–27 threat-led penetration testing (TLPT) for critical functions.

Mythos robustness: Continuous assurance offsets time compression on the defence side. Detection latency is reduced from months and quarters to hours and days. TLPT with Mythos-specific scenarios is the best-known regulatory regime that explicitly addresses the empirical level: 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 protection needs, 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: Mandatory Automation and Resilience Testing

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

Requirement in other frameworks: NIS2 IR No. 3.2.2 mandates automated monitoring where 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 reflecting real threats. In the Glasswing defensive post, Anthropic explicitly recommends tabletop exercises for three to five parallel incidents.

Note on the strictness of the NIS2 IR wording: The IR formulates the automation obligation as a qualified requirement ("where feasible", "subject to operational capacities") and not as an absolute obligation. For Mythos-exposed organisations, feasibility should be interpreted 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 defence 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 built into the design of ISO 27002 as a deficiency – it reflects the framework-by-framework specialisation of the standards landscape.

Under Mythos conditions, this delta becomes an operationally relevant problem. An organisation relying on the pure ISO 27002 base does not have the concrete requirements that C5:2026, NIST, DORA, CRA and NIS2 with Implementing Regulation impose on 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 modernisation of the security architecture (Cluster 1 cryptography, 3 cloud infrastructure, 6 identity) and a temporal shortening of the defence loops (Cluster 2 supply chain, 4 reporting obligations, 5 continuous assurance, 7 automation). Chapter 9 brings these movements together into the prioritised catalogue of Mythos Hardening Controls.

9 Synthesis – Catalogue of the Mythos Hardening Controls

9.1 Purpose and Application of the Catalogue

This chapter compactly consolidates the findings from Chapter 8 into a 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 formulated so that it can be adopted directly into the Statement of Applicability of an existing ISO 27001-certified ISMS.

The thirteen MHC are a prioritised supplementary catalogue for the Mythos threat landscape. They are deployed in addition to the existing ISO controls, but address protection objectives that ISO does not explicitly know or describes only abstractly. Each MHC references an ISO 27002 linkage.

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 limited to the currently most robust and operationally clearly describable supplements. The version notes document the development of this catalogue.

Prioritisation 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). The structural controls remain the most robust defence layer under Mythos. An organisation that has not consistently implemented these controls gains less Mythos resilience by introducing new MHC than by hardening what exists. The assessment language in Chapters 5 and 6 must therefore be read as "selectively degraded" – not as "obsolete". Classic controls are weakened by Mythos in specific effectiveness dimensions, but retain their full protective effect in other dimensions. The correct response is hardening, not replacement.

9.2 Thirteen Mythos Hardening Controls

MHC-01 Post-quantum cryptography strategy and cryptographic inventory

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: 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 finalised PQC standards (13 August 2024); HQC as backup KEM (selected 11 March 2025, standardisation ongoing). NIST SP 1800-38 as the NCCoE migration guide. The EU Commission's roadmap for post-quantum migration.
Hardening recommendation	Implementation: Documented PQC strategy with four elements: (1) a cryptographic inventory of all deployed algorithms, key lengths and implementations, prioritised by impact and likelihood of

	quantum attacks; (2) ongoing observation of the state of the art; (3) use of hybrid cryptography models; (4) defined trigger events, resources, transition paths and success criteria. Annual or event-driven review. Mythos effect: addresses the store-now-decrypt-later threat model and creates structural crypto-agility. Protective effect rests on cryptographic impossibility, not attacker friction.

MHC-02 Software bill of materials and build provenance

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Flanks A.5.21 (ICT supply chain management) and A.8.30 (outsourced development).
Framework comparison	Framework basis: C5:2026 DEV-13. CRA Annex I Part II as the legally binding SBOM obligation (vulnerability reporting 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 benchmark.
Hardening recommendation	Implementation: SBOM generation automatically in the CI pipeline (SPDX, ISO/IEC 5962, or CycloneDX, ECMA-424; CRA-compliant at least SPDX 3.0.1+ or CycloneDX 1.6+) for all self-developed and distributed software artefacts. SBOM requirement contractually with all critical software suppliers. Automated vulnerability correlation against CVE, EUVD and OSV; communication of exploitability separately via VEX or CSAF (the SBOM contains no vulnerability data). SLSA Level 3+ for build provenance of critical artefacts. Mythos effect: creates structural detection capability for supply chain attacks. Detection latency is reduced to the publication time in vulnerability databases.

MHC-03 Phishing-resistant multi-factor authentication

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Flanks A.5.17 (authentication information) and A.8.5 (secure authentication).
Framework comparison	Framework basis: NIST SP 800-63B Revision 4 (final 31 July 2025), AAL2/AAL3 – AAL3 requires device-bound, non-exportable keys; synchronisable passkeys only AAL2. C5:2026 IAM-08. FIDO Alliance standards (WebAuthn, passkeys, CTAP2). NIS2 IR No. 11.2.
Hardening recommendation	Implementation: Phishing-resistant methods under the WebAuthn standard for all privileged accounts, externally reachable services and access to systems with elevated protection needs. Passkeys or hardware tokens as the minimum form. Exclusion of SMS MFA and push notifications without number matching for new deployments. Replace existing legacy MFA methods with a transition plan. Mythos effect: structural answer to AI-quality phishing: origin binding in WebAuthn renders authentication credentials on phishing pages worthless.

MHC-04 Workload identity and Zero Trust network architecture

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Flanks A.5.16 (identity management), A.8.20 (networks security) and A.8.21 (security of network services).
Framework comparison	Framework basis: NIST SP 800-207 Zero Trust Architecture. NIST NCCoE work on agent identity (2026). SPIFFE/SPIRE. Service mesh architectures (Istio, Linkerd).
Hardening recommendation	Implementation: Maturity path in three stages, instead of a big-bang migration. Stage 1 (privileged workloads): SPIFFE/SPIRE identities for service accounts with extended rights, mTLS on the three to five most critical service-to-service paths. Stage 2 (production mainstream): workload identity for all productive microservices, ZTNA for privileged remote access. Stage 3 (full coverage): workload identity for dev and test environments, identity-based micro-segmentation. AI agents as a special case: coding agents (e.g. Claude Code, Cursor, Copilot Workspace) and agentic workflows require dedicated workload identities with time-limited, capability-scoped authorisations per tool call – not personal developer credentials. Complete audit trail of all agent actions with correlation to the initiating human identity. Mythos effect: structural elimination of classic lateral movement paths. Shared secrets are replaced by hardware-anchored, non-reusable identities.

MHC-05 Behaviour-based detection and kill chain correlation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Flanks A.8.16 (monitoring activities) and A.8.12 (data leakage prevention).
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). NIST SP 800-94 only as dated background (2007 edition; the draft of a Revision 1 was withdrawn, no current successor exists).
Hardening recommendation	Implementation: UEBA with ML-supported baseline detection and anomaly alerting. MITRE ATT&CK-based detection rules instead of signature-oriented individual alerts, with measurable coverage (target value: at least 60 % coverage of the top 10 techniques of one's own industry, measured via DeTT&CT or Atomic Red Team validation). Kill chain correlation across multiple events and time windows. Detection focuses against Mythos TTPs: living-off-the-land binaries (PowerShell encoding, schtasks persistence, Office spawning cmd or powershell, unusual curl/wget calls), beacon-less C2 (DNS tunnelling, HTTPS C2 with long sleep intervals, misused cloud API channels). Threat hunting as a standalone function with a documented methodology (PEAK framework or TaHiTI), minimum capacity 0.5 FTE up to 500 employees, 1 or more FTE above; at least twelve hypothesis-based hunts per year with ATT&CK

	mapping. Adversarial robustness of one's own detection AI: test ML models against evasion attacks (adversarial examples) and data poisoning. Practicability note: for organisations without a dedicated SOC, managed detection & response services (CrowdStrike Falcon Complete, Arctic Wolf, SentinelOne Vigilance, Sophos MDR) are a valid alternative. Mythos effect: addresses the aggregation blindness of classic detection. Fragmented micro-step attacks become recognisable through pattern correlation. Realism note: detection is not prevention. MHC-05 shortens detection latency and provides the data basis for downstream containment actions (MHC-11), but does not replace a structural hard barrier. In particular, low-and-slow attacks with action intervals above the correlation time windows, well-camouflaged AI agents within legitimate behavioural baselines, and adversarial ML evasion against one's own detection engine remain residual risks. The effectiveness of MHC-05 must always be assessed in combination with the structural controls from Chapter 4 – not as a replacement.

MHC-06 Container security and confidential computing

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: No direct ISO linkage. Indirectly flanking A.8.31 and A.8.27.
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	Implementation: For containers: signed base images from controlled registries (image signing via Sigstore/cosign), runtime protection with admission controller policies (e.g. OPA Gatekeeper, Kyverno), automated image scans before deployment. For confidential computing: use of TEEs or secure enclaves (e.g. Intel TDX, AMD SEV-SNP, ARM CCA) for workloads with elevated protection needs, with remote attestation. Mythos effect: container signing creates a cryptographic hard barrier against manipulated images. Confidential computing is currently the hardest barrier against attackers with provider infrastructure access.

MHC-07 Multi-tenancy isolation with demonstrable separation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: No direct ISO linkage. Extends A.8.22 with data- and data-centre-specific aspects.
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	Implementation: Documented separation concepts for data (tenant-specific keys, logical or physical separation), for networks (tenant-

	specific VPCs, namespaces, software-defined networking rules) and for compute resources (tenant-specific scheduling policies). Demonstrable separation through regular, ideally automated tests. Mythos effect: structural limitation of the blast radius: an attacker who gains a foothold in one tenant is demonstrably prevented from accessing other tenants.

MHC-08 Immutable backups and recovery validation

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Deepens A.8.13 (information backup), which is fundamentally robust but does not explicitly formulate concrete immutability requirements.
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	Implementation: Implementation of the 3-2-1-1-0 model: three copies on two media types, at least one offsite and at least one immutable or offline, zero errors in the periodic restore test. Separate IAM paths for backup infrastructure. Quarterly restore tests with documentation. Mythos effect: immutable backups are among the most effective hard barriers against ransomware and agentic data integrity attacks.

MHC-09 AI-supported security testing in the pipeline

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Flanks A.8.29 and A.8.25.
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) as a draft (17 December 2025). OWASP ASVS. CRA Annex I Part II.
Hardening recommendation	Implementation: 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). Pre-merge block for high or critical findings with confidence of 80 % or higher; KEV findings immediately blocking. Daily vulnerability scans of running systems. Regular penetration tests with Mythos-specific scenarios (fragmented attacks, AI-supported exploit combinations). Address the secondary threat: AI coding assistants (Copilot, Cursor, Claude Code) regularly generate insecure code patterns – hardcoded credentials, missing input validation, insecure defaults, outdated dependencies. Implement pre-commit hooks with AI code audit

	rules (e.g. Snyk or Semgrep rule sets against AI-typical anti-patterns). Weekly AI code audit reporting with trend tracking. Designate vibe-coded code as a dedicated risk category in the Statement of Applicability. Mythos effect: the shift-left principle removes the attacker's time window between exploit availability and production patch. Structural answer to the collapse of the patch-gap window.

MHC-10 Continuous control monitoring and policy as code

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Largely replaces the effectiveness of A.5.35 and A.5.36 under Mythos conditions.
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	Implementation: Policy as code: formalise security policies as executable rules in OPA, Sentinel or comparable engines. Integration into CI (pre-deploy checks) and runtime enforcement. Automated, continuous compliance checks against baselines. Real-time compliance dashboards with quantified deviation indicators. Deviations automatically trigger incident tickets. Mythos effect: shortening of detection latency for compliance deviations from months or quarters to seconds or minutes.

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

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Largely replaces the effectiveness of A.5.25, supplements A.5.24 and A.5.26.
Framework comparison	Framework basis: C5:2026 SIM-02 and SIM-03. NIS2 IR No. 3.5 (incident response, three phases per 3.5.2). DORA Art. 17. NIST SP 800-61 Revision 3 (incident handling 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	Implementation: SOAR platform (Splunk SOAR, Palo Alto XSOAR, Microsoft Sentinel with Automation Rules) with predefined playbooks for fully automatic containment actions on unambiguous signals. Human escalation targeted at ambiguous cases. Mythos-specific playbooks: mass data exfiltration, parallel multi-vector attack, supply chain compromise, credential stuffing wave, MFA fatigue campaign. Tabletop exercises for three to five parallel incidents, at least every six months. Align the SOC role towards threat hunting and the incident commander function. Hardening of the automation layer itself: the SOAR pipeline becomes an attack

	surface of its own if trigger sources are not authenticated, playbooks become reverse-engineerable, or false-positive floods unintentionally trigger containment actions. Protective measures: signed trigger events with authentication of the detection source, complete audit trail of every playbook execution with execution identity, rate limits per playbook, manual override (human in the loop) for actions with a high blast radius (mass account lock, network isolation of larger segments, certificate revocation), protection of playbook definitions like source code (versioning, code review, signed releases). Practicability note: self-operated SOAR requires a dedicated engineering team for playbook maintenance. For organisations under 1,000 employees without a dedicated security team, managed detection & response services with a 24/7 SOC and a contractually anchored MTTC SLA are the more viable alternative. Mythos effect: structurally closes the time offset between agentic attacker speed and the human response chain. MTTC is reduced from hours to minutes. Effectiveness limit: the SOAR logic itself can become a target – attackers can reverse-engineer trigger conditions, deliberately provoke false positives or turn containment actions against the organisation itself. MHC-11 is therefore not effective without the hardening of the automation layer described in the implementation.

MHC-12 Threat-led penetration testing with Mythos scenarios

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Supplements A.5.35 and A.8.29.
Framework comparison	Framework basis: DORA Art. 26 and 27 as the currently only binding regulatory TLPT regime, concretised 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 Bundesbank/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	Implementation: Conduct threat-led penetration tests with external red teamers operating on real threat scenarios. Include Mythos-specific scenarios: AI-supported spear phishing with deepfake voice, fragmented attack chains via micro-steps, parallel multi-vector attacks, supply chain compromises, cloud API misuse via compromised service accounts. Regular repetition (at least annually for critical functions). Purple team exercises between TLPT cycles with ATT&CK coverage mapping. Practicability note: full TLPTs under the TIBER-EU methodology typically cost 200 to 500 k€ per exercise. For non-DORA addressees, more cost-effective formats are sufficiently effective: purple team exercises with MITRE ATT&CK scenarios, open-source adversary emulation (Caldera, Atomic Red Team) and breach-and-attack simulation platforms (AttackIQ, SafeBreach, Picus Security). Mythos effect: TLPT tests whether controls work against real Mythos TTPs, not whether they are documented. Reduces the risk of false-positive compliance assurances.

MHC-13 AI agent governance and harness security

Category	MYTHOS HARDENING CONTROL
Mythos finding	ISO 27002 linkage: Extends A.5.16 (identity management), A.8.27 (secure system architecture) and MHC-04 with the specific governance and security requirements of agentic AI systems in one's own operations.
Framework comparison	Framework basis: OWASP LLM Top 10 (in particular LLM01 Prompt Injection, LLM06 Excessive Agency). OWASP Agentic Security Initiative (ASI01 to ASI10). MITRE ATLAS (AML.T0051 LLM Prompt Injection, AML.T0053 LLM Plugin Compromise, 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	Implementation: Governance of AI agents in five dimensions. (1) Harness audit: prompts, tool definitions, retrieval pipelines and escalation logic are treated with the same code review discipline as production code. Versioning in the source repository, signed releases, pre-production pen tests against prompt injection and tool confused-deputy attacks. (2) Blast radius limits: capability scoping per tool call (write/read/network with explicitly defined scope), rate limits per agent identity, circuit breakers on unusual action sequences, mandatory human approval thresholds for irreversible actions (production deploy, data delete, spend > X €). (3) Human override mechanisms: every agent session has a human owner with a kill switch; audit trail of all agent actions with correlation to the initiating human identity, retention at least 12 months. (4) Supply chain inventory for agentic components: MCP servers (inventory, source, update channel), VS Code and JetBrains extensions with AI functionality (allowlist, auto-update control), agentic skills and rules files (versioning, code review like source code), external agent frameworks (e.g. raptor, AutoGen, LangGraph). (5) Pre-production checks before agent deployment: documented scope definition, threat model for the agent use case, rollback plan, monitoring plan with defined anomaly thresholds. Audit thresholds: 100 % of production-adjacent coding agents have a documented threat model and a defined blast radius; audit trail coverage 100 % for actions with write rights or network access; mean time to revoke compromised agent identities under 5 minutes. Mythos effect: closes the gap rated CRITICAL in Mythos-Ready (CSA, SANS, OWASP, April 2026), "Unmanaged AI Agent Attack Surface" (privileged AI agents outside existing control frameworks). Addresses both the defensive risk side (insecure, privileged agents in one's own environment) and the supply chain risk side (compromised MCP servers, VS Code extensions, agentic skills). Maturity realism: MHC-13 is conceptually mature and clearly defined in its requirements, but operational implementation readiness in most organisations currently sits at the Initial level. Realistic transition periods: 12 to 18 months to the Defined level, 18 to 24 months to the Managed level. In the first phase, organisations should prioritise inventorying, audit trail and capability scoping – these three building blocks already address the largest part of the risk. Full harness pen tests, adversarial robustness testing of the agent models and automated pre-production checks follow in later phases, depending on market availability of suitable tools.

9.3 Derivation from the Gap Clusters

The thirteen MHC can be traced back to the seven gap clusters from Chapter 8:

- 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 (reporting obligations and root cause): MHC-11 (response side); reporting deadline compliance is not listed as a separate MHC because it must be implemented anyway as a direct legal obligation under NIS2 and DORA.
- Cluster 5 (continuous assurance): MHC-10 and MHC-12.
- Cluster 6 (phishing-resistant identity and Zero Trust): MHC-03, MHC-04 and MHC-13 (identity lifecycle of the AI agents).
- Cluster 7 (automation and resilience testing): MHC-05, MHC-09 and MHC-11.

MHC-08 (immutable backups) is a deepening of A.8.13 and, as a structural resilience measure, sits thematically at the intersection of Clusters 3 and 4. MHC-05 is the operational foundation 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 summarises the thirteen MHC and assigns to each MHC its primary framework basis and its ISO 27002 linkage. The linkages are consistent with the assessment matrix in Annex A and the individual assessments in Chapters 4 to 6.

MHC	Title	Framework basis	ISO linkage
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.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.8.4, 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.8.16
MHC-06	Containers and confidential computing	C5:2026 OPS-32 to OPS-35	A.8.27, A.8.31, A.5.23
MHC-07	Multi-tenancy isolation	C5:2026 OPS-30/31, PSS-10	A.8.22, A.5.23

MHC	Title	Framework basis	ISO linkage
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, ASI01–ASI10, MITRE ATLAS, ISO 42001 Annex A	A.5.16, A.8.27, supplements MHC-04

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

9.5 Maturity Levels per MHC

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

Assessment grid: Initial = ad hoc or documented, Defined = standardised and introduced, Managed = automated, measured and continuously improved.

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	Mandatory FIDO2 for all externally reachable services	Passwordless authentication organisation-wide, SMS MFA switched off
MHC-04	SPIFFE/SPIRE for privileged workloads, mTLS on the 3-5 most critical paths	Workload identity for all productive microservices, ZTNA for privileged access	Workload identity across the board incl. dev/test, AI agents with capability-scoped identities
MHC-05	MITRE ATT&CK as	≥ 60 % coverage of top 10	Adversarially robust ML

MHC	Initial	Defined	Managed
	detection framework, < 30 % coverage	techniques, documented threat hunting	detection, continuous ATT&CK coverage validation
MHC-06	Container images signed, vulnerability scan before deployment	Runtime protection with admission controllers, documented TEE use cases	Confidential computing in production for highly sensitive workloads, remote attestation automated
MHC-07	Logical tenant separation documented	Tenant-specific keys, network/compute isolation	Demonstrable separation through automated tests, regular audits
MHC-08	Backups in place, restore tests annually	3-2-1-1-0 model implemented, quarterly restore tests	Immutable backups with separate IAM paths, automated integrity checking
MHC-09	SAST/DAST/SCA in CI	AI-supported code scanning + pre-merge block on high findings	Daily vulnerability scans, AI code audit for vibe-coded code, automated KEV hotfix channel
MHC-10	Compliance dashboards with manual maintenance	Policy as code in CI, real-time compliance checks against baselines	Runtime enforcement, automated incident tickets on deviation
MHC-11	IR playbooks documented, manual execution	SOAR with partially automated playbooks, MTTC < 60 minutes	Fully automatic tier-1 containment, MTTC < 10 minutes, six-monthly multi-incident tabletops; alternatively MDR service with 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 ≥ 90 % within the SLA window
MHC-13	AI coding agents documented, manual inventory of MCP servers and extensions	Capability-scoped identities + audit trail in production, documented threat model per agent use case, allowlist for extensions and MCP servers	Complete harness code review process, automated pre-production checks, mean time to revoke < 5 minutes, regular prompt injection pen tests

The levels are intended as an audit aid: an external auditor sees at a glance where the organisation stands per MHC and which improvement path is planned. The target level per MHC should derive from the risk assessment of the respective organisation – not every organisation needs the Managed level everywhere.

Part IV – Recommendations and Outlook

Parts I to III have set out the Mythos threat landscape, the assessment of all 93 ISO controls, the gaps relative to other frameworks and the catalogue of the twelve Mythos Hardening Controls. Part IV formulates recommendations derived from this and closes with a reflection on the limitations 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 respective organisation on the basis of its own risk assessment, resource situation and stakeholder structure.

10 Recommendations for ISMS Adaptation

10.1 Scope and Limitations of the Recommendations

The recommendations formulated in this chapter are a prioritised selection, not a definitive catalogue. They cover five thematic areas: immediate reassessment of existing controls, structural hardening, maturing the ISMS process, board communication and metrics work. Other thematic areas (sectoral specifics, special regulatory questions, concrete supplier relationships) are deliberately left out because they cannot be generalised.

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

10.2 Immediate Reassessment of Existing Controls

From the analysis in Part II it follows 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 based on 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 checked whether they are recorded in the SoA in their previous form as the sole mitigation for Mythos-relevant risks. If so, the entry is supplemented or corrected with the respective replacement recommendation from Chapter 6.
- For the 37 partially degraded controls, it is checked which of the five supplementation patterns from Chapter 5.3 has already been implemented in each case and where gaps exist. The result is a gap list per control as input for structural hardening.
- For the 29 robust controls from Chapter 4, consistent implementation is checked. A robust control formally entered in the SoA but not implemented throughout offers no protection under Mythos conditions. Example: an asset inventory without automated updating quickly becomes outdated, leaving Mythos-specific anomaly detection with 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, assignments to the twelve MHC from Chapter 9. The update is the subject of the next management review.

10.3 Structural Hardening: Adopting the MHC into the SoA

Each of the thirteen MHC addresses a Mythos-relevant gap that cannot be closed with ISO 27002 means alone. Adoption into the SoA follows a simple pattern:

- Each MHC is included as an additional entry in the SoA, with reference to the flanking ISO controls from the MHC compact table (Ch. 9.4).
- The implementation status is documented honestly: already implemented, partially implemented, planned or not begun. The category "planned" should carry an explicit justification of why immediate implementation is not possible.
- For each MHC, it is checked whether it is applicable at all in the organisational context. MHC-07 (multi-tenancy isolation), for example, is not relevant for organisations without a multi-tenant architecture. Such non-applicability is likewise noted 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 organisations subject to regulation that requires one of the MHC contents anyway (DORA-regulated financial institutions for MHC-12, NIS2 entities for MHC-11), adoption is regulatorily necessary in any case.

10.4 Maturing the ISMS Process

Several observations from the Mythos threat landscape have direct implications for the operation of the ISMS:

- 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 include the Mythos threat landscape as a recurring agenda item. The version notes of this guide and relevant threat intelligence reports can serve as the basis.
- The internal audit function should be supplemented with continuous audit elements (substantively congruent with MHC-10). The results of automated monitoring feed into the reduced, still annual formal audits, but do not replace them.
- The KPI landscape of the ISMS is supplemented with Mythos-relevant indicators (see Ch. 10.6).
- Establish innovation acceleration governance: a cross-functional body of security, legal, engineering, data protection and, where applicable, procurement with the explicit mandate to evaluate and productionise defensive AI tools and new security controls at accelerated pace. Without this body, every Mythos hardening measure runs into the usual approval friction (procurement cycles, legal review, data protection sign-off), which becomes a structural weakness under Mythos conditions – the NIS2 IR implicitly recognises this in No. 1.1 as a governance obligation. Recommended cadence: fortnightly, with a defined time-to-decision target of at most 30 days for tools with a valid ISO 27001/SOC 2/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, EPSS- and KEV-based prioritisation. Minimum capacity from around 1,000 employees 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 rated 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 trivialisation. The following points have proved useful:

- The four findings from Chapter 2 lend themselves as the structure for a board presentation: collapse of the patch-gap window, timeline compression, fragmentation, capability decoupling. All four points are empirically substantiated.
- The distinction between robust, partially degraded and friction-based controls conveys that not all of security is failing, but that targeted adaptation is required.
- The MHC from Chapter 9 are suitable as a discussion basis for investment decisions. They provide concrete reference points substantiated with framework references.
- The regulatory embedding should be made transparent. Mythos hardening is to a considerable extent a legal obligation, not 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 Ch. 3.3) it applies directly; for other NIS2 entities the respective national transposition is decisive.
- Raise the shifting standard of care. With the entry into force of the EU AI Act (Art. 6 et seq. on high-risk AI systems) and the broad availability of defensive AI tools, the benchmark of what counts as an appropriate security measure is shifting. Boards must increasingly be able to demonstrate to supervisors, insurers and courts which AI tools are used for defensive purposes (code audit, threat hunting, vulnerability discovery) – and justify the non-use of available tools. Recommendation as a recurring board briefing element: "Which defensive AI tools did we introduce in the reporting period? Which are we evaluating? Which are we deliberately not using – and with what justification?" This addresses the risk of regulatory and liability exposure explicitly rated HIGH in Mythos-Ready (CSA, SANS, OWASP).

10.6 Mythos-Relevant Metrics

The following metrics are a selection, not a definitive KPI set. They are oriented towards the twelve MHC and the four assessment criteria from Chapter 3.2.

Metric	Description	Reference
Mean Time to Containment (MTTC)	Average time from detection to automatic or manual containment of an incident. Target value for high-confidence alerts under 10 minutes.	MHC-11, MHC-05
Share of Phishing-resistant MFA	Share of all privileged and externally reachable accounts with phishing-resistant MFA. Target value context-dependent, typically at least 95 % for privileged accounts.	MHC-03
SBOM Coverage	Share of own software artefacts with an automatically generated SBOM and share of critical suppliers with a contractually assured SBOM obligation.	MHC-02
Patch Latency for KEV Listings	Time from inclusion in the CISA Known Exploited Vulnerabilities to roll-out of the patch. Target value under 24 hours.	MHC-09
Restore Test Success Rate	Share of successful quarterly restore tests from	MHC-08

Metric	Description	Reference
	immutable backups. Target value 100 %.	
Continuous Monitoring Coverage	Share of controls whose implementation is checked through automated continuous monitoring, rather than only through periodic audits.	MHC-10
Cryptographic Inventory Coverage	Share of the cryptographic methods in use that are captured in the central inventory and have a priority classification for PQC migration.	MHC-01
TLPT Findings Closure Rate	Share of findings from threat-led penetration tests closed within the time window agreed with the supervisory body.	MHC-12

The eight metrics do not cover all twelve MHC. For structural properties (e.g. MHC-04 workload identity, MHC-07 multi-tenancy isolation), implementation stage gates in project controlling are better suited than continuous metrics.

10.7 Summary of Recommendations

The five recommendation areas – immediate reassessment, structural hardening, ISMS maturation, board communication and metrics work – can be addressed independently of one another. None of the recommendations presupposes that another has already been fully implemented. Organisations with limited resources can prioritise individual areas in a targeted manner.

The unifying bracket of all five areas is the withdrawal of an implicit assumption: that the effort an attacker must invest constitutes a reliable protective quantity. Under Mythos, this is no longer the case. The recommendations aim to anchor protective effect instead in hard barriers: cryptographic, architectural, aggregation-resistant or automation-based.

11 Reflection and Outlook

11.1 Limitations of This Work

First, the assessment is a snapshot as at June 2026. Both the Mythos threat landscape and the cited frameworks continue to evolve. Individual classifications may shift as new empirical evidence becomes available, as new controls are added to frameworks, or as the attackers' model class changes further. The version notes document this development.

Second, the assessment does not replace an organisation-specific risk analysis. Every organisation has its own threat vectors, its own crown jewels and its own dependencies.

Third, the MHC catalogue makes no claim to completeness. Future versions will add further MHC, in particular in the areas of AI system governance (ISO/IEC 42001), AI-specific adversarial testing procedures and the operationalisation of agent identity frameworks.

Fourth, the analysis rests on publicly available primary sources. Internal incident data of individual organisations is not part of the assessment basis.

11.2 What This Guide Does Not Provide

- This guide is not a tool for certification or formal compliance checking. The assessment leads to an effectiveness estimate, not a conformity statement.
- This guide does not replace vendor documentation, a product selection process or an architecture decision. It works at the level of categories and principles.
- This guide is not a legal document. The regulatory references are to be understood as informational and do not replace legal review by qualified lawyers.
- This guide is not a threat intelligence source. It relies on existing reports and makes them usable for control assessment.

11.3 Expected Developments

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

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

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

11.4 Concluding Remarks

Classic controls do not become obsolete. The assessment language of this guide – "friction-only", "partially degraded" – could give the impression that established security measures fundamentally lose their effectiveness under Mythos. That reading would be wrong. Classic controls are weakened by Mythos in specific effectiveness dimensions, but retain their full protective effect in other dimensions. Network segmentation, least privilege, strong authentication, secrets management, immutable backups, cryptographic data separation – all these controls remain the most robust defence layer under Mythos conditions. The correct response to Mythos is the hardening of these structural controls, not their replacement by new detection or automation capabilities. The Mythos Hardening Controls catalogued in Chapter 9 supplement the structural controls and close specific, precisely identified gaps – they are supplement, not replacement.

Information security is in a phase of structural change. The work of CISOs in the coming years will be determined by whether protective effect can be anchored in hard, structural barriers – cryptographic, architectural and automation-based. The path there does not lead through abandoning established frameworks, but through their targeted supplementation.

The ambition of this guide is modest. It delivers no new theoretical framework, no methodology of its own and no compliance statement. It is a working aid that sets existing frameworks against the current threat landscape, makes gaps visible and names concrete starting points for closing them. The actual work lies with the CISOs and their teams.

Part V – Supplementary Layer: Further Frameworks

Parts I to IV so far rest on the primary frameworks defined in Chapter 3.3. Part V supplements this core base with five further frameworks that are useful for deepened Mythos hardening work.

The five supplementary chapters are kept compact. For each framework they answer three questions: What is it? What Mythos relevance does it have? How does it relate to the primary frameworks?

Readers who wish to concentrate on working with the primary frameworks can skip Part V.

12 BSI IT-Grundschutz Compendium

The BSI's IT-Grundschutz Compendium is the most established reference in Germany for building and deepening an ISMS. In early design phases of this guide it was planned as a primary framework, but was demoted in favour of the BSI catalogue C5:2026.

What it is: A comprehensive, modular compendium with building blocks for establishing an ISMS according to BSI Standards 200-1, 200-2 and 200-3. The building blocks are organised into the layers ISMS, ORP, CON, OPS, DER, APP, SYS, IND, NET and INF. Each building block defines basic, standard and high requirements.

Mythos relevance: The Compendium contains practice-oriented, technically concrete requirements that exist only abstractly in ISO 27002:2022. For German organisations in the public sector, in the critical infrastructure context and in companies with IT-Grundschutz certification, the Compendium is the binding basis.

Relationship to the primary frameworks: Complementary to C5:2026. Where C5:2026 provides the cloud-specific audit catalogue, the Grundschutz Compendium offers a broader ISMS construction perspective for organisations whose infrastructure is not purely cloud-based.

The following correspondence list helps map the C5 references named in Parts II and III to Grundschutz building blocks:

C5:2026 domain	IT-Grundschutz building blocks (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

C5:2026 domain	IT-Grundschutz building blocks (typical)
C5 SSO (Service Providers and Suppliers)	OPS.2, CON.14
C5 SIM (Security Incident Management)	DER.2.1 to DER.2.3
C5 BCM (Business Continuity Management)	DER.4
C5 COM (Compliance)	ISMS layer
C5 PSS (Product Safety and Security)	APP building blocks depending on the system

The list is an orientation aid, not a binding one-to-one mapping. In practice, several Grundschutz building blocks fit many controls.

13 MITRE ATT&CK and D3FEND

What it is: ATT&CK is a globally maintained knowledge base of over two hundred attacker techniques, organised into fourteen tactics from initial access to impact. D3FEND is the complementary database on the defence side: it categorises defensive techniques and links them to ATT&CK techniques. Both frameworks are publicly available, continuously updated and the de facto standard in the security operations community.

Mythos relevance: Under Mythos, ATT&CK gains further importance as a detection reference because the kill-chain-oriented structure maps fragmented attacks in coherent tactic steps. That is exactly the perspective that makes individual policy-compliant micro-steps recognisable as part of an overarching attack pattern. D3FEND helps to systematically tie defence investments to concrete attacker techniques.

Relationship to the primary frameworks: ATT&CK and D3FEND have no regulatory character. As an operational working vocabulary they are complementary: C5:2026 OPS-13 (SIEM) and NIS2 IR No. 3.2 can be concretised with ATT&CK as the detection framework; DORA TLPT (Art. 26–27) relies to a considerable extent in practice on ATT&CK-based scenarios.

14 ISO/IEC 42001 – Artificial Intelligence Management System

What it is: A management system standard following the ISO high-level structure, specifically for AI systems. 42001 specifies requirements for context, leadership, planning, support, operation, performance evaluation and improvement of an AI management system. Annex A provides controls for risk management, data quality, lifecycle management, transparency, fairness and robustness.

Mythos relevance: Two dimensions. First, 42001 forms the governance frame for AI systems on the defence side – SIEM, EDR, SAST and SOAR components enriched with AI models. Without documented data lineage, adversarial robustness tests and performance monitoring for these models, new risks arise that a classic ISMS does not cover. Second, 42001 provides vocabulary for AI-specific attacks on one's own models: prompt injection, data poisoning, model extraction.

Relationship to the primary frameworks: Complementary to ISO/IEC 27001. 42001 establishes a standalone management system for the AI components. The EU AI Act does

not reference 42001 normatively, but in practice the industry assumes that a 42001 certification largely covers the requirements for high-risk AI systems.

15 CIS Controls v8

What it is: Eighteen control categories, prioritised by implementation groups IG1 (basic provision), IG2 (moderate maturity) and IG3 (full implementation). The controls are oriented towards operational practice and provide concrete safeguards. The current version v8.1 dates from 2024.

Mythos relevance: The implementation group logic is a pragmatic prioritisation aid for organisations with limited resources. Under Mythos, Control 1 (Inventory and Control of Enterprise Assets), Control 2 (Inventory and Control of Software Assets), Control 6 (Access Control Management), Control 8 (Audit Log Management) and Control 16 (Application Software Security) are particularly relevant.

Relationship to the primary frameworks: The CIS Controls can be mapped to ISO 27002 and C5:2026; the CIS organisation maintains corresponding mapping tables. Not an alternative, but a tactical implementation reference. For organisations in an early maturity phase, the CIS Controls IG1 can serve as a pragmatic entry point.

16 OWASP ASVS and SAMM

What it is: ASVS is a detailed requirements catalogue for the security verification of web applications with three verification levels (L1 surface scan, L2 standard applications, L3 critical applications). SAMM is a maturity model for the security organisation around software development, organised into five business functions (Governance, Design, Implementation, Verification, Operations).

Mythos relevance: Under Mythos, ASVS is the central operational reference for secure software development, because its requirements are formulated at the operational verification level and can therefore be integrated directly into CI pipelines. The requirements provide a concrete basis for automated SAST and DAST rule sets (see MHC-09). SAMM is useful as a maturity model to plan the path from ad hoc security to an institutionalised secure SDLC in a traceable way.

Relationship to the primary frameworks: The OWASP projects concretise the requirements formulated abstractly in ISO 27002 (A.8.25 to A.8.29), in C5:2026 (DEV-01 to DEV-15) and in the CRA (Annex I Part II). For software manufacturers subject to the CRA, ASVS L2 or L3 provides an operational benchmark.

Synthesis of the Supplementary Layer

The five frameworks of this part fall into three groups: the BSI IT-Grundschutz Compendium offers the German ISMS deep-dive perspective. MITRE ATT&CK and D3FEND provide the operational vocabulary for detection and adversarial testing. ISO/IEC 42001, CIS Controls v8 and OWASP ASVS/SAMM are specialisation references for AI system governance, implementation prioritisation and secure software development.

None of the five supplementary frameworks replaces the primary base from Chapter 3.3. The decision to include one or more supplementary frameworks in one's own ISMS depends on the industry, size and strategic orientation of the organisation. A German critical infrastructure company will regularly maintain the BSI IT-Grundschutz Compendium as a second primary reference; a SaaS provider with a high AI share will establish ISO/IEC 42001 as a parallel

work stream; a software manufacturer will anchor OWASP ASVS as the central development reference.

Annexes – Working Materials

The annexes bundle the contents developed in Parts I to V in a form directly usable for practical ISMS work. Annex A provides the complete assessment matrix of all 93 controls with category classification and MHC assignment. 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 provides the RACI model for MHC implementation. Annex F describes the risk assessment bridge to ISO/IEC 27005. Annex G standardises the KPI definitions from Ch. 10.6 for audit-ready measurement reproducible across organisations.

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 the following are stated: 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 assignments 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 associated assets	S	MHC-13
A.5.10	Acceptable use of information	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

ID	Control title	Cat.	Flanking MHC
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 within supplier agreements	T	MHC-02
A.5.21	Managing the ICT supply chain	T	MHC-02
A.5.22	Monitoring of supplier services	T	MHC-02
A.5.23	Cloud services	T	MHC-06, MHC-07
A.5.24	IR management planning	T	MHC-11
A.5.25	Assessment and decision on events	R	MHC-11
A.5.26	Response to incidents	T	MHC-11
A.5.27	Learning from 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 BCM	S	MHC-08
A.5.31	Identification of legal requirements	N	-
A.5.32	Intellectual property rights	N	-
A.5.33	Protection of records	T	MHC-08
A.5.34	Privacy and PII	T	-
A.5.35	Independent review	T	MHC-10, MHC-12
A.5.36	Compliance with policies	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	Awareness, education and training	T	-

ID	Control title	Cat.	Flanking MHC
A.6.4	Disciplinary process	N	-
A.6.5	Responsibilities after termination	S	-
A.6.6	Confidentiality agreements	N	-
A.6.7	Remote working	T	MHC-03, MHC-04
A.6.8	Reporting of security events	T	-
A.7.1	Physical security perimeters	S	-
A.7.2	Physical entry	S	-
A.7.3	Securing offices 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 screen lock	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	-
A.7.13	Equipment maintenance	N	-
A.7.14	Secure disposal or re-use	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

ID	Control title	Cat.	Flanking MHC
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	S	MHC-08
A.8.15	Logging	S	MHC-05
A.8.16	Monitoring activities	T	MHC-05
A.8.17	Clock synchronisation	S	-
A.8.18	Privileged utility programs	S	-
A.8.19	Software installation	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 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 environments	S	MHC-06
A.8.32	Change management	T	MHC-09
A.8.33	Test information	S	-

ID	Control title	Cat.	Flanking MHC
A.8.34	Protection during audit testing	N	-

Legend: S = Robust (Chapter 4), T = Partially degraded (Chapter 5), R = Friction-only (Chapter 6), N = Not affected (Chapter 7). The MHC references point to Mythos Hardening Controls from Chapter 9.

Annex B – Framework Mapping for Core Topics

The following table summarises, for twelve central topics of information security, how the primary frameworks each address them. It is orientation, not a complete mapping.

Topic	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	COM-01 to	GV.OC	Art. 6(5)	Art. 13,	No. 2.2,

Topic	ISO 27002	C5:2026	NIST CSF	DORA	CRA	NIS2 IR
	A.5.36	COM-04			14	2.3
Cloud usage	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 formulates no specific requirement of its own for the respective topic. The NIST CSF column names the respective categories from CSF 2.0; concrete implementation follows via SP 800-53.

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

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 empty and can be filled in organisation-specifically (e.g. already implemented, partially implemented, planned, not applicable).

MHC	Title	Framework basis	ISO linkage	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.8.4, A.8.30	
MHC-03	Phishing-resistant MFA	NIST 800-63B, FIDO2	A.5.17, A.6.7, A.8.1, A.8.4, 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.8.16	

MHC	Title	Framework basis	ISO linkage	Status
MHC-06	Containers and confidential computing	C5:2026 OPS-32 to OPS-35	A.5.23, A.8.27, 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, 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 the detection of Mythos-relevant events. The list is a selection, not a complete catalogue.

Source	Description	URL / access
CVE – Common Vulnerabilities and Exposures	Central vulnerability library from MITRE, published via NVD.	cve.org / nvd.nist.gov
EUVD – European Union Vulnerability Database	ENISA-managed vulnerability database for the EU.	euvd.enisa.europa.eu
KEV – Known Exploited Vulnerabilities Catalog	CISA list of actively exploited vulnerabilities, patch prioritisation.	cisa.gov/known-exploited-vulnerabilities-catalog
EPSS – Exploit Prediction Scoring System	Probability of exploit use of a CVE within 30 days.	first.org/epss

Source	Description	URL / access
OSV – Open Source Vulnerabilities	Vulnerability feed for open-source dependencies.	osv.dev
MITRE ATT&CK	Taxonomy of attacker tactics and techniques.	attack.mitre.org
CISA Advisories	US government notifications on acute threats.	cisa.gov/news-events/cybersecurity-advisories
BSI warnings and situation reports	German security warnings and 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 hub 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 organisation-specifically – particularly in organisations without a dedicated AI governance lead or without a board reporting line for cybersecurity.

Assessment grid: R = Responsible (executes), 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 / officer, AI-Gov = AI governance lead, IT = IT operations / platform engineering, Dev = development, SOC = Security Operations Centre / IR function, Legal = legal and compliance function, Board = executive management / 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 /	A	C	C	R	R	C	I	I

MHC	Title (short)	CISO	ISMS	AI-Gov	IT	Dev	SOC	Legal	Board
	Zero Trust								
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-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 is not delegable and is not split. In the template, the CISO is entered as Accountable throughout; in larger organisations, accountability for individual MHC can be delegated to domain owners (e.g. CTO for MHC-04, Head of Engineering for MHC-09).
- Several R are possible where execution spans several domains. Example MHC-04: IT sets up SPIFFE/SPIRE, Dev migrates the service identities – both are Responsible.
- MHC-13 is the only row in which the Board is listed as Consulted. Rationale: the governance of AI agents going into production concerns risk acceptance questions above the CISO mandate threshold (standard of care, see Ch. 10.5).
- External service providers (MDR, MSSP, TLPT providers, GRC consulting) do not appear as a separate column. They are governed via the respective Responsible role, not via a dedicated RACI column.

Annex F – Risk Assessment Bridge to 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 Ch. 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.
- Identified weaknesses along the four assessment criteria from Ch. 3.2 (attacker patience, time compression, capability decoupling, aggregation resistance).
- Flanking Mythos Hardening Controls from Annex A.
- Current maturity level per flanking MHC per Ch. 9.5 (Initial / Defined / Managed).

F.2 Processing Steps in the Risk Process

Step 1: Risk Reassessment (ISO/IEC 27005:2022 Ch. 7.3)

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

- Raise the likelihood axis by one level for risks whose mitigation rests on a partially degraded control.
- Raise the likelihood axis by two levels for risks whose sole mitigation rests on a friction-only control.
- Consequence axis unchanged, unless the damage scenario itself (e.g. mass exfiltration through agentic tools) suggests a higher consequence level.

Note: the concrete level scale (3, 5 or 10 levels) follows the risk matrix established in the organisation; the principle of raising the likelihood remains unaffected.

Step 2: Treatment Options (ISO/IEC 27005:2022 Ch. 8.1)

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

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

Step 3: Residual Risk Assessment and Acceptance (ISO/IEC 27005:2022 Ch. 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 made.

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	Partially degraded (aggregation blindness, credential compromisability). Flanked by MHC-03.	Likelihood +1; without MHC-03 (phishing MFA, FIDO2) at the Defined maturity level, the control works only to a limited extent.
A.8.4 Source code access	Partially degraded (aggregation blindness). Flanked by MHC-02 and MHC-05.	Likelihood +1; only in combination with repository anomaly detection (MHC-05) and mandatory SBOM (MHC-02) at the Managed maturity level is effectiveness restored.
A.8.16 Monitoring	Partially degraded (aggregation blindness). Flanked by MHC-05.	Likelihood +1; only with behaviour-based detection at the Managed maturity level are fragmented mass clones detected.

Resulting treatment decision: adoption of MHC-02, MHC-03 and MHC-05 into the SoA with a target maturity level of Managed within twelve months. Until that is reached, 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.
- Updated Statement of Applicability with additional MHC entries and a reference to the respective Mythos category of the flanked ISO controls.
- Treatment plan with target maturity levels 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 Standardisation

This annex defines, for each of the eight Mythos metrics from Ch. 10.6, the definition, data source, measurement points (start and stop conditions), calculation formula, reporting frequency and measurement responsibility. Objective: reproducibility between organisations and over time.

Without this standardisation, two organisations can maintain the same KPI with an identical target value and still produce incomparable results – a problem regularly identified in audits.

G.1 Mean Time to Containment (MTTC)

Element	Definition
Definition	Average duration from the first documented detection event of a high-confidence alert to the completed containment action (account lock, network isolation, key rotation or equivalent).
Clock start	Timestamp of the first SIEM or EDR alert with a confidence score ≥ 80 % or with an explicit high-severity flag of the detection tool.
Clock stop	Timestamp of the confirmed execution of the first containment action in the SOAR audit log or in the IR ticket system.
Data source	SOAR platform audit log (primary), SIEM alert log (secondary), IR ticket system (tertiary for non-automated cases).
Calculation	Arithmetic mean of the durations of all high-confidence containment operations in the reporting period. Additionally report the median.
Target value	Mean < 10 minutes; median < 5 minutes.
Reporting	Monthly to the CISO; quarterly in the management review.
Responsibility	SOC lead (measurement); CISO (reporting).
Reference	MHC-11, MHC-05; Mythos-Ready Risk 4.

G.2 Share of Phishing-resistant MFA

Element	Definition
Definition	Share of all privileged accounts and externally reachable service accounts secured by a phishing-resistant authentication method per NIST SP 800-63B AAL2 or higher (FIDO2, WebAuthn, passkeys, hardware tokens per CTAP2).
Calculation	(Number of accounts with phishing-resistant MFA / number of all in-scope accounts) × 100
Scope definition	Privileged accounts: all accounts with domain admin, cloud tenant admin, code repository admin, database admin or comparable roles. Externally reachable: all accounts enabling authentication outside the corporate

Element	Definition
	network.
Data source	Identity provider (IdP) such as Entra ID, Okta, Ping; supplemented by the CMDB list of privileged accounts.
Target value	Privileged accounts ≥ 95 %; externally reachable accounts ≥ 90 %; SMS MFA rate = 0 % for new accounts.
Reporting	Monthly to the CISO; quarterly in the management review.
Responsibility	IAM lead (measurement); ISMS manager (reporting).
Reference	MHC-03; Mythos-Ready Risk 4.

G.3 SBOM Coverage

Element	Definition
Definition	Two sub-metrics: (a) share of own software artefacts with an automatically generated SBOM in CycloneDX or SPDX format, (b) share of critical suppliers with a contractually assured SBOM obligation.
Calculation	(a) (Number of artefacts with an active SBOM in CI / number of all release-relevant artefacts) × 100. (b) (Number of critical suppliers with an SBOM clause in the contract / number of critical suppliers) × 100.
Scope definition	Critical suppliers: all suppliers in tier 1 or 2 per the tiering criteria from A.5.19; software artefacts: all artefacts that go into production compiled or packaged.
Data source	(a) CI pipeline reports (Snyk, Dependency-Track); (b) contract database, reconciled with supplier tiering.
Target value	(a) ≥ 95 % after 12 months; (b) ≥ 80 % after 18 months.
Reporting	Quarterly.
Responsibility	Head of Engineering (a); CISO (b).
Reference	MHC-02; Mythos-Ready Risk 7.

G.4 Patch Latency for KEV Listings

Element	Definition
Definition	Time from the inclusion of a vulnerability in the CISA Known Exploited Vulnerabilities Catalog to the complete roll-out of the correcting patch in the production environment.
Clock start	Publication timestamp in the KEV catalogue (UTC).
Clock stop	Timestamp of the successful patch deployment on 100 % of affected hosts or workloads, reported from the patch management tool.
Data source	KEV catalogue (input), patch management tool (output).
Calculation	Mean and 95th percentile across all KEV-relevant patches in the reporting period.
Target value	Mean < 24 hours; 95th percentile < 72 hours for internet-exposed systems.
Reporting	Per KEV entry in the SOC bridge (immediately); aggregated monthly.
Responsibility	VulnOps lead (measurement); CISO (reporting).
Reference	MHC-09, A.8.8; Mythos-Ready Risks 1 and 9.

G.5 Restore Test Success Rate

Element	Definition
Definition	Share of successful restore tests from immutable backups within the reporting period. Successful = recovery completed within the agreed RTO and delivering consistent, error-free data per the defined acceptance criterion.
Calculation	(Number of successful restore tests / number of restore tests performed) × 100
Minimum frequency	Quarterly per business-critical system; annually a complete disaster recovery test.
Data source	Backup tool audit log, supplemented by documented test protocols of IT operations.
Target value	100 %. Every failed test triggers a major incident in the ISMS process.
Reporting	Quarterly in the management review; upon failure immediately to the CISO and risk owner.
Responsibility	Head of IT operations (execution); ISMS manager (reporting).
Reference	MHC-08, A.8.13.

G.6 Continuous Monitoring Coverage

Element	Definition
Definition	Share of the controls maintained in the SoA whose implementation is checked through automated continuous monitoring (policy as code, automated compliance checks), versus controls checked only through periodic audits.
Calculation	(Number of controls with active automated checking / number of controls in the SoA) × 100
Scope definition	Active automated checking: a check executed at least daily with alerting upon deviation; checking coverage ≥ 80 % of the defined sub-requirements of the control.
Data source	OPA/Sentinel policy repository, GRC platform with automated control tests.
Target value	≥ 60 % after 12 months; ≥ 80 % after 24 months.
Reporting	Quarterly.
Responsibility	ISMS manager.
Reference	MHC-10.

G.7 Cryptographic Inventory Coverage

Element	Definition
Definition	Share of the cryptographic methods used in production systems that are captured in the central crypto inventory and carry a priority classification for the PQC migration.
Calculation	(Number of captured and prioritised cryptographic methods / total number of cryptographic methods in production use) × 100
Minimum capture attributes	Algorithm, key length, implementation (library + version), usage context (in transit / at rest / in use), PQC priority level (1 = long-term confidentiality, 2 = standard, 3 = short-lived).
Data source	CMDB extension with a crypto table; code scanners with crypto discovery (e.g. Cryptosense, Sandbox AQ).
Target value	≥ 80 % after 12 months; ≥ 95 % after 24 months.
Reporting	Every six months.
Responsibility	Architecture lead (maintenance); CISO (reporting).
Reference	MHC-01.

G.8 TLPT Findings Closure Rate

Element	Definition
Definition	Share of the findings from the most recent threat-led penetration test closed within the time window agreed with the supervisory body.
Calculation	(Number of findings closed within the SLA window / number of findings in the reporting period) × 100
Thresholds (standard)	Critical findings: 30 days. High findings: 60 days. Medium findings: 90 days. Low findings: 180 days. Thresholds may be adjusted organisation-specifically and in line with DORA Art. 26.
Closure definition	A finding is considered closed when a re-test by red teamers or internal audit confirms the effectiveness of the mitigation.
Data source	TLPT report, internal findings tracker, re-test report.
Target value	≥ 90 % within the SLA window, aggregated across all severity levels.
Reporting	Quarterly; critical findings immediately after publication of the TLPT report.
Responsibility	CISO (tracking); risk owner (closure).
Reference	MHC-12.

G.9 Standardisation Notes

The eight KPIs are defined so that they can be collected reproducibly across organisations. Prerequisite: the data sources are available in the respective tool stack and the scope definitions are applied consistently.

Recommendation for comparability over time: the KPI definitions are versioned upon material changes, and the version history is documented in the management review. A change in the calculation (e.g. new scope inclusion) must be explicitly flagged as a break in the time series.

Recommendation for comparability between organisations (industry benchmarks, ISAC sharing): when transmitting, the version of the KPI definitions is transmitted together with the value. Sector-specific adjustments (e.g. higher thresholds for DORA-regulated financial institutions) must be declared in the accompanying note.

Glossary

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

Term	Explanation
Aggregation resistance	Property of a control to take effect even when an attack is decomposed into many micro-steps that individually appear legitimate. One of the four assessment criteria (Ch. 3.2).
Attacker patience	Assessment criterion: does the protective effect rest on an attack being too laborious for an adversary? Structurally weakened under Mythos.

Term	Explanation
Crypto-agility	Ability of a system to change cryptographic mechanisms at short notice, for instance upon compromise of an algorithm or during PQC migration.
Capability decoupling	Assessment criterion: does the likelihood assumption depend on a historical correlation between actor and capability that Mythos has removed?
FIDO2 / WebAuthn	Authentication standards of the FIDO Alliance for phishing-resistant MFA. Cryptographically anchor authentication to the origin domain.
Mythos	Designation for Claude Mythos Preview, used in this guide as the reference model for frontier AI capabilities in the hands of attackers.
MHC – Mythos Hardening Control	Category introduced in this guide for a supplementary control that closes a Mythos-relevant gap relative to ISO 27002:2022. Twelve MHC in Chapter 9.
PQC – Post-quantum cryptography	Cryptographic methods resistant to attackers with quantum computing capabilities. NIST standardisation (ML-KEM, ML-DSA, SLH-DSA).
Friction-only	Category for controls whose core effect against Mythos attackers is structurally eliminated, because it rests either on limited attacker capacity or on humanly manageable response times. Four controls in Chapter 6.
SBOM – Software Bill of Materials	Machine-readable list of all software components of an artefact. Formats: SPDX, CycloneDX. Mandatory under CRA Annex I Part II.
SLSA – Supply-chain Levels for Software Artifacts	Framework for build provenance with four maturity levels. Level 3+ recommended for critical build paths.
SOAR	Security Orchestration, Automation and Response. Platform category for automated incident response workflows.
SPIFFE / SPIRE	Standard and reference implementation for workload identities. Basis for Zero Trust architectures in cloud-native environments.
Robust	Category for controls whose protective effect stems from a hard barrier that persists even under Mythos. 29 controls in Chapter 4.
Store-now-decrypt-later	Threat model in which encrypted data exfiltrated today is retained in order to decrypt it once quantum computing becomes available.
Partially degraded	Category for controls whose protective effect persists, but whose originally assumed strength drops under Mythos. 37 controls in Chapter 5.
TEE – Trusted Execution Environment	Hardware-backed, isolated execution area for confidential data processing. Basis for confidential computing.
TLPT – Threat-Led	Penetration testing based on real threat scenarios, anchored in

Term	Explanation
Penetration Testing	DORA Art. 26/27 and in the TIBER-EU framework.
UEBA – User and Entity Behavior Analytics	Detection category based on behavioural baselines and ML-supported anomaly detection.
Time compression	Assessment criterion: does the control presuppose human response time in a chain in which the attacker operates autonomously?
Zero Trust	Architectural principle under which no actor, no device and no network is implicitly trusted. Every transaction is explicitly verified. NIST SP 800-207.

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-H Cyber Resilience Requirements for Software Manufacturers.
- 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).
- GDPR – Regulation (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), February 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 Rev. 2 – Computer Security Incident Handling Guide.

- NIST SP 800-63B – Digital Identity Guidelines, Authentication and Lifecycle 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).

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
- Anthropic – Claude Mythos Preview, April 2026. anthropic.com/news
- 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, Version 0.95, 18 April 2026.
- OWASP Gen AI Security Project – LLM Top 10 (LLM01–LLM10) and Agentic Security Initiative (ASI01–ASI10).
- MITRE ATLAS – Adversarial Threat Landscape for Artificial-Intelligence Systems, atlas.mitre.org.
- NIST AI Risk Management Framework 1.0, AI RMF 1.0.
- MITRE ATT&CK Enterprise Matrix – attack.mitre.org.
- MITRE D3FEND – d3fend.mitre.org.
- OWASP Application Security Verification Standard (ASVS) – owasp.org/www-project-application-security-verification-standard.
- OWASP Software Assurance Maturity Model (SAMM) – owaspsamm.org.
- 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.
- FIDO Alliance – WebAuthn Specifications.
- TIBER-EU – European framework for Threat Intelligence-based Ethical Red Teaming, European Central Bank.

Version Notes

This guide is maintained as a living document. The following table serves as the change log for subsequent versions. The present edition is Version 1.1.

Version	Date	Changes
1.0	April 2026	Initial publication. Assessment of all 93 controls of ISO/IEC 27002:2022 against the Mythos threat landscape. Framework cross-comparison against BSI C5:2026, NIST, DORA, CRA and NIS2 with Implementing Regulation. Catalogue of the 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 point (a); deadlines not in IR 3.3 but in NIS2 Directive Art. 23(4); misinterpretation of 3.3.2; phase count in 3.5 corrected from four to three). Addition of No. 6.7 (network security) for A.8.20–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 Ch. 3.1.3 to resolve the contradiction between definition and application. Marking of the time-factor addition as the author's interpretation of ISO/IEC 27005. Positioning of BSI C5 as an audit catalogue. Consistency Ch. 9.4 vs. Annex A. Addition of Annexes A–G. Standardisation of spelling (ß) and typography.
1.2	April 2026	Concretisation of the hardening recommendations for audit robustness. Tool classes named explicitly (EDR, SAST/DAST/SCA, MDR, SOAR, ZTNA, EASM). Thresholds and quantity structures added (MTTC < 10 min, ATT&CK coverage ≥ 60 %, egress anomaly 2σ/3σ, EDR confidence ≥ 80 %, at least twelve threat hunts per year). Maturity paths in three levels (Initial/Defined/Managed) for all twelve MHC in Ch. 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 focuses in MHC-05. Practicability notes added (MDR alternative, hyperscaler reality, BAS platforms). Methodology references concretised (PEAK framework, TaHiTI, DeTT&CT, STRIDE). Mythos-specific user reporting indicators in A.6.8.
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 the unmanaged AI agent attack surface rated CRITICAL in Mythos-Ready: harness audit, blast radius limits, human override, supply chain inventory for MCP servers/extensions/agentic skills, pre-production checks. A.5.9 extended with a shadow AI inventory (browser plug-ins, IDE extensions, MCP

Version	Date	Changes
		servers). Ch. 10.4 extended with innovation acceleration governance (cross-functional body with a 30-day decision target) and a permanent VulnOps function. Ch. 10.5 supplemented with the standard-of-care shift through the EU AI Act and a board briefing element. Maturity table Ch. 9.5 extended with MHC-13; assessment matrix (Annex A) and MHC standalone worksheet (Annex C) updated accordingly. The Mythos-Ready strategy paper positioned in the Foreword and Ch. 3.4 as the central industry consensus reference (status between threat intelligence report and normative basis). With this, 12 of 13 risks named in Mythos-Ready are mapped at the "Managed" maturity level and 1 at "Defined"; no remaining gaps.
1.4	April 2026	Closure of the structural gaps identified by the audit assessment ("completeness = medium"). New Ch. 1.4 "Position in the ISMS Stack and Limitations of This Guide" explicitly positions MRIS as an effectiveness and reality layer on an ISMS foundation; makes clear what MRIS provides (effectiveness assessment, gap analysis, MHC catalogue, operationalisation) and what it deliberately does not (complete ISMS, own risk management process, PDCA system, compliance mapping tool, organisation-specific policy hierarchy). New Annex E "RACI Model for MHC Implementation" with eight roles (CISO, ISMS, AI-Gov, IT, Dev, SOC, Legal, Board) for all thirteen MHC. New Annex F "Risk Assessment Bridge to ISO/IEC 27005" as an interface definition between MRIS and the ISMS risk process (likelihood raising, treatment options, residual risk acceptance). New Annex G "KPI Definitions and Measurement Standardisation" with complete standardisation of the eight Mythos metrics (data source, clock start/stop, calculation formula, reporting frequency, measurement responsibility). Foreword compactly reduced to one page; focus on facts, less prose.
1.5	April 2026	Sharpening of the classification language and the effectiveness statements following a critical audit assessment. Classic controls are clearly designated as "selectively degraded", not as "obsolete" – corrective paragraphs in Ch. 9.1, Ch. 4.3 and Ch. 11.4 emphasise the priority of structural controls over new MHC. New Ch. 3.1.5 on the context dependency of the classification (the same control can fall into different categories against different attack vectors, example A.5.17). Aggregation resistance operationalised in Ch. 3.2 (SIEM correlation, UEBA, graph analyses, kill chain tracking or structural indivisibility). Threat scope in Ch. 3.5 explicitly limited to Gen-AI-accelerated attacks. MHC-05 supplemented with a realism note (detection is not prevention; low-and-slow, camouflaged agents, adversarial ML evasion as residual risk). MHC-11 extended with 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 with maturity realism (operational implementation readiness currently at the Initial level in most organisations; 12 to 24 months transition period; phased prioritisation inventorying → audit trail → capability scoping).
1.6	June 2026	Reference update without substantive reassessment of the controls. Chapter 9.2: cryptographic standards switched to the finalised FIPS 203/204/205 and HQC (MHC-01); NIST SP 800-94 positioned as the 2007 edition with a withdrawn revision draft (MHC-05); MITRE ATLAS IDs in MHC-13 corrected (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" dropped; covered by LLM06 Excessive Agency). Updated references:

Version	Date	Changes
		NIST SP 800-63B Rev. 4 (MHC-03), CycloneDX/SPDX standardisation, BSI TR-03183-2, CRA timeline and VEX/CSAF (MHC-02), Sigstore/admission controller/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 of the citations: June 2026.

Occasions for updates in future versions will in particular be: 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 guide in practice.