MRIS for TISAX® 2027

Mythos-Resistant Information Security — analysis of the TISAX® question catalogue ISA 2027, module "Information Security"

Part I: analysis and gap assessment. The Implementation Guide (Part II) is available as a separate follow-on publication.

Version 2.0 | September 2026

Author: Richard Peddi (produced with AI assistance)

License and notices

Independence notice

TISAX® is a registered trademark of the ENX Association; the ENX Association administers TISAX® on behalf of the German Association of the Automotive Industry (VDA), which develops and publishes the underlying question catalogue (VDA ISA). This analysis is an independent, privately produced work with no connection to the ENX Association or the VDA. It is neither part of the official TISAX® assessment process nor authorized, reviewed or endorsed by ENX, and it does not replace a TISAX® maturity assessment by an accredited audit provider.

Disclaimer

The classifications in this document are a professional assessment as at July 2026, not legal or certification advice. Application is at the user's own risk and responsibility.

Recommended citation

Peddi, Richard (2026): MRIS for TISAX® 2027 — Part I: analysis and gap assessment. v2.0

About the author

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

Change history

Version	Date	Description
1.0	June/July 2026	First publication.
2.0	August 2026	Methodological alignment with MRIS 2.0: binding reading of the maturity levels (Initial is not evidence of fulfilment, an MHC is deemed implemented from Defined onward), remediation latency instead of time since disclosure for actively exploited vulnerabilities, role reference opened from the CISO to the leadership perspective, spelling of "AI" harmonized, AI assistance in production disclosed.
1.1	July 2026	Revised on the basis of an externally commissioned review report (July 2026, 9 findings F-01–F-09).

Version 1.1 in detail:

Finding	Subject	Implementation
F-01	Inconsistent counting arithmetic and chapter totals	Overall total corrected from 324 to 330 (chapter totals: 116/21/14/37/105/31/6; methodology sharpened: only word-identical catalogue duplicates count as editorial removals); control 1.6.3 reassessed against the primary source (ISA2027-EN.xlsx) (table 15 → 19 rows, overall classification "5 of 15" → "7 of 19" partially degraded); control 5.2.8 "3 of 14" → "3 of 15"; total individual findings 72 → 74 (cascading through the Management Summary, cluster overview, priority table, pattern-formation chapter and Part II).
F-02	Characterization and citation accuracy	Quotation for control 5.2.5 corrected.
F-03	Asymmetric state of sourcing	four detailed claims re-researched (the Cogent study, the NIST AI Agent Standards Initiative, BSI C5:2026 OPS-25.01AC, Mythos benchmark figures); primary sources added to the source apparatus; unconfirmed BSI detail values separated from the confirmed framework statement (footnote 2; likewise in Part II, MHC-09).
F-04	Editorial document control	this change history introduced; date references harmonized to July 2026; a version note added (title page and recommended citation); an orphaned reference "(phase 5, see task list)" corrected.
F-05	Missing economic viability	concerns Part II (see there).

Finding	Subject	Implementation
	consideration	
F-06	SME suitability of individual target values	concerns Part II (see there).
F-07	Missing evidence mapping to the catalogue	concerns Part II (see there).
F-08	Expectation management on the TISAX label and maturity levels	not a document deficiency (usage note); no text change.
F-09	Threat-model dependency of the argument	not a document deficiency (usage note); no text change.

Foreword

With MRIS (Mythos-Resistant Information Security) I previously developed a methodology for testing existing ISMS frameworks specifically for their resilience against AI-accelerated attackers — named after Claude Mythos, the AI model whose publication in April 2026 first documented this shift in the threat landscape concretely and publicly (see “On the name ‘Mythos’”). MRIS-TISAX® 2027 transfers this assessment logic to the TISAX® question catalogue ISA 2027 — as an independent piece of work, not as a derivative of an existing MRIS edition: the threat basis is derived directly from primary sources (see Methodology).

The claim of this document is deliberately limited: it assesses the wording of 330 TISAX® requirements individually against a defined threat model and assigns the findings to recurring patterns. It is a diagnosis. Which concrete supplementary measures follow from it is the subject of the separate follow-on publication (Implementation Guide, MHC). Feedback, contradiction and additions are expressly welcome — the CC BY-NC license is also meant as an invitation to examine, refute or develop this analysis further.

Richard Peddi
Juli 2026

Inhaltsverzeichnis

Management Summary

MRIS-TISAX® 2027 assesses the 45 controls of the module “Information Security” in the TISAX® question catalogue ISA 2027 individually against a criteria grid derived from four documented shifts in the threat landscape (attacker patience, time compression, capability decoupling, aggregation resistance). What is assessed is the wording of the requirements, not the implementation at any particular organization. In its raw text the catalogue comprises 331 requirements marked “+” (321 graded by protection requirement plus 10 for simplified group assessments); 330 distinct requirements are assessed, because control 1.6.3 in the source catalogue contains one requirement that appears twice with identical wording (four further, textually similar pairs of requirements in control 1.6.3 are editorial variations in the source catalogue itself, not genuine duplicates, and are counted as requirements in their own right; see the data quality note in the Methodology). Prototype Protection and Data Protection are not part of this analysis.

At a glance

Metric	Value
Controls assessed	45
Distinct requirements	330
Robust	4 controls (9%)
Partially degraded	26 controls (58%)
Friction-only	1 control (2%)*
Not affected	14 controls (31%)
Recurring pattern categories (phase 2)	8 clusters, 74 individual findings
Comparison frameworks (phase 3)[2]	NIST CSF 2.0, BSI C5:2026, DORA, CRA, NIS2 Directive/IR, ISA/IEC 62443-2-1
Original extension (Part II, beyond TISAX®)[3]	AI agent and GenAI security: MHC-13 (extended) + MHC-16 cover 8 of the 10 OWASP ASI threat categories — TISAX® itself governs only two requirements here (1.3.3, 1.3.4)

*Concerns exclusively the pure password factor without a second factor (4.1.2, one of seven requirements). Classic MFA methods (OTP/push) within the same control are partially degraded — weakened, but not without effect; see Chapter 4.

[2] These six frameworks are the compliance and regulatory frameworks compared in phase 3 against the wording of the TISAX® requirements. The OWASP Top 10 for Agentic Applications (next row) is not a comparison framework in that sense but a threat taxonomy, against which only the original AI agent extension in Part II was tested — the two must not be conflated.

[3] Not derived from the TISAX® diagnosis but a professional judgement of the author on the basis of external threat research (primarily OWASP); marked as original throughout. See the section “Original additions beyond TISAX®”.

Core findings: The TISAX® requirements catalogue does not collapse against a Gen-AI-supported attacker, but it degrades in the majority of its controls: 58% lose part of their protective effect, mostly because wordings such as “regularly” or “appropriate” prescribe no event-driven or time-bound response mechanism. Only a single control (4.1.2, pure password authentication without a second factor) loses its core effect entirely. The robust share is low at 9% and is concentrated almost exclusively in Chapter 5, where TISAX® requires concrete technical mechanisms rather than organizational processes.

The 74 individual requirements affected form eight recurring patterns, headed by periodic rather than event-driven review (19 findings across four chapters). The gap analysis against six external comparison frameworks shows that most of these patterns are already addressed elsewhere in regulation — for example CVSS-bound patch deadlines in BSI C5:2026, or MFA and continuous authentication, which NIS2 Art. 21(2)(j) expressly names as an appropriate measure in the risk-based catalogue of measures — and can be adopted accordingly for the planned hardening control synthesis. One exception: capability decoupling through autonomously acting AI agents that independently integrate external services or software packages is so far governed by none of the six comparison frameworks — even NIST launched a standardization initiative of its own on this only in February 2026.

Part II responds to this with a deliberate, clearly marked exception to the otherwise consistently TISAX®-derived methodology: an extended MHC-13 and a new MHC-16 address AI agent and GenAI security — agent identity, prompt injection, kill switch, multi-agent integrity — and, tested against the OWASP Top 10 for Agentic Applications, cover 8 of the 10 threat categories defined there; TISAX® ISA 2027 itself touches this field in only two requirements (1.3.3, 1.3.4). The extension is justified by the currency of the threat: AI agents are already in production and attackable in today's operating environments, not merely the subject of a future TISAX® release. See the section “Original additions beyond TISAX®” for the full marking of which parts are TISAX®-derived and which are original.

On the name "Mythos"

The name is not an invention. Claude Mythos Preview is a real AI model published by Anthropic on 7 April 2026 — in Anthropic's own description (Frontier Red Team, lead author Nicholas Carlini) a “general-purpose language model that is strikingly capable at computer security tasks”. Its existence had already become known in March 2026 through an accidentally published draft blog post; Anthropic subsequently confirmed to Fortune that the model represented a “step change in capabilities” with considerable cybersecurity risk. Anthropic itself calls Mythos “in a different league” compared with the predecessor model Opus 4.6: in a Firefox 147 exploit benchmark, Opus 4.6 managed two successful attempts in several hundred runs, while Mythos achieved 181 working exploits plus register control in 29 further cases; in an internal OSS-Fuzz benchmark, Opus 4.6 and Sonnet 4.6 reached at best a tier 3 crash, while Mythos achieved 595 tier 1 and 2 crashes and complete control flow takeover (tier 5) on 10 fully patched target systems — among them a 27-year-old, now patched OpenBSD bug, a 16-year-old FFmpeg flaw and a FreeBSD NFS remote code execution developed autonomously without human help (CVE-2026-4747). The New York Times described the disclosure as a “cybersecurity reckoning”.

Anthropic expressly did not release the model publicly — “we do not plan to make Claude Mythos Preview generally available” — but granted access exclusively through “Project Glasswing” to selected industry partners and open source developers (including Microsoft, Apple, AWS, Cisco and Nvidia; more than 40 companies), “to enable defenders to begin securing the most important systems before models with similar capabilities become broadly available”. Its own risk assessment was blunt: “In the short term, this could be attackers, if frontier labs aren't careful about how they release these models.” In the following weeks, independent findings confirmed the significance beyond the original Anthropic publication: roughly two weeks after release, Mozilla reported 271 security vulnerabilities in Firefox found and fixed with Mythos; on 14 May 2026, staff at Calif.io reported a memory corruption exploit for the Apple M5 created with Mythos.

Within a few days, the disclosure and reported cases of unauthorized access triggered crisis responses worldwide: the US Treasury and the Fed warned banks directly, central banks in Canada, the United Kingdom, India, Japan and Australia convened emergency meetings, a US congressional letter demanded a revision of federal cybersecurity policy, and China was denied access. In June 2026 the situation escalated further: following a reported jailbreak that potentially turned Mythos 5 into an “unrestricted cyber tool”, the US Department of Commerce under Howard Lutnick ordered a worldwide suspension of access for all foreign nationals — an intervention unprecedented in this form, which was partially lifted again at the end of June after negotiations with Anthropic.

It is precisely this scale that justifies the choice of name: a single model disclosure that within days simultaneously mobilized central banks, governments and a 40-member industry consortium, followed by worldwide export controls after a real jailbreak incident — that is the moment at which the previously abstract risk of AI-accelerated attacks first materialized

concretely, with public evidence, and on a scale nobody had expected. “Mythos” therefore stands not for a metaphor but for a dated, documented event.

Important for the methodology: The name anchors this analysis at a fixed point in history — it is not a commitment to the proposition that the four shifts in the threat landscape (see Methodology) depend on the capabilities of the Mythos model itself. The actual evidentiary basis of the methodology consists of four independently documented incidents (the Cogent CVE evaluation, the GTG-1002 campaign, the JadePuffer ransomware — see Methodology), not Mythos-specific benchmark results. “Mythos attacker” is, throughout the remainder of this document, a linguistic shorthand for “an attacker with the four evidenced, generation-independent capabilities”, not a reference to a particular product. Should Mythos be technically superseded or withdrawn from the market in a few years, the methodology remains valid so long as the underlying capabilities — automated exploit development, autonomous multi-stage attack execution, fragmented attack patterns, decoupled attacker capability — persist or develop further in any available system. The name is a historical marker for the state of knowledge in July 2026, not a methodological dependency on a single vendor or model.

Quellen:

- Claude Mythos Preview — Anthropic (primary source)
- Frontier Red Team: Claude Mythos Preview — Anthropic (primary report with benchmark details on exploit and crash figures)
- The zero-days are numbered — The Mozilla Blog (primary source on the 271 vulnerabilities found and fixed in Firefox)
- First public kernel memory corruption exploit found by an AI agent — Calif.io (primary source on the Apple M5 exploit)
- Anthropic Claims Its New A.I. Model, Mythos, Is a Cybersecurity 'Reckoning' — The New York Times
- Anthropic Teams Up With Its Rivals to Keep AI From Hacking Everything — Wired
- Anthropic acknowledges testing new AI model representing 'step change' in capabilities — Fortune
- Anthropic Model Scare Sparks Urgent Bessent, Powell Warning to Bank CEOs — Bloomberg
- Anthropic Disabled Fable 5 And Mythos 5 After A U.S. Export-Control Order — Forbes
- White House drops export controls on Anthropic's Mythos and Fable AI models — The Washington Post

Methodology

What this analysis is — and what it is not

This analysis is an additional effectiveness layer on top of the TISAX® maturity assessment, not a substitute for it. A control may reach maturity level 3 ("Established", the TISAX® target value) in the audit and still be classified in this analysis as "partially degraded" — the two are not mutually exclusive. Maturity measures whether a process is documented, introduced and lived. This analysis measures something else: whether the protective mechanism the requirement calls for still holds against a Gen-AI-accelerated attacker. Definitional convention of this analysis: "robust" is deliberately tied here to the TISAX® target maturity level 3 — that is a methodological determination made by MRIS-TISAX® 2027, not a statement of the TISAX® catalogue itself. Under that convention, a requirement counts as "robust" if it sustains maturity level 3 under Mythos conditions without a supplementary measure — not merely "still works somehow".

What is assessed is the wording of the requirement, not the actual implementation at any particular organization. "Robust" means: correctly implemented, this requirement holds. Whether a particular organization has in fact implemented it correctly is answered only by an on-site audit — the regular TISAX® assessment with its maturity levels and evidence (GWP) remains responsible for that.

Limit of this analysis: requirements are assessed individually, not as an attack chain. A classification of "robust" applies to the assessed mechanism in isolation, not to the overall system. Robust encryption (5.1.1) is of little use if the access preceding it takes place through degraded authentication (4.1.2) — the combination of several controls along a realistic attack chain is not assessed here. That is the task of the planned pattern formation (phase 2), which clusters the individual findings along recurring attack patterns.

Proportion figures are a count, not a risk score. "3 of 9 requirements affected" means that three of nine passages of text are weakened — not that the control is one third risky. A single critical gap (for example the absence of phishing resistance in authentication) may weigh more heavily than five administrative individual findings in another control. A risk-weighted prioritization is delivered not by this count but by the later gap analysis (phases 2 and 3).

What this analysis does not (yet) deliver: a prioritization of the findings by urgency or effort, and an assignment to concrete evidence types. This concrete improvement and supplementation layer — what must be changed in a requirement so that it holds again — is deliberately delivered not by this chapter but by the later hardening control synthesis (phase 4) with a subsequent linkage back into Chapters 1 to 7 (phase 5). The categories assigned here are the diagnosis; the MHC synthesis is the prescription.

What this analysis does not assess (scope delimitation): MRIS-TISAX® 2027 is an effectiveness diagnosis at requirement level. Expressly outside the scope:

- Economic viability: no cost-benefit weighing of individual hardening measures — whether a measure pays financially is for the organization to decide.

- Organization-specific risk: no risk analysis for a particular organization (protection requirement, asset criticality, likelihood of occurrence) — the organization's own ISMS risk management remains responsible for that.
- Threat intelligence: no assessment of particular attacker groups, campaigns or target sectors — the threat model below describes capabilities, not actors.
- Likelihood of occurrence: no statement about how likely a particular attack is to strike a particular organization.
- Quantitative risk calculation: no expected values, value at risk or comparable metrics. The categories assigned are ordinal (a ranking), not metric, and the threat priority score used in Part II is likewise a prioritization index, not a risk score (see Part II, section "Assessment logic").

These questions are answered by the organization's own risk analysis; MRIS supplies the diagnosis at requirement level as an input for it, not the finished result.

Controls versus actual security: MRIS assesses TISAX® requirements as text, not the actual security posture of an organization. The category “robust” means that the mechanism called for holds under Mythos conditions when correctly implemented — nothing more. Conversely, “partially degraded” does not mean that a concrete incident is imminent, but that the protective assumption anchored in the wording of the requirement should be sharpened. The categories are a prioritization tool for improving requirements, not a substitute for the security or maturity level of the organization as a whole.

Basis: four shifts in the threat landscape and four assessment criteria

MRIS-TISAX® 2027 is an independent analysis. It has formally nothing in common with any MRIS treatment of ISO 27002 beyond the objective of hardening an existing ISMS effectively and efficiently against Gen-AI-supported threats. The threat model is therefore derived directly from publicly documented incidents and research reports from 2025 and 2026, not adopted from a superordinate framework. Four shifts are evidenced in this way:

1. Collapse of the patch gap window: an evaluation of 69,159 CVE by Cogent Security, published in 2026, shows that the average time from publication of a vulnerability to the first working exploit fell from 125.3 days (January 2025) to 0.5 days (April 2026); for 62% of the critical vulnerabilities with a known exploit, that exploit existed before the scanner detection signature did.
2. Compression of the response time axis: in September 2025, Anthropic uncovered a state-directed campaign (GTG-1002) in which a hijacked AI agent system executed 80 to 90% of the attack steps against roughly 30 target organizations autonomously and at thousands of requests per second — a pace with no human-staffed counterpart.
3. Fragmentation of the threat event: the same campaign first carried out inconspicuous, “low-and-slow” reconnaissance within normal usage patterns, before the remaining steps followed at machine speed — a pattern that detection designed for human correlation of individual anomalies does not capture as a connected whole.
4. Decoupling of capability and actor: in 2026, the security vendor Sysdig documented, under the name “JadePuffer”, the first ransomware campaign executed entirely autonomously by an AI agent. Through an unpatched vulnerability in the open source platform Langflow (CVE-2025-3248, CVSS 9.8) the agent gained entry, moved laterally on its own, captured credentials for several AI and cloud providers, escalated privileges and encrypted 1,342 database and

configuration elements through more than 600 independently executed malware variants — a kill chain that until now presupposed a practised, multi-member attacker team with specialist knowledge of initial access, lateral movement and ransomware operations. The outcome is telling: the agent forgot to secure the key it had generated, and for the ransom demand adopted a Bitcoin address from its own training data — even victims willing to pay would not have got their data back. The real turning point lies precisely there: despite this beginner's error, the complete technical attack chain succeeded through to comprehensive encryption. The capability that until now presupposed an experienced attacker team is already decoupled from the actor today — executed imperfectly, but effectively, and not only in some future mature version.

Together, these four cases mark the turning point that makes hardening measures necessary in the first place: none of the four any longer describes merely a theoretical risk or a gradual intensification of existing threats. All four are events that have already occurred and been documented, in which an AI system performed a task completely or largely without the human expertise, response time or conspicuousness previously required for it. A security concept that tacitly continues to rely on human limits on the attacking side is therefore assessing a threat landscape that no longer exists in that form.

From these four shifts, four assessment criteria are derived, against which every TISAX® requirement is examined:

- Attacker patience: does the protective effect rest on an attack being too laborious or costly for the attacker?
- Time compression: does the control presuppose a human response time even though the attacker now carries out the attack in an automated and autonomous way? Delimitation: the weakness is not “periodic” as such but periodic without documented trigger logic and without risk-based event control. A “periodic review” counts as a time compression weakness only if (1) the object reviewed is a technical or operational security element whose becoming out of date opens an exploitable window (for example software approvals, access rights, backup concepts, log evaluation, supplier security status), AND (2) the wording of the requirement prescribes no supplementary event trigger. If an organization supplements a periodic review with documented, risk-based triggers (for example immediate reassessment on a new critical vulnerability), that closes the gap in practice — the wording of the requirement itself nevertheless remains classified as degraded, because it does not require that trigger. Pure legal-position scans (which laws or contracts apply) are not assessed in this way, because legislation does not change at the pace of an automated attack.
- Capability decoupling: does the assessment “hardly anyone can do that” rest on an empirical assumption that no longer holds because of AI tools?
- Aggregation resistance: does the control detect an attack even when it is decomposed into many individual steps that each appear harmless?

Why exactly four criteria: the four criteria are not an arbitrarily chosen grid width but a direct translation of the four shifts in the threat landscape evidenced above into testable questions — attacker patience translates the collapse of the patch gap window, time compression translates the compression of the response time axis, aggregation resistance translates the fragmentation of the threat event, and capability decoupling translates the decoupling of capability and actor. Each criterion is examined independently of the other three; a requirement may degrade against several criteria at once (see for example 4.1.2, where both capability decoupling and time

compression take effect — the rationale column names all applicable criteria in each case). The number four is therefore a consequence of the threat model, not a category count set in advance.

Reproducibility: every classification in the document follows the same chain: wording of the requirement → comparison against the four shifts in the threat landscape → examination of the four criteria → category (robust / partially degraded / friction-only / not affected) → assignment to one of the eight pattern clusters (phase 2) → MHC allocation (phase 4, Part II). This chain is disclosed per requirement in the rationale column, not merely in the result.

Time reference and currency: MRIS-TISAX® 2027 reflects the state of knowledge as at July 2026 — both for the wording of the TISAX® requirements (ISA 2027) and for the underlying threat model (see above: the Cogent CVE evaluation, GTG-1002, JadePuffer, the Mythos disclosure). Both sides may develop further: TISAX® updates the ISA regularly, and new AI capabilities may sharpen existing classifications or indeed soften them (for instance if phishing-resistant authentication becomes an industry-wide standard and a requirement degraded today thereby becomes robust in practice). This analysis is accordingly to be read as a snapshot, not as a one-off, permanently valid result — a reassessment is sensible as soon as either the TISAX® catalogue or the evidenced attacker capability shifts materially.

The classification of each requirement, the gap analysis and the Implementation Guide are developed independently for the TISAX® context.

Quellen:

- Cogent Research: Exploits Outpace Scanner Detection for 62% of Critical Vulnerabilities as AI Compresses Time-to-Exploit to Under 12 Hours
- Cogent Security: 2026 Q2 Detection Gap Report — Key Findings (primary source)
- AI Shrinks Cyberattack Exploit Time From Years to Days — GovInfoSecurity
- Mapping AI-enabled cyber threats — Anthropic
- AI Agents Drive First Large-Scale Autonomous Cyberattack — Cybersecurity Magazine
- The agentic shift: how autonomous AI is reshaping the global threat landscape — Control Risks
- Jadepuffer: Die erste AI-Ransomware-Attacke verbaselt nicht nur den Schlüssel — heise online
- JadePuffer: The First Successful LLM-Driven Ransomware Attack — Dark Reading
- JadePuffer ransomware used AI agent to automate entire attack — BleepingComputer
- NIST: Announcing the AI Agent Standards Initiative for Interoperable and Secure AI Agents (primary source)
- NIST: Summary and Analysis of Responses to the Request for Information Regarding Security Considerations for AI Agents (primary source)
- BSI: Cloud Computing Compliance Criteria Catalogue C5:2026 (primary document; the full PDF text could not be extracted technically in the course of this research — see footnote 2, section “Gap analysis against six comparison frameworks”)

Categories and their practical meaning

Category	Meaning	What you as CISO do with it
Robust	The security assumption of the requirement does not collapse under the four assessment	Citable as evidence in the equivalent of a Statement of

Category	Meaning	What you as CISO do with it
	criteria — not equivalent to "an attack is impossible", but rather the mechanism rests on a hard barrier (cryptography, physical separation, architecture) that holds under Mythos conditions as well. It thereby reaches the TISAX® target maturity level 3 without a supplementary measure.	Applicability (the ISA catalogue maturity entry). No additional effort from the Mythos perspective — continue auditing as normal.
Partially degraded	Reduced, not absent, protective effect relative to the original security assumption: the mechanism continues to work but is no longer sufficient to sustain maturity level 3 effectively against a Mythos attacker.	A candidate for the risk treatment plan. Needs a supplementary hardening measure (catalogue follows in phase 4). This is where additional budget and attention belong.
Friction-only	Loses its core effect against a Mythos attacker entirely — the security assumption of the requirement no longer holds, irrespective of the maturity level achieved.	Must be replaced or fundamentally redesigned — a mere supplement is not sufficient.
Not affected	Influenced by none of the four assessment criteria (attacker patience, time compression, capability decoupling, aggregation resistance) — not a statement that the requirement is "secure" in general, but that it remains unchanged relative to this specific shift in the threat landscape.	Continue to examine and implement in normal ISMS operation, but deliberately do not prioritize when allocating resources for the Mythos response. No compensation required.

Order of precedence for mixed findings within a control: friction-only > partially degraded > robust > not affected. If a control contains both a friction requirement and several degraded requirements, the overall classification is reported as “friction-only” — the weakest individual finding determines the overall picture, not the average.

Unit of assessment and structure per control

Every TISAX® control (for example 5.2.9) contains several individual requirements marked “+”, graded by protection requirement level: must and should always apply (normal protection requirement), while high and very high are added where the affected asset has a correspondingly higher protection requirement — they do not replace must and should but supplement them. Each individual requirement is assessed on its own, in the original English wording of the source (citable, searchable).

Six controls in Chapter 1 (1.1.1, 1.2.1, 1.2.2, 1.5.1, 1.5.2, 1.6.2) additionally contain SGA requirements (“Additional requirements for Simplified Group Assessments”): these apply irrespective of the protection requirement, and only where the organization uses the simplified TISAX® group assessment (several similar sites assessed together). They are marked in the tables as a separate level “SGA” and are NOT included in the must/should/high/very high counts — the proportion figure per control reports SGA requirements separately.

At the end of each control there is an overall classification: for most controls it is identical across all three protection requirement levels and is summarized in a single sentence with a proportion figure (“3 of 9 requirements affected”). For a few controls the picture changes as the protection requirement rises — there it is spelled out which additional requirement triggers the difference and what that means for you in practice.

Responsibility column: for degraded, friction and robust requirements, the function typically responsible is additionally stated (IT operations, CISO function/ISMS, development/engineering, IAM team, procurement/supplier management, human resources, facility/site security). This is an organizational assignment, not a binding RACI — it serves as a first indication of who typically works on the finding.

Data quality: control 1.6.3, column “high”, contains five strikingly similar pairs of requirements in the source catalogue. Checked against the original ISA 2027 table, exactly one of those pairs is word-identical (counted once); the remaining four differ editorially from one another (among other things the typographical error “outage of” versus “outage oaf”, and the wording “Methods” versus “Instruments to monitor communication”) and are counted as distinct requirements in their own right (see control 1.6.3 below). Control 3.1.2 contains no requirements of its own (“Superseded by 1.6.3, 5.2.8 and 5.2.9”).

Chapter 1 — IS Policies and Organization (14 Controls, 116 distinct requirements including 10 SGA-requirements)

1.1.1 — To what extent are information security policies available?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for information security have been determined, documented and adapted to the organization's goals."	Not affected	Governance basis of the policy, no attack relevance.	CISO function
must	"A policy exists, and is approved by the organization's management."	Not affected	Act of approval, no defensive function.	CISO function
must	"The policy includes objectives and the significance of information security within the organization."	Not affected	Minimum content requirement for the document itself.	CISO function
must	"The policies are made available to employees in a suitable form (e.g., intranet)."	Not affected	Distribution obligation, no attack relevance.	CISO function
must	"Employees and external business partners are informed of any changes relevant to them."	Not affected	Communication obligation, no attack relevance.	CISO function
should	"The information security requirements are based on the strategy of the organization. Legislation and contracts are considered in the policy."	Not affected	Substantive orientation of the policy, no technical protective effect.	CISO function
should	"The policy indicates consequences in case of non-conformance."	Not affected	Sanction clause, organizational.	CISO function
should	"Other relevant security policies are established."	Not affected	Completeness requirement for the	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
			body of policies.	
should	"Periodic review and, if required, revision of the policies are established."	Partially degraded	Time compression: “periodic” without a concrete review trigger. If new attack patterns emerge within a few weeks while a fixed review cycle runs, say, annually, the policy remains out of date in the meantime.	CISO function
SGA	"Policies are published and implemented in the entire assessment scope."	Not affected	Coverage requirement for group assessments (all sites within scope), no defensive function.	CISO function

Overall classification: Partially degraded (1 of 10 requirements affected), irrespective of protection requirement. Under a Simplified Group Assessment this additionally applies to all sites within scope.
Hardening reference: MHC-10 (see Part II)

1.2.1 — To what extent is information security managed within the organization?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The scope of the ISMS (the organization managed by the ISMS) is defined."	Not affected	Determination of scope, governance.	CISO function
must	"The organization's requirements for the ISMS are determined."	Not affected	Governance.	CISO function
must	"The organizational management has commissioned and approved the ISMS."	Not affected	Act of approval by management.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The ISMS provides the organization's management with suitable monitoring and control means (e.g., management review)."	Not affected	Requires only the existence of a steering instrument, no frequency specification — the frequency question follows separately.	CISO function
must	"Applicable controls have been determined (e.g., ISO 27001 Statement of Applicability, or completed ISA catalogue)."	Not affected	Scoping governance.	CISO function
must	"The effectiveness of the ISMS is regularly reviewed by the management."	Partially degraded	Time compression: an annual management review rhythm, as is usual in practice, responds structurally too slowly to a threat landscape that changes within weeks.	CISO function
SGA	"The management system is approved by an entity that has the necessary authority for the entire assessment scope (i.e., all locations within the scope)."	Not affected	Approval governance for the group-wide scope.	CISO function

Overall classification: Partially degraded (1 of 7 requirements affected), irrespective of protection requirement. Under a Simplified Group Assessment this additionally applies to all sites within scope.
Hardening reference: MHC-10 (see Part II)

1.2.2 — To what extent are information security responsibilities organized?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Responsibilities for information security within the organization are defined, documented, and assigned."	Not affected	Governance.	CISO function
must	"The responsible employees are defined, qualified, and empowered for their task."	Not affected	Governance and personnel.	CISO function
must	"The required resources are available."	Not affected	Resource commitment, no defensive function.	CISO function
must	"The contact persons are known within the organization and to relevant business partners."	Not affected	Awareness of points of contact, organizational.	CISO function
should	"There is a definition and documentation of an adequate information security structure within the organization."	Not affected	Governance.	CISO function
should	"Security-relevant roles that are not part of the ISMS and are relevant for information security are considered."	Not affected	Completeness of scoping; a possible point of connection for “shadow AI roles” in Chapters 4 and 5, but as yet no sufficiently concrete textual basis for a different classification.	CISO function
high	"An appropriate organizational separation of responsibilities should be established in order to avoid	Partially degraded	Capability decoupling: separation of duties has	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
	conflict of interests (segregation of duties). (C, I, A)"		protected until now because one person can hardly master two specialist roles at once. With AI assistance a single person can convincingly carry out the tasks of a second role as well, where the separation is only organizational and not additionally enforced through system rights.	
SGA	"A named person with overall responsibility for the management system exists and is available."	Not affected	Governance and accountability, group-wide.	CISO function

Overall classification: Unremarkable at normal protection requirement (Not affected, 0 of 7 including SGA) — no action needed. From high protection requirement onward, the segregation-of-duties obligation is added, which shows a weakness without technical enforcement: Partially degraded (1 of 8 requirements affected). The SGA requirement applies additionally irrespective of protection requirement and does not change this. In practice: relevant to you only if this control applies to assets classified as high or very high.
Hardening reference: standalone measure (see Part II)

1.2.3 — To what extent are information security requirements considered in projects?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Projects are classified while taking into account the information security	Not affected	Governance and project control.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	requirements."			
should	"The procedure and criteria for the classification of projects are documented."	Not affected	Documentation obligation.	CISO function
should	"During an early stage of the project, risk assessment is conducted based on the defined procedure and repeated in case of changes to the project."	Not affected	An explicit double trigger (“early in the project AND repeated on changes”) dispels the time compression concern.	CISO function
should	"For identified information security risks, measures are derived and considered in the project."	Not affected	Governance.	CISO function
high	"The measures thus derived are reviewed regularly during the project and reassessed in case of changes to the assessment criteria. (C, I, A)"	Not affected	Likewise a double trigger (“regularly AND on changes”) is present.	CISO function

Overall classification: Not affected (0 of 5 requirements affected), irrespective of protection requirement. Consistently double triggers — a good example of wording that structurally dispels the time compression concern.

1.3.1 — To what extent are information assets identified and recorded?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Information assets and other assets where security is relevant to the organization are identified and recorded."	Not affected	Inventory governance.	CISO function
must	"The supporting assets processing the information assets are identified and recorded."	Not affected	Inventory governance.	CISO function
should	"A catalogue of the relevant information assets exists. The following aspects are considered: the corresponding supporting assets are assigned to each relevant information asset, a person responsible for these information assets is assigned, the catalogue is subject to regular review."	Partially degraded	Time compression: without a fixed review trigger, the inventory remains incomplete between reviews. Newly created, unrecorded systems (for instance independently set up cloud or AI services) fall precisely into that window — which is where automated attacks strike first.	IT operations

Overall classification: Partially degraded (1 of 3 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-10 (see Part II)

1.3.2 — To what extent are information assets classified and managed in terms of their protection needs?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A consistent scheme for the classification of information assets regarding the protection goal of confidentiality is available."	Not affected	A pure classification scheme; the effect arises only in downstream technical controls.	CISO function
must	"Evaluation of the identified information assets is carried out according to the defined criteria, and assigned to the existing classification scheme."	Not affected	Application of the scheme, governance.	CISO function
must	"Specifications for the handling of supporting assets (e.g., identification, usage, transport, storage, return, deletion/disposal) depending on the classification of information assets are in place and implemented."	Not affected	Governance requirement; the technical effect lies in the referenced controls themselves.	CISO function
should	"The protection goals of integrity and availability are taken into consideration."	Not affected	Completeness of the scheme.	CISO function

Overall classification: Not affected (0 of 4 requirements affected), irrespective of protection requirement.

1.3.3 — To what extent is it ensured that only evaluated and approved external IT services are used for processing the organization's information assets?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"External IT services are not used without explicit assessment and	Partially degraded	Capability decoupling: the approval obligation	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	implementation of the information security requirements. The following aspects are considered: availability of a risk assessment of the external IT services, legal, regulatory, and contractual requirements."		presupposes that a person deliberately selects a particular external service and submits it for review. AI assistants with tool access can, however, call external services on their own in the course of a task, without anyone ever having submitted that particular service for approval. No technical block (for example network or access control) is named in the wording.	
must	"The external IT services have been harmonized with the protection need of the processed information assets."	Not affected	Checks whether the classification of the service matches the protection requirement of the data — an allocation decision, not a defensive measure.	CISO function
should	"Requirements regarding the procurement, commissioning and release associated with the use of external IT services are determined and fulfilled."	Not affected	Procurement requirement.	Procurement/supplier management
should	"A procedure for release in consideration of the	Not affected	Procedural requirement for the	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	protection need is established."		approval.	
should	"External IT services and their approval are documented."	Not affected	Documentation obligation.	CISO function
should	"It is verified at regular intervals that only approved external IT services are used."	Partially degraded	Time compression: a sample-based check at fixed intervals leaves a window in which an unapproved service remains undetected — without continuous automated detection, larger than the response time of an automated attack.	IT operations

Overall classification: Partially degraded (2 of 6 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-10, MHC-13 (see Part II)

1.3.4 — To what extent is it ensured that only evaluated and approved software is used for processing the organization's information assets?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Software is approved before installation or use. The following aspects are considered: limited approval for specific use-cases or roles, conformance to the information security requirements, software use rights and licensing, source/reputation of the software."	Partially degraded	Capability decoupling: an AI-supported coding assistant can load software packages on its own in the course of a task, without a human deliberately submitting that particular package for approval.	Development/engineering

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Software approval also applies to special purpose software such as maintenance tools."	Partially degraded	Same rationale as above, here merely a clarification of scope.	Development/engineering
should	"The types of software such as firmware, operating systems, applications, libraries, device drivers to be managed are determined."	Not affected	Scoping governance.	IT operations
should	"Repositories of managed software exist."	Not affected	Existence of infrastructure.	IT operations
should	"The software repositories are protected against unauthorized manipulation."	Partially degraded	The strength of the protection is left open: a mere access list is unlikely to withstand a sufficiently capable automated attack, whereas a cryptographic signature check would — the wording does not commit itself.	IT operations
should	"Approval of software is regularly reviewed."	Partially degraded	Time compression: “regularly” without a concrete review trigger.	IT operations
should	"Software versions and patch levels are known."	Partially degraded	Mere knowledge of version and patch status says nothing about the speed of response. Where only hours now lie between publication of a vulnerability and a working attack,	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
			knowledge alone is not sufficient.	
very high	"Additional requirements for software use (e.g., need for control or monitoring of usage) are determined (if any) (C, I, A)"	Not affected	Worded too indefinitely (“if applicable”) for a concrete assessment.	IT operations

Overall classification: Partially degraded (5 of 8 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-10, MHC-09, MHC-13 (see Part II)

1.4.1 — To what extent are information security risks managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Risk assessments are carried out both at regular intervals and in response to events."	Not affected	A double trigger dispels the time compression concern.	CISO function
must	"Information security risks are appropriately assessed (e.g., for probability of occurrence and potential damage)."	Partially degraded	Capability decoupling: classic estimates of likelihood of occurrence often assume that certain attacks are accessible only to a few highly specialized actors. That assumption no longer holds where AI tools make the same capability broadly accessible. No updated assessment methodology is prescribed.	CISO function
must	"Information security risks are	Not affected	Documentation	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	documented."		obligation.	
must	"A responsible person (risk owner) is assigned to each information security risk."	Not affected	Governance.	CISO function
should	"A procedure is in place defining how to identify, assess and address security risks within the organization."	Not affected	Governance.	CISO function
should	"Criteria for the assessment and handling of security risks exist."	Partially degraded	Same rationale as above, here at the level of the assessment criteria rather than the result.	CISO function
should	"Measures for handling security risks and the persons responsible for these are specified and documented."	Not affected	Governance and tracking.	CISO function
should	"In case of changes to the environment (e.g., organizational structure, location, changes to regulations), reassessment is carried out in a timely manner."	Not affected	Event-driven trigger present.	CISO function

Overall classification: Partially degraded (2 of 8 requirements affected), irrespective of protection requirement.
Hardening reference: standalone measure (see Part II)

1.5.1 — To what extent is compliance with information security ensured in procedures and processes?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Observation of policies is verified throughout the organization."	Not affected	Existence of the verification, governance.	CISO function
must	"Information security policies and procedures are reviewed at regular intervals."	Partially degraded	Time compression: “regularly” without a concrete review	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
			trigger.	
must	"Measures for correcting potential non-conformities (deviations) are initiated and pursued."	Not affected	Corrective process, governance.	CISO function
must	"Compliance with information security requirements (e.g., technical specifications) is verified at regular intervals."	Partially degraded	Same rationale as above.	CISO function
must	"The results of the conducted reviews are recorded and retained."	Not affected	Documentation obligation.	CISO function
should	"A plan for content and framework conditions (time schedule, scope, controls) of the reviews to be conducted is provided."	Not affected	Planning artefact.	CISO function
SGA	"Internal controls are implemented for all entities within the assessment scope."	Not affected	Coverage requirement, governance.	CISO function
SGA	"The implementation of all applicable security requirements and control objectives are included."	Not affected	Completeness requirement, governance.	CISO function
SGA	"Suitably qualified and appropriate information security audit responsibilities and resources are defined."	Not affected	Resource and qualification governance.	CISO function
SGA	"A detailed internal audit plan and schedule is defined and followed. [Entire assessment scope covered, Internal audits conducted regularly, Results tracked within ISMS structures]"	Partially degraded	Time compression: the same “regularly without an event trigger” finding as for the must element above, here applied to the group-wide audit plan — newly added sites are not necessarily brought	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
			into the audit plan promptly without a trigger.	

Overall classification: Partially degraded (3 of 10 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-10 (see Part II)

1.5.2 — To what extent is the ISMS reviewed by an independent authority?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Information security reviews are carried out by an independent and competent body at regular intervals and after fundamental changes."	Not affected	A double trigger (“regularly AND after material changes”) — a counter-example to 1.5.1, showing that the concrete wording decides, not the subject as such.	CISO function
must	"Measures for correcting potential deviations are initiated and pursued."	Not affected	Corrective process.	CISO function
should	"The results of conducted reviews are documented and reported to the management of the organization."	Not affected	Documentation and reporting.	CISO function
SGA	"In case of fundamental changes, audits are conducted by an independent and appropriately qualified entity (i.e., external entities or special internal departments which are objective, competent and free from undue influence (independent))."	Not affected	Event-triggered (“in case of fundamental changes”), well worded — same rationale as for the must element above.	CISO function
SGA	"Findings and implementation of corrective actions is tracked by the independent entity."	Not affected	Tracking process.	CISO function

Overall classification: Not affected (0 of 5 requirements affected), irrespective of protection requirement.

1.6.1 — To what extent are information security relevant events or observations reported?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A definition for a reportable security event or observation exists and is known by employees and relevant stakeholders."	Not affected	Definitional governance.	CISO function
must	"Adequate mechanisms based on perceived risks to report security events are defined, implemented, and known to all relevant potential reporters."	Partially degraded	Aggregation resistance: a reporting system that depends on a person noticing something conspicuous does not recognize an attack deliberately decomposed into many small steps that each appear harmless — no single step stands out.	IT operations
must	"Adequate channels for communication with event reporters exist."	Not affected	Infrastructure and process.	CISO function
should	"A common point of contact for event reporting exists."	Not affected	Organizational.	CISO function
should	"Different reporting channels according to perceived severity ... are available."	Not affected	Process and infrastructure.	CISO function
should	"Employees are obligated and trained to report relevant events."	Not affected	Training obligation.	Human resources
should	"Security event reports can also be submitted by external	Not affected	Process and channel.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	parties."			
should	"Mechanism to, and information how to, report incidents is accessible by all relevant reporters."	Not affected	Accessibility.	CISO function
should	"A feedback procedure to reporters is established."	Not affected	Process.	CISO function
very high	"Tests and exercises of event and observation reporting are conducted regularly. (C, I, A)"	Not affected	Testing activity — consistently not treated like live detection (see 1.6.2).	CISO function

Overall classification: Partially degraded (1 of 9 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-05 (see Part II)

1.6.2 — To what extent are reported security events managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Reported events are processed without undue delay."	Partially degraded	Time compression: presupposes that people can handle the processing at the necessary pace. Where an attacker automates the greater part of their approach and carries it out without human involvement, purely human-borne processing structurally cannot keep up — automation is not required.	IT operations
must	"An adequate reaction to reported security events is ensured."	Partially degraded	Same rationale as above.	IT operations
must	"Lessons learned are incorporated into continuous improvement."	Not affected	Improvement process.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
should	"During processing, reported events are categorized, qualified, and prioritized."	Not affected	Process and methodology.	IT operations
should	"Responsibilities for handling of events based on their category are defined and assigned."	Not affected	Governance.	CISO function
should	"A strategy for reporting potentially criminal aspects of security incidents to relevant authorities when necessary or appropriate."	Not affected	Legal process.	Legal
high	"Maximum response times based on class, category and severity are defined."	Partially degraded	Defining a target does not mean that it can also be met against an attacker acting in an automated and rapid way — an obligation to automate is missing for that.	IT operations
high	"Events not processed appropriately according to their priority are escalated."	Partially degraded	A sensible safety net, but itself still human-borne, and it does not resolve the time compression gap of the previous row.	IT operations
high	"Lawful, regulatory, and contractual reporting obligations, and respective contact information are known."	Not affected	Legal knowledge.	Legal
high	"A communication strategy for security related events exists."	Not affected	Communication planning.	CISO function
high	"Procedures for response to supplier	Not affected	Process, examined in depth in Chapter 6.	Procurement/supplier

Level	Requirement (original wording)	Category	Rationale	Responsible
	security incidents are established."			management
very high	"Handling of events in different categories and priorities is regularly tested."	Not affected	Testing activity.	CISO function
SGA	"Standard mechanisms to report and track relevant security events are established."	Not affected	Standardization requirement for group-wide consistency; the speed of detection and response itself is already assessed above.	CISO function

Overall classification: Partially degraded (4 of 13 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-11 (see Part II)

1.6.3 — To what extent is the organization prepared to handle crisis situations?

Data quality note: in the source catalogue, the “high” column contains five strikingly similar pairs of requirements. Only one pair is word-identical (counted once, 19 rather than 20 raw rows); the other four differ editorially from one another (typographical errors or divergent wording in the source catalogue itself, among them “outage of” versus “outage oaf”, “Methods” versus “Instruments to monitor communication”) and are carried below as distinct rows.

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"An appropriate plan to react to and recover from crisis situations exists and the required resources are available."	Not affected	Existence of planning.	CISO function
must	"Responsibilities and authority for crisis management within the organization are defined, documented, and assigned."	Not affected	Governance.	CISO function
must	"The responsible employees are defined and qualified for	Not affected	Governance and personnel.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	their task."			
should	"Methods to detect crisis situations are established. General indications for the existence or imminence of a crisis situation and specific predictable crisis are identified."	Partially degraded	Aggregation resistance and time compression: classic crisis detection is designed for gradually escalating situations. An IT incident that grows from many small, unremarkable steps into a genuine crisis within hours is likely to be recognized too late.	IT operations
should	"A procedure to invoke and/or escalate crisis management is in place."	Not affected	Process.	CISO function
should	"Strategic goals and their priority in crisis situations are defined and known to relevant personnel."	Not affected	Governance.	CISO function
should	"A crisis management team is defined and approved."	Not affected	Governance.	CISO function
should	"Crisis policies and procedures are defined and approved."	Not affected	Governance.	CISO function
should	"Crisis planning is reviewed and updated regularly."	Partially degraded	Time compression: “regularly” without a concrete review trigger.	CISO function
high	"Relevant different potential crisis scenarios are identified."	Partially degraded	The example scenarios named (failure of personnel, resources or infrastructure) are classic outage scenarios; an AI-orchestrated, multi-track attack with rapid escalation is not expressly provided for	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
			among them.	
high	"Relevant different potential crisis scenarios are identified." (second catalogue version, see data quality note)	Partially degraded	A substantively identical requirement, carried twice in the source catalogue with editorially divergent wording (“outage of” versus “outage oaf”) — same rationale as the preceding row.	CISO function
high	"Necessary resources and information to handle crisis ... are identified."	Not affected	Resource planning.	CISO function
high	"Necessary resources and information to handle crisis ... are identified." (second catalogue version, see data quality note)	Not affected	Resource planning; editorial duplicate of the preceding row.	CISO function
high	"A communication strategy for crisis situations exists."	Partially degraded	The same time compression concern as for crisis detection: no reference to a sharply shortened response window.	CISO function
high	"A communication strategy for crisis situations exists." (second catalogue version, see data quality note)	Partially degraded	A substantively identical requirement, carried twice in the source catalogue with editorially divergent wording (“Methods” versus “Instruments to monitor communication”) — same rationale as the preceding row.	CISO function
high	"The efficiency, feasibility, and appropriateness of the crisis planning is evaluated	Partially degraded	Time compression: “regularly” without a concrete review trigger.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	regularly."			
high	"Spot based testing of crisis planning is conducted."	Not affected	Testing activity.	CISO function
high	"Spot based testing of crisis planning is conducted." (second catalogue version, see data quality note)	Not affected	Testing activity; editorial duplicate of the preceding row.	CISO function
very high	"Crisis exercises and simulations involving all relevant persons, including decision makers are conducted regularly."	Not affected	Testing activity.	CISO function

Overall classification: Partially degraded (7 of 19 requirements affected), irrespective of protection requirement.
Hardening reference: MHC-10, MHC-05, standalone measure (see Part II)

Chapter 1 summary

Control	Overall classification	Share affected
1.1.1	Partially degraded	1/10
1.2.1	Partially degraded	1/7
1.2.2	Not affected → from high onward: Partially degraded	0/7 → 1/8
1.2.3	Not affected	0/5
1.3.1	Partially degraded	1/3
1.3.2	Not affected	0/4
1.3.3	Partially degraded	2/6
1.3.4	Partially degraded	5/8
1.4.1	Partially degraded	2/8
1.5.1	Partially degraded	3/10
1.5.2	Not affected	0/5
1.6.1	Partially degraded	1/9
1.6.2	Partially degraded	4/13
1.6.3	Partially degraded	7/19

10 of 14 controls consistently degraded, 3 consistently not affected, 1 dependent on protection requirement. No robust, no friction-only — to be expected, since Chapter 1 is a purely governance chapter with no cryptographic or architectural hard barriers. Two recurring patterns carry the high degraded share: (1) “regularly” or “periodically” without a concrete review trigger, (2) human-borne reporting or response speed against an attacker operating in automated form. Responsibility predominantly the CISO function and ISMS processes — consistent with this being the governance chapter. SGA requirements (6 controls affected, 10 additional requirements, of which 1 degraded) do not change the overall classification of any of the 14 controls — they extend the scope to group level without introducing new weaknesses.

Chapter 2 — Human Resources (4 Controls, 21 requirements)

2.1.1 — To what extent is the qualification of employees for sensitive work fields ensured?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Sensitive work fields and jobs are determined."	Not affected	Scoping governance.	Human resources
must	"The requirements for employees with respect to their job profiles are determined and fulfilled."	Not affected	Governance.	Human resources
must	"The identity of potential employees is verified (e.g., checking identity documents)."	Partially degraded	Capability decoupling: convincing forged identity documents and — in combination with video interviews — deepfake identities were until now regarded as specialist knowledge. AI-generated document forgery and real-time video and audio deepfakes make deceptively realistic identity pretence in the hiring process broadly accessible	Human resources

Level	Requirement (original wording)	Category	Rationale	Responsible
			today (a known current pattern: fictitious remote employees).	
should	"The personal suitability of potential employees is verified by means of simple methods (e.g., job interview)."	Partially degraded	Same rationale: a live interview as a verification method is precisely the target of real-time deepfake impersonation.	Human resources
should	"An extended suitability verification depending on the respective work field and job is conducted (e.g., assessment centre, psychological analysis, checking of references, certificates and diploma, checking of certificates of conduct, checking of professional and private background)."	Partially degraded	Capability decoupling: references, testimonials and diplomas can be forged more convincingly with AI tools than before; fictitious referees and companies are also easier to present.	Human resources

Overall classification: Partially degraded (3 of 5 requirements affected), irrespective of protection requirement. Recurring theme: all three identity and eligibility checks rely on human verification based on “authenticity markers” that are no longer reliable today.
Hardening reference: MHC-15 (see Part II)

2.1.2 — To what extent is all staff contractually bound to comply with information security policies?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A non-disclosure obligation is in effect."	Not affected	Contract law, no attack relevance.	Legal
must	"An obligation to comply with the information security policies is in effect."	Not affected	Contract law.	Legal

Level	Requirement (original wording)	Category	Rationale	Responsible
should	"A non-disclosure obligation beyond the employment contract is in effect."	Not affected	Contract law.	Legal
should	"Information security aspects are considered in the employment contracts of the staff."	Not affected	Contract law.	Legal
should	"A procedure for handling violations of said obligations is described."	Not affected	Process and governance.	Human resources

Overall classification: Not affected (0 of 5 requirements affected), irrespective of protection requirement. Pure contract law with no technical or behavioural protective effect — untouched by Mythos.

2.1.3 — To what extent is staff made aware of and trained with respect to the risks arising from the handling of information?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Employees are trained and made aware."	Partially degraded	Capability decoupling: classic awareness content trains people to recognize spelling mistakes, impersonal salutations and other “tells” of phishing and social engineering. AI-generated attacks are error-free, personalized and linguistically inconspicuous — the trained recognition features no longer work reliably.	Human resources
should	"A concept for awareness and training of employees	Not affected	Existence of a structured concept, governance.	CISO function

Level	Requirement (original wording)	Category	Rationale	Responsible
	is prepared. [Policy, Event reporting, Malware reaction, Password policy, Compliance, NDA use, External IT services use]"			
should	"Target groups for training and awareness measures... are identified and considered in a training concept."	Not affected	Scoping governance.	CISO function
should	"The concept has been approved by the responsible management."	Not affected	Act of approval.	CISO function
should	"Training and awareness measures are carried out both at regular intervals and in response to events."	Not affected	A double trigger dispels the time compression concern.	Human resources
should	"Participation in training and awareness measures is documented."	Not affected	Documentation obligation.	Human resources
should	"Contact persons for information security are known to employees."	Not affected	Communication and governance.	CISO function

Overall classification: Partially degraded (1 of 7 requirements affected), irrespective of protection requirement. The finding concerns not the existence or organization of training but its content: training concepts that do not expressly convey that classic recognition markers of deception attempts no longer work leave staff with a false sense of security.
Hardening reference: MHC-15 (see Part II)

2.1.4 — To what extent is mobile work regulated?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for mobile work are determined and fulfilled. [Behavior in private surroundings, Behavior in public surroundings, Secure handling of and access to information, Protection from theft, Measures to ensure a secure access to organization's IT-services]"	Not affected	Physical and behavioural requirements; technical access protection is assessed separately in 4.1.2 and 5.1.2, no double counting here.	IT operations
should	"[Measures for travelling e.g., viewing by authorities, Measures for travelling to security-critical countries]"	Not affected	Travel and physical security, no specific Mythos relevance.	CISO function
should	"Employee awareness."	Not affected	In the context of this control this refers to physical travel risks (shoulder surfing, theft), not to the general phishing and social engineering training in 2.1.3.	Human resources
high	"Protective measures against overhearing and viewing are implemented. (C)"	Robust	Physical and technical countermeasure (privacy screens, shielded rooms) — a hard barrier that works irrespective of attacker capability.	IT operations

Overall classification: Not affected at normal protection requirement (0 of 3 requirements affected) — no action needed. From high protection requirement onward, a physical robust requirement is added (visual and acoustic protection); since no degraded requirement exists, the overall classification from high onward becomes Robust. In practice: for mobile-work scenarios with high or very high protection requirements, physical shielding is demonstrably effective and citable as SoA evidence.

Chapter 2 summary

Control	Overall classification	Share affected
2.1.1	Partially degraded	3/5
2.1.2	Not affected	0/5
2.1.3	Partially degraded	1/7
2.1.4	Not affected → from high onward: Robust	0/3 → Standfest

Chapter 2 shows a new pattern compared with Chapter 1: the degraded findings arise not from vague review cycles but from human identity and eligibility verification (the hiring process, awareness content) that rests on recognition markers no longer reliable today — the same capability decoupling evidenced in the Methodology through the JadePuffer case appears here at the level of identity verification rather than technical attack execution. The catalogue's first robust example (physical visual and acoustic protection).

Chapter 3 — Physical Security (4 Controls, 14 requirements, of which 1 Control with no requirements of its own)

3.1.1 — To what extent are security zones managed to protect information assets?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A security zone concept including the associated protective measures based on the requirements for the handling of information assets is in place. [Physical conditions in the definition of security zones, Delivery and shipping areas]"	Not affected	Existence of a concept, governance; the effectiveness of the measures themselves is assessed in the following rows.	Facility/site security
must	"The defined protective measures are implemented."	Partially degraded	The wording leaves open WHICH measures are implemented. In so far as access control rests on human checking (reception, security service), it is vulnerable to AI-supported pretence (forged badges, convincing cover stories by telephone) — capability decoupling.	Facility/site security
must	"The code of conduct for security zones is known to all persons involved."	Not affected	Communication of rules of conduct.	Facility/site security
should	"Procedures for allocation and revocation of access	Not affected	Administrative process for	Facility/site security

Level	Requirement (original wording)	Category	Rationale	Responsible
	rights are established."		physical access rights.	
should	"Visitor management policies (including registration and escorting of visitors) are defined."	Partially degraded	Capability decoupling: visitor registration and escorting rest on examining a pretended identity and appointment (a forged supplier credential, a forged appointment confirmation by telephone or email) — easier to stage today.	Facility/site security
should	"Policies for carrying along and using mobile IT devices and mobile data storage devices... are defined and implemented."	Not affected	Device policy, examined technically in 3.1.4.	IT operations
should	"Network/infrastructure components (own or customer networks) are protected against unauthorized access."	Robust	Physical locking or enclosure of network components (server room, distribution cabinet) is a hard barrier irrespective of attacker capability; the human access decision is already assessed in the preceding rows (no double counting).	Facility/site security
should	"External properties used for storing and processing information assets are	Not affected	Completeness of scoping.	Facility/site security

Level	Requirement (original wording)	Category	Rationale	Responsible
	considered in the security zone concept..."			
high	"Protective measures against simple overhearing and viewing are implemented. (C)"	Robust	Physical and architectural measure, hard and independent of attacker capability.	Facility/site security

Overall classification: Partially degraded (2 of 9 requirements affected), irrespective of protection requirement. Two robust elements (enclosure of network components, visual and acoustic protection from high onward) exist in parallel but do not change the overall classification, because human-mediated access and visitor checking remains a weakness.
Hardening reference: MHC-15 (see Part II)

3.1.2 — Superseded by 1.6.3, 5.2.8 and 5.2.9

No requirements of its own — the content is fully covered in the controls named. No classification required.

3.1.3 — To what extent is the handling of supporting assets managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for the handling of supporting assets (e.g., transport, storage, repair, loss, return, disposal) are determined and fulfilled."	Not affected	Governance requirement for the life cycle of physical assets; the technical strength of disposal is assessed separately below.	IT operations
high	"Disposal of supporting assets is conducted in accordance with one of the relevant standards (e.g. ISO 21964, at least security level 4). (C)"	Robust	A physical destruction standard (for example a defined shredder particle size) is a hard, verifiable process — AI-	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
			supported attacker capability does not help in reconstructing properly destroyed data.	

Overall classification: Not affected at normal protection requirement (0 of 1) — only a general process requirement. From high protection requirement onward, a physical destruction standard is added, and the overall classification becomes Robust. In practice: a positive example — TISAX® expressly requires a robust, Mythos-resistant disposal method at higher protection requirements.

3.1.4 — To what extent is the handling of mobile IT devices and mobile data storage devices managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for mobile IT devices and mobile data storage devices are determined and fulfilled. [Encryption, Access protection e.g. PIN/password, Marking]"	Partially degraded	"Encryption" is only one of several aspects named as equally ranked, not mandatory and comprehensively prescribed (that happens only at high). Where encryption is absent or only partial, PIN or password protection remains the effective barrier — vulnerable to automated or AI-supported	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
			guessing and to social engineering.	
should	"Registration of the IT devices."	Not affected	Inventorying.	IT operations
high	"General encryption of mobile data storage devices or the information assets stored thereon. Where this is technically not feasible, information is protected by similarly effective measures. (C, I)"	Robust	Requires general (comprehensive) encryption — a hard cryptographic barrier that closes the gap left open at the must level.	IT operations

Overall classification: Partially degraded at normal protection requirement (1 of 2), because only an incomplete encryption obligation exists. From high protection requirement onward, the obligation of general encryption closes precisely this gap — the overall classification becomes Robust. In practice: devices at normal protection requirement remain structurally less well protected than those at high or very high protection requirement; check whether “normal” devices in your estate are in fact fully encrypted regardless.
Hardening reference: standalone measure (see Part II)

Chapter 3 summary

Control	Overall classification	Share affected
3.1.1	Partially degraded	2/9
3.1.2	(no requirements)	–
3.1.3	Not affected → from high onward: Robust	0/1 → Standfest
3.1.4	Partially degraded → from high onward: Robust	1/2 → Standfest

Physical Security delivers the densest concentration of robust findings to date: physical destruction, general encryption, visual and acoustic protection, and enclosure of network equipment are all hard barriers independent of attacker capability. The remaining degraded findings arise consistently at the same point: human-mediated access and visitor checking.

Chapter 4 — Identity and Access Management (4 Controls, 37 requirements)

4.1.1 — To what extent is the use of identification means managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for the handling of identification means over the entire lifecycle are determined and fulfilled. [Creation/handover/return/destruction, Validity periods, Traceability, Handling of loss]"	Not affected	Life cycle governance for badges and tokens; actual forgery resistance is the subject of other controls (cryptography: 5.1.1).	IAM team
should	"Identification means can be produced under controlled conditions only."	Not affected	Controls the legitimate issuing channel (only authorized bodies produce badges) — a governance control over the production process, not a statement about the detection of forgeries by third	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
			parties.	
high	"A strategy of blocking or invalidation of identification means in case of loss is prepared and implemented as far as possible. (C, I, A)"	Partially degraded	Time compression: “as far as possible” without a mandatory immediate or automated revocation leaves a window in which a lost or copied credential remains valid — an attacker acting in an automated way exploits precisely that window faster than a manual revocation can take effect.	IAM team

Overall classification: Not affected at normal protection requirement (0 of 2 requirements affected). From high protection requirement onward, the lock and invalidation strategy is added, whose non-binding wording (“as far as possible”) opens a time compression gap — the overall classification becomes Partially degraded (1 of 3). In practice: check whether your lock processes for lost badges or tokens actually take effect automatically and immediately, rather than relying on the discretionary wording.
Hardening reference: MHC-03 (see Part II)

4.1.2 — To what extent is the user access to IT services and IT systems secured?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The procedures for user authentication have been selected based on a risk assessment. Possible attack scenarios have been considered (e.g., direct accessibility via the internet)."	Not affected	Selection methodology; the strength of the method chosen is assessed separately in the following row.	IAM team
must	"State of the art procedures for user authentication are applied."	Partially degraded	Capability decoupling and time compression: “state of the art” does not expressly require phishing-resistant methods (for example FIDO2/WebAuthn). Automated real-time phishing relays (adversary-in-the-middle) defeat OTP- and push-based methods that many organizations still regard as “state of the art”. Target state: a mandatory use of phishing-resistant methods would raise this requirement to robust level, because the protective assumption would then hold structurally rather than merely gradually.	IAM team
should	"The user authentication procedures are defined and implemented based on the	Not affected	Process for deriving requirements.	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
	business-related and security-relevant requirements"			
should	"Users are authenticated at least by means of strong passwords according to proven and widely accepted practices."	Friction-only	The protective assumption of this requirement — that only the authorized person knows the password — no longer holds against automated, AI-scaled phishing, irrespective of password strength: the attacker obtains the correct password directly from the user rather than by guessing. Unlike the row below (privileged accounts, where PAM at least offers session monitoring as partial protection), there is here no second factor and no compensating mechanism in the wording — the requirement loses its core effect entirely, not merely in part.	IAM team
should	"Superior procedures are used for the authentication of privileged user accounts (e.g., privileged access management, two-factor authentication)."	Partially degraded	The “e.g.” wording allows simple two-factor authentication without phishing resistance to suffice — the same adversary-in-the-middle relay gap as above, here applied to privileged accounts. It	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
			remains “only” degraded rather than friction-only, because PAM as the alternative named can offer additional session monitoring and time limitation, which partly compensates for a compromised factor.	
high	"Depending on the risk assessment, authentication procedure and access control have been enhanced by supplementary measures (e.g., continuous access monitoring with respect to irregularities or use of strong authentication, automatic logout, disabling in case of inactivity, or brute force prevention). (C, I, A)"	Partially degraded	The list of options permits fulfilment through weaker measures (automatic logout, brute-force protection) that do not address the adversary-in-the-middle and phishing relay gap; only the option “continuous access monitoring” would actually help, but it is not mandatory.	IAM team
very high	"Before gaining access to data of very high protection needs, users are authenticated by means of strong authentication (e.g., two-factor authentication) according to the state of the art. (C, I)"	Partially degraded	Same gap as above: the example names classic two-factor authentication rather than expressly phishing-resistant methods.	IAM team

Overall classification: Friction-only (1 of 7 requirements loses its core effect entirely; 5 of 7 affected in total, of which 4 remain effective but weakened), irrespective of protection requirement. Under the precedence convention (friction-only > partially degraded > robust > not affected), the pure password factor determines the overall classification, even though the other four findings in this control are only weakened, not without effect. The first friction-only finding in the entire catalogue — the recurring pattern in the control is the missing obligation of phishing-resistant (FIDO2/WebAuthn-type) authentication. Target state: with a mandatory

phishing-resistant method (FIDO2/WebAuthn/passkey) for all seven requirements, the control would be robust throughout, not merely the pure password factor resolved — this target state is the subject of MHC-03 (see Part II).
Hardening reference: MHC-03 (see Part II)

4.1.3 — To what extent are user accounts and login information securely managed and applied?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The creating, changing, and deleting of user accounts is conducted."	Not affected	Basic process.	IAM team
must	"Unique and personalized user accounts are used."	Not affected	Traceability principle, structural.	IAM team
must	"The use of \"shared accounts\" is regulated (e.g., restricted to cases where traceability of actions is dispensable)."	Not affected	Governance.	IAM team
must	"User accounts are disabled immediately after the user has resigned from or left the organization (e.g., upon termination of the employment contract)."	Not affected	“Immediately” is a clear immediate trigger with no time compression gap in the wording.	IAM team
must	"User accounts are regularly reviewed."	Partially degraded	Time compression: periodic review without an event trigger leaves orphaned or superfluous accounts undetected between cycles.	IAM team
must	"The login information is provided to the user in a secure manner."	Not affected	Transmission process.	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A policy for the handling of login information is defined and implemented. [No disclosure, No writing down/unencrypted storage, Immediate change on suspected compromise, No reuse, Change of temporary credentials after first login, Quality requirements]"	Not affected	Hygiene and conduct policy; the strength of the authentication mechanism itself is assessed separately in 4.1.2 (no double counting).	IAM team
should	"A basic user account with minimum access rights and functionalities is existent and used."	Not affected	Least-privilege baseline, structural.	IAM team
should	"Default accounts and passwords pre-configured by manufacturers are disabled (e.g., by blocking or changing of password)."	Not affected	A one-off hardening step; once implemented, closed irrespective of attacker capability.	IT operations
should	"User accounts are created or authorized by the responsible body."	Not affected	Governance.	IAM team
should	"Creating user accounts is subject to an approval process (four-eyes principle)."	Not affected	Four-eyes principle, organizationally robust.	IAM team
should	"User accounts of service providers are disabled upon completion of their task."	Not affected	Clearly event-triggered.	IAM team
should	"Deadlines for disabling and deleting user accounts are defined."	Not affected	Governance and deadline setting.	IAM team
should	"The use of default	Not affected	Technical	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	passwords is technically prevented."		enforcement, binary, effective once applied.	
should	"Where strong authentication is applied, the use of the medium (e.g., ownership factor) is secure."	Partially degraded	Capability decoupling: the possession factor (for example a telephone number for SMS OTP) is vulnerable to SIM swapping; convincingly deceiving telecommunications support staff has become easier with AI assistance.	IAM team
should	"User accounts are reviewed at regular intervals. This also includes user accounts in customers' IT systems."	Partially degraded	Same time compression rationale as for the must element above, here additionally applied to customer systems.	IAM team
should	"Interactive login for service accounts (technical accounts) is technically prevented."	Not affected	Technical, structural enforcement.	IT operations

Overall classification: Partially degraded (3 of 17 requirements affected), irrespective of protection requirement (the control has no high or very high requirements).
Hardening reference: MHC-10, MHC-03 (see Part II)

4.2.1 — To what extent are access rights assigned and managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The requirements for the management of access rights (authorization) are determined and fulfilled. [Procedure for application/verification/approval, Need-to-know/least-privilege, Access rights revoked when no longer needed]"	Not affected	Event-triggered revocation logic (“when no longer needed”) instead of vague periodicity — well worded.	IAM team
must	"The access rights granted for normal and privileged user accounts and technical accounts are reviewed at regular intervals also within IT systems of customers."	Partially degraded	Time compression: the classic periodic review pattern without an event trigger.	IAM team
should	"Strategies for authorizing access to information are prepared."	Not affected	Governance.	IAM team
should	"Authorization roles are used."	Not affected	RBAC design, structural.	IAM team
should	"Rights are allocated on a need-to-use basis and according to the role and/or area of responsibility."	Not affected	Design principle.	IAM team
should	"Normal user accounts are not granted privileged access rights."	Not affected	Structural separation.	IAM team
should	"The user's access rights are updated after the user's responsibilities have changed (e.g., to another field of responsibility)."	Not affected	Event-triggered, well worded.	IAM team
high	"The access rights are approved by the responsible internal Information Officer. (C, I, A)"	Not affected	Additional approval governance.	IAM team
very	"Information is stored in	Robust	File-level content	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
high	encrypted form at content level (e.g., file level) to prevent unauthorized persons from gaining access and information (privileged users). Where encryption is not feasible, information shall be protected by similarly effective measures. (C)"		encryption is a hard cryptographic barrier that works even against privileged insiders or circumvented access controls.	
very high	"Existing access rights are regularly reviewed at shorter intervals (e.g., quarterly). (C)"	Partially degraded	Shortens the review interval but remains periodic without an event trigger — the structural time compression gap persists, only the window becomes smaller.	IAM team

Overall classification: Partially degraded (1 of 7 requirements affected at normal protection requirement; 2 of 10 including high and very high), irrespective of protection requirement. Notably, at very high protection requirement, file-level encryption adds a genuine robust element, which nevertheless does not tip the overall classification, because periodic access review remains a weakness in its own right.
Hardening reference: MHC-10 (see Part II)

Chapter 4 summary

Control	Overall classification	Share affected
4.1.1	Not affected → from high onward: Partially degraded	0/2 → 1/3
4.1.2	Friction-only	1/7 (friction) + 4/7 (degraded)
4.1.3	Partially degraded	3/17
4.2.1	Partially degraded	1/7 → 2/10

Identity and Access Management delivers the first and, to date, only friction-only finding in the entire catalogue: the pure password factor in 4.1.2 loses its core effect against automated phishing entirely, not merely in part. The control also remains the most heavily affected overall (5 of 7 requirements) — the recurring pattern is the absence of an explicit obligation of phishing-resistant authentication. Alongside this, the familiar periodic review pattern recurs (4.1.1, 4.1.3,

4.2.1). A robust bright spot at very high protection requirement: content-based encryption (4.2.1).

Chapter 5 — IT Security/Cyber Security (15 Controls, 105 requirements, largest chapter)

5.1.1 — To what extent is the use of cryptographic procedures managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"All cryptographic procedures used (e.g., encryption, signature, and hash algorithms, protocols) provide the security required by the respective application field according to the recognized industry standard to the extent legally feasible."	Robust	Requires industry-standard cryptography. Correctly implemented modern cryptography (AES-256, current TLS and so on) is mathematically hard against a Gen-AI-supported attacker — AI shortens neither factorization nor brute force against modern key lengths. (Post-quantum risk is a separate field of threat, deliberately not addressed here — no Mythos or Gen-AI relevance.)	CISO function
should	"A concept for the application of cryptography is defined and implemented. [Methods, Key strengths, Full lifecycle: generation/storage/archiving/retrieval/distribution/deactivation/renewal/deletion]"	Not affected	Key management process, governance; the cryptographic strength itself is already assessed above.	IT operations
high	"Key sovereignty requirements (particularly in case of external processing) are determined and fulfilled. (C, I)"	Not affected	A question of legal and contractual sovereignty (jurisdiction for cloud processing), no Mythos relevance.	CISO function

Overall classification: Robust (no degraded or friction element present), irrespective of protection requirement. Properly implemented cryptography is one of the clearest examples of a genuine hard barrier in the entire catalogue — directly citable as SoA evidence.

5.1.2 — To what extent is information protected during transfer?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The network services used to transfer information are identified and documented."	Not affected	Inventory and governance.	IT operations
must	"Policies and procedures in accordance with the classification requirements for the use of network services are defined and implemented."	Not affected	Governance.	IT operations
must	"Measures for the protection of transferred contents against unauthorized access are implemented."	Partially degraded	Unspecific “measures” without a binding minimum standard (for example encryption) leave room for weaker protection that an AI-supported attacker positioned within the network can defeat. Only high and very high close the gap expressly (see below).	IT operations
should	"Measures for ensuring correct addressing and correct transfer of information are implemented."	Not affected	Correctness of transmission, not an attacker barrier in the narrower sense.	IT operations
should	"Electronic data exchange is conducted using content or transport encryption according to the respective classification."	Robust	Expressly requires content or transport encryption — closes the must-level gap for organizations that implement this should	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
			element.	
should	"Remote access connections to the organization's network possess adequate security features. [Encryption, Authentication strength, Granting/termination of access, Secured connection e.g. VPN]"	Partially degraded	Remote access is the preferred target of real-time phishing relay attacks (see 4.1.2) — VPN encryption protects the transport path but does not resolve the authentication weakness if that remains conventional (password or OTP).	IT operations
high	"Information is transported or transferred in encrypted form (at least transport-encryption) or protected by similarly effective measures. (C)"	Robust	Explicit encryption obligation for a high protection requirement.	IT operations
very high	"Information is transported or transferred in content-encrypted form. (C)"	Robust	Content encryption, the strongest form — end-to-end-style protection.	IT operations

Overall classification: Partially degraded (2 of 6 requirements affected), irrespective of protection requirement. Two separate weaknesses: the vague must wording is largely closed by the should element on encryption; the remote access authentication gap (identical to the finding in 4.1.2), however, persists across all levels independently of that.
Hardening reference: MHC-03, standalone measure (see Part II)

5.2.1 — To what extent are changes managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Information security requirements for changes to the organization, business processes, IT systems are determined and fulfilled."	Not affected	Governance.	IT operations
should	"A formal approval	Not affected	Governance.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	procedure is established."			
should	"The potential impact on the information security of changes is assessed."	Not affected	Process assessment.	IT operations
should	"Changes affecting the information security are subjected to planning and testing."	Not affected	Process.	IT operations
should	"Procedures for fallback in fault cases are considered."	Not affected	Resilience governance.	IT operations
high	"Compliance with the information security requirements is verified during and after the changes are applied. (C, I, A)"	Not affected	A double, clearly defined review point (“during and afterwards”) — well worded, no time compression relevance.	IT operations

Overall classification: Not affected (0 of 6 requirements affected), irrespective of protection requirement. A cleanly specified process control with no Mythos-relevant gap — a good example of a consistently unremarkable governance control.

5.2.2 — To what extent are development and testing environments separated from operational environments?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The IT systems have been subjected to risk assessment in order to determine the necessity of their separation into development, testing and operational systems."	Not affected	Assessment process.	Development/engineering
must	"A segmentation is implemented based on the results of risk analysis."	Robust	Architectural separation is a structural barrier: once correctly implemented, it works irrespective of how an attacker automates the initial access attempt.	IT operations
should	"The requirements for development and testing environments are determined and fulfilled. [Separation of dev/test/prod, No dev tools on operational systems, Different user profiles]"	Robust	Deepens the same structural separation (separate profiles, no development tools in production) — architecturally excludes lateral movement and tool abuse in the production environment.	IT operations

Overall classification: Robust (no degraded or friction element), irrespective of protection requirement. Note: the effectiveness applies to correctly implemented segmentation; an incomplete or faulty implementation is a matter of implementation, not a weakness inherent in the wording.

5.2.3 — To what extent are IT systems protected against malware?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Requirements for protection against malware are determined."	Not affected	Governance.	IT operations
must	"Technical and organizational measures for protection against malware are defined and implemented."	Partially degraded	Capability decoupling: classic signature-based malware protection is structurally weaker against AI-generated and polymorphic malware, because signatures take effect only once a variant has become known.	IT operations
should	"Unnecessary network services are disabled."	Not affected	Attack surface reduction, structural and binary.	IT operations
should	"Access to network services is restricted to necessary access through suitable protective measures (see examples)."	Not affected	Access control, structural.	IT operations
should	"Malware protection software is installed and updated automatically at regular intervals (e.g., virus scanner)."	Partially degraded	The example given, “virus scanner”, signals signature-based detection; automatic but periodic updates structurally lag behind AI-generated variants.	IT operations
should	"Received files and software are automatically inspected for malware prior to their execution (on-access scan)."	Partially degraded	Same signature detection limit.	IT operations
should	"The entire data contents of all systems is regularly inspected for malware."	Partially degraded	A double weakness: periodic review (time compression) and signature-based	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
			detection (capability decoupling).	
should	"Data transferred by central gateways (e.g., e-mail, internet, third-party networks) is automatically inspected by means of protection software. [Encrypted connections, Encrypted content]"	Partially degraded	Encrypted content structurally bypasses gateway inspection — a known circumvention path deliberately used by AI-supported attackers.	IT operations
should	"Measures to prevent protection software from being deactivated or altered by users are defined and implemented."	Not affected	Tamper protection for the protection agent itself, structural and binary.	IT operations
should	"For IT systems operated without the use of malware protection software, alternative measures... are implemented."	Not affected	Compensating control, generic.	IT operations

Overall classification: Partially degraded (5 of 10 requirements affected), irrespective of protection requirement. One of the most heavily degraded controls in the catalogue — signature-based malware detection is a classic example of capability decoupling: AI-generated, polymorphic malware systematically evades signature matching.
Hardening reference: MHC-05 (see Part II)

5.2.4 — To what extent are event logs recorded and analysed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Information security requirements regarding the handling of event logs are determined and fulfilled."	Not affected	Governance.	IT operations
must	"Security-relevant requirements regarding the logging of activities of system administrators and	Not affected	Scoping.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	users are determined and fulfilled."			
must	"The IT systems used are assessed regarding the necessity of logging."	Not affected	Assessment process.	IT operations
must	"When using external IT services, information on the monitoring options is obtained and considered in the assessment."	Not affected	Supplier and governance aspect.	IT operations
must	"Event logs are checked regularly for policy violations and noticeable problems in compliance with the permissible legal and organizational provisions."	Partially degraded	Aggregation resistance and time compression: a periodic manual review geared to “conspicuous” events structurally fails to recognize an attack deliberately fragmented into many unremarkable individual steps.	IT operations
should	"A procedure for the escalation of relevant events to the responsible body... is defined and established."	Not affected	Downstream escalation process.	IT operations
should	"Event logs (contents and meta data) are protected against alteration. (e.g., by a dedicated environment)."	Robust	Tamper protection for log data (for example WORM storage, a separate log system) is a structural barrier, independent of attacker capability.	IT operations
should	"Adequate monitoring and recording of any actions on the network that are relevant to information security are established."	Not affected	Existence of the record; the quality of evaluation is already assessed above.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
high	"Information security requirements relevant to the security during the handling of event logs, e.g., contractual requirements, are determined and implemented. (C, I, A)"	Not affected	Contractual and governance-related.	IT operations
high	"Events related to establishing and terminating remote access sessions (e.g., for remote maintenance). (C, I, A)"	Not affected	Scoping (what is logged), no statement on evaluation.	IT operations
very high	"Logging of any access to data of very high protection needs as far as technically feasible and as permissible according to legal and organizational provisions. (C, I)"	Not affected	Scoping.	IT operations

Overall classification: Partially degraded (1 of 8 requirements affected), irrespective of protection requirement. The weakness lies specifically in the evaluation (“checked regularly”), not in the logging itself — log integrity is in fact well secured through a robust element of its own. Hardening reference: MHC-05 (see Part II)

5.2.5 — To what extent are vulnerabilities identified and addressed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Information on technical vulnerabilities for the IT systems in use is gathered (e.g., information from the manufacturer, system audits, CVS database) and evaluated (e.g., Common Vulnerability Scoring System CVSS)."	Not affected	Gathering of information; the time-critical assessment follows in the next row.	IT operations
must	"Potentially affected IT	Partially	Collapse of the patch	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	systems and software are identified and the risk caused by the vulnerability is assessed."	degraded	gap window (see Methodology, the Cogent Security evaluation 2026): classic risk assessments usually calculate in weeks up to a working exploit. Demonstrably this window shortens through AI-supported exploit development to in some cases less than a day — an assessment methodology that does not take this into account systematically underestimates the urgency.	
must	"Risks arising from vulnerabilities are treated."	Not affected	Treatment process, downstream.	IT operations
should	"An adequate patch management is defined and implemented (e.g., patch testing and installation)."	Partially degraded	Same patch gap rationale: “adequate” without a binding speed SLA guarantees no pace that keeps up with AI-accelerated exploit availability.	IT operations
should	"Risk minimizing measures are implemented, as necessary."	Not affected	Generic compensating measure.	IT operations
should	"Successful installation of patches is verified in an appropriate manner."	Not affected	Quality assurance, downstream.	IT operations

Overall classification: Partially degraded (3 of 6 requirements affected), irrespective of protection requirement. A central finding of the entire catalogue: this control directly reflects the evidenced shift “collapse of the patch gap window” from the Methodology — a priority for the later hardening control synthesis.

Hardening reference: MHC-09 (see Part II)

5.2.6 — To what extent are IT systems and services technically checked (system and service audit)?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Requirements for auditing IT systems or services are determined."	Not affected	Governance.	IT operations
must	"The scope of the system audit is specified in a timely manner."	Not affected	Planning.	IT operations
must	"System or service audits are coordinated with the operator and users of the IT systems or services."	Not affected	Coordination.	IT operations
must	"The results of system or service audits are stored in a traceable manner and reported to the relevant management."	Not affected	Documentation and reporting.	IT operations
must	"Measures are derived from the results and are implemented in an appropriate timeframe."	Partially degraded	Time compression: “appropriate timeframe” is indefinite and carries no binding upper limit relative to an accelerated threat pace.	IT operations
should	"System and service audits are planned taking into account any security risks they might cause (e.g., disturbances)."	Not affected	Planning certainty.	IT operations
should	"Regular system or service audits are performed. [Qualified personnel, Suitable tools e.g. vulnerability scanners, From internet and internal network]"	Partially degraded	Periodic rather than continuous review leaves the same patch gap open between audit cycles as in 5.2.5.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
should	"Within a reasonable period following completion of the audit, a report is prepared."	Not affected	Documentation pace, downstream.	IT operations
high	"For critical IT systems or services, additional system or service audit requirements have been identified and are fulfilled (e.g., service specific tests and tools and/or human penetration tests, risk-based time intervals). (A)"	Not affected	“Risk-based time intervals” is an adaptive, well-worded requirement without rigid periodicity.	IT operations
very high	"IT systems and services are regularly scanned for vulnerabilities. Suitable protective measures must be implemented for systems and services that may not be scanned. (C, I, A)"	Partially degraded	Even at the highest protection requirement level, scanning is only “regular” rather than continuous — the opportunity for a binding near-real-time check at very high protection requirements is left unused.	IT operations

Overall classification: Partially degraded (2 of 8 requirements affected at normal protection requirement; 3 of 10 including very high), irrespective of protection requirement. Recurring theme: periodic rather than continuous review, consistent with the patch gap finding in 5.2.5. Hardening reference: MHC-09 (see Part II)

5.2.7 — To what extent is the network of the organization managed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Requirements for the management and control of networks are determined and fulfilled."	Not affected	Governance.	IT operations
must	"Requirements regarding network segmentation are	Robust	Segmentation obligation — consistent with the	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	determined and fulfilled."		architectural robust assessment in 5.2.2.	
should	"Procedures for the management and control of networks are defined."	Not affected	Governance.	IT operations
should	"For a risk-based network segmentation. [Limitations for connecting IT systems, Security technologies, Performance/trust/availability/security/safety, Limitation of impact if compromised, Detection of potential attacks and lateral movement, Separation of networks with different purpose, Internet-exposure risk, Technology-specific separation for external IT services, Separation from customer networks, Detection/prevention of data loss]"	Partially degraded	The detection sub-aspects (lateral movement, data outflow) are the weaker component: generic “detection” without any requirement of behaviour-based or anomaly-based methods is geared more to obvious than to fragmented, slowly unfolding attack patterns.	IT operations
high	"Extended requirements for the management and control of networks are determined and implemented. [Authentication of IT systems on the network, Restricted access to management interfaces, Specific risks e.g. wireless and remote access] (C, I, A)"	Partially degraded	The reference to remote access risks imports the same authentication gap as in 4.1.2 and 5.1.2; the remaining sub-aspects (device authentication, protection of the management interface) are structurally sound in themselves.	IT operations

Overall classification: Partially degraded (1 of 4 requirements affected at normal protection requirement; 2 of 5 including high), irrespective of protection requirement. The segmentation obligation itself is robust; the weakness lies in the detection and remote access dimension.
Hardening reference: MHC-03, MHC-05 (see Part II)

5.2.8 — To what extent is continuity planning for IT services in place?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Critical IT services are identified, and business impact is considered."	Not affected	Business impact analysis, governance.	IT operations
must	"Requirements and responsibilities for continuity and recovery of those IT services are known to relevant stakeholders and fulfilled."	Not affected	Governance.	IT operations
should	"Critical IT systems are identified. [Classification for appropriate protection need, Implementation of adequate security measures]"	Not affected	Governance and classification.	IT operations
should	"Continuity planning exists and is regularly reviewed and updated. [Alternative communication, Alternative storage, Alternative power/network]"	Partially degraded	Time compression: periodic updating without an event trigger for scenarios that escalate faster today (for example the speed of ransomware propagation).	IT operations
should	"Continuity planning includes at least (Distributed) Denial of Service attacks, successful ransomware attacks and other sabotage activities, system failure scenarios, and natural disaster scenarios affecting critical IT systems."	Not affected	Modern attack scenarios (DDoS, ransomware, sabotage) are expressly named — a well-secured wording, irrespective of how an attack is orchestrated in detail.	IT operations
high	"Continuity planning includes predefined time frames (Recovery Time Objective) for resumption of	Not affected	Defined RTO parameter.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	operation in line with requirements. (A)"			
high	"Appropriate SLAs with external service providers according to continuity planning are in place. (A)"	Not affected	Contractual.	Procurement/supplier management
high	"Continuity plans include coordination of contractually agreed communication with business partners. (A)"	Not affected	Contractual and process.	CISO function
high	"Continuity planning is regularly tested including a full recovery and reconstitution of the system to a known state and compliance with defined target times. (A)"	Not affected	Testing activity.	IT operations
high	"A backup and recovery strategy for critical IT services and information is defined and implemented. (C, I, A)"	Not affected	Existence of the strategy; the concrete backup mechanics are assessed separately in 5.2.9 (no double counting).	IT operations
high	"Backups of critical IT services and information are sufficiently protected against unauthorized modification or deletion by malicious software, (I, A)"	Partially degraded	“Sufficiently protected” without an immutability requirement allows ransomware with compromised backup administrator rights to delete or encrypt backups nonetheless. Only 5.2.9 (very high) closes this gap expressly through immutable or offline backups.	IT operations
high	"Backups of critical IT	Partially	Same rationale, here	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	services and information are sufficiently protected against unauthorized access by malicious software or operators. (C, I)"	degraded	the confidentiality side.	
very high	"Continuity planning is coordinated with the continuity plans of relevant external service providers. (A)"	Not affected	Coordination and contractual.	Procurement/supplier management
very high	"Continuance of essential mission and business functions with minimal or no loss of operational continuity is possible. [Alternate operation strategies/separated standby systems, Alternate/backup storage sites with equivalent controls]"	Robust	Spatially separated standby systems or sites with equivalent controls are an architectural redundancy — effective irrespective of attacker capability.	IT operations
very high	"Continuity planning is tested regularly. Test scenarios, results, and lessons learned are recorded. (I, A)"	Not affected	Testing activity.	IT operations

Overall classification: Partially degraded (1 of 5 requirements affected at normal protection requirement; 3 of 15 including high and very high), irrespective of protection requirement. Cross-reference: the backup immutability gap at high is expressly closed by control 5.2.9 (very high) — see there for the full assessment.
Hardening reference: MHC-08 (see Part II)

5.2.9 — To what extent is the backup and recovery of data and IT services ensured?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Backup concepts exist for	Not	Existence of a concept;	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	relevant IT systems. Appropriate protective measures to ensure confidentiality, integrity, and availability for data backups are considered."	affected	the concrete mechanics are assessed specifically at high and very high.	
must	"Recovery concepts exist for relevant IT services."	Not affected	Existence of a concept.	IT operations
should	"A backup and recovery concept exists for each relevant IT service. Dependencies between IT services and the sequence for recovery are considered."	Not affected	Planning and sequencing.	IT operations
high	"Backup and recovery concepts are methodically reviewed at regular intervals. (A)"	Partially degraded	Time compression: periodic review of the concept without an event trigger — a concept can become out of date even where the underlying storage medium (see very high) is immutable.	IT operations
high	"General restore capability is considered and tested (e.g., sample testing, test systems). (I, A)"	Not affected	Sample-based testing activity.	IT operations
high	"Backup and recovery concepts exist. [RPO, RTO, Required resources for recovery, Avoidance of overload scenarios, Appropriate spatial redundancy e.g. separate room/fire section/data center/site] (A)"	Robust	Spatial redundancy (a separate server room, fire compartment or site) is an architectural barrier that works irrespective of attacker capability.	IT operations
very	"(Additional) Backups are	Robust	Offline or immutable	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
high	performed via offline procedures, immutable backups or by using an isolated IAM solution. (I, A)"		backups, or an isolated IAM solution, are a hard architectural barrier against ransomware and credential-based attacks — one of the strongest robust examples in the entire catalogue.	
very high	"Restore procedures are technically tested in a methodical way at regular intervals. (I, A)"	Not affected	Testing activity.	IT operations
very high	"Geographical redundancy is considered in data backup and recovery concepts. (A)"	Robust	Geographic redundancy, structural.	IT operations

Overall classification: Not affected at normal protection requirement (0 of 3 requirements affected) — only the existence of a concept is required. From high protection requirement onward, periodic review of the concept without an event trigger is added — the overall classification becomes Partially degraded (1 of 6), while at the same time spatial redundancy also appears as a robust element. At very high protection requirement, the robust side strengthens considerably (immutable and offline backups, geographic redundancy) — the overall classification nevertheless remains Partially degraded (1 of 9), because the periodic-review gap from the high level persists unchanged. In practice: even an organization with exemplary immutable backups is classified here as “partially degraded”, because the two requirements concern different aspects — the backup medium (immutable, strong) versus the backup concept review cycle (periodic, weak). Cross-reference: supplements the backup immutability gap from control 5.2.8 (high).
Hardening reference: MHC-10 (see Part II)

Summary (Chapter 5, part 1: 5.1.1–5.2.9)

Control	Overall classification	Share affected
5.1.1	Robust	0/3
5.1.2	Partially degraded	2/6
5.2.1	Not affected	0/6
5.2.2	Robust	0/3

Control	Overall classification	Share affected
5.2.3	Partially degraded	5/10
5.2.4	Partially degraded	1/8
5.2.5	Partially degraded	3/6
5.2.6	Partially degraded	2/8 → 3/10
5.2.7	Partially degraded	1/4 → 2/5
5.2.8	Partially degraded	1/5 → 3/15
5.2.9	Not affected → from high onward: Partially degraded	0/3 → 1/6 → 1/9

5.3.1 — To what extent is information security considered in new or further developed IT service?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The information security requirements associated with the design and development of IT service are determined and considered."	Not affected	Governance.	Development/engineering
must	"The information security requirements associated with the acquisition or extension of IT service and IT components are determined and considered."	Not affected	Governance.	Procurement/supplier management
must	"Information security requirements associated with changes to developed IT services are considered."	Not affected	Governance, thematically related to 5.2.1.	Development/engineering
must	"System approval tests are carried out under consideration of the information security requirements."	Not affected	Approval gate, existence level.	Development/engineering
should	"Requirement specifications are prepared. [IS requirements, Vendor recommendations/best practices, Best practices/security guidelines, Fail safe design]"	Not affected	Documentation and design input.	Development/engineering
should	"Requirement specifications are reviewed against the information security requirements."	Not affected	Review governance.	Development/engineering
should	"The IT service is reviewed for compliance with specifications prior to	Not affected	Approval gate.	Development/engineering

Level	Requirement (original wording)	Category	Rationale	Responsible
	productive use."			
should	"The use of production data for testing purposes is avoided as far as possible (if applicable, anonymization or pseudonymization). [Lifecycle requirements for test data, Case-related generation specs]"	Partially degraded	Capability decoupling: re-identifying pseudonymized or anonymized data by correlation with additional information has until now required specialist statistical knowledge. AI-supported correlation methods make this more broadly accessible, while the requirement prescribes no re-identification-resistant methodology.	Development/engineering
should	"Test-systems are provided with protective measures comparable to those in the operational environment where productive data are used for testing purposes."	Not affected	Compensating control, structurally comparable.	Development/engineering
very high	"The security of purpose built software or significantly customized software is tested (e.g., penetration testing) during commissioning, in case of significant changes, or at regular intervals. (C, I, A)"	Not affected	A triple trigger (commissioning, material changes, regular interval) — a well-secured wording with no isolated time compression gap.	Development/engineering

Overall classification: Partially degraded (1 of 10 requirements affected), irrespective of protection requirement.
Hardening reference: standalone measure (see Part II)

5.3.2 — To what extent are requirements for network services defined?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Requirements regarding the information security of network services are determined and fulfilled."	Not affected	Governance.	IT operations
should	"A procedure for securing and using network services is defined and implemented."	Not affected	Governance and process.	IT operations
should	"The requirements are agreed in the form of SLAs."	Not affected	Contractual.	Procurement/supplier management
should	"Adequate redundancy solutions are implemented."	Robust	Redundancy is a structural availability measure, consistent with other redundancy assessments in the catalogue.	IT operations
high	"Procedures for monitoring the quality of network traffic (e.g., traffic flow analyses, availability measurements) are defined and carried out. (A)"	Not affected	Quality and availability monitoring, no security detection relevance (assessed separately from that in 5.2.7).	IT operations

Overall classification: Robust (no degraded or friction element), irrespective of protection requirement.

5.3.3 — To what extent is the return and secure removal of information assets from external IT services regulated?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"A procedure for the return and secure removal of information assets from each external IT service is defined and implemented."	Not affected	Governance and contract process.	Procurement/supplier management
should	"A description of the termination process is	Not	Contractual.	Procurement/supplier

Level	Requirement (original wording)	Category	Rationale	Responsible
	given, adapted to any changes, and contractually regulated."	affected		management

Overall classification: Not affected (0 of 2 requirements affected), irrespective of protection requirement.

5.3.4 — To what extent is information protected in shared external IT services?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Effective segregation (e.g., segregation of customers) prevents access to own information by unauthorized users of other organizations."	Robust	Tenant separation, correctly implemented (cryptographic isolation, strict access controls), is a structural barrier — consistent with the segmentation assessment in 5.2.2 and 5.2.7.	IT operations
should	"The provider's segregation concept is documented and adapted to any changes. [Separation of data/functions/software/OS /storage/network, Risk assessment for shared environment]"	Not affected	Documentation of the concept; the strength of the mechanism is already assessed above.	Procurement/supplier management

Overall classification: Robust (no degraded or friction element), irrespective of protection requirement.

Chapter 5 summary (overall, 15 controls)

Control	Overall classification	Share affected
5.1.1	Robust	0/3
5.1.2	Partially degraded	2/6
5.2.1	Not affected	0/6
5.2.2	Robust	0/3
5.2.3	Partially degraded	5/10
5.2.4	Partially degraded	1/8
5.2.5	Partially degraded	3/6
5.2.6	Partially degraded	2/8 → 3/10
5.2.7	Partially degraded	1/4 → 2/5
5.2.8	Partially degraded	1/5 → 3/15
5.2.9	Not affected → from high onward: Partially degraded	0/3 → 1/6 → 1/9
5.3.1	Partially degraded	1/10
5.3.2	Robust	0/5
5.3.3	Not affected	0/2
5.3.4	Robust	0/2

Chapter 5 is the technically densest chapter and shows the greatest range: 4 consistently robust (cryptography, dev/test/prod separation, network redundancy, tenant separation — all architectural or cryptographic hard barriers), 9 partially degraded, 2 not affected, 1 dependent on protection requirement (5.2.9, with the fully reasoned exception that the very strong robust supplement at very high protection requirement does not tip the degraded overall classification). Two market-relevant core Mythos findings in this chapter: the collapse of the patch gap window (5.2.5, 5.2.6) and the consistently missing obligation of phishing-resistant authentication (5.1.2, 5.2.7, congruent with 4.1.2).

Chapter 6 — Supplier Relationships (3 Controls, 31 requirements)

6.1.1 — To what extent is information security ensured among contractors and cooperation partners?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Contractors and partners are subjected to a security risk assessment."	Not affected	Governance and process.	Procurement/supplier management
must	"An appropriate level of information security is ensured by contractual agreements with contractors and partners."	Not affected	Contractual.	Procurement/supplier management
must	"Where applicable, contractual agreements with clients and customers are passed on to contractors and partners."	Not affected	Contractual, pass-through.	Procurement/supplier management
should	"Contractors and partners are contractually obligated to pass on any requirements regarding an appropriate level of information security to their subcontractors."	Not affected	Contractual cascading.	Procurement/supplier management
should	"Service reports and documents by contractors and partners are reviewed."	Partially degraded	Capability decoupling: the review rests on documents submitted by the supplier. AI tools make convincingly forged or embellished reports and evidence easier to produce, while no independent verification is required.	Procurement/supplier management
high	"Evidence is provided that the information security level of the supplier is	Partially degraded	The list of examples allows mere self-declaration or	Procurement/supplier management

Level	Requirement (original wording)	Category	Rationale	Responsible
	adequate for the protection needs of the information (e.g., checked questionnaire/self-assessment, attestation, certificate, supplier audit). (C, I, A)"		attestation to suffice — the same document forgery gap as above, here applied to formal evidence. Only very high closes this through the obligation of a third-party audit (see below).	
high	"The supplier's level of compliance with the required evidence is documented, regularly and in case of changes... reviewed and monitored. (C, I, A)"	Partially degraded	The review frequency is well worded with a double trigger; the underlying problem of evidence authenticity nevertheless remains unaffected.	Procurement/supplier management
high	"The supplier's compliance with contractual agreements is checked, documented, regularly and in case of changes... reviewed and monitored. (C, I, A)"	Partially degraded	Same rationale.	Procurement/supplier management
very high	"The adequate level of information security should be demonstrated by a third party audit (an adequate TISAX® label or equivalent) or an adequate supplier audit covering agreed and applicable customer requirements. If an audit was not conducted, a risk-based decision must be made by the organization's management to continue doing business with the supplier. A record of this decision exists. (C, I, A)"	Robust	Requires an independent third-party audit (a TISAX® label or similar) instead of self-declaration — an on-site audit with interviews and sample checks is far harder to falsify at scale than a submitted document. Closes the gap identified at high in a targeted way.	Procurement/supplier management

Level	Requirement (original wording)	Category	Rationale	Responsible
very high	"Contractual obligations with customers regarding transparency of supply chain risks are fulfilled. (C, I, A)"	Not affected	Contractual.	Procurement/supplier management

Overall classification: Partially degraded (1 of 5 requirements affected at normal protection requirement; 4 of 8 from high onward; 4 of 10 including very high), irrespective of protection requirement. At very high protection requirement, the self-declaration weakness is expressly and effectively closed by the third-party audit obligation (a TISAX® label) — the overall classification nevertheless remains “partially degraded”, because the two monitoring requirements at high must, independently of the type of evidence, only be checked “regularly”. In practice: actively require a TISAX® label or an equivalent third-party audit from critical suppliers rather than relying on self-declarations — this is the single most effective lever in this control.
Hardening reference: MHC-14 (see Part II)

6.1.2 — To what extent is non-disclosure regarding the exchange of information contractually agreed?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The non-disclosure requirements are determined and fulfilled."	Not affected	Legal.	Legal
must	"Requirements and procedures for applying non-disclosure agreements are known to all persons passing on information in need of protection."	Not affected	Communication and awareness.	Legal
must	"Valid non-disclosure agreements are concluded prior to forwarding sensitive information."	Not affected	Legal gate.	Legal
must	"The requirements and procedures for the use of non-disclosure agreements and the handling of	Partially degraded	Time compression: expressly covers the operational procedures for handling	Legal

Level	Requirement (original wording)	Category	Rationale	Responsible
	information requiring protection are reviewed at regular intervals."		information requiring protection (not merely the contractual text) — an out-of-date handling instruction without an event trigger is a technical and operational risk, not a purely legal one.	
should	"Non-disclosure agreement templates are available and checked for legal applicability."	Not affected	Legal.	Legal
should	"Non-disclosure agreements include the persons/organizations involved, the type of information covered by the agreement, the subject of the agreement, the validity period of the agreement and the responsibilities of the obliged party."	Not affected	Legal minimum content.	Legal
should	"Non-disclosure agreements include provisions for the handling of sensitive information beyond the contractual relationship."	Not affected	Legal.	Legal
should	"Options of demonstrating compliance with specifications (e.g., review by an independent third-party or audit rights) are defined."	Not affected	Contractual mechanism.	Legal
should	"A process for monitoring the validity period of temporary non-disclosure	Not affected	“In due time” refers to a known expiry date — a calendar-driven rather	Legal

Level	Requirement (original wording)	Category	Rationale	Responsible
	agreements and initiating their extension in due time is defined and implemented."		than a vague periodic trigger.	

Overall classification: Partially degraded (1 of 9 requirements affected), irrespective of protection requirement. The only finding in an otherwise purely legally framed control — it concerns specifically the operational handling instruction, not the contractual text itself.
Hardening reference: MHC-10 (see Part II)

6.1.3 — To what extent are the responsibilities between external IT service providers and the own organization defined?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"The IT services concerned are identified."	Not affected	Scoping.	Procurement/supplier management
must	"The security requirements relevant to the IT service are determined."	Not affected	Governance.	Procurement/supplier management
must	"The organization responsible for implementing the requirement is defined and aware of its responsibility."	Not affected	Governance and accountability.	Procurement/supplier management
must	"Mechanisms for shared responsibilities are specified and implemented."	Not affected	Governance.	Procurement/supplier management
must	"The responsible organization fulfils its respective responsibilities."	Not affected	Level of fulfilment, no specific mechanism named for the assessment.	Procurement/supplier management
should	"In case of IT services, configuration has been conceptualized, implemented, and documented based on the necessary security	Not affected	Documentation.	IT operations

Level	Requirement (original wording)	Category	Rationale	Responsible
	requirements."			
should	"The responsible staff is adequately trained."	Not affected	Training on responsibilities, thematically separate from the phishing awareness in 2.1.3.	Human resources
high	"A list exists indicating the IT services concerned and the respective responsible IT service providers. (C, I, A)"	Not affected	Inventory.	Procurement/supplier management
high	"The applicability of the ISA controls has been evaluated and documented. (C, I, A)"	Not affected	Governance and scoping (the TISAX® applicability check itself).	CISO function
high	"The service configuration is included in the regular security assessments. (C, I, A)"	Partially degraded	Time compression: periodic rather than continuous inclusion in security assessments leaves configuration deviations at externally operated services undetected between cycles — consistent with the audit cycle finding in 5.2.6.	Procurement/supplier management
high	"Proof is provided that the IT service providers fulfil their responsibility. (C, I, A)"	Partially degraded	Capability decoupling: an unspecified “evidence” without any requirement of independent verification imports the same document authenticity gap as in 6.1.1.	Procurement/supplier management
high	"Integration into local protective measures (such	Partially degraded	Refers to “secure authentication	IAM team

Level	Requirement (original wording)	Category	Rationale	Responsible
	as secure authentication mechanisms) is established and documented. (C, I, A)"		mechanisms" without an explicit phishing resistance obligation — imports the same gap as 4.1.2.	

Overall classification: Not affected at normal protection requirement (0 of 7 requirements affected). From high protection requirement onward, three additional requirements are added (regular inclusion in security assessments, independently verifiable evidence of responsibility, secure authentication integration) — the overall classification becomes Partially degraded (3 of 12). In practice: relevant only to service provider relationships for assets with high or very high protection requirements; at normal protection requirement there is no Mythos-specific additional need.
Hardening reference: MHC-10, MHC-03, MHC-14 (see Part II)

Chapter 6 summary

Control	Overall classification	Share affected
6.1.1	Partially degraded	1/5 → 4/10
6.1.2	Partially degraded	1/9
6.1.3	Not affected → from high onward: Partially degraded	0/7 → 3/12

Consistent chapter pattern: supplier security relies heavily on submitted documents and self-declarations, whose resistance to forgery falls with AI tools. TISAX® expressly closes this at very high protection requirement (6.1.1) through the obligation of an independent third-party audit — the single most effective lever in the entire chapter.

Chapter 7 — Compliance (2 Controls, 6 requirements)

7.1.1 — To what extent is compliance with regulatory and contractual provisions ensured?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Legal, regulatory, and contractual provisions of relevance to information security (see examples) are determined at regular intervals."	Not affected	A pure scan of the legal position: legislation does not change at the pace of an automated attack — periodic consideration is appropriate here and not relevant to time compression.	Legal
must	"Policies regarding compliance with the provisions are defined, implemented, and communicated to the responsible persons."	Not affected	Governance and communication.	Legal
should	"The integrity of records in accordance with legal, regulatory, and contractual provisions and business requirements is considered."	Not affected	Governance level (taking into account), no concrete technical mechanism named.	Legal

Overall classification: Not affected (0 of 3 requirements affected), irrespective of protection requirement.

7.1.2 — To what extent is the protection of personal data considered when implementing information security?

Level	Requirement (original wording)	Category	Rationale	Responsible
must	"Legal and contractual information security requirements regarding the procedures and processes	Not affected	Legal.	Legal

Level	Requirement (original wording)	Category	Rationale	Responsible
	in the processing of personal data are determined."			
must	"Regulations regarding the compliance with legal and contractual requirements for the protection of personal data are defined and known to the involved persons."	Not affected	Legal and communication.	Legal
must	"Processes and procedures for the protection of personally identifiable information are considered in the information security management system."	Not affected	Governance linkage between ISMS and data protection.	Legal/CISO function

Overall classification: Not affected (0 of 3 requirements affected), irrespective of protection requirement.

Chapter 7 summary

Control	Overall classification	Share affected
7.1.1	Not affected	0/3
7.1.2	Not affected	0/3

Chapter 7 is entirely Not affected — a clean example of controls originally framed in purely legal or regulatory terms, outside the Mythos zone of effect.

Phase 1 complete: overall overview of all 7 chapters

Distribution per chapter (dominant overall classification, normal protection requirement)

Chapter	Controls	Robust	Partially degraded	Friction-only	Not affected
1. IS Policies and Organization	14	0	10	0	4 (of which 1 becomes degraded from high onward)
2. Human Resources	4	0	2	0	2 (of which 1 becomes robust from high onward)
3. Physical Security	3 (+1 with no requirements)	0	2	0	1 (becomes robust from high onward; in addition, 1 of the 2 degraded controls [3.1.4] likewise becomes robust from high onward, see there)
4. Identity and Access Management	4	0	2	1	1 (becomes degraded from high onward)
5. IT Security/Cyber Security	15	4	8	0	3 (of which 1 becomes degraded from high onward)
6. Supplier Relationships	3	0	2	0	1 (becomes degraded from high onward)
7. Compliance	2	0	0	0	2
Total (45 controls)	45	4 (9%)	26 (58%)	1 (2%)	14 (31%)

Assessment of the distribution

TISAX® contains a single friction-only case (4.1.2, password-only authentication with no second factor) — practically all other heavily affected controls do not lose their core effect entirely; they degrade to “partially degraded”. The robust share, at 9%, is comparatively low — the “Information Security” module is predominantly worded in process and organizational terms (many governance requirements with no concrete technical mechanism). The robust share is concentrated almost entirely in Chapter 5 (cryptography, segmentation, tenant separation, redundancy) — precisely where TISAX® imposes explicit technical or architectural requirements.

Recurring pattern categories (preview of phase 2)

Across all chapters, the 26 partially-degraded findings and the one friction-only finding can be assigned to four or five recurring causes, which are consolidated into clusters in phase 2:

1. Periodic rather than event-driven review (time compression) — affects policy reviews, software approvals, access rights reviews, log evaluation, vulnerability and patch management, system and service audits, review of the backup concept, supplier monitoring. By far the most frequent individual finding.
2. Collapse of the patch gap window — risk assessment and patch speed without regard to AI-accelerated exploit development (5.2.5, 5.2.6).
3. Missing obligation of phishing-resistant authentication — consistently present in authentication and remote access controls (4.1.2 — here as friction-only for the pure password factor, elsewhere as partially degraded —, 5.1.2, 5.2.7, 6.1.3).
4. Document and evidence forgery through AI (capability decoupling) — the hiring process (2.1.1), supplier evidence (6.1.1, 6.1.3), compliance documentation.
5. Signature-based detection (capability decoupling) — malware protection (5.2.3).

These five clusters cover the most frequent but not all affected requirements.

Pattern formation and gap analysis (phases 2 and 3)

Unit of analysis: the chapter summaries count 26 controls with an overall classification of "partially degraded" and 1 control with "friction-only" (27 of 45) — this figure refers to the overall classification per control at normal protection requirement and does not reflect several concurrent finding causes per control (for example 4.1.2: both the missing phishing resistance and pure time compression, depending on the protection requirement level). 4 further controls (1.2.2, 4.1.1, 5.2.9, 6.1.3) are unremarkable at normal protection requirement but degrade at high or very high, and accordingly do not appear in the figure of 27, although they do contain individual affected requirements. Pattern formation therefore works at the level of the 74 individual affected requirements (73 partially degraded plus 1 friction-only) across 31 distinct controls, parsed from the document text by the same extraction procedure as the priority table in the appendix. This figure was corrected during revision from an original 71 to 72, and with version 1.1 to 74: the earlier compact index had failed to capture one of the two affected requirements in control 1.4.1 (the should level, assessment criteria), even though Chapter 1 itself already correctly showed it as partially degraded; version 1.1 adds two further individual findings from control 1.6.3, resulting from the primary-source-based reassessment of the editorial catalogue duplicates there (see Methodology, data quality note).
Methodological approach: the eight clusters were formed inductively from the 74 individual findings, not derived from a category grid fixed in advance. Each individual finding was first considered on its own on the basis of its rationale in Chapters 1 to 7; only then were recurring causal patterns (for example "periodic rather than event-driven") consolidated into a named cluster. The number eight is accordingly an empirical result of this particular analysis, not a methodological given — a different TISAX® module or a different requirements catalogue could produce a different number of clusters.

Cluster overview

Cluster	Findings	Controls	Chapter	Dominant criterion
1. Periodic rather than event-driven review	19	12 (1.1.1, 1.2.1, 1.3.1, 1.3.3, 1.3.4, 1.5.1, 1.6.3, 4.1.3, 4.2.1, 5.2.9, 6.1.2, 6.1.3)	1, 4, 5, 6	Time compression
2. Collapse of the patch gap window / ransomware resilience	9	4 (1.3.4, 5.2.5, 5.2.6, 5.2.8)	1, 5	Time compression
3. Missing phishing-resistant authentication	10	6 (4.1.1, 4.1.2, 4.1.3, 5.1.2, 5.2.7, 6.1.3)	4, 5, 6	Capability decoupling / time compression

Cluster	Findings	Controls	Chapter	Dominant criterion
4. AI-supported identity and document forgery	11	5 (2.1.1, 2.1.3, 3.1.1, 6.1.1, 6.1.3)	2, 3, 6	Capability decoupling
5. Signature-based detection versus AI-generated malware	5	1 (5.2.3)	5	Capability decoupling
6. Aggregation resistance	4	4 (1.6.1, 1.6.3, 5.2.4, 5.2.7)	1, 5	Aggregation resistance
7. Autonomous AI agents bypass approval processes	4	2 (1.3.3, 1.3.4)	1	Capability decoupling
8. Missing automation obligation in response	4	1 (1.6.2)	1	Time compression
Individual findings (no recurring pattern)	8	6 (1.2.2, 1.4.1, 1.6.3, 3.1.4, 5.1.2, 5.3.1)	1, 3, 5	—

Total: 74. Cluster 1 (periodic review) is by far the largest and covers four of the seven chapters — the time compression gap is not an isolated phenomenon of individual functional areas but a structural pattern spanning the CISO function, IT operations, IAM and procurement.

Gap analysis against six comparison frameworks

Methodology: for each cluster it is examined whether NIST CSF 2.0, BSI C5:2026, DORA, the CRA, the NIS2 Directive and Implementing Regulation, and ISA/IEC 62443-2-1:2024 already close the same structural gap. This is a comparison at the level of regulatory and normative baseline stance (the maturity of the best practice), not a statement about whether a given company is formally subject to these six frameworks — DORA addresses financial entities and their critical ICT third-party providers, the CRA addresses manufacturers of products with digital elements, NIS2 addresses essential and important entities in certain sectors. The assessment rests on web research from July 2026 into the publicly available texts and secondary sources; it does not replace a clause-precise certification review.

Legend: ✓ = expressly addresses the gap with a binding deadline or method · ~ = mitigates the gap but does not close it fully · ✗ = shares the same structural gap · n/a = outside the substantive scope of the framework

Cluster	NIST CSF 2.0	BSI C5:2026	DORA	CRA	NIS2 Directive/IR	ISA/IEC 62443-2-1:2024
1. Periodic review	~	~	✓ (reporting)	n/a	~	~

Cluster	NIST CSF 2.0	BSI C5:2026	DORA	CRA	NIS2 Directive/IR	ISA/IEC 62443-2-1:2024
2. Patch gap / ransomware	~	~[2]	~	✓ (disclosure)	~	~
3. Phishing-resistant auth	~	n/a[1]	~	n/a	~	~
4. Document forgery	✗	n/a[1]	✓ (suppliers)	n/a	~	n/a[1]
5. Signature-based detection	~[1]	n/a[1]	n/a[1]	n/a	n/a[1]	~[1]
6. Aggregation resistance	~	n/a[1]	n/a[1]	n/a	n/a[1]	~
7. Autonomous AI agents	✗	✗	✗	✗	✗	✗
8. Automation obligation in response	~	~	✓ (de facto)	~ (de facto)	~	~

[1] For these cells, the research could not verify a sufficiently concrete, citable individual requirement; the classification is cautious and indicative rather than clause-precise.

[2] Applies only to the vulnerability-deadline half of the cluster: C5:2026 OPS-25.01AC requires a defined, monitored, CVSS-bound deadline framework for remediating vulnerabilities, binding for review periods from 1 June 2027 — this structural requirement is confirmed across several secondary sources. The concrete CVSS-tiered example deadline values in the criteria catalogue itself could not, however, be verified directly against the BSI primary document (the full PDF text was not technically extractable); one checked secondary source names divergent values for this. They are therefore deliberately not stated in figures here. The ransomware and backup immutability half (control 5.2.8) remains open even under this more cautious classification: criteria OPS-06 to OPS-09 require encryption, RTO/RPO alignment, monitoring and regular restore tests, but nowhere immutability or WORM — the term “immutable” occurs in the entire criteria catalogue only in the context of container images (OPS-26), not for backups.

Die wichtigsten Einzelbefunde:

Cluster 2 (patch gap and ransomware) is the most clearly closed case on the vulnerability-deadline aspect, but not on backup immutability: BSI C5:2026 (binding for review periods from 1 June 2027) requires, in criterion OPS-25.01AC, a defined and monitored, CVSS-bound deadline framework for remediating vulnerabilities (see footnote 2 on the state of sourcing for the exact deadline values). That is precisely the precision the vague TISAX® wordings lack — “adequate

patch management” with no speed SLA (5.2.5, should) and “appropriate timeframe” (5.2.6, must) — and independently confirms that the patch gap collapse diagnosed here is not a special case of this analysis but was already treated as a regulatory gap by a German supervisory authority in 2026. The same does not hold for backup immutability: C5:2026 criteria OPS-06 to OPS-09 (Data Backup and Recovery) require encryption, RTO/RPO alignment, monitoring and restore tests at least annually, daily where sharpened — the term “immutable” occurs in the entire criteria catalogue only in the context of container images (OPS-26), not for backups. The gap diagnosed in control 5.2.8 accordingly also remains open against C5:2026. The CRA additionally obliges manufacturers to report actively exploited vulnerabilities to ENISA within 24 hours, with a follow-up report after 72 hours and a final report 14 days after a remedy becomes available — binding from 11 September 2026. That concerns the disclosure obligation for the manufacturer's own product, not internal patch management of the IT environment, but it is a precedent for clock-bound rather than vaguely worded deadlines.

Cluster 3 (phishing resistance) shows a finer nuance: NIS2 Art. 21(2)(j) is the first EU law of any kind to name multi-factor or continuous authentication expressly — thereby closing the “no second factor at all” gap. The adversary-in-the-middle relay gap (real-time phishing defeats OTP and push), however, persists, because neither NIS2 nor the remaining five frameworks expressly prescribe phishing-resistant methods (FIDO2/WebAuthn) in the core text. That level of precision exists to date only in secondary technical guidance (for example NIST SP 800-63B, AAL3) — TISAX® ISA 2027 accordingly shares this gap with the majority of the comparison frameworks, not merely with itself.

Cluster 4 (document forgery) shows a robust ranking among the comparison frameworks themselves: DORA Art. 30(3) — a sector-specific reference framework for financial entities and their critical ICT third-party providers, see the Methodology above — requires an unrestricted right of audit and access and expressly clarifies that submitted certifications “support but do not replace” the entity's own review. NIS2 Art. 21(2)(d)/(3) in conjunction with NIS2 IR No. 5.1.4 likewise requires a documented supplier assessment with contractually anchored audit rights or access to audit reports — with proportionality latitude (“where appropriate”) rather than DORA's unrestricted claim, but without standard certificates (ISO 27001, SOC 2 Type II, a TISAX® assessment, CSA STAR) being expressly permitted to replace those audit rights. TISAX® ISA 2027, with mere self-declaration or attestation, remains the weakest of the three requirements. For the personnel process (2.1.1: deepfake identity in the hiring process), none of the six frameworks offers an equivalent — that part of the cluster remains entirely TISAX®-original.

Cluster 7 (autonomous AI agents) is the only cluster in which all six comparison frameworks share the same gap as TISAX® ISA 2027 — and that is evidenced by current research, not merely assumed: neither the NIST AI RMF nor ISO/IEC 42001 nor the EU AI Act to date contains any reference to “agent”, “agentic” or autonomous AI systems. NIST has itself recognized this and, in February 2026, launched the “AI Agent Standards Initiative” (937 responses received to a Request for Information); an “AI Agent Interoperability Profile” is announced for Q4 2026 at the earliest — not yet available, therefore, as at the state of this analysis. The finding diagnosed here (1.3.3, 1.3.4: AI assistants call external services or packages on their own, without a human ever

having submitted that particular tool for approval) is accordingly not a TISAX® peculiarity but a gap at the frontier of regulation as a whole. A second, narrower NIST approach is likewise still at draft stage: the NIST Cyber AI Profile (NIST IR 8596) published an initial draft in December 2025 with a comment period until the end of January 2026; according to NIST's own project page, its processing status at the time of this analysis is “Reviewing Comments” — no final text exists, and an express agent reference claimed by secondary sources could not be verified against the NIST primary text (see Part II, MHC-13, Section 2). The MHC catalogue in Part II deliberately goes beyond the pure TISAX® diagnosis at this point (see the section “Original additions beyond TISAX®” below).

Cluster 5 (signature-based detection) and cluster 6 (aggregation resistance) could not be verified with comparable precision: the research found no sufficiently concrete, citable individual requirement on behaviour-based or AI-supported malware detection, or on an explicit obligation to detect deliberately fragmented attack chains, in any of the six frameworks. ISA/IEC 62443-2-1:2024 mentions, in secondary sources, “behavioral analytics” and “continuous monitoring for anomaly detection” as an extension relative to the previous edition, without a sharp, citable obligation being derivable from that. These two clusters are therefore classified more cautiously (~ rather than ✓) than the others.

Individual findings outside the clusters

8 requirements across 6 controls could not be assigned to any recurring pattern: 1.2.2 (segregation of duties undermined by an AI-enabled dual role), 1.4.1 (outdated likelihood-of-occurrence assumptions in the risk methodology — at both the level of the assessment itself and of the underlying assessment criteria, 2 requirements), 1.6.3 high (AI-orchestrated scenarios not provided for in crisis planning, including the editorially duplicated second catalogue version, 2 requirements), 3.1.4 (PIN or password as a fallback option where no encryption obligation exists), 5.1.2 must (unspecific measures for remote access with no minimum standard), 5.3.1 (re-identification of pseudonymized data through AI-supported correlation). For individual findings with no recurring pattern, a six-framework gap analysis is not informative; they remain anchored with their full rationale in Chapters 1 to 7 and feed directly into the MHC catalogue (phase 4) as standalone measures.

Positioning for phase 4

Cluster	Assessment	Recommendation for the MHC
1. Periodic review	Partially addressed industry-wide	Align with the CSF continuity principle (DE.CM) and NIS2 wording (“away from periodic snapshots”), but for TISAX® condense to a single event-driven core principle — 12 affected controls do not need 12 different wordings
2. Patch gap / ransomware	Patch deadlines resolved	For patch deadlines, align closely with the BSI C5:2026 principle (a defined, CVSS-bound

Cluster	Assessment	Recommendation for the MHC
	industry-wide, backup immutability still open	deadline framework instead of “adequate” or “appropriate”); backup immutability requires a standalone MHC requirement despite the C5:2026 review
3. Phishing-resistant auth	Partially addressed industry-wide	NIS2 already names MFA and continuous authentication expressly as an appropriate measure; the phishing-resistance precision (FIDO2/WebAuthn) must be supplied by the MHC itself — no framework yet prescribes it bindingly
4. Document forgery (suppliers)	Partially addressed industry-wide	The DORA principle (sector-specific: financial sector) “a certificate does not replace one’s own review” can be adopted
4. Document forgery (personnel and access, incl. 3.1.1)	TISAX®-original	No comparison framework applicable — a standalone MHC development is needed (deepfake verification in the hiring and access process)
5. Signature-based detection	Predominantly original	A point of reference in IEC 62443-2-1 behavioural analytics, but at its core a standalone MHC wording
6. Aggregation resistance	Predominantly original	No framework found with a sharp requirement — a standalone MHC development
7. Autonomous AI agents	Universally unresolved	The greatest opportunity for a genuinely forward-leaning MHC contribution — even NIST is only at the beginning here (a result at the earliest in Q4 2026)
8. Automation obligation in response	Addressed de facto industry-wide	The DORA principle (sector-specific: financial sector) can be adopted: a hard deadline instead of a soft obligation forces automation de facto without literally requiring it

Cluster 8 and the patch-deadline half of cluster 2 are the “lightest” cases for phase 4 — the solution already exists outside TISAX® and need only be adopted; the backup immutability half of cluster 2, by contrast, requires a wording of its own despite the review undertaken. Cluster 7 is the most valuable individual finding of this analysis: a gap that no relevant framework currently closes, so that an MHC contribution here genuinely breaks new ground rather than merely mirroring existing practice.

Original additions beyond TISAX®

The MHC catalogue in Part II is, with one exception, derived throughout from a concrete TISAX® requirement finding in this document (see Methodology, “What this analysis is — and what it is not”). On one subject this boundary is deliberately and expressly crossed: the security of AI agents.

Occasion: cluster 7 shows that TISAX® ISA 2027 touches AI agents at only two points (1.3.3, 1.3.4 — unreviewed tool or service integration) and that none of the six comparison frameworks from phase 3 closes this gap. A comparison against the OWASP Top 10 for Agentic Applications (ASI01–ASI10, as at June 2026) — a community-based threat taxonomy developed specifically for AI agents — shows that the MHC-13 originally derived from cluster 7 fully covered only one of these ten threat categories and partially covered two more. Seven categories remained unaddressed — not because they are irrelevant to organizations, but because TISAX® ISA 2027 nowhere addresses them concretely enough to derive a finding of its own.

Approach: Part II extends MHC-13 by four additional minimum standards — (a) agent identity and NHI governance, (b) baseline hardening against prompt injection, (c) a kill switch, (d) continuous reassessment of approved components — and adds a new, standalone MHC-16 (multi-agent integrity) for threats that presuppose an architecture with several communicating agents. This brings the catalogue's coverage to eight of the ten OWASP ASI categories (ASI01, ASI02, ASI03, ASI04, ASI05 via MHC-13; ASI06, ASI07, ASI08 via MHC-16). Two categories remain unaddressed even after this extension: ASI09 (human-agent trust exploitation) and ASI10 (rogue agents) — named here deliberately and openly rather than treated tacitly as solved.

Why this marking matters: the methodological value of MRIS lies in the fact that every classification can be traced back to a concrete, citable TISAX® requirement text (see Methodology, “Reproducibility”). The additions described here do not meet that criterion — they are the author's professional judgement on the basis of external threat research (primarily OWASP), not a diagnosis distilled from a TISAX® requirement text. This section, together with MHC-13 and MHC-16 (each in Section 2) in Part II, therefore consistently marks which part is TISAX®-derived and which is original, so that the two categories are not conflated. The “74 individual findings across 31 controls” of the core analysis (see Management Summary) remain unaffected by this extension.

Regulatory outlook: the NIST Cyber AI Profile (NIST IR 8596) is, at the time of this analysis, still at draft stage (status “Reviewing Comments”, see cluster 7 above) and was at no point citable enough to derive a TISAX®-comparable requirement from it. Should NIST IR 8596, or a future TISAX® ISA release, expressly address agent security, the content added here as original should be reviewed and, where appropriate, converted into a TISAX®-derived classification.

Appendix: priority table and standard cross-references

This appendix contains two supplementary overviews for Chapters 1 to 7: a priority table of all individual findings requiring action, together with the MHC reference for implementation, and the assignment of each control to ISO 27001/27017, ISA/IEC 62443-2-1 and NIST CSF 2.0 as groundwork for the gap analysis in phases 2 and 3.

SGA addendum: 10 of the 330 requirements shown (level “SGA”) were added subsequently and were not part of the original count of 320 requirements from phase 1 — affected: 1.1.1, 1.2.1, 1.2.2, 1.5.1 (4), 1.5.2 (2), 1.6.2. They apply irrespective of protection requirement, exclusively under a Simplified Group Assessment.

Standard cross-references per control

Control	Referenz zu anderen Standards (ISO 27001/27017, ISA/IEC 62443-2-1, NIST CSF 2.0)
1.1.1	ISO 27001:2022: A.5.1 ; ISA/IEC 62443-2-1: 6.2.1 ; NIST CSF 2.0: GV.PO-1, GV.PO-2
1.2.1	ISO 27001:2022: 4, 9.3 ; NIST CSF 2.0: GV.OC-01, GV.OC-02, GV.OV-01, GV.OV-02, GV.OV-03, GV.RM-01, ID.IM-01, ID.IM-03, ID.IM-04
1.2.2	ISO 27001:2022: A.5.2, A.5.3 ; ISA/IEC 62443-2-1: 6.2.3, 6.2.5 ; NIST CSF 2.0: GV.RR-02, GV.RR-03
1.2.3	ISO 27001:2022: A.5.8 ; NIST CSF 2.0: GV.OC-04
1.3.1	ISO27001:2022: A.5.9 ; ISA/IEC 62443-2-1: 7.2.1 ; NIST CSF 2.0: ID.AM-01, ID.AM-07
1.3.2	ISO 27001:2022: A.5.10, A.5.12, A.5.13, A.8.10 ; ISA/IEC 62443-2-1: 10.2.1, 10.2.2, 13.3.4 ; NIST CSF 2.0: ID.AM-05, GV.OC-02, PR.PS-03, PR.DS-01, PR.DS-02, PR.DS-10
1.3.3	ISO 27001:2022: A.5.23 ISO 27017: 14.1.1 ; NIST CSF 2.0: ID.RA-10, ID.AM-02, ID.AM-04, GV.SC-07, GV.SC-09, GV.SC-10, PR.PS-02, PR.PS-05
1.3.4	ISO 27001:2022: A.5.19 ; ISA/IEC 62443-2-1: 9.4.x ; NIST CSF 2.0: ID.AM-02
1.4.1	ISO 27001:2022: 6.1, 6.2, 8.2, 8.3 ; ISA/IEC 62443-2-1: 6.3.1 ; NIST CSF 2.0: GV.RM-01, GV.RM-02, GV.RM-03, GV.RM-04, GV.RM-06, ID.RA-03, ID.RA-04, ID.RA-05, ID.RA-06
1.5.1	ISO 27001:2022: 9.2 ISO 27001:2022: A.5.36 ; ISA/IEC 62443-2-1: 6.3.4 ; NIST CSF 2.0: ID.RA-07, ID.IM-01
1.5.2	ISO 27001:2022: A.5.35 ; ISA/IEC 62443-2-1: 6.3.4
1.6.1	ISO 27001:2022: A.5.24, A.6.8 ; ISA/IEC 62443-2-1: 6.2.6, 12.2.1, 12.2.2, 12.2.3 ; NIST CSF 2.0: GV.RM-05, DE.CM-06
1.6.2	ISO 27001:2022: A.5.24-27 ; ISA/IEC 62443-2-1: 12.2.7, 12.2.8 ; NIST CSF 2.0: RS.MA-02, RS.MA-03, RS.MA-04, GV.SC-08, ID.IM-02, ID.IM-04, RS.MA-01, RS.MA-05, RS.CO-02, RS.CO-03, RS.MI-01, RS.MI-02, RC.RP-01, RC.RP-02, RC.RP-04, RC.RP-05, RC.RP-06, RC.CO-03, RC.CO-04
1.6.3	ISO 27001:2022: A.5.29 ; ISA/IEC 62443-2-1: 13.2.1
2.1.1	ISO 27001:2022: A.6.1 ; ISA/IEC 62443-2-1: 6.2.2 ; NIST CSF 2.0: GV.RR-04 PR.AT-02
2.1.2	ISO 27001:2022: A.6.2, A.6.5
2.1.3	ISO 27002: A.7.2.1, A.7.2.2 ; ISA/IEC 62443-2-1: 6.2.4 ; NIST CSF 2.0: PR.AT-01, PR.AT-02, GV.RR-01, GV.RR-02, GV.OC-02, GV.RR-04
2.1.4	ISO 27001:2022: A.6.7, A.8.1 ; ISA/IEC 62443-2-1: 8.4.2

Control	Referenz zu anderen Standards (ISO 27001/27017, ISA/IEC 62443-2-1, NIST CSF 2.0)
3.1.1	ISO 27001:2022: A.7.1 - A.7.6 ; ISA/IEC 62443-2-1: 6.4.1 ; NIST CSF 2.0: GV.OC-04, PR.AA-06, DE.CM-02, DE.CM-03
3.1.2	—
3.1.3	ISO 27001:2022: A.5.10, A.5.11 ; ISA/IEC 62443-2-1: 9.2.2
3.1.4	ISO 27001:2022: A.6.7, A.7.10, A.8.1 ; ISA/IEC 62443-2-1: 9.2.2
4.1.1	ISO 27001:2022: A.5.18 ; NIST CSF 2.0: PR.AA-01
4.1.2	ISO 27001:2022: A.5.15, A.8.5 ; NIST CSF 2.0: PR.AA-02, PR.AA-03
4.1.3	ISO 27001:2022: A.5.16, A.5.17, A.5.18, A.8.9 ; NIST CSF 2.0: PR.AA-04
4.2.1	ISO 27001:2022 A.5.18 ; ISA/IEC 62443-2-1: 11.2.4, 11.3.1, 11.3.4 ; NIST CSF 2.0: PR.AA-05
5.1.1	ISO 27001:2022: A.8.24 ; ISA/IEC 62443-2-1: 10.2.8
5.1.2	ISO 27001:2022: A.5.14 ; NIST CSF 2.0: PR.DS-02, PR.IR-01
5.2.1	ISO 27001:2022: A.8.32 ; ISA/IEC 62443-2-1: 7.2.4 ; NIST CSF 2.0: PR.PS-01, PR.PS-04
5.2.2	ISO 27001:2022: A.8.31 ; ISA/IEC 62443-2-1: 6.3.3
5.2.3	ISO 27001:2022: A.8.1 ; ISA/IEC 62443-2-1: 9.2.1, 9.3.1, 9.3.2, 9.3.3
5.2.4	ISO 27001:2022: A.5.17, A.8.15 ; ISA/IEC 62443-2-1: 10.2.5, 12.2.4, 12.2.5 ; NIST CSF 2.0: PR.PS-04, DE.AE-02, DE.AE-03, DE.AE-04, DE-AE-08, DE.CM-01, DE.CM-09, DE.AE-06
5.2.5	ISO 27001:2022: A.8.8, A.8.19 ; ISA/IEC 62443-2-1: 9.4.2, 12.2.9 ; NIST CSF 2.0: ID.RA-01, ID.RA-08
5.2.6	ISO 27001:2022: A.5.36, A.8.34, A.8.8 ; ISA/IEC 62443-2-1: 6.3.2 ; NIST CSF 2.0: ID.IM-02, DE.CM-09, IR.RA-09
5.2.7	ISO 27001:2022: A.8.20, A.8.22 ; ISA/IEC 62443-2-1: 6.3.2, 8.2.1, 8.2.2, 8.2.6, 8.2.8 ; NIST CSF 2.0: ID.AM-03, PR.IR-01, PR.IR-03
5.2.8	ISO 27001:2022: A.5.30 ; ISA/IEC 62443-2-1: 13.2.1 ; NIST CSF 2.0: PR.IR-02
5.2.9	ISO 27001:2022: A.8.13 ; ISA/IEC 62443-2-1: 13.3.1, 13.3.2, 13.3.3, 13.3.4, 13.3.5 ; NIST CSF 2.0: PR.DS-11, RC.RP-03
5.3.1	ISO 27001:2022: A.5.8, A.8.25 ; NIST CSF 2.0: PR.PS-01, PR.PS-06
5.3.2	ISO 27001:2022: A.8.21 ; NIST CSF 2.0: PR.IR-04
5.3.3	ISO 27001:2022: A.5.23 ; ISO 27017: CLD.8.1.5 ISA/IEC 62443-2-1: 10.2.6 NIST CSF 2.0: ID.AM-08, GV.RR-04, PR.PS-03
5.3.4	ISO 27017: CLD.9.5.1, CLD.9.5.2 ; NIST CSF 2.0: DE.CM-06, PR.IR-01, GV.SC-07
6.1.1	ISO 27001:2022: A.5.19 - A.5.22 ; NIST CSF 2.0: GV.SC-02, GV.SC-03, GV.SC-04, GV.SC-05, GV.SC-06
6.1.2	ISO 27001:2022: A.5.14, A.6.6
6.1.3	ISO 27001:2022: A.5.23 ISO 27001:2022: A.8.9 ; ISO 27017: CLD.6.3.1 ; ISA/IEC

Control	Referenz zu anderen Standards (ISO 27001/27017, ISA/IEC 62443-2-1, NIST CSF 2.0)
	62443-2-1: 6.2.3 ; NIST CSF 2.0: GV.OC-05, GV.SC-01, GV.SC-02
7.1.1	ISO 27001:2022: A.5.31, A.5.32, A.5.33 ; NIST CSF 2.0: GV.SC-02, GV.OC-03
7.1.2	ISO 27001:2022: A.5.34

Priority table: individual findings requiring action

This table replaces the earlier “compact index per requirement”. Rather than all 330 assessed requirements, it shows only the 74 individual findings requiring action (category partially degraded or friction-only), each with a short label and a direct reference to the relevant MHC entry in Part II. Unremarkable requirements (not affected, robust) are not listed here; they appear with full wording and rationale in Chapters 1 to 7.

Control	Short label	Category	Hardening reference
1. IS Policies and Organization			
1.1.1	Policy review not event-driven (“periodic”)	Partially degraded	MHC-10
1.2.1	ISMS effectiveness review not event-driven	Partially degraded	MHC-10
1.2.2	Segregation of duties organizational only, not technically enforced	Partially degraded	standalone measure
1.3.1	Asset inventory maintenance not event-driven	Partially degraded	MHC-10
1.3.3	AI agent can call external services without approval	Partially degraded	MHC-13
1.3.3	Approval sample check with no ongoing detection	Partially degraded	MHC-10
1.3.4	AI coding assistant loads software packages without approval	Partially degraded	MHC-13
1.3.4	Same gap for specialist software (maintenance tools)	Partially degraded	MHC-13
1.3.4	Repository protection: signature check not mandatory	Partially degraded	MHC-13
1.3.4	Software approval review not event-driven	Partially degraded	MHC-10
1.3.4	Patch status known, but no response speed required	Partially degraded	MHC-09

Control	Short label	Category	Hardening reference
1.4.1	Risk assessment still assumes a narrow attacker circle	Partially degraded	standalone measure
1.4.1	Same gap at the level of the assessment criteria	Partially degraded	standalone measure
1.5.1	Policy compliance review not event-driven	Partially degraded	MHC-10
1.5.1	Requirements compliance review not event-driven	Partially degraded	MHC-10
1.5.1	Group-wide audit plan with no trigger for new sites	Partially degraded	MHC-10
1.6.1	Reporting system does not detect fragmented attack steps	Partially degraded	MHC-05
1.6.2	Incident handling purely human, no automation obligation	Partially degraded	MHC-11
1.6.2	Same gap in responding to reports	Partially degraded	MHC-11
1.6.2	Response time targets with no automation obligation to meet them	Partially degraded	MHC-11
1.6.2	Escalation mechanism itself remains human-borne	Partially degraded	MHC-11
1.6.3	Crisis detection designed for gradual, not rapid, escalation	Partially degraded	MHC-05
1.6.3	Crisis plan updating not event-driven	Partially degraded	MHC-10
1.6.3	Crisis scenarios do not include an AI-orchestrated multi-track attack	Partially degraded	standalone measure
1.6.3	Crisis scenarios do not include an AI-orchestrated multi-track attack (editorial catalogue duplicate, see data quality note)	Partially degraded	standalone measure
1.6.3	Communication strategy has no shortened response window	Partially degraded	MHC-05
1.6.3	Communication strategy has no shortened response window (editorial catalogue duplicate, see data quality note)	Partially degraded	MHC-05
1.6.3	Crisis plan assessment not event-driven	Partially	MHC-10

Control	Short label	Category	Hardening reference
		degraded	
2. Human Resources			
2.1.1	Identity check vulnerable to a deepfake applicant identity	Partially degraded	MHC-15
2.1.1	Live interview vulnerable to a real-time deepfake	Partially degraded	MHC-15
2.1.1	References and testimonials more easily forged	Partially degraded	MHC-15
2.1.3	Awareness training designed around outdated phishing recognition markers	Partially degraded	MHC-15
3. Physical Security			
3.1.1	Access control vulnerable to AI-supported pretence	Partially degraded	MHC-15
3.1.1	Visitor registration vulnerable to forged credentials	Partially degraded	MHC-15
3.1.4	Encryption of mobile devices not mandatorily prescribed	Partially degraded	standalone measure
4. Identity and Access Management			
4.1.1	Locking of lost credentials not immediate or automated	Partially degraded	MHC-03
4.1.2	“State of the art” does not enforce phishing resistance	Partially degraded	MHC-03
4.1.2	Password strength ineffective against real-time phishing	Friction-only	MHC-03
4.1.2	Privileged accounts: simple 2FA without phishing resistance suffices	Partially degraded	MHC-03
4.1.2	List of options allows fulfilment without phishing resistance	Partially degraded	MHC-03
4.1.2	Two-factor example carries no phishing-resistance obligation	Partially degraded	MHC-03
4.1.3	Account review not event-driven	Partially degraded	MHC-10
4.1.3	Possession factor (SMS OTP) vulnerable to SIM swapping	Partially degraded	MHC-03

Control	Short label	Category	Hardening reference
4.1.3	Account review in customer systems not event-driven	Partially degraded	MHC-10
4.2.1	Access rights review not event-driven	Partially degraded	MHC-10
4.2.1	Shorter review interval remains periodic, not event-driven	Partially degraded	MHC-10
5. IT Security/Cyber Security			
5.1.2	Transport protection with no mandatory minimum encryption standard	Partially degraded	standalone measure
5.1.2	Remote access: VPN protects the transport path, not authentication	Partially degraded	MHC-03
5.2.3	Signature-based malware protection weak against AI-generated variants	Partially degraded	MHC-05
5.2.3	Virus scanner example signals signature-based detection	Partially degraded	MHC-05
5.2.3	Same signature detection limit in file scanning	Partially degraded	MHC-05
5.2.3	Full scan periodic and signature-based — a double weakness	Partially degraded	MHC-05
5.2.3	Encrypted content bypasses gateway inspection	Partially degraded	MHC-05
5.2.4	Log review does not detect fragmented attack patterns	Partially degraded	MHC-05
5.2.5	Vulnerability assessment underestimates the shortened exploit window	Partially degraded	MHC-09
5.2.5	Patch management with no mandatory speed SLA	Partially degraded	MHC-09
5.2.6	Remediation deadline (“appropriate timeframe”) indefinite	Partially degraded	MHC-09
5.2.6	Audit cycle leaves the patch gap open between reviews	Partially degraded	MHC-09
5.2.6	Scan frequency only periodic even at very high	Partially degraded	MHC-09
5.2.7	Lateral movement detection not behaviour-based	Partially degraded	MHC-05

Control	Short label	Category	Hardening reference
5.2.7	Remote access risk imports the authentication gap from 4.1.2	Partially degraded	MHC-03
5.2.8	Contingency plan updating not event-driven	Partially degraded	MHC-08
5.2.8	Backup protection with no immutability requirement (integrity)	Partially degraded	MHC-08
5.2.8	Same gap, confidentiality side	Partially degraded	MHC-08
5.2.9	Review of the backup concept not event-driven	Partially degraded	MHC-10
5.3.1	Test data: no re-identification-resistant anonymization required	Partially degraded	standalone measure
6. Supplier Relationships			
6.1.1	Supplier reports with no independent counter-check	Partially degraded	MHC-14
6.1.1	Self-declaration or attestation sufficient as evidence	Partially degraded	MHC-14
6.1.1	Review frequency good, evidence authenticity problem remains	Partially degraded	MHC-14
6.1.1	Same evidence authenticity gap in contract review	Partially degraded	MHC-14
6.1.2	NDA handling procedure not event-driven	Partially degraded	MHC-10
6.1.3	Service provider configuration checked periodically rather than continuously	Partially degraded	MHC-10
6.1.3	Unspecified evidence with no independent verification	Partially degraded	MHC-14
6.1.3	Authentication mechanisms carry no phishing-resistance obligation	Partially degraded	MHC-03