FRAMEWORK · COMPANION · V1.0

The control-objective chain

Bearing west · Second piece · PDCA+ v2.0 · Public Review

Why the substrate pivots on the control objective rather than the control implementation, and how the chain from regime intent to stakeholder entitlement is structured around that fulcrum.

AUTHOR

Joacim Brandell

Written for practitioners working with regulatory regimes · worked example draws on defence supply-chain assurance across CMMC, DEFSTAN, and Canada Protected B

How to read this piece

“W.1 – Supporting assurance towards a regime” established that PDCA+ supports assurance toward a regime by holding evidence shapes, classifications, translation rules, and shared capabilities — and that the framework composes across regimes because operational reality is held once and regime-specific addressing accumulates. This is the answer to the leading question. It is also, deliberately, not the whole picture.

What was left implicit is the interior structure of how a regime's intent connects, through a sequence of records, to the evidence that demonstrates a stakeholder's entitlement is being met. The structure matters because the choice of where the chain pivots — where regime requirements meet operational reality — is what determines whether the multi-regime composition argument actually works. The pivot has to be at the right level. If it is at the wrong level, the composition fails and the framework collapses into parallel programmes with extra steps.

This piece elaborates that interior structure. It develops the conceptual distinction between control objectives and control implementations, walks through the chain of records from regime intent down to stakeholder entitlement, shows why the chain's fulcrum sits at the Control Objective Record, and works one example through the full chain across the three regimes the W.1's matrix used.

A practitioner who has read W.1 will find here the answer to the question they were probably already asking: "yes, but how, exactly, does the framework hold the regime's requirements without rebuilding them per regime?" The answer is: by pivoting the chain at the control objective, where regimes converge, rather than at the control implementation, where they diverge.

Control objective versus control implementation

The two terms are routinely used interchangeably and routinely confused. The distinction is structural and has direct consequences for whether multi-regime assurance scales.

A control implementation is the specific way an organisation has chosen to address a particular requirement, in the particular operational environment the organisation runs. "We require FIDO2 hardware keys for administrative access to production systems" is an implementation. So is "we run authenticated vulnerability scans weekly using Tenable across all hosts in the production environment." Implementations are concrete, context-dependent, and rich — they reflect tooling choices, operational maturity, cost-benefit decisions, and the specific risks the organisation is responding to in its own environment.

A control objective is what the implementation is supposed to achieve. "Administrative access to production systems is authenticated using phishing-resistant factors" is an objective. So is "the organisation maintains current awareness of exploitable vulnerabilities in its information systems." Objectives are abstract relative to implementations — they describe

the end state rather than the means of achieving it. Multiple implementations can satisfy a single objective, and a single implementation often satisfies multiple objectives.

Control objective	Control implementation
Administrative access is authenticated using phishing-resistant factors	FIDO2 hardware keys for production admin access; backup TOTP for break-glass
The organisation maintains current awareness of exploitable vulnerabilities	Weekly authenticated Tenable scans; daily CVE feed ingestion; CISA KEV monitoring
Configuration deviations from baseline are detected and addressed	Hourly drift detection via configuration-management tooling; standing change-review board for accepted deviations
Incidents are recognised, classified, and responded to within defined timelines	24×7 SOC with SIEM-driven detection; defined IR runbooks with severity-based SLAs; quarterly tabletop exercises

The distinction matters because regimes converge much more strongly on objectives than on implementations. CMMC, DEFSTAN, Canada Protected B, and most credible cybersecurity regimes all care that administrative access is phishing-resistant — they share that objective. They differ on which implementations they accept as evidence (some accept TOTP, some require hardware tokens, some specify particular FIPS validation levels), on how often evidence must be refreshed, on what specific factors qualify as phishing-resistant. The objective is shared substantially across regimes; the implementations vary substantially within regimes.

If the framework's substrate is structured around control implementations, the substrate has to hold a near-complete record per regime, because the implementation-level details a regime cares about are largely regime-specific. Multi-regime reuse is poor at the implementation layer. **If the framework's substrate is structured around control objectives**, the substrate holds one objective record that points to one implementation record, and multiple regimes can address the same objective from their own requirement records. Multi-regime reuse is high at the objective layer. The framework's structural-reuse claim depends on choosing the right layer to pivot on, and the right layer is the objective.

> **THE CONCEPTUAL MOVE, IN ONE SENTENCE**
>
> *Regimes converge on what the organisation must achieve; they diverge on how the organisation may achieve it. Pivoting the substrate at the objective layer captures the convergence; pivoting at the implementation layer captures the divergence and forfeits the reuse.*

The chain, as a sequence of records

The substrate represents the connection between a regime's intent and a stakeholder's demonstrated entitlement as a chain of records. Each record holds a specific kind of content; each is owned by an identifiable party; each connects to its neighbours through declared relationships rather than through implicit understanding held in someone's head.

Reading the chain from top to bottom traces the path from regime to evidence. Reading it from bottom to top traces the path from evidence to entitlement. The chain runs in both directions, and the substrate uses both. Surfacing a finding at any record propagates to the records above and below it, which is what lets the conjunction observe assurance-relevant inadequacies as they arise.

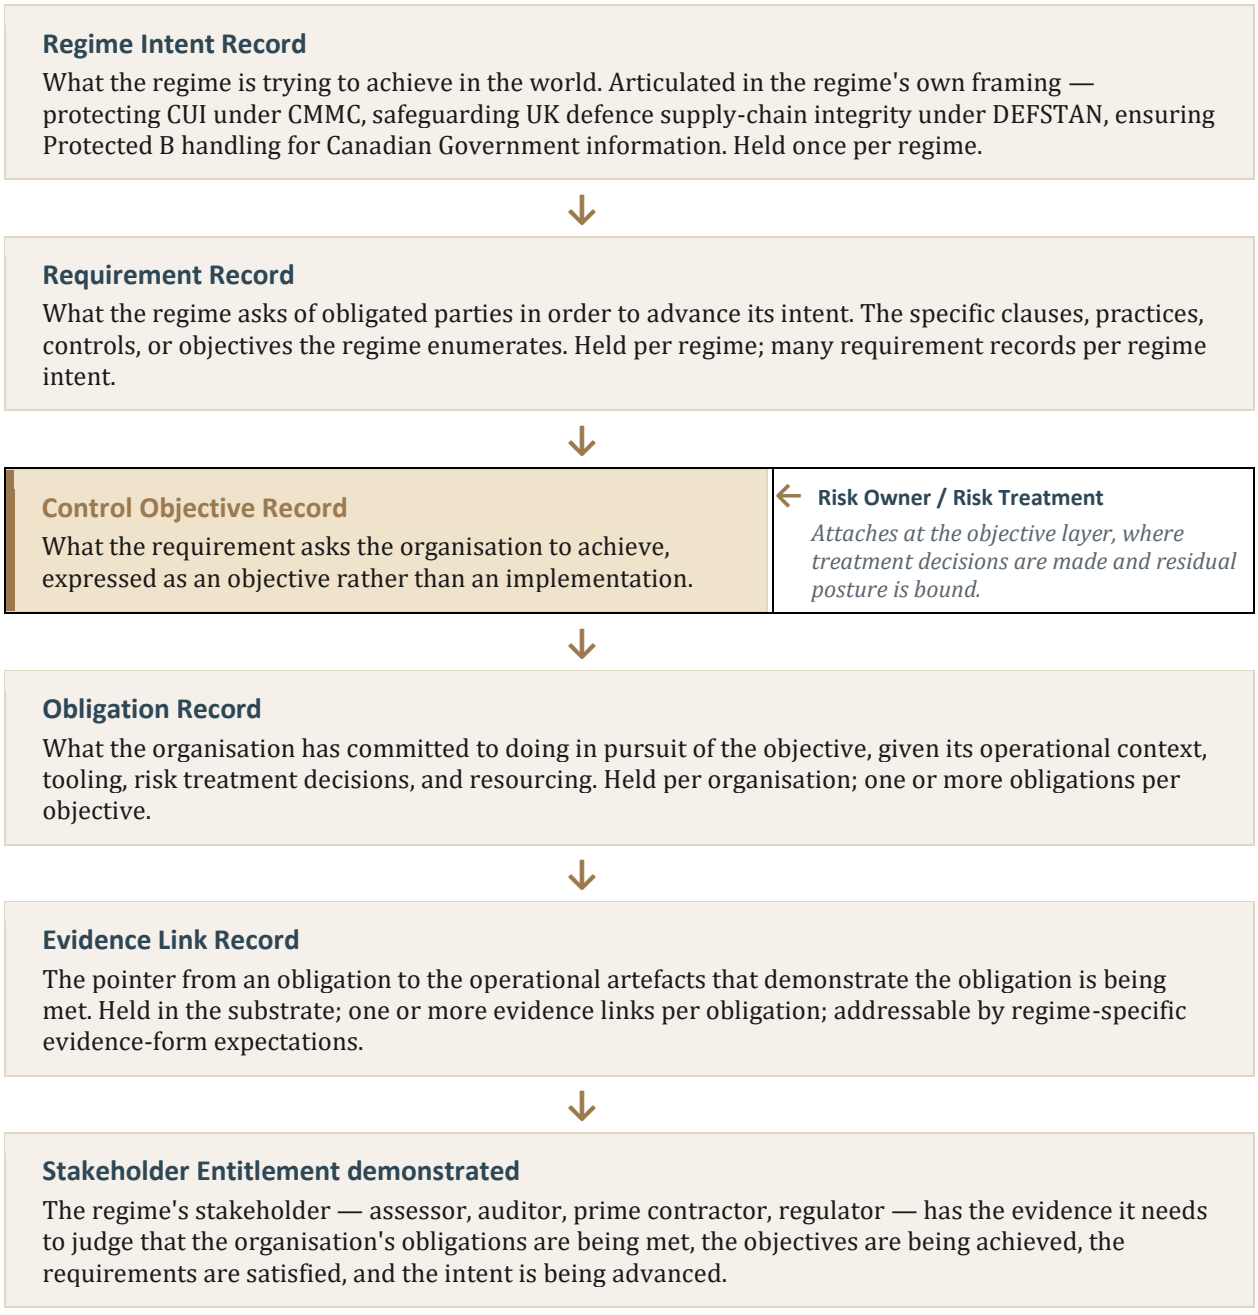

How the chain works in practice

Reading top-to-bottom, the chain decomposes the regime's intent into something the organisation can act on. Intent decomposes into requirements; requirements decompose into objectives; objectives generate obligations the organisation commits to; obligations link to evidence; evidence demonstrates entitlement. Each step is a refinement, and each refinement is held as a record the substrate can address.

Reading bottom-to-top, the chain composes the operational reality back into something a regime stakeholder can consume. Operational artefacts link as evidence to obligations the organisation has made; obligations connect to the objectives they pursue; objectives roll up to the requirements that ask for them; requirements aggregate into the intent the regime is advancing. The stakeholder, examining evidence, can trace any evidence element back up to the intent it ultimately advances — and back down to the operational artefact that underwrites it.

The substrate holds all the records and all the links. Adding a new regime adds Regime Intent Record, Requirement Records, and Obligation Records specific to that regime — but the Control Objective Records, where the regime's requirements connect to operational reality, are typically shared with the objectives already held for other regimes. The structural reuse the W.1's §4 named lives precisely here, in this layer of the chain.

Why the fulcrum sits at the Control Objective Record

The chain has a privileged position — the place where regime-side records and organisation-side records meet. That position is the Control Objective Record, and the choice to place the fulcrum there is what makes the rest of the architecture work.

Above the Control Objective Record, the records are regime-side. They are held by the regime, articulated in its vocabulary, owned by its authorities. The organisation reads them but does not write them. New regime versions, new flow-downs, and new regimes themselves all enter the substrate through this side of the chain.

Below the Control Objective Record, the records are organisation-side. They are held by the organisation, articulated in operational vocabulary, owned by operational authorities. The regime reads them through evidence presentation but does not write them. Operational reality enters the substrate through this side of the chain.

At the Control Objective Record, the two sides meet. The objective is articulated in vocabulary that both regime and organisation can read — abstract enough to be regime-portable, concrete enough to constrain operational obligation. This is what makes the objective the right place to pivot. Pivoting higher (at the Requirement Record) would lock the organisation into regime-specific articulation and forfeit reuse across regimes. Pivoting lower (at the Obligation or Evidence Link Record) would force the regime to address operational specifics it neither owns nor cares about, and would make every regime addressing the same operational reality have to re-do the connection work.

Why risk treatment attaches here

The lateral input to the Control Objective Record — Risk Owner / Risk Treatment — is structurally significant. Risk treatment is the activity by which an organisation decides what posture to commit to in pursuit of an objective. It is not regime-side work (the regime does not own the risk treatment decision; it owns only the requirement that some treatment be made). It is not implementation-side work (the implementation realises the treatment decision but does not make it). It is objective-level work.

Risk treatment decisions belong at the objective layer because that is where the question "what posture will we accept in pursuit of this objective?" is meaningfully answerable. A risk owner can examine the objective, examine the operational context, examine the obligations the organisation is willing to commit to, and decide whether the residual posture is acceptable. The decision binds the objective to a treatment commitment, which the obligations realise. The Control Objective Record holds the binding; the Risk Owner record holds the authority; the substrate connects them so that risk-treatment changes propagate to obligations and evidence below, and risk-treatment authority is traceable to the objective above.

Without this lateral connection at the objective layer, risk treatment would either disappear from the chain (decisions made implicitly, untraceable) or be misplaced (attached to implementations, which fragments treatment across many implementations that should share one treatment posture). The Control Objective Record is the right place for risk treatment to attach because it is the right level at which risk acceptance can be meaningfully bound.

> **THE FULCRUM, STRUCTURALLY**
> *The Control Objective Record is where regime-side records meet organisation-side records, where multiple regimes' requirements converge on a shared objective, and where risk treatment decisions are meaningfully made. The substrate's whole architecture rests on this layer being where the chain pivots.*

A worked example: risk acceptance, named

W.1's matrix in §4 used six evidence elements. Walking one of them — risk acceptance, named — through the full chain across the three regimes makes the architecture concrete.

Regime intent records

Three Regime Intent Records exist in the substrate, one per regime. The CMMC intent record holds the US Department of Defense's intent to protect Controlled Unclassified Information across the defence industrial base. The DEFSTAN intent record holds the UK Ministry of Defence's intent to maintain cyber security across its supply chain. The Canada Protected B intent record holds the Government of Canada's intent to protect Protected B information across its information systems. The intent records differ substantively — different

jurisdictions, different information categories, different policy contexts. They are held separately because they are separate intents.

Requirement records

Each regime decomposes its intent into requirements. For risk acceptance, named, the relevant requirement records are: CMMC practice RA.L2-3.11.1 ("Periodically assess the risk to organizational operations…resulting from the operation of organizational systems and the associated processing, storage, or transmission of CUI"), DEFSTAN 05-138 clause 4.4 (risk treatment with named accountability), and ITSG-33 control RA-3 (risk assessment with documented and accepted residuals). Three Requirement Records, one per regime, each held in the regime's own vocabulary.

Control objective record

Where these three Requirement Records connect down into the substrate, they meet a single Control Objective Record. The objective: "the organisation deliberately and authoritatively binds the residual risk posture for systems within scope, with named accountability for the binding decision and traceable evidence of the basis on which the decision was made." This objective is regime-portable — all three regimes can read it as articulating what they require, even though their requirement records phrase the ask differently. The objective is also operationally meaningful — the organisation can act on it, and its risk owners can attach treatment decisions to it.

This is the structural reuse working in practice. Three regime-side requirement records, three different vocabularies, three different control reference numbers — addressing one objective. The substrate holds one Control Objective Record, not three; the regime mappings address that one record from three sides.

Risk Owner / Risk Treatment, attached at the objective layer

The named risk owner for systems handling CUI (which is also covered DEFSTAN-side and Canada-Protected-B-side) attaches their risk treatment decision to the Control Objective Record. The treatment specifies the posture: which risks are accepted, which are mitigated, which are transferred. The named accountability is the risk owner's; the supporting evidence is referenced; the conditions under which acceptance was given are recorded.

Because the treatment attaches at the objective rather than at an implementation, the same treatment binding serves all three regimes. If the treatment changes — if the risk owner accepts a different residual, or different evidence becomes available — the change propagates to all three regimes through the substrate. No regime-specific treatment reconciliation is needed.

Obligation records

The Control Objective Record generates Obligation Records that describe what the organisation has committed to doing in pursuit of the objective. For risk acceptance, these obligations include: maintaining a current risk register for systems in scope, performing risk analysis on a defined cadence, recording risk acceptance decisions with named accountability, retaining evidence of the decision basis for the duration each regime requires.

Some obligations are regime-portable (the obligation to maintain a current risk register satisfies all three regimes). Some are regime-specific (CMMC may expect particular evidence forms; DEFSTAN may require specific retention periods; Canada Protected B may require specific authorisation signatures). The substrate holds the obligations as a coherent set, addresses regime-portable obligations once, and addresses regime-specific obligations per regime.

Evidence link records and stakeholder entitlement

Each obligation links to the operational artefacts that demonstrate it is being met. The risk register itself is an artefact; risk-acceptance memoranda are artefacts; the records of risk analyses performed on cadence are artefacts. Evidence Link Records connect each obligation to the artefacts that underwrite it, with the artefacts addressable by regime-specific evidence-form expectations.

When the assessor or auditor arrives — the C3PAO for CMMC, the DEFSTAN auditor, the Canadian Communications Security Establishment authoriser — they consume the evidence through the regime's expected forms. The Stakeholder Entitlement demonstrated is the regime's stakeholder having what they need to make their assurance judgement. The same operational reality, evidenced once, demonstrates entitlement to three different stakeholders, each addressing it through their own regime's lens.

> **THE CHAIN DOING ITS WORK**
>
> *Three regimes, three intents, three sets of requirements — meeting at one Control Objective Record, one risk treatment binding, one operational reality, and one set of evidence artefacts. The structural reuse is not asserted; it is held in the records, traceable end-to-end, addressable from any of the three regime sides.*

What this changes for the practitioner

Three operational changes follow from this chain structure, each worth naming directly.

First, control-objective work becomes the highest-leverage assurance activity. The objective is where multiple regimes' requirements converge, where risk treatment attaches, where the substrate's structural reuse is realised. Practitioners who invest in articulating control objectives precisely — abstract enough to be regime-portable, concrete enough to constrain obligation — get disproportionate value from the work. Practitioners who skip this layer and jump from regime requirements directly to implementations lose the convergence, and with it the reuse.

Second, risk treatment becomes substrate-attached rather than register-attached. The risk register continues to exist as an operational artefact, but the risk treatment binding lives at the Control Objective Record. This is a subtle but consequential shift: it means the risk owner's decision about residual posture is connected to the objectives the residual applies to, not just listed in a register. Changes to objectives surface to risk owners; changes to treatment

surface to obligations and evidence; the substrate connects what was previously connected only through human memory.

Third, evidence becomes a property of obligations linked to objectives, not a property of regimes. The same evidence artefact can underwrite multiple obligations across multiple regimes, because the link is from obligation to evidence, not from regime to evidence. The audit-readiness question shifts from "do we have CMMC evidence?" to "are our obligations linked to current evidence?" — and the answer to the second question simultaneously answers the assurance-readiness question for every regime.

None of these changes eliminates work, but each one moves work to where it is leveraged. The total assurance burden across multiple regimes becomes much smaller than the sum of the per-regime burdens — which is, recursively, W.1's central claim, viewed from inside the chain that makes it true.

Closing

W.1 answered the leading question — in what way does PDCA+ support assurance toward a regime — at the level of the framework's overall contribution. This piece has answered the more specific question that the answer from W.1 raises: how does the framework actually hold the connection between regime requirements and operational evidence, in a way that scales across multiple regimes rather than fragmenting under them?

The answer is the chain. Regime intent decomposes through requirements to control objectives; the chain pivots at the objective layer where regimes converge and risk treatment attaches; objectives generate obligations; obligations link to evidence; evidence demonstrates entitlement. Each record is held in the substrate, addressable, traceable, and connected to its neighbours through declared rather than implicit links.

The choice to pivot the chain at the Control Objective Record is what makes the multi-regime composition work. It is also what makes risk treatment structurally well-placed, what makes evidence reusable across regimes, and what makes the substrate's continuous surfacing of inadequacies operationally meaningful. The chain is small — six record types — but each is doing specific work that the substrate cannot do without it. Taking any one out would compromise the architecture.

Practitioners working with this architecture in mind will recognise that most of the high-leverage assurance work is at the objective layer. Most of the high-frequency operational work is at the obligation and evidence layers. Most of the high-stakes governance work — risk treatment, named acceptance — also attaches at the objective layer. The chain tells the practitioner where to direct attention, and the substrate makes the attention productive.

References

- U.S. Department of Defense. Cybersecurity Maturity Model Certification (CMMC) 2.0 Programme documentation, 32 CFR Part 170.
- NIST Special Publication 800-171 Rev. 3. Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations.
- NIST Special Publication 800-171A. Assessing Security Requirements for Controlled Unclassified Information.
- UK Ministry of Defence. Defence Standard 05-138 (Cyber Security for Defence Suppliers).
- Government of Canada. ITSG-33 IT Security Risk Management: A Lifecycle Approach.
- Government of Canada, Treasury Board Secretariat. Directive on Security Management — Protected B handling requirements.
- U.S. Department of State. International Traffic in Arms Regulations (ITAR), 22 CFR §§120–130.
- Directive (EU) 2022/2555 — NIS2 Directive, and national transpositions thereof.