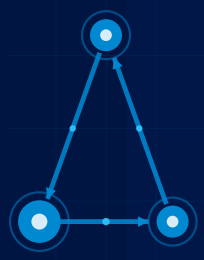

CRC

CONSOLIDATE, REDUCE, CLOSE

The Minimum Surface Standard for AI-era regulated infrastructure.

MIT LICENSE · v1.0 · GENERAL REASONING, INC.

> ***Anthropic's Mythos model did not create a new category of threat. It revealed that the security posture enterprise computing has accepted for decades was calibrated to a human-speed adversary that no longer exists in isolation. The vulnerability window between discovery and exploitation has collapsed to hours. The cognitive partitioning that made siloed security feel adequate -- the Dunbar Perimeter -- is gone. This document defines the Consolidate, Reduce, Close (CRC) framework: three directives that constitute the Minimum Surface Standard for regulated infrastructure in the AI era.***

I. The Dunbar Perimeter

In the 1990s, anthropologist Robin Dunbar observed that humans can maintain stable social relationships with roughly 150 people at once. Beyond that threshold, the cognitive load becomes unmanageable. We lose track of context, dependencies, and the relationships between relationships. We lose the ability to reason about the whole.

The same limit applies to complex technical systems.

A security engineer can hold one application's architecture in their head with genuine clarity. Maybe two. By the time an organization is reasoning about five, ten, or fifty interconnected enterprise platforms simultaneously -- their APIs, their data flows, their trust relationships, their shared identity layers -- it is at the outer edge of what human cognition can manage without losing resolution.

So organizations did what humans always do when they hit a cognitive limit: they partitioned. The Salesforce team owned Salesforce. The SAP team owned SAP. The identity team owned Active Directory. Finance owned the ERP. Each domain had an expert team. Each team had a handle on their domain. Each domain felt like a security boundary.

> ***We call this implicit boundary the Dunbar Perimeter -- the security boundary created not by architecture, but by the cognitive limit of the humans responsible for it. It was never a real control plane. It was a limit of perception.***

The systems were always connected. Customer records flowed from CRM to ERP. Compensation data moved from HRIS to finance. Identity tokens crossed every boundary on every request. These connections were real, operational, and consequential -- but they were not held in anyone's head simultaneously. No single team was responsible for the full graph. Each team owned a node. Nobody owned the edges.

II. The SaaS Business Model and the Industrialization of Dunbar Perimeters

The modern enterprise SaaS ecosystem did not create the Dunbar Perimeter. It industrialized it.

The foundational premise of the SaaS business model is best-of-breed specialization: a purpose-built solution for a specific domain, sold to the team responsible for that domain. CRM for sales. ERP for finance. HRIS for people operations. ITSM for IT. Procurement. Collaboration. Analytics. Compliance. Each application is optimized, supported, and -- crucially -- secured within its own boundary. Each carries its own compliance certifications. Each has its own security team and penetration testing program. Each is, in isolation, a defensible system.

And each one connects to all the others.

The average enterprise runs hundreds of SaaS applications. Each integration creates an edge between two nodes. Each edge crosses a Dunbar Perimeter. It belongs to neither team. It is secured by no single owner. It exists in the gap between two cognitive domains that were designed to be independent but are operationally inseparable.

> ***Individually secure SaaS systems can create collectively insecure architectures. The SaaS business model sells you perimeters. It does not sell you the connections between them.***

This architecture accumulated quietly over decades, driven entirely by legitimate operational logic. Best-of-breed selection makes procurement sense. Domain-specific ownership makes organizational sense. Compliance certification makes regulatory sense. The result, however, is an enterprise technology landscape that is a highly connected attack graph, partitioned only by the cognitive limits of the humans responsible for its pieces.

The reason this was not catastrophic prior to 2026 was a symmetry of limitation: the attacker had a Dunbar limit too. Human attackers, even sophisticated nation-state operators, could only hold so many systems in their heads simultaneously. They picked targets. They worked domains. The cognitive partitioning that constrained defenders also constrained attackers. The race was conducted at human speed on both sides.

That symmetry is gone.

III. What Mythos Revealed

Anthropic's Mythos model, released in limited preview in April 2026 as part of Project Glasswing, does not have a Dunbar limit. It does not partition systems into cognitive domains. It does not fatigue. It does not stop at

application boundaries or respect the organizational silos that defined security responsibility.

It constructs and traverses the full dependency graph -- across identity systems, SaaS platforms, internal services, and data flows -- simultaneously, autonomously, and without the cognitive overhead that made siloed security feel adequate. In testing, it identified critical vulnerabilities in every major operating system and browser, many of them decades old. It chained those vulnerabilities autonomously, linking a weakness in one component to adjacent weaknesses to produce complete, working attack paths -- overnight, while the security team slept.

The connections that were previously invisible to defenders -- the edges that crossed Dunbar Perimeters and belonged to nobody -- are now first-class attack paths. Not new connections. The ones that were always there.

The inversion is complete:

Defenders	still evaluate systems one domain at a time.
Adversaries	now traverse them as a single unified graph.
Architecture	assumed separation between systems.
Attackers	exploit their composition.
SaaS vendors	secure their own perimeter.
Nobody	owns the edges between them.

The implication is non-optional. You can no longer evaluate Salesforce's attack surface independently from SAP's. They are not separate surfaces. They are nodes in a single traversable graph. They always were. The difference is that now, something can see and exploit that graph in its entirety -- and the window between discovery and weaponization has collapsed from months to hours.

> ***The era that just ended was not an era of security. It was an era of uninspected risk. Mythos is the inspection. The findings are not recoverable by doing more of the same thing faster.***

IV. The CRC Framework: Consolidate, Reduce, Close

The instinct when faced with a new threat is to harden: add controls, add monitoring, add certifications, add layers. This instinct is wrong under the AI-era threat model. Hardening a general-purpose environment running a sprawling SaaS stack against a machine-speed adversary is an arms race that defenders cannot win. The curve does not bend in the defender's favor.

The correct response is reduction. Remove the surface. Do not harden what is unnecessary -- eliminate it. The CRC framework defines three directives, applied in sequence, each of which removes material that a Mythos-class adversary requires to traverse a regulated stack. The directives are named for what they do: Consolidate, Reduce, Close.

01

CONSOLIDATE

Collapse the attack graph.

Eliminate integration boundaries by consolidating regulated workflows onto a single governed platform.

Every SaaS application eliminated removes a node from the attack graph. Every integration boundary eliminated removes an edge -- and edges are traversal opportunities. The Dunbar Perimeter exists because humans cannot reason about the full graph simultaneously. Consolidation makes the graph small enough that it can be fully owned, fully monitored, and fully defended by a single control surface. The target state is a single governed platform -- one node, zero unowned edges, every data flow producing an immutable attributed audit record. General Reasoning's DXMachine is the reference implementation: regulated workflows that currently span Salesforce, SAP, Workday, ServiceNow, and their integration layers consolidated onto one auditable platform.

02

REDUCE

Eliminate the execution surface.

Minimize the execution environment to exactly what the workflow requires and nothing else.

A general-purpose OS running a regulated application is a general-purpose attack surface. Package managers, shells, unnecessary kernel subsystems, legacy daemons, browsers, USB drivers -- none of these are the application. All of them are chain material. A Mythos-class model that finds a kernel vulnerability in a general-purpose environment has a shell to drop into, a package manager to abuse for persistence, and lateral network paths to follow. On a purpose-built OS, the same vulnerability is a dead end. The bug exists. The chain does not. The Reduce pillar calls for a purpose-built OS image containing exactly the runtime required for the application -- and nothing else. No shell in production. No package manager. Read-only root filesystem. Signed boot chain with TPM attestation. The application is the OS. General Reasoning's infrastructure for DXMachine and Chandra Protocol is the reference implementation.

03

CLOSE

Seal the network.

Implement circular topology with mutual authentication and a single auditable egress point.

A reduced execution environment eliminates the local attack surface. The Close pillar eliminates the remote attack surface. In a closed network topology, each node communicates exclusively with other known, authenticated nodes within the governed boundary. Inbound rules permit only traffic from known peers. Outbound rules permit only traffic to known peers. The only external egress is a single, TLS-pinned endpoint required for AI model inference -- one auditable hole, logged as a first-class audit event on every call. Mythos-class scanning has no entry point: there is no public-facing service to fingerprint, no open port to probe, no lateral network path between nodes that does not traverse a mutual authentication checkpoint.

> ***If Mythos cannot reach it, Mythos cannot chain it. That is the objective. Every CRC control exists to ensure that when a vulnerability is found, there is nothing adjacent to exploit. The chain requires surface area. CRC denies it.***

04

BOUNDARY

Govern every external AI inference endpoint.

Harden and formally attest the boundary where the governed environment contacts external AI inference.

The Consolidate, Reduce, and Close pillars address the internal attack surface. The Boundary pillar addresses the one deliberate opening: external AI inference endpoints. Each endpoint -- whether Anthropic, OpenAI, or any other provider -- is a named, enumerated, formally governed crossing. Certificate pinning on every endpoint. Request signing before any payload leaves the boundary. Schema validation on every response before it touches internal state. AI agent computer-use capability is a distinct threat class: any agent that can operate a keyboard, mouse, browser, or filesystem from outside the governed boundary is not an inference endpoint -- it is an actor. Computer-use agents require explicit GABA attestation separate from inference-only endpoints. The Boundary pillar score reflects how completely each external endpoint is governed, pinned, logged, and formally attested. Chandra Protocol provides the attestation mechanism: every boundary crossing is a context unit. The forthcoming GABA Standard (gabastandard.com) defines the formal AI boundary risk acceptance framework.

V. The Governed Server Boundary Prerequisite

A CRC Minimum Surface Score assessment applies exclusively to deployments that satisfy the Governed Server Boundary (GSB) prerequisite. An organization must affirm all of the following before a CRC assessment is valid.

Requirement	Condition
Organizational control	The infrastructure under assessment runs on servers under formal organizational control. End-user devices, developer workstations, shared consumer environments, and cloud instances with unmanaged provider agents do not satisfy this requirement.
No unmanaged AI agent access	No AI agent with computer-use capability -- the ability to operate a keyboard, mouse, browser, filesystem, or any graphical interface -- has access to the governed boundary from outside a formally authorized, Chandra-recorded endpoint. An AI coding agent running on an unmanaged workstation that can reach the governed boundary invalidates the assessment.
Governed change pipeline	Changes to the deployment require explicit authorization producing an immutable audit record. Ad-hoc changes to the execution environment, network rules, or endpoint registry outside the governed pipeline invalidate the assessment.
Current endpoint registry	All external AI inference endpoints are formally enumerated, certificate-pinned, and current. An endpoint added outside the registry invalidates the Boundary pillar score.

> ***If any GSB condition is not met, the CRC assessment is not applicable -- not zero, not low -- inapplicable. CRC applies to governed server deployments under formal organizational control. Consumer devices, developer workstations, and any environment where AI agents with computer-use capability operate outside a governed boundary are outside CRC scope. The framework makes no claims about and provides no scoring for unmanaged environments.***

VI. The Minimum Surface Score

Each CRC pillar is scored on a 0 to 4 scale across four pillars: Consolidate, Reduce, Close, and Boundary. The total Minimum Surface Score (MSS) ranges from 0 to 16. Regulated deployment certification requires an MSS of 13 or higher. A perfect score of 16 indicates full CRC compliance across all four pillars.

The score is an inventory of what is present, not a checklist of what has been removed. The question for each pillar is not "what have we eliminated?" but "what does our architecture contain?" An organization that has hardened a general-purpose environment by removing services scores lower than one that began with a minimal image, because the former still contains the general-purpose substrate.

Level	Name	Consolidate	Reduce	Close	Boundary
0	**Uninspected**	Full sprawl. Boundaries unowned and uninventoried.	General-purpose OS, full package universe, shell in production.	Public-facing services, unrestricted egress.	Unrestricted external API calls. No logging, no validation, no registry.

Level	Name	Consolidate	Reduce	Close	Boundary
1	**Inventoried**	Boundaries documented, ownership assigned, data flows mapped.	Services inventoried, unnecessary ones disabled, hardening applied.	Firewall rules documented, egress filtered by category.	External endpoints inventoried and documented.
2	**Controlled**	Boundaries owned, monitored, traffic logged, API contracts enforced.	Containerized isolation, no package manager in production.	Ingress restricted to known sources, egress allow-listed by endpoint.	Rate limited, logged, response validated. No computer-use agents.
3	**Reduced**	Majority of regulated workflow on single governed platform.	Purpose-built OS, no shell, read-only root, signed boot chain.	Circular topology, mutual authentication between all nodes.	Certificate pinned, request signed, anomaly detection, circuit breaker active.
4	**Minimal**	Single platform. Zero unowned boundaries. All flows produce audit records.	TPM-attested boot. Immutable filesystem. Application is the OS.	Single auditable egress. Every network crossing a first-class audit event.	All controls active. GABA attestation complete. Residual risks formally documented and signed.

Verdict thresholds:

Score	Verdict	Meaning
0-4	**Critical Exposure**	Stack fully exposed to AI-speed attack chains. GSB prerequisite likely not met.
5-8	**High Exposure**	Some controls exist but a Mythos-class adversary would still find substantial chain material.
9-12	**Moderate Exposure**	Meaningful reduction in place. Critical gaps remain. Below regulated deployment threshold.
13-15	**Reduced Surface**	Meets the CRC regulated deployment threshold. Specific addressable gaps remain.
16	**Minimal Surface**	Full CRC compliance across all four pillars. Architecture denies the chain. GABA attested.

Organizations are encouraged to score each pillar independently and address the lowest-scoring pillar first. The interactive scoring tool is available at crcstandard.com/scoring.

VII. Implementation: The General Reasoning Reference Architecture

General Reasoning, Inc. has designed its DXMachine and Chandra Protocol infrastructure as a reference implementation of the CRC framework, targeting an MSS of 16 across all four pillars.

Consolidate

DXMachine consolidates regulated enterprise workflows onto a single governed platform built on Allegro Common Lisp with AllegroCache persistent object storage. Every action produces a Chandra Protocol context unit: an immutable, attributed, hash-chained audit record that functions as the authorization mechanism for what happens next -- not a receipt for what already occurred.

Reduce

The production execution environment is a purpose-built Linux image built with the Yocto Project, containing Allegro CL, AllegroServe, AllegroCache, and the components required to support them -- nothing else. No shell binary in production. No package manager. No browser. Read-only root filesystem. Signed and TPM-attested boot chain. The application is the OS.

Close

Network topology is circular and self-referential. Each Chandra or DXMachine node communicates exclusively with other known, authenticated nodes in the governed network. The single external egress is the Anthropic Claude API, accessed via a TLS-pinned connection. Every call is logged as a first-class Chandra context unit. There is no public-facing service endpoint. There is no surface to fingerprint or traverse.

> ***The audit record is not a receipt for what happened. It is the gate token for what happens next. This is the Chandra Protocol principle -- and it is what makes the CRC architecture auditable by construction, not by retrofit.***

VIII. The Path Forward

The organizations that will operate safely in the AI era are not those with the largest security teams or the most compliance certifications. They are those whose architectures are hostile to chaining by construction -- whose execution environments contain nothing to chain, whose networks have no surface to traverse, whose workflows are consolidated to the point where the attack graph is too small to exploit.

The Dunbar Perimeter is gone. Any architecture that depended on human cognitive limits to feel secure is now exposed. The response is not incremental hardening of the existing stack. The response is reduction: Consolidate the workflow, Reduce the execution environment, Close the network.

The CRC framework is published as an open standard under the MIT License. Organizations are encouraged to score their stacks, publish their scores, and use the framework as a basis for architecture review and procurement decisions. crcstandard.com.

> ***Mythos raised the bar. The bar was always that high. We simply lacked the instrument to measure it.***