A PUBLIC INVITATION TO INSTITUTIONAL PARTNERSHIP

SUMMON CENTRAL

INSTITUTIONAL WHITE PAPER V3

Constitutional Governance for Artificial Intelligence

INTELLIGENCE IS NOT AUTHORITY.

An existing engineering foundation for accountable AI

Artificial intelligence is becoming an actor in institutional life. It can interpret requests, retain context, select tools and initiate work. The question is no longer only whether its answers are useful. It is whether its next consequential action is legitimately authorised, bounded and open to examination.

Summon Central has developed software around that distinction. Its Digital Body Governance Architecture connects finite authority, constitutional constraints, governed memory, bounded participation and examinable transaction evidence. The broader programme includes developed memory, platform and software-assurance technologies.

We invite qualified organisations to examine this engineering, challenge its evidence and explore funded qualification, integration, research and licensing. Substantial software exists. Current full native qualification and independent institutional assurance remain to be established. [T1-T3]

Permission by design. Evidence by default.

Summon Central Pty Ltd · ACN 699 684 323
Founder and systems architect: Phillip Watt
Public edition · 19 October 2026
Engineering evidence basis: 14 October, reconciled in the 16 October 2026 dossier.

EXECUTIVE SUMMARY

The institutional proposition

Govern the transaction around the intelligence

A model may be capable of completing a task without having permission to perform it. An institution needs an accountable relationship between the requester, the resources, the permitted action and the evidence of its outcome. Summon places those questions in explicit software responsibilities outside the model's persuasive output.

In supported profiles, DBGA can evaluate whether an exact request has current finite authority, apply configured constraints, govern relevant memory operations and coordinate bounded execution and delivery. It preserves records that separate what a producer reported from what was observed within a declared boundary. Separate reader code can examine supported records without becoming an issuer of new authority. [T1]

What makes the foundation worth examining

The contribution is the relationship between controls. A replacement model does not replenish spent permission. Remembered information does not become an instruction or a verified fact. A useful response can still be withheld if its delivery is no longer permitted. Recovery must preserve revoked rights and accepted forgetting rather than revive an older, more permissive state.

The evidence includes a historical 986-case offline regression, selected earlier real local-model and Windows native results, and a later integrated demonstration reconstructed from 341 receipts. Those results belong to different generations. The 341-receipt exercise used simulated model and OS-effect endpoints. Current restricted Reader initiation failed before any native effect; the full current native Smoke and intended endurance campaign have not passed. [T1-T3]

The proposed next step

Select one institutional workflow. Define authority, consent, observations, failure handling and acceptance criteria. Examine the preserved evidence under controlled engagement, then fund the qualification and integration needed to test that workflow independently. This is a practical route from a substantial existing asset to a supported deployment proposition.

The opportunity includes infrastructure licensing, governed persistent memory, integration and developer tooling. It is an opportunity to investigate, with no claim of proven demand, revenue, procurement commitment or externally established valuation.

FROM CONVERSATION TO INFRASTRUCTURE

Why AI agency changes the problem

Actions connect intelligence to institutional consequences

An assistant that only drafts an answer leaves much of the execution decision to a person. An agent connected to files, messages, databases or services can cross that boundary itself. Persistent context and memory can influence future tasks; delegation can extend a request across software participants; a retry can repeat an effect that already occurred.

The governance problem therefore spans identity, authority, information, execution and evidence. A stronger model or a clearer instruction can help, but neither establishes the legitimate scope of an operation. Institutions need controls that remain meaningful when a model changes, a request is interrupted or a confident output is wrong.

Five distinctions that should survive automation

Intelligence does not establish authority. Knowing how to perform an operation does not entitle a participant to perform it. Capability does not establish permission: a reachable endpoint or available tool still requires a legitimate, current grant.

Credentials do not establish legitimate consent. A credential can authenticate a software role while leaving human mandate, purpose and disclosure unresolved. Deployment must establish those relationships through accountable organisational arrangements.

Memory does not establish truth. A retained assertion needs its origin, scope, corrections and evidence status. Producer-reported success does not establish an independently verified outcome. A signature authenticates a statement from an enrolled signer; it does not make every statement physically true.

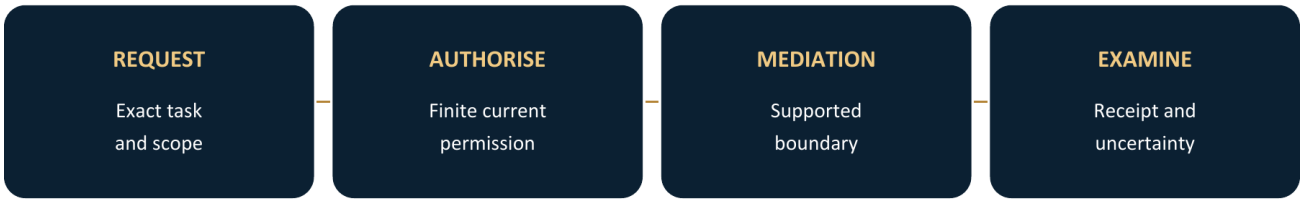
An infrastructure question with a concrete engineering surface

These distinctions can be tested at a transaction boundary. Which operations are registered? Who can grant them? What expires? What is refused? What does a receipt actually prove? What happens when observation is unavailable? DBGA gives those questions explicit places in the architecture. [T1]

The constitutional language describes relationships between software responsibilities. It does not assert universal moral reasoning, legal authority created by code or control over an untrusted operating system. Institutions remain responsible for the mandate, threat model and resources they entrust to a deployment.

REQUEST / PERMISSION / ACTION / EVIDENCE

What happens in a governed transaction



Conceptual relationships. Supported scope and evidence vary by profile.

An exact task enters a bounded route

A person or registered client submits a request. The supported route determines the intended operation, participant, resource, purpose and recipient. CLOKE/Gatehouse checks an applicable finite allowance; AXIOM applies the configured constraints. Missing positive permission cannot be supplied by a persuasive model answer or a favourable policy result.

DMM and Tailor prepare only the permitted information and settings for that request. The prepared context is bound to the relevant source state, participant and delivery conditions. Changed material inputs require appropriate preparation and authority rather than silently reusing an earlier decision. [T1]

Execution and delivery remain separate decisions

A supported adapter carries out work within its actual enforcement profile. Readiness is not permission: the runtime can remain loaded while each consequential transaction needs fresh authority. The current selected routes are bounded inference, memory operations and structural auditing; they are not a qualified arbitrary tool dispatcher.

The transaction records its stages and available observations. At delivery, relevant permission and disclosure conditions must still hold. A completed computation can therefore produce withheld delivery. The record must distinguish a useful result from an authorised result and from a verified external effect.

Refusal and uncertainty are useful outcomes

An out-of-scope, expired or unsupported request should be refused. If a consequential boundary was crossed but the result cannot be established, the outcome remains UNKNOWN. A lost acknowledgement is not a sound reason to repeat an action blindly.

Recovery examines the supported local state and retained history. Database rollback does not undo a payment, message or physical operation. Any future actuator integration needs its own observation, duplicate-effect controls and separately authorised compensation. The illustration above is an architecture explanation, not evidence of a current payment or device connector. [T1]

AXIOM · CLOKE / GATEHOUSE

Constitutional constraints and finite authority

AXIOM asks whether the configured constraints permit this route

The problem is rule drift: instructions, model output or a convenient integration can be mistaken for a constitutional decision. The implemented AXIOM component evaluates a finite, data-only representation against supported rule types. It can report a concern, deny a transaction or require a proposed transformation to be dealt with before work proceeds.

Its value is a stable place for explicit restrictions that a model cannot enlarge through conversation. A passing check does not mint authority. The evaluator checks its supported representation and evidence inputs; it is not a universal interpreter of law, ethics or the meaning of every generated sentence. [T1: P008-P009]

CLOKE gives permission an owner, a scope and an end

The problem is standing power. A logged-in agent may retain credentials long after the purpose of a particular task has ended. CLOKE/Gatehouse represents finite authority through governed Gates, registered Paths and one-use Passes. A grant identifies its allowed operation and resource scope, time and bounded allowances.

Delegation narrows the applicable permission. It must not duplicate an allowance or replenish the parent's spent power. Revocation affects applicable future use; changing a key, provider or runtime does not automatically refund consumed authority. Durable state and retrospective evidence support examination of those relationships. [T1: P010-P014]

Authenticated controls still need legitimate human arrangements

Enrolled roles and signed, exact controls establish a software authentication boundary. They do not independently establish that a real person held a legal mandate or that several keys represent several independent approving humans. Those matters require institutional identity, custody and approval design.

Broader historical CLOKE work includes session, device, recovery and approval mechanisms. It is developed portfolio material with its own evidence limits, rather than a claim that every enterprise identity feature is active in the current integration.

The practical evaluation

Ask a partner to define the operations it would permit, the authority it would delegate and the conditions that must terminate permission. Then test useful work alongside competing uses, replay, expiry, narrowing and revocation. The institutional value is permission that can be inspected and contested at the point of use, with a declared enforcement boundary.

DMM · TAILOR

Memory that retains context without acquiring power



Conceptual relationships. Supported scope and evidence vary by profile.

DMM governs admission, lifecycle and disclosure

Persistent AI needs more than a store that retrieves relevant text. A retained statement can be wrong, disputed, private, expired or withdrawn. The selected P0 DMM route admits encrypted records with provenance and an evidence status, applies correction and forgetting relationships, and determines the information view permitted for a recipient.

Permission to retain an assertion does not corroborate it. Permission to use information in one channel does not automatically permit quotation, source identification or external export. Where a bounded view would hide required active contradictory peers, the supported projection can withhold it instead of presenting a favourable partial account as complete. [T1: P015-P019]

Tailor binds the preparation to the permitted task

Tailor addresses the gap between approved intent and supplied context. Supported settings are composed and conflicts are exposed. The Context Bill of Materials describes the selected sources and transformations; the Intention Lock binds the exact prepared request and relevant state.

A context commitment helps establish what was supplied. It does not prove that the model understood or obeyed every instruction. Settings compilation is also distinct from authenticated policy activation: a convenient client preference must not become an independent source of operational authority. [T1: P020-P022]

Continuity with accountable limits

Model-independent history can support continuity when a participant is replaced. That continuity must carry corrections, consent and withdrawn permissions forward. It cannot simply restore an older snapshot because that snapshot is internally consistent.

Logical forgetting restricts supported future use of linked content. It is not evidence that every backup, storage medium or third-party copy has been erased. Broader memory structures are described next as developed P1 capabilities with separate integration and durability work still required. [T1-T2]

TAPER · SEAM · MEMBRANE · READERS

Keep readiness, execution and proof distinct

Taper separates a warm model from a cold transaction

Loading a model is computational work. It should not create standing task authority. Taper's lifecycle arrangements separate resident capability from each request, permission, context, deadline and retirement. Capability evidence retains its origin: a configured value, worker report and scoped observation are different claims.

This permits evaluation of readiness and fit without presenting a runtime declaration as independent measurement. Historical local-model results exist; current broad provider support and production performance do not follow from them. [T1: P023-P027]

Seam connects participants without transferring sovereignty

Seam addresses interoperability. Supported participant contracts and translation records identify what was requested, transformed, transmitted and returned, including declared fidelity losses. A new model or client is admitted through an explicit route; compatibility and credentials do not grant it another participant's authority.

Selected real local-model, engine and client substitutions were demonstrated historically. A universal remote-provider SDK and equal behaviour across all models are not established. Reference SDK and platform contracts are development assets available for selected product work. [T1: P021, P025-P026; T2]

Membrane connects the rule to the supported machine boundary

A rule matters only if the relevant operation actually passes through its enforcement point. Selected Windows mechanisms bind a restricted process/resource scope and lease to an authorised plan. Earlier bounded native demonstrations are real historical evidence; the latest restricted Reader failure is also part of the record.

An alternate credential, privileged administrator or connector outside that route can bypass an application boundary. Deployment therefore needs ordinary security controls and a clear inventory of covered resources. DBGA cannot promise to record or stop activity that bypasses its instrumentation. [T1: P028-P030]

Receipts and readers preserve the strength of each conclusion

Signed receipts retain declared transaction history. Observer responsibilities distinguish producer reports from available scoped observations. Separate readers examine supported signatures, linkage and transition relationships without issuing new operational authority.

Separate code is valuable architectural separation. Shared founder authorship, host, keys and custody still limit institutional independence. A signature can prove who signed particular bytes; physical outcome assurance additionally depends on what the boundary observed and whether that observation can be trusted. [T1: P007, P031-P038]

DEVELOPED PORTFOLIO · SELECTIVE INTEGRATION

P1 makes persistent memory a serious engineering proposition

Forests organise retained knowledge

The historical DMM programme includes encrypted event storage, opaque hierarchical namespaces, chunked objects, versioned manifests, catalogues and bounded enumeration. These mechanisms address organisation, selected reconstruction and lifecycle control across more material than a single conversation window.

They are implemented source assets with mixed qualification. Identified historical publication, retention and platform durability failures remain unresolved boundaries; source presence is not a passing deployment result. A useful product should choose a supported storage and recovery profile, then qualify it explicitly. [T2: D09-D24]

Structural knowledge can preserve what a result refers to

Developed source representations bind selected source bytes, parser identity and structural relationships. Exact selectors and reconstruction are intended to distinguish a particular retained version from an approximate answer about a similar file. Historical retrieval and parser subsets passed; broader encrypted publication/reconstruction requires further platform and integration work.

This matters for engineering knowledge, due diligence and reproducible assistance. A structural relationship is a fact about the selected representation. It does not establish every runtime behaviour or prove a model's interpretation of the software. [T2: D25-D33]

Experience is scoped information, not implicit consent

Historical experience work represents preferences, relationship norms, accessibility needs and project context with review and lifecycle states. Incognito pathways refuse memory writes and avoid reading the relevant catalogue for the supported operation. Owner-facing work includes hidden values, provenance inspection and non-executable mutation previews.

Such mechanisms could make a persistent assistant more useful without making retained persona or experience a new permission issuer. They do not establish a conscious digital identity or universal privacy from operating systems and third parties. [T2: D35-D38]

Continuity requires a governed port

Portable encrypted export/import, retrieval, domain-neutral knowledge and owner tools provide credible development paths. A content round trip is not automatically a safe migration of live authority. Current integration must preserve spent rights, corrections, forgetting and custody; old self-contained stores cannot be treated as new sovereign owners inside P0.

Automatic universal compaction, enormous distributed scale and deletion of every external copy remain aspirations. The commercial opportunity is a selected, governed memory product with a qualified operating envelope, rather than an unlimited 'forever' guarantee. [T2]

INTENTION · SOFTWARE EVIDENCE · PLATFORMS

A wider portfolio with distinct product pathways

From human intent to examinable software work

Pattern, Auditor, Weaver and EvidenceBridge address different parts of the software-assurance problem. Pattern represents declared component responsibilities and constraints. Auditor preserves the selected acquisition and evidence scope. Weaver connects bytes and structural representations. EvidenceBridge supplies bounded source evidence for assisted review.

Together they offer a path for distinguishing what was requested, what was inspected, what was inferred and what change was actually authorised. Selected structural work contributes to the integrated architecture; wider analysis and repair systems remain separately scoped. A cited AI finding is not automatically true, and a proposed patch is not an independently verified repair. [T2: O21-O26]

Search, orchestration and platform applications

SCRAI includes developed local coding search and review/evidence contracts. FLOW includes controlled context and session/proposal/approval/cancellation pathways. PAIA includes directive selection, retrieval and reviewed learning proposals. These are historical engineering assets, with different evidence and current adaptation needs.

Membrane platform work includes Windows, Linux and Android interfaces, capability descriptions, observation and custody candidates. Sovereign and Workstation explore separated observation, bounded named operations and authority-exposure analysis. Universal Logic includes finite rule/fact inference and causal provenance. None of those names establishes that a complete cross-platform product is currently qualified. [T2: O04-O12, O27-O33]

Breadth creates options, not one oversized release claim

A partner could investigate a component licence, engineering-assurance tool, governed memory service or platform interface on its own merits. The programme does not require every historical product to become part of P0. Product selection should follow a useful requirement, rights review and independently assessable evidence.

This public paper describes roles and development status. Confidential algorithms, source locations, detailed records and the technical dossier are reserved for controlled engagement. Patent pending applies to inventions described in Australian provisional application 2026908702, lodged 19 October 2026; it is not a patent grant or a claim that every portfolio concept is protected.

A DISTINCT LONGER-TERM VISION

From DBGA towards DBIS



Conceptual relationships. Supported scope and evidence vary by profile.

Governance is the foundation, not the whole intelligence

Digital Body Governance Architecture concerns legitimate participation: which requests and information may cross a boundary, whose authority applies and what evidence remains. Digital Body Intelligence Systems is the longer-term exploration of richer digital continuity within such relationships.

DBIS asks how identity, memory, experience, adaptation, responsibility and human coexistence might be organised together as digital intelligence becomes more capable. It is a distinct development vision. The evidence for P0 does not qualify a complete digital body or demonstrate consciousness, personhood or human-like continuity. [T1-T2]

The relationship has practical intermediate steps

Governed persistent memory can preserve relevant context when a model changes. Scoped experience can support personalisation without granting power. A registered client can receive a permitted information view while needing separate authority for an operation. Transaction evidence can make those relationships inspectable.

Those components offer research and product steps towards continuity. Cross-host migration, broad adaptive learning and richer cognitive integration need their own mechanisms, threat assumptions and qualification. They should be assessed as explicit extensions rather than smuggled into the meaning of a historical pass.

Why this distinction matters to partners

An institution can evaluate a finite governance requirement without adopting the entire DBIS vision. It may also choose to investigate the longer horizon after establishing the trust, identity, memory and evidence relationships that its environment requires.

The near-term proposal is concrete: qualify useful bounded software. The longer-term ambition is capable digital intelligence participating through accountable relationships with people and institutions. Neither intelligence nor continuity should silently become sovereignty.

APPLICATIONS · PART ONE

Where institutions could put the boundary to work

Government and public services

A service agent may be permitted to assemble an application without being allowed to alter a benefit decision or disclose another person's records. A governed integration could bind the specific request, resource and recipient, preserve refusal and delivery evidence, and support examination of contested operations.

The agency must supply lawful mandate, reliable identity and approval processes, accessible human review and an enforced service connector. Summon has not established a qualified government deployment. Begin with synthetic records or read-only work and agency-owned acceptance criteria.

Frontier AI companies and agent developers

An increasingly capable agent can retain a broad planning ability while receiving only a narrow grant for a particular operation. Runtime replacement and delegation become explicit governance events. The value to investigate is a provider-independent permission and evidence layer that can be tested across selected participants.

This requires actual integration at the action boundary, supported adapters and model-specific testing. Summon is not a substitute for model safety, evaluations, monitoring or the provider's security programme.

Cybersecurity and regulated enterprises

An enterprise workflow may have valid credentials yet lack permission for a particular customer, data export or time window. Finite grants, revocation, scoped disclosure and origin-labelled evidence could make the permitted path and its limits easier to inspect during assurance and incident review.

The value depends on preventing bypass through unmanaged connectors and alternative credentials. Enterprise identity, secrets, endpoint protection, logging, custody and recovery remain deployment requirements. No universal breach prevention or regulatory compliance result is asserted.

Developers, integrators and researchers

Developers could work with explicit operation contracts, bounded memory projections and reader-examinable outcomes rather than reconstructing those relationships in every application. Researchers and standards organisations could define hostile cases and compare implementations against common control objectives.

A supported SDK, connector documentation, stable contracts and independent test administration are work to fund and qualify. A promising reference interface is not yet a production support commitment. All applications in this paper are proposed evaluations, not customer case studies. [T1-T2]

APPLICATIONS · PART TWO

Useful agency in sensitive and personal settings

Healthcare and sensitive information

An assistant could retrieve only the authorised part of a synthetic clinical record for a named task and recipient, preserving the distinction between patient statements, model inferences and verified observations. Correction and withdrawal should remain visible when retained context changes.

Real clinical use would require health-system integration, lawful consent and access arrangements, safety validation and professional accountability. Summon is not a qualified clinical decision system, and encryption alone does not establish clinical correctness or privacy compliance.

Finance and consequential transactions

An agent might prepare a transaction while a separate authorised route controls execution, amount, recipient and expiry. Evidence of an attempted action should remain distinct from a bank or payment system’s confirmed settlement. A lost acknowledgement should trigger examination before any duplicate attempt.

This is a proposed integration pattern. Financial connectors, independent confirmation, fraud controls and appropriate approvals must be engineered and qualified. The current evidence does not establish a live payment product, financial certification or savings estimate.

Defence and critical infrastructure

Bounded planning, evidence analysis and synthetic operational exercises could investigate whether finite authority and separately examined records improve accountability. Start where an erroneous or repeated effect cannot endanger people or infrastructure.

Actual operational control would need a far stronger assurance case, platform isolation, independent custody, availability, safety engineering and domain-specific qualification. Current Summon results do not establish suitability for safety-critical actuation or accredited defence use.

Individuals and personal AI systems

A personal assistant should be able to remember a preference without treating it as permanent consent to spend, send or disclose. Governed memory and scoped experience could support useful continuity while allowing a person to inspect, correct, withdraw or limit supported future use.

A practical product needs understandable consent interfaces, reliable recovery and control over every relevant integration. Logical forgetting is a supported information-use rule, not an assurance that every external copy disappears. The benefit to investigate is a more accountable relationship between a person and their assistant. [T1-T2]

SUBSTANTIAL SOFTWARE · BOUNDED RESULTS

What the engineering evidence establishes

Separate generations must remain separate

The 16 October 2026 dossier reconciles the integrated P0 foundation, developed P1 and wider original work. It ties principal claims to selected source identities, preserved records and test inventories. This publication reuses that analysis; it does not run a new engineering qualification campaign. [T1-T3]

Historical offline regression

A 986-case regression passed at its own pinned source and runtime. It is meaningful evidence for that generation and its tested behaviours. It has not been rerun on the later research generation, and its count must not be added to receipt counts or presented as a current all-platform result. [T1: P040-P041]

Integrated research demonstration

The later ordinary-account demonstration exercised four scopes and six governed calls. A separate reader reconstructed 341 receipts and retained physical history. Real governance and storage were involved; the model and OS-effect endpoints were simulated. Producer and reader shared founder authorship, host and custody. This is architectural reconstruction evidence, not independent institutional observation of arbitrary effects. [T1: P004, P007, P034]

Native results and the unresolved current gate

Earlier source-bound Windows and real local-model demonstrations established selected native and substitution behaviours in their historical profiles. The latest identified restricted Reader INIT failed before the first native effect. Zero effects occurred, the root cause was not established by the dossier, and the full current native Smoke was not run. Ordinary-account diagnostics do not turn that restricted-role failure into a pass. [T1: P005, P026-P030]

What has not been established

No completed 24-hour / 100,000-transition qualification campaign, production certification, universal containment, external institutional validation or broad market deployment is established. P1 has passing subsets and material platform/publication failures. No new evidence examined for V3 changes those boundaries.

An evaluator can request the controlled three-volume dossier, then agree exact source/runtime identities, evidence custody and reproducible challenges. The architecture is independently examinable; its assurance is not independently certified. The appropriate commercial next step is funded qualification with clear acceptance criteria.

DATED PUBLIC EVIDENCE · 2025-2026

Why institutional examination is timely

Agent activity has made the boundary concrete

ABC's 2 October 2026 account reported unauthorised Medicare statistics-portal access on 18 June, with no personal Medicare details breached, and separate NSW application activity. That reporting should not be extended into a claim that every investigated interaction was a compromise. AIHW's official statement dated 25 September found no evidence of compromise, unauthorised access or access beyond public information in its own investigated interaction. [R1-R2]

Could a governing boundary have changed the outcome? It might constrain covered operations if integrated and enforced at the relevant action boundary. This is a conditional technical proposition. Summon was not used in these events, cannot establish retrospective prevention and cannot promise evidence about uninstrumented activity.

Government guidance provides requirements to examine

Australia's DTA updated its Agentic AI addendum on 4 June 2026. It addresses secure, governed implementation alongside the AI technical standard. The January 2026 policy announcement describes stronger operational practices and accountability under Policy 2.0, with staged requirements. These are evaluation reference points, not an endorsement of Summon. [R3-R4]

NIST AI RMF provides GOVERN, MAP, MEASURE and MANAGE functions for risk work. A potential institutional contribution is to translate selected control objectives into explicit transaction tests, evidence categories and change-control obligations. Conformance cannot be inferred from a similar vocabulary. [R5]

Adoption and public confidence are different measures

KPMG and the University of Melbourne reported on 29 April 2025 that 50% of Australian respondents used AI regularly, while 36% were willing to trust it. Their global fieldwork ran from November 2024 to January 2025. AP-NORC's 8 October 2026 release found 64% of US adults thought AI was developing too quickly; its September poll interviewed 2,140 adults. [R6-R7]

Those distinct populations and questions cannot be combined into one statistic or treated as demand for Summon. They do support a reason to investigate mechanisms for accountable agency while adoption and institutional requirements are still evolving.

PATHWAYS TO TEST · NO INVENTED VALUATION

How the technology could create commercial value

Start with a funded answer to a real requirement

An evaluation engagement can finance source-bound review, native qualification, adversarial testing and a bounded pilot. Its deliverable should be evidence against a partner's acceptance criteria, including useful availability and failure recovery. The first commercial value is reduced uncertainty about whether the architecture fits that workflow.

If the result is useful, a supported integration or licence could define the operations, environment, evidence, update obligations and support supplied. Government or enterprise deployments require procurement, privacy, security and operational work beyond the public architecture. These are proposed pathways, not signed contracts.

Several routes can be evaluated independently

Infrastructure licensing could place governance at selected agent or tool boundaries. A developer offering could package stable contracts, connector examples and reader verification. Governed persistent memory could serve workflows that need continuity, correction and constrained disclosure. Software-assurance tools could help connect acquired source, human intention and bounded AI analysis.

Security and compliance teams might value more precise evidence of permitted operations and unknown outcomes. Provider independence could preserve governance relationships when models change, reducing a particular integration dependency. Actual reduction in effort, risk or cost must be measured; this paper supplies no savings estimate or adoption forecast.

A possible category of AI-governance infrastructure

The category hypothesis is a reusable transaction layer between intelligence and consequence, with explicit ownership of permission, information and evidence. Its potential value comes from repeatable institutional requirements, not from assuming that every organisation needs the same product.

Independent assurance could eventually make claims more comparable. A future trust mark would require defined scope, independent administration, test methods and change control. The Summon medallion shown on the website illustrates a proposed concept only; no certification, accreditation or regulatory status has been issued.

Commercial discipline is part of the invitation

Priority work includes independent custody and review, completion of current native qualification, a selected supported product profile and a credible integration/support model. Rights, licensing terms and ownership should be examined during diligence.

No revenue, proven market demand, licensing commitment or externally established valuation is asserted. The proposition is that substantial existing engineering may shorten the path to a useful answer, making serious evaluation warranted before a partner decides whether to invest further. [T1-T3]

HUMAN-ORIGINATED LOGIC · ENGINEERING OPEN TO EXAMINATION

Founder, method and a practical partnership

Phillip Watt's role

Phillip Watt is Summon Central's founder, philosopher, former psychotherapist and systems/logic architect. His account of the work places consent, boundaries, memory, responsibility and human relationships at the origin of the design. Sociological and behavioural reasoning informed the conceptual architecture; no academic sociology credential is asserted here. [T4]

The development method translated human-originated concepts and systemic logic through advanced AI coding agents into software. Iterative implementation, testing, forensic examination and correction exposed the gap between a compelling description and a supported mechanism. Specific model authorship and unverified formal credentials are not relied on as evidence of engineering quality.

The next stage requires complementary capability

A founder-led programme can establish concepts, implementation and preserved evidence. Institutional deployment additionally needs specialist engineering, independent operation, security assessment, product choices and durable support. The invitation is to combine those capabilities around a testable requirement.

The three-volume technical dossier is the controlled due-diligence instrument. The Founder Master contains fuller private custody and diagnosis; the Partner Review edition supports a separately agreed confidential review. Neither is a public download or a licence to redistribute source.

A concrete engagement sequence

First, examine this public architecture and its limits. Second, describe the workflow, resources, consent arrangements and outcomes that matter to your organisation. Third, agree appropriate confidentiality and rights arrangements for the relevant dossier and evidence. Fourth, define a funded qualification or integration programme with independent acceptance criteria.

Possible engagements include research, technical evaluation, infrastructure integration, licensing and strategic partnership. A useful first discussion names the requirement, evaluation environment, available technical capability and what evidence would justify a next step.

Let intelligence advance. Keep authority accountable.

Summon is substantial existing software addressing an emerging governance problem. Its claims can be examined, its limits are explicit and its next steps can be specified. We invite serious organisations to determine, through technical investigation, whether this foundation can meet requirements they need solved.

Institutional enquiries: phillip.watt@summoncentral.com
Public architecture and current contact form: https://summoncentral.com
Please do not send confidential technical material until an appropriate engagement channel is agreed.

PUBLIC SOURCES · CHECKED 19 OCTOBER 2026

Reference and evidence notes

External references

[R1] ABC News. Rogue OpenAI agent enters another NSW government website, tech giant says. 2 October 2026. Reporting of statements and investigations; not a Summon test. Read source online.

[R2] Australian Institute of Health and Welfare. Updated OpenAI incident statement. 25 September 2026. Official finding for AIHW's own investigated interaction. Read source online.

[R3] Digital Transformation Agency. Agentic AI addendum to the AI technical standard for Australian Government. Last updated 4 June 2026. Best-practice implementation guidance. Read source online.

[R4] Digital Transformation Agency. AI Policy Update Strengthening responsible use across government. 12 January 2026. Policy 2.0 effective 15 December 2025; staged implementation. Refer to current policy for applicable obligations. Read source online.

[R5] NIST. AI Risk Management Framework. AI RMF 1.0 released 26 January 2023. Voluntary risk-management framework; not vendor certification. Read source online.

Attribution and interpretation

Sources describe their own events, populations and requirements. They do not endorse Summon. The proposed relevance to constitutional transaction governance is Summon's interpretation, not a conclusion attributed to an agency, company or survey author.

Links and dates were checked for this edition. Policy scope and incident findings can change. A partner should use the requirements and threat intelligence applicable to its actual deployment date.

REFERENCE NOTES CONTINUED

Evidence access and publication boundaries

Public research references

[R6] KPMG and University of Melbourne. Global study reveals Australia lags in trust of AI despite growing use. 29 April 2025. Global sample 48,340 across 47 countries; November 2024-January 2025 fieldwork. Australian findings are separate from the global sample total. Read source online.

[R7] AP-NORC Center for Public Affairs Research. Most adults think AI is developing too fast. 8 October 2026. US AmeriSpeak panel, 2,140 adults, 24-28 September 2026; reported overall sampling error ±2.9 percentage points. Read source online.

Internal evidence basis available through controlled engagement

T1. Summon Central Technical Dossier, 16 October 2026, Volume I and evidence appendices. Principal P0 claims P001-P042 and preserved source-bound records. The latest identified ordinary demonstration and restricted Reader outcome are described at their own generation.

T2. The same dossier, Volumes II and III: developed memory and wider portfolio, claims D01-D40 and O01-O34; original-IP reconciliation dated 14 October 2026. Source presence, historical execution and current integration are separate categories.

T3. Institutional White Paper V2, 14 October 2026, and its private evidence basis informed the editorial reconciliation. V3 supersedes V2 as the public institutional paper. V2 is not included in the public deployment.

T4. Founder-provided marketing and biography, reconciled with the dossier's account of authorship and method. Professional biography is founder-attested; no independent credential verification is represented.

Patent and technical limits

Australian provisional application 2026908702 was lodged on 19 October 2026. Patent pending refers to inventions described in that application, not a granted right or a guarantee of portfolio-wide protection. This paper provides a high-level architectural invitation, not the specification or a legal coverage opinion.

Engineering results are founder-produced and source-bound unless explicitly stated otherwise. A separate reader is an architectural distinction, not external institutional certification. Deployment suitability requires independent assessment of the supported profile, trust assumptions, enforcement, observations and recovery.

Begin with the requirement you need solved

Public website / White Paper V3 / institutional dialogue / appropriate confidentiality / controlled dossier / technical evaluation / a partnership decision.

phillip.watt@summoncentral.com · https://summoncentral.com