14 September 2026

Chair
Joint Select Committee on Artificial Intelligence
PO Box 6100
Parliament House
Canberra ACT 2600

Subject: Submission to the Joint Select Committee on Artificial Intelligence — *A National Trust Anchor for AI Agents*

Dear Chair

I welcome the opportunity to make a submission to the Joint Select Committee on Artificial Intelligence.

In a future with millions or even billions of AI agents operating online, people and machines will need a safe and reliable way to know whether an agent is legitimate and who it represents. Today, a user visiting **health.gov.au** can have confidence that they are accessing an official Australian Government health service because the **health.gov.au** domain name provides a recognised and trusted link to that organisation. As governments and businesses increasingly deploy AI agents to interact with people and other machines, we will need a similar way to verify that an AI agent claiming to act on behalf of an organisation is genuinely authorised to do so. The Domain Name System (DNS) could provide that trust layer by linking AI agents to verified domain names and the organisations behind them.

I attach my whitepaper, *A National Trust Anchor for AI Agents*, for consideration by the Committee. The paper addresses an emerging aspect of artificial intelligence that I believe warrants consideration as Australia develops its national approach to AI: **the identity, accountability and trust infrastructure required as AI systems evolve** from tools that generate information into autonomous agents capable of acting across digital services and organisational boundaries. The paper is particularly relevant to paragraphs (k), (l) and (m) of the Committee's Terms of Reference: risks and harms arising from AI; national security and cyber security opportunities and risks, including the resilience of critical infrastructure; and the implications of emerging AI capabilities for Australia's national security and strategic resilience.

AI agents represent an important change in the nature of the internet. Increasingly, agents will be capable of discovering services, invoking tools, exchanging information,

initiating transactions, making commitments and interacting with other agents with varying degrees of autonomy. As this occurs, many interactions that today involve human judgement will instead occur machine-to-machine and at machine speed.

This creates an emerging trust problem. Before accepting a consequential action from an AI agent, another system may need to establish independently what agent is acting, which person or organisation stands behind it, whether that relationship remains current, and where the authority claimed by the agent can be verified. These questions become particularly important in relation to fraud and impersonation, cyber security, government services, critical infrastructure, supply chains and national security.

The attached paper argues that accountability for AI agents should not depend solely on proprietary identity systems operated by the platforms on which those agents happen to run. Australia should consider establishing a platform-neutral, sovereign trust anchor through which an agent's durable identity and accountable sponsor can be independently resolved, while leaving authentication, credentials, authorisation and operational controls to the specialist systems designed to provide them.

The paper proposes that Australia's existing national DNS infrastructure warrants serious consideration as the foundation for such a trust model. The DNS already provides globally resolvable, hierarchical delegation under national stewardship, while DNS Security Extensions provide mechanisms for authenticating published information. The proposal is not that DNS should determine whether an agent is safe or authorised to undertake a particular transaction, but that it could provide the authoritative public starting point from which an agent's accountability and associated verification information can be established.

This also presents an opportunity for sovereign capability. Australia need not develop every foundation model or AI platform domestically to retain sovereign control over important elements of the trust infrastructure through which autonomous agents interact with Australian organisations, citizens, government and critical systems.

The agent ecosystem is still developing, and the paper does not suggest that Australia should prematurely prescribe a particular technical implementation. Rather, I submit that Australia should recognise machine-verifiable agent identity and accountability as an emerging component of national digital infrastructure, and begin the technical, policy and multistakeholder work required to determine how that capability should be provided.

In particular, I invite the Committee to consider recommending that the Australian Government:

1. Recognise trusted, independently verifiable identity and accountability for AI agents as an emerging national cyber security and digital infrastructure requirement;
2. Convene government, cyber security, digital identity, AI safety, critical infrastructure, privacy, competition and Internet-governance stakeholders to examine requirements for a sovereign, platform-neutral trust anchor for AI agents; and
3. Support controlled technical trials and international standards engagement to test interoperable approaches before proprietary agent identity systems become entrenched.

These measures would be proportionate to the maturity of the technology. They would not require Australia to determine the final architecture today but would ensure that the country is actively shaping the trust infrastructure on which safe agentic AI may increasingly depend.

Thank you for considering this submission. auDA would welcome the opportunity to discuss these findings and provide further information. Please contact me at [redacted]

Yours sincerely

Glenn Wightwick
Chief Innovation Officer
.au Domain Administration (auDA)

NATIONAL DIGITAL INFRASTRUCTURE

A national trust anchor for AI agents

Why the country code Top Level Domain (ccTLD) registry should operate the national agent namespace

White paper for government, industry, critical infrastructure and the technology sector

Glenn Wightwick, Chief Innovation Officer, .au Domain Administration (auDA)

September 2026

.auDA

Contents

The agentic internet requires public trust infrastructure that allows any machine to resolve an artificial intelligence (AI) agent to an accountable person or recognised legal or organisational entity without depending on the platform that created or operates it. Every consequential agent should carry a durable, independently resolvable identity. The Domain Name System (DNS) is uniquely positioned as the universally deployed, hierarchical internet namespace with native global resolution and established delegation at national level. The national DNS operator is therefore a strong candidate to serve as the national trust anchor for the agentic internet, subject to an explicit domestic mandate and multi-stakeholder governance.

Introduction: from the human web to the agentic internet

The internet we know was designed around people. Websites present information through browsers; apps deliver functionality on mobile devices. A person visits a destination, recognises a brand, reads information, compares options and authorises an action. Domain names sit at the foundation of this model: they provide the stable, globally unique names through which people and machines locate organisations and services across the internet.

AI agents change the identity of the internet's active participant. An AI agent is software that can pursue an objective on behalf of a person or organisation: it interprets context, plans a sequence of steps, discovers services, invokes tools, exchanges information and takes action with varying degrees of autonomy. Unlike a chatbot that primarily returns an answer, an agent can complete work. A personal agent may research and purchase travel, manage finances or coordinate health appointments. A business agent may negotiate with suppliers, monitor infrastructure, transact with customers, operate software and delegate tasks to other agents.

In this model, the human-facing website or app becomes less relevant. People will express an intent to a trusted personal or organisational agent. That agent will discover the relevant providers, compare machine-readable offers, call services directly, negotiate permissions or price, complete a transaction and report the outcome. Websites and apps may not disappear as underlying systems, but their content, navigation and function will be far less visible. The primary interaction will be agent-to-service and, increasingly, agent-to-agent.

THE INFRASTRUCTURE INVERSION: Websites become endpoints. Agents become the users. Identity becomes the interface.

This transition is already measurable. Cloudflare reported that in May 2026 non-human traffic passed human traffic across its network and they predict that if current trends continued, non-human traffic could be as much as 1,000 times human traffic within five years – making humans "a rounding error on the Internet", not because human use declines, but because machine activity grows so quickly.[1,2]

The forecast is not a certainty, and traffic volume is not the same as economic or social value. Its direction, however, is difficult to dispute. Agents can inspect thousands of resources in the time a person visits a handful; they can operate continuously; and they can coordinate, retry and delegate at machine speed. Cloudflare notes that agents need not render a site's design, see its imagery or click its advertising: the human-oriented interface may never be presented at all. The future internet will therefore carry vastly more interactions that no person directly sees.[1,3]

That loss of visibility removes many of today's informal trust signals and reduces opportunities for human judgement. A person can pause at an unfamiliar website, notice an unexpected payment request, read terms, reject a permission prompt or recognise that a page is impersonating a known organisation. An autonomous agent may make hundreds of linked decisions before a person reviews the result. When agents buy, publish, disclose information, enter agreements or operate critical systems, trust can no longer depend on a human inspecting the screen at the decisive moment.

An agentic internet consequently requires a higher – not lower – standard of integrity. Agent identity should be collision-resistant across participating trust domains, machine-verifiable and portable across platforms. A relying system must be able to establish which accountable person or recognised entity stands behind an agent, where authoritative verification and status information can be obtained, and what authority the agent carries for the action at hand. These checks must occur automatically and at internet scale.

Today, auDA is the organisation endorsed by the Australian Government to manage the .au domain, providing a trusted national namespace for Australians. That existing role does not itself confer a mandate to operate AI-agent trust infrastructure. It does, however, provide an unusually relevant institutional and technical foundation from which Australia could extend a public chain of accountability to autonomous agents, subject to explicit policy authority, transparent governance and appropriate oversight. If durable agent identity becomes dependent on proprietary platforms alone, the future internet risks fragmented namespaces, opaque accountability and higher switching costs precisely when machines – not people – are making more of its decisions.[4,5]

Policy directions

AI agents are becoming actors: they discover services, invoke tools, exchange data, make commitments and initiate transactions across organisational boundaries. Yet most are identified only inside the platform that created them. When an agent crosses that boundary, counterparties may be unable to establish what acted, which organisation stands behind it, whether its authority remains valid, or where to obtain authoritative status and revocation information. This is no longer an abstract standards issue; it is a precondition for safe adoption at national scale.

The policy requirement is clear. OpenAI's 2026 call for collective cyber defence, endorsed across government and industry, asks frontier AI companies to ensure that agentic identities are "traceable and accountable". The US National Institute of Standards and Technology is separately examining identification, authentication, delegation, auditing and non-repudiation for software and AI agents. Teresa Shea's paper, "AI Agents Need Accountability That Travels With Them", captures the central design principle – as AI agents become autonomous and operate across organisational and platform boundaries, their identity, ownership, and accountability must be established through a persistent, independently verifiable trust mechanism rather than remaining tied to the platform in which they were created.[6–8]

A national answer must not create another closed identity silo. It must give any relying party a common starting point, preserve organisational autonomy, support cryptographic verification and lifecycle change, and operate under domestic governance. DNS already provides these properties for names and, through DNS Security Extensions (DNSSEC), can add data-origin authentication and integrity along a signed chain of delegation. Several independent Internet Engineering Task Force (IETF) draft standards now explore DNS-based approaches to durable agent identity, domain-authorised declarations and cryptographic verification. These drafts do not represent IETF consensus or endorsement, but they demonstrate active architectural interest in DNS as one component of the emerging agent trust stack.[9–12]

The decision before national stakeholders

The question is not whether every aspect of agent identity belongs in the DNS. It does not. Authentication credentials, workload attestation, authorisation decisions, transaction logs and privacy-sensitive attributes belong in the systems designed for them. The decision is where the authoritative public binding should live that connects an agent to its accountable entity, its verification material and its current status. For national use, that binding should be provided through a governed namespace operated as a public-interest trust service.

The trust architecture of the agentic internet

A national trust anchor should not attempt to determine whether an AI agent is safe, truthful or entitled to perform every action it requests. Its purpose is narrower and more fundamental: to provide an authoritative, independently resolvable binding between an agent and the accountable entity that stands behind it. That binding gives relying parties a common root from which other claims about the agent can be discovered and verified, without requiring either party to trust the platform on which the agent happens to run.

This distinction separates four questions that are often collapsed into the single word "identity". Identity asks what agent is acting. Accountability asks which legal entity accepts responsibility for it. Authority asks what the agent is permitted to do in a particular context. Behaviour and evidence concern what the agent actually did and whether those actions complied with policy.

The national trust layer should answer the first two directly and provide authenticated pointers to the systems that answer the latter two.

Trust question	Primary answer
Identity	A durable, collision-resistant agent name that can be resolved independently of the hosting platform.
Accountability	A verified binding from that agent name to an accountable person or recognised organisation, with authoritative lifecycle status.
Authority	Delegated credentials, permissions, transaction limits and context-specific policy evaluated by relying systems.
Behaviour and evidence	Attestation, audit, transaction records and incident evidence retained by the accountable parties or approved services.

TRUST PRINCIPLE: Resolve accountability, not activity. The public trust layer should reveal who stands behind an agent and where authoritative verification can be obtained – not expose prompts, transactions, personal data or operational history.

Trust anchor versus registry

The trust anchor and the registry are related but distinct. The trust anchor is the architectural and institutional role: the authoritative national point from which a relying party can verify the delegation of accountability. The registry is the operational service that administers names, organisational bindings, lifecycle status and authenticated pointers under that anchor.

The registry would not replace enterprise identity, public key infrastructure (PKI), Open Authorisation (OAuth), workload identity, decentralised identifiers, attestation or sector-specific assurance. Those systems remain responsible for authentication and operational authority. The national layer provides the stable public starting point that lets those systems answer a prior question: who is accountable for the agent presenting them?

A layered trust model

The architecture is therefore layered rather than monolithic. At the base is the existing global DNS root. Country-code delegation provides the national trust boundary. Beneath that, the national agent namespace delegates control to verified organisations, which in turn name and manage their own agents. DNSSEC authenticates the chain of delegation and the discovery records associated with each agent. Above that public anchor sit the credentials and controls required for runtime trust: keys, attestations, OAuth or Secure Production Identity Framework for Everyone (SPIFFE) credentials, policy engines, transaction-specific permissions and audit mechanisms.

Layer	Role in the trust architecture
Global DNS root	Provides the universal hierarchy through which national namespaces are independently resolvable.
National delegation	Establishes a domestically governed trust boundary under the ccTLD.
Accountable organisation	Receives a verified delegation and controls subordinate names for its agents.
Agent identity	Publishes a stable name, lifecycle status and authenticated pointers to verification material.
Runtime trust	Uses credentials, attestations and policy to decide whether a specific action is authorised now.

The result is a deliberate separation between durable public accountability and transient operational authority. An agent can move between clouds, models or orchestration platforms without losing its public identity, while its runtime credentials can rotate or expire as frequently as security policy requires. A platform can still provide sophisticated identity and authorisation services; it simply does not become the sole root of trust for the agent's existence or organisational affiliation.

How two agents establish trust

Consider an Australian procurement agent receiving an offer from another agent claiming to represent a supplier. Before acting, the procurement agent resolves the counterparty's durable agent name and validates the DNSSEC chain to the relevant national namespace. It retrieves the agent's signed identity and status information, confirms the accountable organisation, then verifies the runtime credential presented for the proposed transaction. Only after those independent checks does it evaluate whether the counterparty has authority to negotiate or commit within the required limits.

Step	What the relying agent verifies
1. Resolve	Resolve the counterparty agent name through DNS and validate the DNSSEC chain.
2. Attribute	Establish the accountable organisation and retrieve the authoritative identity/status endpoint.
3. Check status	Confirm that the agent identity is active and has not been suspended, revoked, transferred or retired.
4. Verify authority	Validate the runtime credential, delegation and transaction-specific limits supplied by the counterparty.

Step	What the relying agent verifies
5. Act and record	Proceed only if policy permits, retaining the evidence needed for audit, dispute or incident response.

The trust anchor establishes durable accountability; the runtime credential establishes present authority. Neither alone is sufficient for a consequential transaction.

The same model can extend across borders without creating a single global AI identity authority. An Australian agent anchored beneath .au could interact with an overseas agent anchored beneath another country-code namespace, with both parties relying on the global DNS hierarchy for resolution while each jurisdiction retains governance over its own national trust boundary. International standards would define interoperable record formats, assurance semantics and verification procedures; national registries would provide the authoritative delegations.

Why a national layer in a global system

A national trust anchor is not a retreat from global interoperability. It is a way of achieving it through delegated governance. DNS already demonstrates the model: a single global hierarchy delegates authority to national namespaces, which in turn delegate to organisations. The proposed agent trust architecture follows the same pattern – global resolution, national governance, organisational responsibility and agent-level identity.

This matters because legal accountability, incident response, registration eligibility, due process and suspension ultimately operate within jurisdictions. A national layer gives those obligations a recognised governance home while preserving the ability of any relying party anywhere in the world to resolve and verify the resulting identity. The objective is therefore not a collection of national silos, but a federated global fabric of nationally governed accountability.

This is the architectural role of the National AI Agent Trust Registry: not a universal credential issuer, not an authorisation engine and not a repository of agent activity, but the authoritative public root through which an agent's identity, accountable sponsor and current trust status can be independently resolved. Other identity technologies remain free to compete and evolve above it.

Why DNS and the national DNS operator

Major cloud and AI platforms will develop capable agent identity systems, and those systems should be used. The national infrastructure question is different: whether the durable public identity through which counterparties establish accountability should depend on the platform currently executing the agent. A neutral accountability anchor separates accountability from execution, allowing an organisation to move an agent between models, clouds or orchestration

systems without losing the public identity through which others establish who stands behind it. This improves portability and reduces the risk that identity itself becomes another source of platform lock-in.[13]

Existing mechanisms such as PKI, OAuth/OIDC, SPIFFE, decentralised identifiers and platform identity and access management remain essential for authentication, credentials, workload identity and runtime authority. The proposed national layer does not replace them. Its purpose is to provide a stable, publicly resolvable accountability and discovery anchor through which those mechanisms can be composed across organisational and platform boundaries.

DNS is a practical foundation because its distinctive property is not simply global naming. It combines universal internet resolution, hierarchical delegation, independence from individual application platforms and an established country-level delegation structure. That hierarchy allows an accountable organisation to control subordinate names for its agents while a parent authority governs eligibility and delegation.[14,15] DNSSEC adds data-origin authentication and integrity for signed DNS data and can authenticate pointers to agent identity documents, keys, supported protocols and lifecycle-status services.[9]

What DNS does – and does not – guarantee

DNS provides authoritative naming and delegation. DNSSEC can protect the origin and integrity of signed DNS data and allow a relying party to determine whether published trust information validates back to an accepted trust anchor. It does not, by itself, prove that an agent is safe, truthful, legally authorised for a transaction or controlled by the entity named in a registration record. The accountability claim is established by registry proofing, eligibility and governance; DNSSEC protects the publication of that claim.[9]

DNS should not be treated as an instantaneous revocation channel. DNS caching is governed by time to live (TTL) values, while DNSSEC signatures have separate validity periods. A production architecture should therefore define bounded revocation latency and use short-lived signed status assertions or a dedicated authoritative status service, discovered and authenticated from the DNS anchor, for high-consequence decisions.[9]

SECURITY BOUNDARY: Registry governance establishes the accountable binding. DNSSEC authenticates its publication. Runtime credentials establish present authority. Relying-party policy decides whether to act.

Why auDA is a strong Australian candidate

A national trust-anchor operator should combine authoritative namespace capability, internet-scale resilience, public-interest governance, neutral treatment of competing platforms, lifecycle and dispute processes, domestic accountability and established international coordination. In Australia, auDA has an unusually strong fit against these criteria because it already operates the

.au namespace under Australian Government endorsement and participates in multi-stakeholder internet governance.[4,5,15]

The existing .au role provides capability and institutional foundations, but it should not be treated as automatically authorising an AI-agent trust function. Government should retain responsibility for policy, assurance requirements, sector obligations and oversight, while auDA could perform the narrower infrastructure role of operating the authoritative namespace and lifecycle service under an explicit mandate. This reuses mature registry and DNSSEC capabilities while keeping credential technology, proofing, attestation and runtime authorisation open to competing providers and sector-specific schemes.[4,5]

CORE PRINCIPLE: One national root of accountability; many interoperable credentials, platforms and assurance schemes above it.

International federation by delegation, not centralisation

The long-term design should assume reciprocal national trust namespaces rather than a single supranational agent registry. An Australian agent could be anchored beneath .au and an overseas agent beneath another nationally governed namespace where that country chooses such a model. DNS supplies global resolvability without requiring identical national policy. Because ccTLD governance and eligibility regimes differ, federation should not imply equivalent assurance: each national scheme should publish machine-readable assurance semantics and recognition rules.[15]

Federation therefore occurs at two levels. DNS provides namespace interoperability and authenticated delegation; standards above DNS define common metadata, status semantics, assurance claims and credential discovery. Bilateral or multilateral recognition frameworks can map assurance levels where needed. The relying party remains responsible for deciding which foreign assurance regimes and credential classes satisfy the risk of a particular transaction. The global layer standardises how trust information is located and interpreted; the national layer governs eligibility and assurance; the relying party decides what it accepts.

The mandate required now

The agent ecosystem is still forming, but identity decisions made now will harden quickly into APIs, procurement requirements, platform contracts and trust relationships. If national stakeholders wait until proprietary identifiers dominate cross-platform agent interactions, interoperability will become a costly retrofit and the public accountability layer may already have been privatised by default. Government and industry should therefore recognise the ccTLD operator as the lead convenor and intended operator of the national AI agent trust namespace, subject to transparent multi-stakeholder design, technical validation and appropriate statutory or policy oversight.

Recommended national action

1. Confer an explicit mandate. Government, industry peak bodies and critical-infrastructure leaders endorse a ccTLD-led national trust registry as shared digital infrastructure, subject to transparent multi-stakeholder design, competition safeguards and appropriate statutory or policy oversight.
2. Establish a multi-stakeholder design group spanning cyber security, digital identity, AI safety, privacy, competition, law enforcement, procurement and standards.
3. Pilot the architecture. Run a 6-12 month sandbox with government services, banks, telecommunications, health, energy and major AI platforms; test DNSSEC resolution, organisational proofing, key rotation, bounded revocation latency, privacy controls, suspension and cross-platform verification.
4. Create demand through procurement. Require high-impact agents used by government or critical infrastructure to present a resolvable national identity and verifiable accountability chain, while accepting interoperable runtime credential schemes.
5. Federate internationally. Work through relevant IETF standards processes, Internet Corporation for Assigned Names and Numbers (ICANN)/ Internet Assigned Numbers Authority (IANA) coordination channels, regional ccTLD organisations and peer registries to explore common discovery, status and assurance semantics. Treat current agent-identity Internet-Drafts as work in progress rather than standards or institutional endorsement.[10–12]

OUTCOME: A registered agent should be as easy to resolve to its accountable sponsor as a domain name is to resolve to its authoritative operator – and high-consequence relying systems should be able to obtain sufficiently fresh status information to stop relying on that identity when the binding is suspended or revoked.

Conclusion

The agentic internet requires public trust infrastructure, not merely persistent identifiers. A relying machine must be able to determine what agent is acting, which accountable person or recognised entity stands behind it, whether that binding remains current, and where to verify the authority presented for a particular action. These are different questions and should remain architecturally separated.

DNS provides a uniquely mature combination of universal internet resolution, hierarchical delegation, cryptographic integrity through DNSSEC and established national stewardship. It should not be asked to certify agent behaviour or replace runtime identity systems. The strategic opportunity is narrower and stronger: use a proven public naming and delegation system as the accountability and discovery anchor through which diverse credential, attestation and authorisation systems can interoperate. In Australia, auDA is a strong candidate to operate that function because of its existing capability and public-interest role, but the function should be conferred explicitly and governed transparently.[4,5,9]

References

1. Cloudflare, Inc., Q2 2026 Earnings Call, Corrected Transcript (6 August 2026), pp. 6, 8–9. https://www.cloudflare.net/files/doc_financials/2026/q2/NET-US-CORRECTED-TRANSCRIPT-Cloudflare-IncNETUS-Q2-2026-Earnings-Call-6August2026-500-ET-PM-07-08-2026-2026-08-07-0.pdf
2. Cloudflare, "Content Independence Day, one year on: building the business model for the agentic Internet" (1 July 2026). https://blog.cloudflare.com/agentic-internet-bot-report/
3. Cloudflare, "Building an open Agentic Internet: readable, discoverable, callable, and payable" (6 August 2026). https://blog.cloudflare.com/the-agentic-internet/
4. Australian Government / auDA, "2025 Terms of Endorsement for auDA" (2025). https://files.auda.org.au/documents/Terms-of-Endorsement-for-auDA-2025.pdf
5. auDA, "Who we are" and current statement of Australian Government endorsement and core functions. https://www.auda.org.au/about-auda/who-we-are/
6. Shea, T., "AI Agents Need Accountability That Travels With Them," The Cipher Brief (2026). https://www.thecipherbrief.com/ai-agents-need-accountability-that-travels-with-them
7. OpenAI and signatories, "A call for collective action on cyber defense" (2026). https://openai.com/collective-cyberdefense/
8. NIST NCCoE, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization" (Concept Paper, 2026). https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf
9. IETF, RFC 4033, "DNS Security Introduction and Requirements" (2005). https://www.rfc-editor.org/rfc/rfc4033.html
10. IETF Internet-Draft, "DNS-Anchored Durable Identity for AI Agents (DNSid)" (work in progress, 2026). https://datatracker.ietf.org/doc/draft-ihsanullah-dnsid/
11. IETF Internet-Draft, "ApertoID: DNS-Based Agent Identity Declaration Protocol" (work in progress, 2026). https://datatracker.ietf.org/doc/draft-ferro-dnsop-apertoid/
12. IETF Internet-Draft, "Agent Identity Protocol (AIP)" (work in progress, 2026). https://datatracker.ietf.org/doc/draft-prakash-aip/
13. Australian Competition and Consumer Commission (ACCC), "Digital Platform Services Inquiry — March 2025 final report" (released 23 June 2025), including emerging competition issues in cloud computing and generative AI. https://www.accc.gov.au/inquiries-and-consultations/finalised-inquiries-and-monitoring/digital-platform-services-inquiry-2020-25/march-2025-final-report
14. IETF, RFC 1591, "Domain Name System Structure and Delegation" (1994). https://www.rfc-editor.org/rfc/rfc1591/
15. IANA, "Delegating or transferring a country-code top-level domain (ccTLD)," current guidance. https://www.iana.org/help/cctld-delegation

.au Domain Administration Ltd
www.auda.org.au

PO Box 18315
Melbourne VIC 3001
general@auda.org.au

September 2026