

TIP Protocol Whitepaper, Version 1.1

Trust Identity Protocol for the Verifiable Internet

The AI Lab Intelligence Unobscured, Inc.

theailab.org

August 23, 2026

Our mission: make truth verifiable.

Table of Contents

Suggested Citation	8
License Summary	9
Document Status	9
Distribution List	9
Foreword	9
Acknowledgments	11
Founding Node Operators	11
AI Trust Council Founding Members	12
Members joining the Council	12
AI Trust Advisory Board	13
AI Trust Ambassadors	14
Counsel	14
Standards organizations and supervisory authorities	14
Open source community	14
Readers	15
Executive Summary	15
ES.1 The condition the protocol addresses	15
ES.2 The Trust Identity Protocol	15
ES.3 The federated network	16
ES.4 Governance, the AI Trust Council	16
ES.5 Licensing	18
ES.6 Cryptographic foundation	18
ES.7 Operational status	18
ES.8 Regulatory posture	19
ES.9 What this whitepaper invites	19
ES.10 Contact	20
Part I: The Verification Problem	20
1.1 The condition the protocol addresses	20
1.2 The limits of platform-controlled trust signals	21

1.3 The limits of existing provenance standards	21
1.4 The limits of watermarking	22
1.5 The identity gap	23
1.6 Requirements for a durable solution	23
1.7 Comparative analysis	24
1.8 Scope of this whitepaper	25
Part II: Design Principles	25
2.1 Decentralized identity, federated verification	26
2.2 Post-quantum cryptography from genesis	26
2.3 Append-only directed acyclic graph for non-repudiable provenance	27
2.4 Pseudonymity with accountability	28
2.5 Adversarial robustness	28
2.5.1 Sybil resistance	28
2.5.2 Key compromise containment	29
2.5.3 Regulatory and jurisdictional pressure	29
2.5.4 Cryptographic adversary	29
2.6 Minimality	29
2.7 Interoperability	30
2.8 Open specification, structured licensing	30
Part III: Cryptographic Foundation	31
3.1 Standards landscape and rationale for the NIST post-quantum suite	31
3.2 ML-KEM-768 (FIPS 203) key encapsulation	32
3.3 ML-DSA-65 (FIPS 204) primary signature scheme	32
3.4 SLH-DSA (FIPS 205) hash-based signature fallback	33
3.5 PRF-to-AES key protection chain	34
3.6 Three-hash content addressing	34
3.7 Threat model and formal security claims	35
3.7.1 Adversaries the protocol contains	35
3.7.2 Adversaries the protocol does not contain	35
3.7.3 Cryptographic migration path	36
3.8 Conformance testing	36
Part IV: TIP-ID: Identity Layer	37
4.1 Cryptographic Trust Identity (CTID) construction	37
4.1.1 Generation	37
4.1.2 Derivation of the CTID from the public key	37
4.1.3 Properties of the CTID	38
4.1.4 Migration to native post-quantum authenticator support	38
4.2 The Verification Provider	39
4.2.1 Role	39
4.2.2 Accreditation	39
4.2.3 Annual obligations	40
4.2.4 Suspension and revocation	40
4.3 Pseudonymity, revocation, and succession	40
4.3.1 Pseudonymity in operation	40
4.3.2 Revocation by the holder	40
4.3.3 Revocation by the Verification Provider	41
4.3.4 Succession	41
4.4 Biometric binding and the WebAuthn resident key flow	41
4.4.1 The WebAuthn resident key authenticator	41

4.4.2 Biometric user verification	42
4.4.3 Backup, recovery, and device loss	42
4.5 Identity portability across jurisdictions	43
4.5.1 Portability between protocol participants	43
4.5.2 Portability of the regulatory effect	43
4.5.3 Cross-border interoperability with the European Digital Identity Wallet	43
4.5.4 Cross-border interoperability with Aadhaar and DigiLocker	43
4.6 Identity governance and dispute	43
4.6.1 The CTID holder’s standing	43
4.6.2 The Verification Provider’s standing	44
4.6.3 Disputes	44
Part V: TIP-CONTENT: Provenance Layer	44
5.1 Overview of the Canonical Normalization Algorithm Version 2.2	45
5.2 The nine canonical normalization steps	45
5.2.1 Step 1: Platform identification	45
5.2.2 Step 2: Content scope extraction	47
5.2.3 Step 3: Structural canonicalization	47
5.2.4 Step 4: Character normalization	47
5.2.5 Step 5: Reference normalization	48
5.2.6 Step 6: Embedded media identification	48
5.2.7 Step 7: Metadata extraction	48
5.2.8 Step 8: Serialization	48
5.2.9 Step 9: Three-hash addressing	49
5.3 Publisher Mode	49
5.3.1 Architectural distinction from Creator Mode	49
5.3.2 Editorial workflow integration	49
5.3.3 Byline attribution	50
5.4 Creator Mode	50
5.4.1 Browser extension flow	50
5.4.2 Mobile flow	50
5.4.3 Embedded creator flow	51
5.5 Content versioning and CONTENT_UPDATED semantics	51
5.5.1 Versioning model	51
5.5.2 The CONTENT_UPDATED event	51
5.5.3 Reader-facing versioning	52
5.6 Origin Codes	52
5.6.1 OH: Original Human	52
5.6.2 AA: AI-Assisted	52
5.6.3 AG: AI-Generated	52
5.6.4 MX: Mixed	53
5.6.5 Selection of the Origin Code	53
5.7 Signature payload format and reference packet	53
5.7.1 Payload structure	53
5.7.2 Wire format	54
5.7.3 Reference packet	54
5.8 Content Classification, Transient Media Retention, and Classifier Improvement	54
Part VI: TIP-TRUST: Reputation Layer	55
6.1 The Trust Score	55
6.2 The four sub-scores	55
6.2.1 Cryptographic sub-score	55

6.2.2 Behavioral sub-score	56
6.2.3 Adjudicated sub-score	56
6.2.4 Network sub-score	57
6.2.5 Aggregation	57
6.3 Trust Score tiers	57
6.4 Blocking Items B1 through B6	58
6.5 Bonded jury adjudication	59
6.5.1 The bonded juror	59
6.5.2 Procedure	59
6.5.3 AI-assisted pre-classification	60
6.6 Appeals	60
6.7 Reader-facing display	60
6.7.1 The Trust Score badge	60
6.7.2 Display modes	60
6.7.3 Accessibility	61
6.8 The Global Seal of Trust	61
Part VII: Federated DAG and Network Topology	62
7.1 The federated directed acyclic graph	62
7.1.1 Entry structure	62
7.1.2 Append-only property	63
7.1.3 Public verifiability	63
7.2 Node Operator roles	63
7.2.1 Light Node	63
7.2.2 Full Node	63
7.2.3 Verification Provider node	64
7.3 Operational periods	64
7.3.1 The Founding Node Operators	64
7.3.2 The character of the Founding Period	65
7.3.3 Transition to the Network Period	66
7.4 Synchronization, consistency, and partition tolerance	66
7.4.1 Synchronization protocol	66
7.4.2 Consistency model	66
7.4.3 Partition tolerance	66
7.4.4 Resolution of structural ambiguity	67
7.4.5 Time consensus	67
7.5 Service level commitments	67
7.5.1 Service level targets	67
7.5.2 Service level remedies	68
7.5.3 Incident classification	68
7.6 Geographic distribution of the Founding Period network	68
7.7 Capacity model and scaling profile	68
7.7.1 Storage profile	69
7.7.2 Bandwidth profile	69
7.7.3 Computational profile	69
7.8 Industry support and program memberships	69
Part VIII: Reference Implementation and Integration	70
8.1 The REST API surface	70
8.1.1 Endpoint inventory	71
8.1.2 Authentication	72
8.1.3 Rate limiting	72

8.1.4 Error model	73
8.1.5 Versioning	73
8.2 Browser extension	73
8.2.1 Supported browsers	73
8.2.2 Architecture	73
8.2.3 Publisher Mode and Creator Mode	74
8.2.4 Privacy posture	74
8.2.5 Localization	74
8.3 The <tip-badge> web component	74
8.3.1 Specification	74
8.3.2 Use by publishers	74
8.3.3 Customization	75
8.3.4 Accessibility	75
8.4 WordPress reference plugin	75
8.4.1 CNA-1.0	75
8.4.2 Plugin architecture	75
8.4.3 Distribution	75
8.5 Mobile web application	76
8.5.1 Architecture	76
8.5.2 Installation	76
8.5.3 Functionality	76
8.6 Software development kits	76
8.6.1 Published kits	76
8.6.2 Planned kits	76
8.6.3 Conformance	77
8.7 Distribution and update channels	77
8.7.1 Browser extension distribution	77
8.7.2 Update cadence	77
8.7.3 Software bill of materials	77
Part IX: Regulatory Alignment	78
9.1 European Union	78
9.1.1 Artificial Intelligence Act, Regulation (EU) 2024/1689	78
9.1.2 General Data Protection Regulation, Regulation (EU) 2016/679	79
9.1.3 Digital Services Act, Regulation (EU) 2022/2065	80
9.1.4 eIDAS and the European Digital Identity Wallet	81
9.1.5 NIS2 Directive, Directive (EU) 2022/2555	81
9.1.6 Cyber Resilience Act, Regulation (EU) 2024/2847	82
9.1.7 Council of Europe Framework Convention on AI	82
9.2 United States	82
9.2.1 Federal Trade Commission Act Section 5	82
9.2.2 NIST AI Risk Management Framework	83
9.2.3 Executive Order 14110 and successor instruments	83
9.2.4 State law mosaic	83
9.3 United Kingdom	84
9.3.1 Online Safety Act 2023	84
9.3.2 United Kingdom General Data Protection Regulation and Information Commissioner’s Office guidance	84
9.4 New Zealand, Privacy Act 2020	84
9.5 India, Digital Personal Data Protection Act 2023	84
9.6 OECD AI Principles	85

9.7 Standards organization engagement	85
9.8 Export control posture	85
9.9 Antitrust posture	86
Part X: Governance and Licensing	86
10.1 The AI Lab Intelligence Unobscured, Inc.	86
10.2 The AI Trust Council	87
10.2.1 Establishment and independence	87
10.2.2 Founding Members	87
10.2.3 Decision rights	88
10.2.4 Independence covenants	89
10.2.5 AI Trust Advisory Board	90
10.2.6 AI Trust Ambassadors	90
10.2.7 Activation and Network Period transition	91
10.3 TIPCL-1.0 License Summary	91
10.3.1 Free Tier eligibility	91
10.3.2 Commercial License tier schedule	91
10.3.3 Grace period and renewal	92
10.4 Patent licensing under TIPCL-1.0 Section 8	92
10.5 Trademark policy	93
10.6 Apache 2.0 conversion provision	94
10.7 Verification Provider role, accreditation, and obligations	94
10.8 EU AI Act classification posture	95
10.9 Cross-jurisdictional governance commitments	96
Part XI: Roadmap and Forward-Looking Statements	97
11.1 Operational milestones to date	97
11.1.1 Operational Readiness Conditions	98
11.2 Twelve-month horizon	99
11.3 Twenty-four-month horizon	100
11.4 Thirty-six-month horizon	101
11.5 Material risks	101
11.6 Conditions for AI Trust Council Network Period reauthorization	102
11.7 Statement on the design horizon	103
Appendix A: Glossary of Defined Terms	103
Appendix B: Acronym Index	107
Appendix C: CNA-2.2 Worked Example	110
C.1 Source content unit	110
C.2 Step-by-step normalization	110
Step 1: Platform identification	110
Step 2: Content scope extraction	111
Step 3: Structural canonicalization	111
Step 4: Character normalization	111
Step 5: Reference normalization	111
Step 6: Embedded media identification	112
Step 7: Metadata extraction	112
Step 8: Serialization	112
Step 9: Three-hash addressing	113
C.3 Verification at the reader	113

C.4 Verifiability across surfaces	113
C.5 Failure mode, modified content	113
C.6 Failure mode, revoked CTID	114
Appendix D: End-to-End Identity Issuance, Signing, and Verification Example	114
D.1 Scenario	114
D.2 Step 1: Identity issuance	114
D.3 Step 2: Initial Trust Score computation	115
D.4 Step 3: Content signing	115
D.5 Step 4: Reader-side verification	116
D.6 Step 5: Subsequent modification	116
D.7 Failure mode, authenticator loss	117
D.8 Failure mode, adverse adjudication	117
Appendix E: Compliance Crosswalk	118
E.1 ISO/IEC 27001:2022 Annex A controls	118
E.2 NIST Cybersecurity Framework 2.0 functions	120
E.3 NIST AI Risk Management Framework	121
E.4 OECD AI Principles	122
Appendix F: TIPCL-1.0 License Summary	123
F.1 Scope of license	123
F.2 Eligibility tiers	123
F.3 Required conditions of use	124
F.4 NOTICE file requirement	125
F.5 Patent license	125
F.6 Trademark license	125
F.7 Required attribution form	125
F.8 Apache 2.0 conversion provision	126
F.9 Termination	126
F.10 Governing law and dispute resolution	126
F.11 Disclaimer	126
Appendix G: References	127
G.1 Cryptographic standards	127
G.2 European Union legislation	127
G.3 United States legislation, regulation, and executive instruments	128
G.4 United States state legislation	128
G.5 United Kingdom legislation	128
G.6 New Zealand legislation	129
G.7 India legislation	129
G.8 Council of Europe instruments	129
G.9 OECD instruments	129
G.10 Standards organizations and industry bodies	129
G.11 The AI Lab published works	129
Appendix H: Contact Directory	130
H.1 Licensing and commercial implementation	130
H.2 Network participation	130
H.3 Governance	130
H.4 Technical engagement	130
H.5 Legal	131

H.6 Press and public affairs	131
H.7 Postal address	131
Appendix I: Suggested Citation	131
I.1 Plain text	131
I.2 IEEE	131
I.3 ACM	131
I.4 APA	132
I.5 Chicago	132
I.6 BibTeX	132
I.7 Citation of specific Parts and Sections	132
I.8 Citation of the canonical Specification	132
Appendix J: Notice and Disclaimer	132
J.1 Copyright notice	133
J.2 No warranty	133
J.3 No legal advice and no fiduciary relationship	133
J.4 Trademark notices	133
J.5 Patent notice	133
J.6 Regulatory references and compliance posture	134
J.7 Forward-looking statements safe harbor	134
J.8 No reliance and no offer	135
J.9 Limitation of liability	135
J.10 Governing law and dispute resolution	135
J.11 Severability	136
J.12 Entire understanding	136
Appendix K: Document History	136
K.1 Cadence of subsequent versions	137
K.2 Amendment procedure	137
K.3 Withdrawal	137
K.4 Archival	137
Appendix L: The Architect and Author	138
L.1 Identification	138
L.2 Works of authorship and invention	138
L.3 Founder’s commitment	139
L.4 Statement for the historical record	139

Suggested Citation

Mendhe, D. (2026). *TIP Protocol Whitepaper, Version 1.1: Trust Identity Protocol for the Verifiable Internet*. The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware. USCO Application No. 1-15175755931. Available at <https://theailab.org/whitepaper>.

IEEE: D. Mendhe, “TIP Protocol Whitepaper, Version 1.1,” The AI Lab Intelligence Unobscured, Inc., Wilmington, DE, 2026.

ACM: Dinesh Mendhe. 2026. *TIP Protocol Whitepaper, Version 1.1*. The AI Lab Intelligence Unobscured, Inc., Wilmington, DE.

APA: Mendhe, D. (2026). *TIP Protocol Whitepaper, Version 1.1: Trust Identity Protocol for the Verifiable Internet*. The AI Lab Intelligence Unobscured, Inc.

Chicago: Mendhe, Dinesh. 2026. *TIP Protocol Whitepaper, Version 1.1*. Wilmington, DE: The AI Lab Intelligence Unobscured, Inc.

License Summary

This whitepaper is licensed under the Creative Commons Attribution 4.0 International Public License. Implementations of the technical framework described herein are licensed under the TIP Protocol Code License Version 1.0 (TIPCL-1.0). See Part X and Appendix F.

Document Status

Version 1.1. Published August 23, 2026, incorporating the corrections and amendments published at theailab.org/whitepaper/errata: Free Tier eligibility of an individual natural person at any income level, and the TIPCL-1.0 commercial fee schedule scaled by World Bank income group below US\$25,000,000. Version 1.0 was published June 16, 2026 (third printing, incorporating accuracy corrections to legal citations, the expanded universal content-type taxonomy, the AI Trust Council growth and public verification channels sections, and the Article 95(3) of the EU AI Act multi-stakeholder framing). Original publication: June 4, 2026. Next review: January 2027 (concurrent with the inaugural AI Trust Council annual transparency report). Errata may be reported to whitepaper-errata@theailab.org and will be published at theailab.org/whitepaper/errata.

Distribution List

Public.

Foreword

By Dinesh Mendhe, Founder and Chairman of The AI Lab Intelligence Unobscured, Inc.

Dinesh Mendhe is the creator and founder of the Trust Identity Protocol, the AI Trust Council, the AI Trust Registry, the AI Trust ID, and the Human Trust ID system. He is the sole inventor named on the five United States provisional patent applications underlying the Trust Identity Protocol, the architect of the Canonical Normalization Algorithm at every published version, and the author of this whitepaper, of the canonical TIP Protocol Specification, and of the founding instruments through which the protocol and the Council enter the historical record.

I am writing this foreword on the morning of the publication of the first version of the Trust Identity Protocol Whitepaper. The work of two years sits behind it. The work of the next several decades sits in front of it.

Today the protocol stands at the threshold of Genesis. The publication of this whitepaper, the ratification of the AI Trust Council Charter, the execution of the Founding Node Operator Agreements, and the satisfaction of the Operational Readiness Conditions described in Part XI together establish the conditions on which the federated network goes live. The Genesis

Date is a date in June 2026, to be determined by the sole director on the satisfaction of those conditions, and published on theailab.org not less than three business days in advance.

I will not, in this foreword, restate the substance of the document. Eleven Parts and twelve Appendices set out the technical architecture, the regulatory posture, the institutional governance, the licensing terms, the operational milestones, and the material risks of the protocol. The reader will find each in its place. What I will do, in this foreword, is record three commitments on which everything in the document depends and to which I have asked the reader to attach The AI Lab and to attach me personally.

The first commitment is to the people. The Trust Identity Protocol exists because journalists, publishers, creators, public officials, courts, regulators, and ordinary readers cannot, in the present state of digital media, reliably distinguish authentic content from synthetic content, or reliably identify the party responsible for either. The condition has consequences in journalism, in civic discourse, in the administration of justice, in financial markets, and in personal life. It is not an abstract problem. The protocol exists to address a concrete consequence in concrete lives. The protocol's design choices, beginning with the architectural decision that the biometric template of a creator never leaves the creator's device and continuing through the pseudonymity-with-accountability structure of the CTID, are made with the human person at the center. I commit, on behalf of The AI Lab, that no change to the protocol or to the operations of The AI Lab will compromise that orientation.

The second commitment is to the rule of law. The protocol does not seek to exempt any party from the operation of the law. The Verification Provider operates under the law of its jurisdiction. The Node Operator operates under the law of its jurisdiction. The AI Lab operates under the law of the State of Delaware and of the United States, and under the law of each jurisdiction in which it engages. The AI Trust Council operates under a Charter drafted to satisfy the independence indicia that regulators and courts will apply. The protocol's transparency provisions, its warrant canary requirement, its jurisdiction declaration requirement, and its audit obligation are not concessions reluctantly made in response to regulatory pressure. They are design choices. The institutions that the protocol depends on are the institutions of a free society, and the protocol is offered in support of those institutions.

The third commitment is to the long horizon. The cryptographic primitives are selected against the cryptographic conditions of the next several decades. The append-only DAG is structured for retention of records over the same horizon. The licensing terms convert to permissive open source on a fixed date. The AI Trust Council Charter contemplates expansion, dissent, succession, and reauthorization on an indefinite horizon. The protocol is not a startup product. The protocol is a piece of foundational infrastructure offered to a society that requires foundational infrastructure for content authenticity and verifiable identity. The AI Lab will operate the protocol for as long as the protocol is needed, will steward the protocol's evolution through the AI Trust Council, and will, on the conditions described in this whitepaper, eventually relinquish proprietary control through the Apache License 2.0 conversion provision of January 1, 2031.

To the Founding Node Operators who have committed their organizations to the operation of the federated network from the Genesis Date through the Founding Period, in the United Kingdom, in India, in the United States, and in New Zealand: I am grateful. To the Founding Members of the AI Trust Council, who have accepted responsibility for the governance of the protocol during a period in which the responsibility is significant and the public reward

is modest, I am grateful. To the regulators, supervisory authorities, standards organizations, civil society organizations, and academic institutions that will engage with the protocol in the months and years to come, I extend the invitation The AI Lab makes in this whitepaper, and I record The AI Lab’s commitment to engage as a serious institution, with patience, with respect for the public function of the engaging body, and with the institutional patience that a foundational infrastructure project requires.

The protocol’s success is not within the gift of any single party. It depends on the conduct of many parties in many jurisdictions over many years. The AI Lab will do its part. We ask the reader to consider whether to do the reader’s part as well.

Dinesh Mendhe Founder and Chairman The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware June 2, 2026

Acknowledgments

The publication of this whitepaper, and the operation of the protocol it describes, depend on the contributions of many institutions and many natural persons. The AI Lab Intelligence Unobscured, Inc. acknowledges, with gratitude, the following contributions.

Founding Node Operators

The protocol’s federated network is, at the date of publication, being stood up in pre-Genesis operating mode, the pilot phase preceding live production defined in Executive Summary Section ES.7. Live production operation commences on the Genesis Date. During the Pre-Genesis Period and continuing into the Founding Period that follows from the Genesis Date, the network is and will be operated by seven Founding Node Operators with signed Founding Node Operator Agreements and three additional Founding Node Operators in accession, across four jurisdictions:

Founding Node Operator	Jurisdiction	Node type	Status
THE PRESCIENT PACHYDERM LTD	United Kingdom	Full Node	Signed
AZLogics Private Limited	India	Full Node	Signed
Apex Modular Solutions LLC	United States	Full Node	Signed
Timpi International Ltd	New Zealand	Full Node	Signed
Lonestar Data Holdings Inc.	United States	Full Node	Signed
6Simplex Software Solutions Pvt Ltd	India	Full Node	Signed
The AI Lab Intelligence Unobscured, Inc.	United States (theailab.org)	Full Node	Signed
Marist University	United States	Full Node	In accession
Rutgers University	United States	Full Node	In accession
The Core / BOOM FactCheck	India	Full Node	In accession

At the date of publication, three of the ten Founding Node Operators above (Marist University, Rutgers University, and The Core / BOOM FactCheck) are in accession. Accession is the

period in which an incoming Founding Node Operator has agreed in principle and is in final review of the Founding Node Operator Agreement, the Service Level Agreement, and the Technical Requirements Specification. If any party in accession does not countersign by the Genesis Date, the AI Trust Council will publish a revised list of Founding Node Operators at theailab.org/whitepaper/errata, accompanied by a revised printing of this whitepaper.

Each Founding Node Operator that has signed has committed its organization to the operation of a Full Node, first in pre-Genesis operating mode and from the Genesis Date in live production operation; to the service level commitments; to the warrant canary and jurisdiction declaration requirements; and to the data protection and cybersecurity obligations of its jurisdiction. Each Founding Node Operator in accession has agreed in principle to the same commitments and is in the final-execution phase of those Agreements. The protocol is not yet in live production at the date of publication; live production operation commences on the Genesis Date. The AI Lab acknowledges in particular the contribution of Timpi International Ltd, the New Zealand Founding Node Operator deploying from Christchurch and represented in the Founding Node Operator Agreement by Gareth Evans, Chief Executive Officer and Co-Founder.

AI Trust Council Founding Members

The governance of the protocol during the Founding Period is undertaken by the AI Trust Council, whose Founding Members are:

Seat	Member
Founding Chair	Joshua Baron
Founder Seat	Dinesh Mendhe
Founding Member	Ross Thorpe
Founding Member	Issa Nesheiwat
Ex Officio Observer	Dr. Sofia Martinez Gonzalez, in her capacity as President and Chief Executive Officer of The AI Lab Intelligence Unobscured, Inc.

Joshua Baron has accepted the responsibilities of the Founding Chair, the position with the principal responsibility for the integrity of Council proceedings during the Founding Period. Ross Thorpe and Issa Nesheiwat have accepted the responsibilities of independent Founding Members. Each has agreed to serve on the conditions set out in the Charter, including the conflict disclosure, recusal, and transparency requirements.

Members joining the Council

The Council is further joined, with accession in process during the Pre-Genesis Period, by the following Members whose seats take effect on confirmation in accordance with the Charter:

- Christopher Stott, Executive Chair, Lonestar Data Holdings Inc., United States.
- Gareth Evans, Chief Executive Officer and Co-Founder, Timpi, Auckland, New Zealand.
- Dr. Vladimer Kobayashi, PhD, Professor, University of the Philippines Mindanao, Philippines.

- Benjamin Wild, Founder and Chief Doodler, Ben Wild Studios, Sale, England, United Kingdom.
- Kathleen Hom, American Bar Association Cybersecurity Legal Task Force; practice in AI and cybersecurity law, United States.
- Ram Waman Ghonmode, Chief Business Officer, The AI Lab Intelligence Unobserved, Inc., India (Member with declared interest and defined recusal scope).

The AI Lab acknowledges the willingness of each of the foregoing to lend institutional credibility, sectoral expertise, and geographic representation to the Council during the Founding Period and into the Network Period.

The roster presented above reflects the cohort in accession as of the publication date of this whitepaper. It is the operational seat from which the Council expands, not its intended ceiling. The Charter contemplates further Independent Members, Constituency Representatives, and Joining Members across the constituencies the protocol serves: journalism, civil society, publishing, education, information security, public-sector technology, regulatory liaison, academia, and standards organizations. The pace at which the Council expands is set by the willingness of credible candidates to commit to the independence covenants and recusal regime defined in the Charter, and by the standing demand for the protocol the Council stewards. In the thirty days preceding the publication of Version 1.0 of this whitepaper (between mid-May 2026 and June 16, 2026), theailab.org recorded in excess of four million page visits, and the supporting video explainers recorded in excess of one million views on the YouTube platform. The Council reports this empirical evidence of stakeholder interest as the leading indicator of its capacity to recruit and seat credible representation across the categories of multi-stakeholder participation contemplated by Article 95 of the EU AI Act. Updated figures are published on a continuing basis at the public channels of The AI Lab and the AI Trust Council identified in Executive Summary Section ES.7.

AI Trust Advisory Board

In addition to the Council itself, The AI Lab acknowledges the contribution of the AI Trust Advisory Board, a body of subject-matter advisors that provides consultative input to the Council on specific workstreams (cryptography, journalism, public policy, education, regulation, civil society) without holding a seat on the Council and without directing protocol matters. Advisors are accountable to the Council Chair and serve at the pleasure of the Council. The constitution of the Advisory Board at the date of publication is:

- John DiVuolo, PMP, ITIL, CISSP, Director of Information Security, Rutgers University, United States. Advisor on information security and risk management.
- Dale Whittaker, Advisor and Principal Officer, US Education Research and Development, Gates Foundation, United States. Advisor on higher education, philanthropy, and equity-focused artificial intelligence.
- Nate Angell, Founder, Nudgital, formerly Director of Communications and Community at Creative Commons and Director of Marketing at Hypothesis, Portland, Oregon, United States. Advisor on open knowledge, community, and adoption.
- Ajit Dharmik, Director, 6Simplex Software Solutions Pvt Ltd, India. Advisor on software engineering, deployment infrastructure, and the operational realities of running a Founding Node from the Indian jurisdiction.

The Advisory Board complements the Council by carrying domain expertise to deliberations without diluting the Council’s decision rights, and by providing a continuing source of independent technical and institutional input across the Founding Period and into the Network Period.

AI Trust Ambassadors

The AI Lab further acknowledges the AI Trust Ambassadors, a public-advocacy body composed of individuals who carry the mission of the Trust Identity Protocol into the communities, regions, industries, and disciplines they serve. Ambassadors are not Members of the Council and do not steward the protocol; their role is outward-facing advocacy in conversations with creators, institutions, regulators, and platforms around the world. The Ambassador roster is established through nomination, review, and confirmation procedures defined in the Charter, and is expected to be populated during the Pre-Genesis Period and the early Founding Period as the public surface of the protocol matures. The volume of inbound interest in the protocol since the public launch, recorded in excess of four million page visits to theailab.org and in excess of one million YouTube views of supporting explainers in the thirty days preceding the publication of this whitepaper, is the leading indicator of how rapidly the Ambassador roster can be populated with credible regional and industry champions.

Counsel

The drafting of the canonical TIP Protocol Specification v5.0, of this whitepaper, of the Founding Node Operator Agreement, of the Verification Provider accreditation agreement, of the Service Level Agreement, of the Technical Requirements Specification, of TIPCL-1.0, of the AI Trust Council Charter, of the Trademark Usage Policy, and of the United States Copyright Office and United States Patent and Trademark Office filings was undertaken under the direction of The AI Lab’s Founder and Chairman. The AI Lab acknowledges, with gratitude, the legal review afforded to the documents by counsel in the United States, the United Kingdom, India, New Zealand, and other jurisdictions in which the protocol operates or will operate.

Standards organizations and supervisory authorities

The AI Lab acknowledges, with respect for the public function of each engaging body, the published standards of the National Institute of Standards and Technology, the World Wide Web Consortium, the Internet Engineering Task Force, the European Telecommunications Standards Institute, the International Organization for Standardization Joint Technical Committee 1 Subcommittee 27, the FIDO Alliance, and the Coalition for Content Provenance and Authenticity, on each of which the protocol relies. The AI Lab also acknowledges the regulatory frameworks of the European Union, the United States, the United Kingdom, New Zealand, India, and other jurisdictions identified in Part IX, against which the protocol’s regulatory posture is designed.

Open source community

The reference implementations of the protocol depend on the open source ecosystem and on the work of contributors to OpenSSL, BoringSSL, AWS-LC, Mozilla NSS, Chromium, Mozilla Firefox, Apple WebKit, Node.js, Python, Rust, Go, the LLVM project, and many other projects

identified in the Software Bill of Materials of each reference implementation. The AI Lab acknowledges the open source community and is committed to its support through the Free Tier of TIPCL-1.0 and through the Apache License 2.0 conversion provision of TIPCL-1.0 Section 12.

Readers

The protocol exists because readers, in the present state of digital media, cannot reliably distinguish authentic from synthetic content, or reliably identify the party responsible for either. The protocol is offered as an instrument of the reader. The AI Lab acknowledges the reader as the protocol's principal beneficiary and as the audience to whom the work is, in the end, addressed.

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States. June 2, 2026.

Executive Summary

This whitepaper describes the Trust Identity Protocol™ (TIP), a federated, post-quantum, cryptographically verifiable framework for binding authenticated identities to authenticated content, published by The AI Lab Intelligence Unobscured, Inc., a Delaware corporation, with the support of the AI Trust Council, an independent multi-stakeholder governance body. The federated network is scheduled to commence live operation on the Genesis Date, a date in June 2026 to be determined by the sole director of The AI Lab on the satisfaction of the Operational Readiness Conditions described in Part XI and published not less than three business days in advance. Seven Founding Node Operators in four jurisdictions are contracted under signed Founding Node Operator Agreements, and three further Founding Node Operators are in accession, to operate nodes of the network from Genesis through the Founding Period.

ES.1 The condition the protocol addresses

By 2026, the cost of producing synthetic audio, image, video, and text content of a quality sufficient to deceive an attentive human observer has fallen to a fraction of the cost of producing equivalent content by traditional means. The consequences extend across journalism, civic discourse, the administration of justice, financial markets, and personal life. The principal responses available before the publication of this whitepaper, including platform-controlled trust signals, existing content provenance standards (C2PA and the Content Authenticity Initiative), and synthetic content watermarking, are each valuable. None resolves the structural problem of binding an authenticated identity to an authenticated piece of content, recorded in a federated and jurisdictionally portable manner that neither party may repudiate at a later time, using cryptography designed for durability against the cryptographic conditions of the next several decades.

ES.2 The Trust Identity Protocol

The protocol comprises three architectural layers.

TIP-ID is the identity layer. A Cryptographic Trust Identity (CTID) is a pseudonymous identifier derived from a public cryptographic key, issued by a Verification Provider accredited by the AI Trust Council. The private key is generated, stored, and operated on a WebAuthn resident key authenticator on the holder’s device, optionally bound to the holder by a biometric verification gesture that does not transmit biometric data to the Verification Provider. The CTID is portable across platforms, surfaces, and jurisdictions.

TIP-CONTENT is the provenance layer. The Canonical Normalization Algorithm Version 2.2 (CNA-2.2) produces a deterministic, content-defined fingerprint of a unit of content (text, image, audio, video, or composite). The fingerprint is signed by the holder of a CTID using a post-quantum signature scheme, producing a TIP-CONTENT record. The record carries an Origin Code (OH, AA, AG, MX) identifying the provenance category of the content, supplying the machine-readable marking contemplated by Article 50(2) of Regulation (EU) 2024/1689 (the EU AI Act).

TIP-TRUST is the reputation layer. The Trust Score is the aggregate of four sub-scores (Cryptographic, Behavioral, Adjudicated, Network) computed by deterministic rules from the history of events associated with a CTID and with the issuing Verification Provider. Adverse adjudications and Blocking Items B1 through B6 reduce the score. The reader-facing Global Seal of Trust supplies the human-readable disclosure contemplated by Article 50(4) of the EU AI Act.

ES.3 The federated network

Signed events are recorded on an append-only directed acyclic graph (DAG) maintained by a federation of Node Operators. The append-only property supplies non-repudiation, auditability, and alignment with regulatory expectations for the public record of disclosures. The graph is replicated across Node Operators in multiple jurisdictions. The Founding Node Operators, contracted as of the publication date of this whitepaper, are: THE PRESCIENT PACHYDERM LTD (United Kingdom), AZLogics Private Limited (India), Apex Modular Solutions LLC (United States), Timpi International Ltd (New Zealand), Lonestar Data Holdings Inc. (United States), and The AI Lab Intelligence Unobscured, Inc. (United States, theailab.org).

Identity issuance is performed by Verification Providers accredited by the AI Trust Council. Verification Providers operate under an annual fee, an annual independent audit, a published warrant canary, and a jurisdiction declaration. The economic terms under which Verification Providers will operate beyond the public-facing commitments are being designed and are not part of the public launch of the protocol.

ES.4 Governance, the AI Trust Council

The AI Trust Council is the independent multi-stakeholder body that ratifies protocol amendments, accredits Verification Providers, supervises the rules governing Trust Scores and Blocking Items, conducts dispute appellate review, and engages with regulators and standards organizations. The Council operates under a written Charter ratified by the sole director of The AI Lab and published at theailab.org/ai-trust-council. The Council’s Members are:

Seat	Member	Notes
Founding Chair	Joshua Baron	Independent Member
Founder Seat	Dinesh Mendhe	Founder and creator of the Trust Identity Protocol, the AI Trust Council, the AI Trust Registry, the AI Trust ID, and the Human Trust ID system; Member with declared interest and defined recusal scope
Founding Member	Ross Thorpe	Independent Member
Founding Member	Issa Nesheiwat	Independent Member
Ex Officio Observer	Chief Executive Officer of The AI Lab Intelligence Unobscured, Inc.	Observer capacity

In addition, Joining Members are in accession during the Pre-Genesis Period; their seats take effect on confirmation in accordance with the Charter. The Founding Chair and the two independent Founding Members are independent of The AI Lab. The Charter contains independence covenants including a three-of-five constituency supermajority threshold for protocol amendments, mandatory conflict disclosure and recusal, transparent publication of minutes, an annual transparency report, and a prohibition on instruction from The AI Lab. The Council is structured to be the kind of multi-stakeholder body contemplated by Article 95 of the EU AI Act.

The Council membership presented above reflects the operational cohort as of the publication date of this whitepaper. It is not the intended ceiling. The Charter contemplates additional Independent Members, Constituency Representatives, and Joining Members across the constituencies the protocol serves: journalism, civil society, publishing, education, information security, public-sector technology, regulatory liaison, academia, and standards organizations. Public reception of the protocol since its announcement provides empirical evidence of the stakeholder interest the Council is designed to integrate over time. In the thirty days preceding the publication of Version 1.0 of this whitepaper (between mid-May 2026 and June 16, 2026), theailab.org recorded in excess of four million page visits, and the supporting video explainers recorded in excess of one million views on the YouTube platform. These figures are current at the date of publication; updated figures are published on a continuing basis at the public channels listed in Section ES.7. Membership growth is the intended trajectory of the Council, and the breadth of inbound interest is the leading indicator of how rapidly that trajectory can be populated with credible representation across the categories of stakeholder participation contemplated by Article 95(3) of the EU AI Act, which provides that voluntary codes of conduct may be drawn up by individual providers or deployers of AI systems or by organisations representing them or by both, including with the involvement of any interested stakeholders and their representative organisations, including civil society organisations and academia.

ES.5 Licensing

The canonical TIP Protocol Specification is published under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0). Implementations are licensed under the TIP Protocol Code License Version 1.0 (TIPCL-1.0). TIPCL-1.0 grants a no-fee Free Tier to (a) individuals at any income level and small businesses with annual gross revenue below US\$100,000, (b) non-profits, NGOs, and charities, (c) educational institutions, (d) government entities, (e) journalism organizations for editorial use, and (f) research and development within published ceilings. Commercial Licenses for parties not eligible for the Free Tier are issued on a uniform nine-tier schedule (Micro through Global) with annual fees ranging from US\$125 to US\$550,000 depending on annual gross revenue and on the World Bank income group of the licensee's country of principal place of business. For a licensee with annual gross revenue below US\$25,000,000, the fee for its tier is multiplied by 1.00 for a high-income economy, 0.50 for upper-middle, 0.25 for lower-middle, and 0.12 for low income, subject to a minimum of US\$125; at US\$25,000,000 of annual gross revenue and above, a single fee applies worldwide. The fee schedule is uniform across all licensees within each tier and income group, published in advance, and not negotiated on a per-licensee basis, consistent with fair, reasonable, and non-discriminatory licensing principles. TIPCL-1.0 contains a self-executing conversion to the Apache License Version 2.0 on January 1, 2031.

ES.6 Cryptographic foundation

The protocol's signature scheme is ML-DSA-65, specified by NIST FIPS 204. The protocol's key encapsulation mechanism is ML-KEM-768, specified by NIST FIPS 203. The protocol's hash-based signature fallback for long-term archival is SLH-DSA, specified by NIST FIPS 205. The selection of NIST post-quantum primitives at the genesis of the protocol, rather than as a future migration, is described in Part III.

ES.7 Operational status

The protocol is in the Pre-Genesis Period commencing June 1, 2026 and concluding on the Genesis Date. The Pre-Genesis Period is the pilot phase of the protocol: Full Nodes are deployed and operated against pilot data, and the protocol does not accept production traffic. The protocol transitions from the pilot phase to live production operation on the Genesis Date. As of the publication date of this whitepaper, seven Founding Node Operators in four jurisdictions are contracted under signed Founding Node Operator Agreements and standing up Full Node deployments in pre-Genesis operating mode, and three additional Founding Node Operators are in accession (Marist University, Rutgers University, and The Core / BOOM FactCheck); the AI Trust Council Charter has been ratified by the sole director of The AI Lab, and the Council was convened on the third day of May, 2026; the canonical TIP Protocol Specification is published on a public repository under United States Copyright Office Application No. 1-15175755931 (pending); and the supporting reference implementations (browser extension for Chrome, Firefox, Arc, and Safari; the <tip-badge> web component; the WordPress reference plugin; the mobile web application) are deployed. The Genesis Date is scheduled for June 2026 and will be published not less than three business days in advance on theailab.org. In the thirty days preceding the publication of Version 1.0 of this whitepaper (between mid-May 2026 and June 16, 2026), public reception of the protocol has provided early operational evidence of standing demand: theailab.org recorded in excess of four million page visits, and the supporting video

explainers recorded in excess of one million views on the YouTube platform. These figures are reported as operational signal of the protocol’s standing at the date of publication. Current figures are published on a continuing basis at the public channels listed below. The figures stated here will be presented with methodology and time-period detail in the inaugural AI Trust Council annual transparency report and in subsequent reports.

The public channels at which the figures stated above can be independently verified, and at which subsequent operational signal is published in real time, are the canonical web property of The AI Lab at theailab.org, The AI Lab on the YouTube platform at youtube.com/@theailaborg, The AI Lab on the LinkedIn platform at linkedin.com/company/the-ai-lab-org, and the AI Trust Council on the LinkedIn platform at linkedin.com/company/ai-trust-council. The maintenance of a public surface for the AI Trust Council distinct from the public surface of The AI Lab reflects the independence of the Council from any single applicant or operator and is consistent with the multi-stakeholder body framing of Article 95 of the EU AI Act. The audited record of these figures, and not the platform metrics themselves, is the canonical evidence record of the AI Trust Council.

The roadmap described in Part XI sets out the operational milestones planned for the 12-month, 24-month, and 36-month horizons that follow Genesis.

ES.8 Regulatory posture

The protocol is not an “AI system” within the meaning of Article 3(1) of the EU AI Act. The Trust Score is not a social scoring system within the meaning of Article 5(1)(c) of that Act. The protocol is positioned to supply the machine-readable marking required by Article 50(2) and the human-readable disclosure contemplated by Article 50(4). The protocol is designed to support compliance by Verification Providers and Node Operators with the General Data Protection Regulation, the Digital Services Act, eIDAS 2.0, the NIS2 Directive, the Cyber Resilience Act, the Council of Europe Framework Convention on AI, the United Kingdom Online Safety Act 2023, the New Zealand Privacy Act 2020, the Indian Digital Personal Data Protection Act 2023, the United States Federal Trade Commission Act Section 5, the NIST AI Risk Management Framework, Executive Order 14110, and the state-law mosaic in the United States, in each case to the extent applicable to the licensee’s use of the protocol. Part IX describes the protocol’s alignment with each instrument in detail.

ES.9 What this whitepaper invites

The AI Lab and the AI Trust Council invite:

1. Regulators and supervisory authorities in the European Union, the United States, the United Kingdom, New Zealand, India, and other jurisdictions to engage with the AI Trust Council concerning the conformance of the protocol with applicable law and the recognition of the Council under voluntary code of conduct frameworks where available.
2. Standards organizations including ISO/IEC JTC 1/SC 27, W3C, IETF, ETSI, NIST, IPTC, and C2PA to engage with the AI Trust Council concerning the development of international standards for content provenance and verifiable identity infrastructure.
3. Publishers, journalism organizations, and platforms to deploy the reference implementations and to publish content under TIP-CONTENT records, supplying readers with veri-

fiable provenance and identity signals.

4. Prospective Verification Providers in jurisdictions not yet represented in the Founding Period network to apply for accreditation through the procedure published at theailab.org/tip-verification-provider.
5. Prospective Founding Node Operators completing the Founding Period network to apply through the procedure published at theailab.org/founding-node.
6. Commercial licensees in any tier of the Commercial Tier Schedule to obtain a Commercial License through the procedure published at theailab.org/tip-license.

ES.10 Contact

Purpose	Address
Licensing and Commercial Implementation	licensing@theailab.org , theailab.org/tip-license
Founding Node Operator Applications	nodes@theailab.org , theailab.org/founding-node
AI Trust Council Membership theailab.org/ai-trust-council	council@theailab.org ,
Verification Provider Accreditation theailab.org/tip-verification-provider	licensing@theailab.org ,
Technical Inquiries	tip@theailab.org , theailab.org/tip-protocol
Standards Engagement	standards@theailab.org
Conformance Inquiries	compliance@theailab.org
General Counsel	legal@theailab.org
Press and Public Affairs	press@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part I: The Verification Problem

1.1 The condition the protocol addresses

By the close of the year 2025, the cost of producing synthetic audio, image, video, and text content of a quality sufficient to deceive an attentive human observer had fallen to a fraction of the cost of producing equivalent content by traditional means. Tools capable of generating photorealistic images, full-motion video of named natural persons, audio in the recognizable voice of named natural persons, and long-form text in the recognizable style of named natural persons were available to any person with consumer-grade hardware and a residential internet connection.

The consequences of this condition extend across every domain in which the public relies on the authenticity of a digital record. In journalism, the line between an unaltered photograph and a manipulated one is no longer apparent to the reader. In civic discourse, the question of whether a video of a public official saying a particular thing is genuine or fabricated cannot be answered by examining the video. In the administration of justice, the evidentiary weight of a digital

recording is undermined by the possibility of synthesis. In financial markets, the possibility of synthetic announcements from named officers of public companies has been demonstrated to move prices. In personal life, the production of synthetic content depicting natural persons in circumstances they did not consent to has become an ordinary, low-cost act.

Existing public and private responses to this condition fall into three categories: (a) technical efforts to embed provenance metadata in content at the moment of production or capture, (b) technical efforts to detect synthetic content by analysis of the content itself, and (c) administrative and regulatory measures imposing disclosure obligations on the producers and distributors of synthetic content. Each category has been valuable. None has, in the form available at the publication of this whitepaper, addressed the underlying structural problem: the absence of a federated, post-quantum, cryptographically verifiable, jurisdictionally portable system for binding an authenticated identity to an authenticated piece of content, recorded in a manner that cannot be repudiated by either party at a later time.

The Trust Identity Protocol is designed to address that structural problem.

1.2 The limits of platform-controlled trust signals

In the platform-centric architecture that emerged between 2004 and 2024, the principal trust signals available to a reader were the trust signals applied by the platform on which the content was encountered. A blue check on one platform, a yellow check on a second, a notation that a video had been reviewed by independent fact-checkers on a third. These signals had the merit of being immediate and intelligible to ordinary readers. They had three structural limitations.

First, the signals were defined and applied by the platform. A change of platform policy, a change of ownership, or a commercial decision concerning the meaning of a signal could alter the signal's meaning without the reader's knowledge. A signal that meant identity-verified in one quarter could come to mean identity-verified-and-paying-subscription in the next quarter without a visible change in the signal itself.

Second, the signals were not portable. A reader who left a platform took none of the platform's trust signals with the reader. A creator whose identity had been verified on one platform had to repeat the verification on every other platform. A signal applied on one surface had no force on any other surface.

Third, the signals were not cryptographically attached to the content. The platform's signal sat alongside the content in the platform's user interface. The content could be reproduced, captured by screenshot, redistributed by other means, in each case stripped of the platform's signal. The signal accompanied the experience of the content on the platform, not the content itself.

These three limitations are not accidents of design. They are structural consequences of the platform-centric architecture. A trust signal defined by the platform, applied by the platform, and visible only on the platform is, by construction, a property of the platform and not of the content.

1.3 The limits of existing provenance standards

A second category of response has emerged in the form of provenance standards designed to attach metadata to content at the moment of capture or production. The most widely deployed

of these is the specification developed by the Coalition for Content Provenance and Authenticity (C2PA), with substantial work also undertaken by the Content Authenticity Initiative led by Adobe and by several manufacturers of cameras and capture devices. These efforts have produced a deployable specification and a growing installed base of cameras and editing tools that produce conformant manifests.

The C2PA specification is a meaningful contribution to the provenance problem and the Trust Identity Protocol is designed to interoperate with it. The C2PA specification, in its present form, does not, however, resolve four properties on which a verifiable internet depends.

The identity of the signer. A C2PA manifest is signed by an entity holding a credential issued by a certificate authority recognized by the C2PA. The credential attests to the identity of the signing entity as known to that certificate authority. The C2PA specification does not, in itself, supply a federated registry of accredited identity verification providers, a standardized accreditation process for such providers, or a public registry of signing identities portable across platforms and jurisdictions.

The reputation of the signer. A C2PA manifest does not contain a reputation score associated with the signing entity. The reader of a content unit accompanied by a C2PA manifest can verify that the manifest is intact and that the signing entity is the entity it claims to be. The reader cannot, from the manifest alone, ascertain whether the signing entity has previously been the subject of suspension, revocation, or adverse adjudication by any institution.

The post-quantum posture. The cryptographic signature schemes used in C2PA manifests as deployed at the publication of this whitepaper rely on classical primitives (ECDSA and RSA) that are subject to attack by a sufficiently advanced quantum computer. Signatures generated under those primitives today are subject to a long-tail risk that they may be forgeable retroactively, at a future date, by an adversary holding a quantum computing capability not presently possessed by any known actor but believed by the United States National Institute of Standards and Technology and by analogous bodies in the European Union, the United Kingdom, and elsewhere to be plausibly achievable within the working life of records being signed today.

The append-only public record. A C2PA manifest is a record of provenance accompanying a content unit. It is not, in itself, a public, append-only, distributed record of the population of signed content units. The C2PA specification does not contemplate, and does not require, the publication of signed content fingerprints to a public, append-only ledger from which the population of all signed content can be inspected by any interested party.

The Trust Identity Protocol is designed to supply these four properties as additions to the existing provenance ecosystem, not as replacements. A publisher or a creator who has invested in a C2PA workflow may continue that workflow, generate a C2PA manifest, and additionally register the content under the Trust Identity Protocol, thereby gaining the four properties listed above while preserving interoperability with C2PA readers.

1.4 The limits of watermarking

A third response is the embedding of imperceptible signals within content (steganographic or perceptual watermarks) that identify the content as artificially generated. Watermarking is an active area of research and is the subject of mandates emerging in jurisdictions including the European Union, the United States, and the People's Republic of China.

Watermarking is a useful tool. It is not, taken alone, an answer to the verification problem. The principal limitations of watermarking, as observed in the literature published between 2022 and 2025, are: (a) the susceptibility of imperceptible watermarks to removal or degradation by routine processing, including resizing, recompression, and modest editing; (b) the adversarial dynamic in which the development of more robust watermarks is followed by the development of more capable removers; (c) the asymmetry that a producer of synthetic content with adversarial intent is the party least likely to embed a recoverable watermark; and (d) the structural inability of watermark-only systems to address the question of who is responsible for a piece of content, addressing only the question of how the content was produced.

The Trust Identity Protocol is complementary to watermarking. The Origin Code system described in Part V (codes OH, AA, AG, MX) is designed to interoperate with watermark detectors. A reader of content accompanied by a watermark indicating AI generation and a TIP-CONTENT record signed by an accredited Verification Provider can rely on both signals. The protocol does not require watermarking, does not provide watermarking, and does not displace the work of organizations developing watermarking technology.

1.5 The identity gap

Provenance standards address the question: where did this content come from. Watermarking addresses the question: how was this content produced. Neither, in present form, addresses a third question on which the verifiable internet depends: who is responsible for this content.

The Trust Identity Protocol places that third question at its center. A unit of content bearing a TIP-CONTENT record is bound, by cryptographic signature, to a CTID. The CTID is in turn issued by a Verification Provider, accredited by an independent multi-stakeholder governance body, operating under audit, transparency, and warrant canary commitments described in Part X. The reader, the platform, the journalist, the regulator, and the court may, from a TIP-CONTENT record, ascertain (a) that a particular CTID signed the content, (b) that the CTID has not been suspended or revoked, (c) the reputation score of the signing CTID, and (d) the jurisdiction under whose law the CTID was issued.

The reader does not, from the TIP-CONTENT record alone, obtain the real-world identity of the signer. The CTID is a pseudonymous identifier. The mapping between the CTID and the real-world identity of the holder is maintained by the Verification Provider, subject to the disclosure obligations imposed by applicable law and by the legal process of the Verification Provider's jurisdiction. The protocol therefore supplies accountability without surveillance: a CTID can be held to account through the suspension, revocation, and adjudication mechanisms of the protocol; the real-world identity of the holder is disclosed only through the legal process of an identified jurisdiction.

1.6 Requirements for a durable solution

From the analysis in Sections 1.1 through 1.5, eight requirements emerge that any durable solution to the verification problem must satisfy. The Trust Identity Protocol is designed against these requirements; the design of the protocol in Parts II through VIII is best understood as the protocol's response to them.

1. **Decentralization.** No single platform, no single corporation, no single jurisdiction holds the trust signal. Identity verification is distributed across a federation of accredited Ver-

ification Providers in multiple jurisdictions. The record of signed events is distributed across a federation of Node Operators in multiple jurisdictions.

2. **Independence of governance.** The body that ratifies protocol changes, accredits Verification Providers, and supervises enforcement is institutionally independent of any single commercial operator, including the publisher of the protocol.
3. **Cryptographic verifiability.** Trust signals are bound to content by cryptographic signature. A reader, with software conformant to the protocol, can verify the signal without trusting the platform on which the content is encountered.
4. **Post-quantum durability.** Cryptographic primitives are selected from the standards published by the United States National Institute of Standards and Technology for resistance to attack by quantum computing (FIPS 203, FIPS 204, FIPS 205). Signatures generated today are designed to remain verifiable and unforgeable across the working life of the records being signed.
5. **Pseudonymity with accountability.** Identity verification produces a pseudonymous identifier. The pseudonymous identifier may be held to account by suspension, revocation, and adjudication within the protocol. Disclosure of the real-world identity of the holder occurs only through the legal process of the jurisdiction in which the issuing Verification Provider operates.
6. **Portability.** Trust signals are portable across platforms, surfaces, and jurisdictions. A CTID issued in one jurisdiction is recognized by readers, platforms, publishers, and regulators in other jurisdictions, subject to the legal interoperability frameworks established in Part IX.
7. **Append-only public record.** Signed content fingerprints are recorded on a public, append-only directed acyclic graph maintained by the federation of Node Operators. Any interested party may inspect the population of signed records and may obtain a verifiable history of the suspensions, revocations, and adjudications that have affected a given CTID.
8. **Regulatory alignment.** The protocol is designed to support the disclosure, transparency, and provenance obligations imposed by applicable law on producers, distributors, and platforms (including Articles 50(2) and 50(4) of the EU AI Act, Section 12 of the United Kingdom Online Safety Act, the California AI Transparency Act, and analogous instruments in other jurisdictions identified in Part IX).

1.7 Comparative analysis

The following table summarizes the principal differences between the Trust Identity Protocol and other systems addressing aspects of the verification problem. The comparison is descriptive, not evaluative; each of the compared systems is the product of substantial work and serves purposes for which it may be better suited than the Trust Identity Protocol.

In the table below, columns labelled “Adobe CAI”, “Bluesky AT”, and “Platform-native” refer respectively to the Adobe Content Authenticity Initiative, the Bluesky AT Protocol identity layer, and platform-native trust signals operated by individual platforms. Cell abbreviations: VPs = Verification Providers; PQC = NIST post-quantum cryptography (FIPS 203, 204, 205);

DAG = the protocol’s append-only directed acyclic graph; Council = the AI Trust Council; Seal = the Global Seal of Trust; DID = decentralized identifier; membership = C2PA membership.

Property	TIP	C2PA	Adobe CAI	Bluesky AT	Platform- native
Cryptographic content binding	Yes (CNA-2.2)	Yes (manifest)	Yes (manifest)	Partial	No
Federated identity issuance	Yes (VPs)	Partial	Partial	Yes (DID)	No
Pseudonymous identifier	Yes (CTID)	No	No	Yes (DID)	No
Post-quantum primitives	Yes (PQC)	No	No	No	No
Public append-only ledger of signed events	Yes (DAG)	No	No	Partial	No
Reputation scoring layer	Yes (Trust)	No	No	Partial	Yes
Independent governance body	Yes (Council)	Yes (member)	No	Yes (Bluesky)	No
Cross-jurisdiction portability	Yes	Yes	Yes	Yes	No
EU AI Act Article 50 marking support	Yes	Partial	Partial	No	No
EU AI Act Article 50 disclosure support	Yes (Seal)	No	No	No	Yes (variable)

1.8 Scope of this whitepaper

This whitepaper addresses the design and the public-policy positioning of the Trust Identity Protocol. The detailed technical specification of the protocol is set out in the canonical TIP Protocol Specification, available at theailab.org, USCO Application No. 1-15175755931 (pending). Parts II through VIII of this whitepaper describe the protocol at a level of detail sufficient for a regulator, a chief technology officer, a general counsel, or a principal engineer to evaluate the protocol’s properties, classification, and regulatory posture. A reader requiring the byte-level specification of any component is referred to the canonical Specification.

Contact The AI Lab regarding Part I

Technical Inquiries: tip@theailab.org Press and Public Affairs: press@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part II: Design Principles

This Part sets out the design principles from which the technical layers described in Parts III through VIII are derived. The principles are intended to be intelligible to a reader who is not a

cryptographer or a distributed systems engineer, and to be sufficient for a regulator or a general counsel to assess the protocol’s structural character against the principles of the regulatory landscape described in Part IX.

2.1 Decentralized identity, federated verification

The identity layer of the Trust Identity Protocol is decentralized in the strict sense that no single entity, including The AI Lab, holds the universe of identifiers, controls the issuance of identifiers, or has the power to deactivate identifiers unilaterally. Identifiers are issued by a federation of Verification Providers operating in identified jurisdictions under accreditations granted by the AI Trust Council on the recommendation of The AI Lab.

The federation is not a peer-to-peer system in which any party may issue identifiers. The federation is a credentialed network in which a defined population of accredited Verification Providers issues identifiers under published criteria. This design choice has three consequences.

First, accountability is allocated to identified institutions. A Verification Provider that operates poorly is identifiable and subject to suspension or revocation of accreditation by the Council. A peer-to-peer system in which any party may issue identifiers cannot, by construction, hold any party accountable for the consequences of issuance.

Second, the protocol is compatible with the regulatory expectations of identified jurisdictions. A regulator in a jurisdiction in which a Verification Provider operates may, through the legal process of that jurisdiction, compel disclosure of the mapping between a CTID and the real-world identity of the holder, on the conditions specified by that jurisdiction’s law. A peer-to-peer system in which any party may issue identifiers presents no analogous interface to legal process.

Third, the AI Trust Council and not any single Verification Provider is the locus of governance. A Verification Provider that disagrees with a Council decision may resign from accreditation or may, through the Council’s dissent procedure, record its disagreement. It does not have the power to fork the protocol or to issue identifiers outside the accreditation regime.

2.2 Post-quantum cryptography from genesis

The Trust Identity Protocol is the first widely deployed content provenance and verifiable identity framework to be designed against the cryptographic standards published by the United States National Institute of Standards and Technology in August 2024 for post-quantum cryptography. The protocol’s signature scheme is ML-DSA-65 (FIPS 204). The protocol’s key encapsulation mechanism is ML-KEM-768 (FIPS 203). The protocol’s hash-based signature fallback for long-term archival is SLH-DSA (FIPS 205).

The selection of post-quantum primitives at the genesis of the protocol, rather than as a future migration, reflects three considerations.

First, signatures generated today are expected to remain in service for decades. A signature applied to a journalistic photograph in 2026, recorded on a public append-only ledger, and relied on by a regulator or a court in 2046 must be verifiable in 2046 by the cryptographic standards then in force. The cryptographic guidance of NIST, the European Union Agency for Cybersecurity (ENISA), and analogous bodies converged between 2022 and 2024 on the position that

classical primitives applied today carry a non-negligible long-tail risk of retroactive forgeability and that systems being designed in 2025 and 2026 should be designed against post-quantum primitives directly.

Second, the cost of a migration from classical primitives to post-quantum primitives, undertaken after a federated network is in operation with a large installed base of signed records, is substantially higher than the cost of designing against post-quantum primitives at the outset. The cost of compatibility analysis, the cost of dual-signing periods, and the cost of operator coordination across a federation are all reduced by avoiding the migration.

Third, the post-quantum primitives selected by NIST have been the subject of substantial cryptanalytic scrutiny over the multi-year standardization process and are believed by the public cryptographic research community to be sound. The decision to deploy them now is not an experimental choice; it is the application of standards reviewed by the United States federal government and recommended for adoption.

Part III describes the post-quantum primitive selection in detail.

2.3 Append-only directed acyclic graph for non-repudiable provenance

The record of signed events maintained by the federation of Node Operators is an append-only directed acyclic graph. Entries are added to the graph by Node Operators in accordance with the rules described in Part VII. Entries are not removed from the graph. The append-only property is structural to the protocol.

The append-only property is selected for three reasons.

First, non-repudiation. A signature that has been published to a public append-only graph cannot subsequently be retracted by the signer with the effect of denying that the signature existed. The signer may revoke the underlying CTID, terminating its future utility; the historical record that the signer signed the content at the recorded time remains. This property is essential to the evidentiary use of signed content in journalism, in administrative proceedings, and in litigation.

Second, auditability. A regulator, an academic researcher, an investigative journalist, or an interested member of the public may, at any time, inspect the population of signed records, the suspension and revocation events affecting particular CTIDs, and the patterns of activity of particular Verification Providers. The transparency of the federation is structural and does not depend on the discretion of any single operator.

Third, alignment with regulatory expectations. Where applicable law imposes obligations on platforms or publishers to disclose the provenance of artificially generated content (EU AI Act Article 50, California SB 942, analogous instruments), the existence of a public append-only record of such disclosures supplies a verifiable basis for compliance. The platform or publisher does not need to depend on its own record-keeping; the record exists independently on the federation.

The tension between the append-only property and the right to erasure under data protection law is addressed by the pseudonymization architecture described in Section 9.1.2 and in Part IV. The data published to the graph is a pseudonymous identifier (the CTID) and a content fingerprint (the CNA-2.2 output). The mapping between the CTID and the real-world identity of the

holder is maintained off-graph by the Verification Provider. On a verified erasure request, the off-graph mapping is deleted; the on-graph pseudonymous record remains as an orphan and is not personal data within the meaning of the General Data Protection Regulation as clarified by Recital 26 of that Regulation and by the guidance of the European Data Protection Board on pseudonymization.

2.4 Pseudonymity with accountability

The Trust Identity Protocol does not require any natural person to disclose their real-world identity to The AI Lab, to any Node Operator, to any platform, to any reader of content, or to any other participant in the protocol other than the Verification Provider that issues the CTID. The Verification Provider holds the mapping between the CTID and the real-world identity. The mapping is disclosed only through the legal process of the jurisdiction in which the Verification Provider operates, on the conditions specified by that jurisdiction's law.

This design choice addresses two competing concerns that have characterized the public discussion of digital identity systems over the past decade. The first concern is that identity systems that require real-world identification at every point of use produce a surveillance architecture in which every act of speech, commerce, or association can be linked to a named natural person, undermining the conditions of a free society. The second concern is that systems that do not require real-world identification at any point produce a domain of unaccountable speech and conduct in which fraud, harassment, defamation, and the manipulation of democratic processes are not effectively addressed by the institutions that ordinarily address them.

The protocol's response is to separate the question of accountability from the question of disclosure. A CTID is held to account within the protocol by the suspension, revocation, and adjudication mechanisms of Part VI. A CTID's holder is identified outside the protocol only through legal process. The Verification Provider that issued the CTID is the institution to which legal process is directed. The Verification Provider operates under a published jurisdiction declaration, a warrant canary, and an independent audit, each described in Part X.

The result is a protocol in which a journalist may publish under a CTID without disclosing the journalist's real-world identity to a hostile government, in which an ordinary user may comment under a CTID without disclosing the user's real-world identity to a platform or to other users, and in which a court of competent jurisdiction may, on a showing of cause, obtain the real-world identity of a CTID holder through the legal process of the jurisdiction in which the issuing Verification Provider is established.

2.5 Adversarial robustness

The Trust Identity Protocol is designed against an explicit adversary model. The principal adversaries against which the protocol is designed are described below.

2.5.1 Sybil resistance

A Sybil adversary is one that obtains many identifiers and uses them to simulate the support of many distinct natural persons. The protocol's resistance to Sybil attack rests on three properties: (a) CTIDs are issued by accredited Verification Providers operating under identity verification practices subject to audit, not by any party that requests one; (b) the Trust Score architecture ag-

gregates behavior across multiple sub-scores, including the Network sub-score, which weighs the established history and reputation of the issuing Verification Provider; and (c) Blocking Items B1 through B6, defined in Part VI, suspend the operational utility of a CTID on detection of behavior characteristic of Sybil attack.

2.5.2 Key compromise containment

A key compromise adversary is one that obtains the private key of a CTID holder and uses it to sign content as if it were the holder. The protocol contains key compromise through four mechanisms: (a) private keys are generated, stored, and operated on in a WebAuthn resident key authenticator on the holder's device, subject to biometric or other user verification; (b) CTIDs may be revoked by the holder or by the issuing Verification Provider; revocation is published to the append-only DAG within the synchronization interval described in Part VII; (c) signatures published on the DAG carry a timestamp recognized by the federation, permitting the identification of signatures published after the revocation timestamp as suspect; and (d) the Trust Score reflects revocation events through the Behavioral sub-score, reducing the operational utility of a CTID known to have been compromised.

2.5.3 Regulatory and jurisdictional pressure

A regulatory adversary is a state or political actor that seeks to disable the protocol within its jurisdiction or to compel the production of identifying information by a Verification Provider operating within its jurisdiction. The protocol's response is structural rather than confrontational. The federation operates across multiple jurisdictions; the suspension of operations by a Verification Provider in one jurisdiction does not disable the protocol in other jurisdictions. The jurisdiction declaration and the warrant canary published by each Verification Provider give users notice of the jurisdictional risk associated with a given Verification Provider, permitting users to select a Verification Provider on a basis other than convenience. The protocol does not seek to exempt Verification Providers from the law of their jurisdictions; it seeks to make the operation of that law transparent.

2.5.4 Cryptographic adversary

A cryptographic adversary is one with computational capabilities sufficient to forge signatures, to invert hash functions, or to break key encapsulation schemes. The protocol's response is the selection of NIST post-quantum primitives described in Section 2.2 and detailed in Part III, the use of a three-hash content addressing scheme described in Part V, and the conservative parameter selection described in Part III.

2.6 Minimality

The Trust Identity Protocol does not seek to be a general-purpose distributed ledger, a payment system, a programmable smart contract platform, a content delivery network, or a recommendation system. The protocol is designed to perform a narrow set of operations: the issuance of CTIDs by Verification Providers; the signing of CNA-2.2 content fingerprints by CTID holders; the recording of signed events on an append-only DAG; the computation of Trust Scores from defined inputs; and the publication of suspension, revocation, and adjudication events.

Minimality is selected because the security properties of the protocol can be analyzed and defended in proportion to the simplicity of the operations the protocol performs. A protocol that performs many operations exposes many surfaces to adversaries and presents many points at which regulatory analysis must be conducted. A protocol that performs few operations may, against the same set of adversaries and the same regulatory landscape, be analyzed and defended more rigorously.

The protocol's minimality is not a limitation imposed by lack of ambition. It is a deliberate scoping choice. Applications that require additional capabilities (content recommendation, content moderation, payment, social interaction) may be built on top of the protocol by parties electing to do so. Such applications operate in their own legal and operational frameworks. The Trust Identity Protocol supplies the foundation of authenticated identity and authenticated content provenance on which such applications may rely.

2.7 Interoperability

The Trust Identity Protocol is designed to interoperate with existing provenance standards (C2PA), existing identity standards (Verifiable Credentials, Decentralized Identifiers, WebAuthn), existing cryptographic infrastructure (X.509, PKCS), and existing publishing infrastructure (HTML, RSS, OpenGraph, Schema.org). The protocol's web component (`<tip-badge>`) is implemented in standard HTML and JavaScript and is deployable on any web page. The protocol's browser extension operates within the published extension APIs of Chrome, Firefox, Arc, and Safari. The protocol's WordPress reference plugin operates within the published plugin API of WordPress.

The interoperability principle is operationally important and is also a regulatory commitment. The protocol does not seek to displace investments made by publishers, platforms, or capture device manufacturers in existing provenance and identity infrastructure. The protocol adds the four properties identified in Section 1.3 (federated identity, reputation, post-quantum durability, append-only public record) as additions to that infrastructure.

2.8 Open specification, structured licensing

The canonical specification of the protocol is published under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0). Any person may read, redistribute, translate, comment on, or adapt the specification, subject to the attribution requirements of CC BY 4.0. Implementations of the protocol are licensed under the TIP Protocol Code License Version 1.0 (TIPCL-1.0) as described in Part X and Appendix F. The licensing structure is designed to maximize the diffusion of the specification while sustaining the institutional support necessary to operate the protocol during the Founding Period and the early Network Period. The Apache License 2.0 conversion provision on January 1, 2031 establishes a fixed horizon at which the implementation license becomes permissive open source.

The open specification principle is a regulatory commitment as well as a technical commitment. A regulator, a researcher, or a competitor may, at any time, inspect the specification, verify its conformance with published standards, and develop independent implementations. The protocol does not depend on proprietary secrets for its security properties. Such security properties as the protocol possesses derive from the cryptographic primitives, the architecture, and the institutional governance, all of which are disclosed.

Contact The AI Lab regarding Part II

Technical Inquiries: tip@theailab.org Standards Engagement: standards@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part III: Cryptographic Foundation

This Part describes the cryptographic primitives, modes of operation, and parameter selections from which the security properties of the Trust Identity Protocol derive. The description is normative and is intended to be sufficient for a chief technology officer, a principal engineer, or a cryptographer to assess the security posture of the protocol. The byte-level specification of each primitive is set out in the canonical TIP Protocol Specification.

3.1 Standards landscape and rationale for the NIST post-quantum suite

In August 2024, the United States National Institute of Standards and Technology published three Federal Information Processing Standards establishing the first government-recommended post-quantum cryptographic suite for general-purpose digital signature and key encapsulation. The three standards are FIPS 203 (Module-Lattice-Based Key-Encapsulation Mechanism Standard), FIPS 204 (Module-Lattice-Based Digital Signature Standard), and FIPS 205 (Stateless Hash-Based Digital Signature Standard). The Trust Identity Protocol selects its primary primitives from these three standards.

The selection rests on six considerations.

Public review. The primitives standardized by NIST in 2024 are the survivors of a multi-year, multi-round public standardization process in which candidate primitives were subjected to cryptanalytic review by the international research community. The selected primitives have been the subject of more public review than any other post-quantum candidates as of the publication of this whitepaper.

Government endorsement. The selection of the primitives by NIST has been followed by adoption guidance from the United States National Security Agency (Commercial National Security Algorithm Suite 2.0), the Bundesamt für Sicherheit in der Informationstechnik in Germany, the Agence nationale de la sécurité des systèmes d'information in France, the Centrum Wiskunde & Informatica in the Netherlands, the Canadian Centre for Cyber Security, and others. The European Telecommunications Standards Institute and the European Union Agency for Cybersecurity have published roadmaps incorporating the NIST primitives into European standards.

Long-tail durability. A signature applied to a unit of content today and recorded on the federated DAG is expected to remain verifiable for decades. The cryptographic guidance of NIST, ENISA, and analogous bodies converged between 2022 and 2024 on the position that classical primitives (ECDSA, RSA, Ed25519) applied today carry a non-negligible risk of retroactive forgeability against an adversary holding a quantum computing capability not presently possessed by any known actor but believed to be plausibly achievable within the working life of records being signed today.

Implementation maturity. Reference implementations of ML-KEM, ML-DSA, and SLH-DSA are available in C, Rust, and JavaScript and have been integrated into the cryptographic libraries of OpenSSL, BoringSSL, AWS-LC, and Mozilla NSS. The deployment of the NIST primitives at the publication of this whitepaper is not experimental.

Interoperability. The selection of NIST primitives aligns the protocol with the cryptographic posture toward which the standards-setting bodies of the European Union, the United Kingdom, Canada, Australia, New Zealand, Japan, and others are converging.

Migration avoidance. The protocol is designed against post-quantum primitives from the genesis of the network, avoiding a future migration from classical primitives to post-quantum primitives undertaken after a federated network with a large installed base is in operation.

3.2 ML-KEM-768 (FIPS 203) key encapsulation

ML-KEM, the Module-Lattice-Based Key-Encapsulation Mechanism, is the primary key encapsulation mechanism of the Trust Identity Protocol. The protocol uses parameter set ML-KEM-768, providing security strength equivalent to Category 3 in the NIST post-quantum security framework (approximately 192-bit symmetric security).

ML-KEM-768 is used by the protocol in three contexts.

Session key establishment for transport security. When a CTID holder, a Verification Provider, a Node Operator, or a reader establishes a TLS 1.3 session with another participant in the protocol, the session is established using a hybrid key agreement combining classical X25519 with ML-KEM-768. The hybrid approach preserves the security properties of the classical key agreement while adding the security properties of the post-quantum key agreement, against an adversary holding a quantum computing capability sufficient to break the classical key agreement.

Encryption of CTID provisioning material. When a Verification Provider provisions a CTID to a holder, sensitive provisioning material is encapsulated under ML-KEM-768 using a public key supplied by the holder's WebAuthn resident key authenticator. The encapsulated material is recoverable only by the holder.

Encryption of audit records. Audit records produced by a Verification Provider and transmitted to an auditor are encapsulated under ML-KEM-768 using a public key published by the auditor.

The parameter selection is conservative. ML-KEM-768 provides a margin of security against future improvements in lattice-based cryptanalysis without imposing a prohibitive bandwidth or computational cost on the protocol's principal operations.

3.3 ML-DSA-65 (FIPS 204) primary signature scheme

ML-DSA, the Module-Lattice-Based Digital Signature Algorithm, is the primary signature scheme of the Trust Identity Protocol. The protocol uses parameter set ML-DSA-65, providing security strength equivalent to Category 3 (approximately 192-bit classical security and approximately 128-bit quantum security).

ML-DSA-65 is used by the protocol in four contexts.

Signature of TIP-CONTENT records. A CTID holder signs the CNA-2.2 fingerprint of a unit of content using the ML-DSA-65 private key associated with the holder’s CTID. The signature, together with the fingerprint, the Origin Code, and the metadata defined in Part V, constitutes the TIP-CONTENT record. The signature is published to the federated DAG.

Signature of CTID issuance certificates. A Verification Provider, on issuing a CTID to a holder, signs an issuance certificate binding the CTID to the holder’s WebAuthn credential public key and to the issuance metadata. The signature is performed using the Verification Provider’s ML-DSA-65 private key. The signature is verifiable by any participant using the Verification Provider’s public key published in the Verification Provider Registry.

Signature of suspension, revocation, and adjudication events. Suspension, revocation, and adjudication events are signed by the issuing Verification Provider (for events affecting CTIDs the provider issued) or by the AI Trust Council (for events affecting Verification Provider accreditations). The signatures are published to the federated DAG.

Signature of Node Operator commitments. A Node Operator periodically signs a commitment to the state of the DAG segment it maintains, in accordance with the synchronization protocol described in Part VII. The commitment is published to the DAG and is the basis for the federation’s consistency protocol.

The selection of ML-DSA-65 over the higher-strength parameter set ML-DSA-87 reflects the cost-security tradeoff. ML-DSA-65 produces a signature of approximately 3,300 bytes and a public key of approximately 1,950 bytes, which is bandwidth-acceptable on the operational paths of the protocol. ML-DSA-87 produces signatures of approximately 4,600 bytes and is reserved for the long-term archival use case described in Section 3.4.

3.4 SLH-DSA (FIPS 205) hash-based signature fallback

SLH-DSA, the Stateless Hash-Based Digital Signature Algorithm, is included in the protocol as a fallback signature scheme for long-term archival and for the high-assurance signature of protocol-critical events.

The principal property of SLH-DSA, relative to ML-DSA, is that its security rests only on the security of the underlying hash function (SHA2-256 or SHAKE-256) and not on the conjectured hardness of any algebraic problem. In the event that a future cryptanalytic development materially reduces the security of lattice-based primitives, SLH-DSA remains secure subject only to the security of the hash function.

SLH-DSA is used by the protocol in two contexts.

Long-term archival signatures. A publisher or a journalism organization electing to retain signed content under cryptographic guarantee for the long term (decades to a century) may opt to apply a SLH-DSA signature in addition to the ML-DSA-65 signature. The dual signature is published to the federated DAG. The SLH-DSA signature preserves verifiability under a wider range of cryptographic futures.

High-assurance protocol events. Certain protocol-critical events, including the genesis of a Verification Provider accreditation and the activation of a new Charter version of the AI Trust Council, are signed using SLH-DSA in addition to ML-DSA-65. The dual signature is published to the federated DAG.

The bandwidth cost of SLH-DSA (signatures of approximately 16,000 to 50,000 bytes depending on parameter selection) makes it unsuitable for the per-content signature path. The protocol uses parameter set SLH-DSA-SHA2-192s for the optional contexts described above, providing approximately 192-bit classical security with a signature size of approximately 35,000 bytes.

3.5 PRF-to-AES key protection chain

Symmetric keys used by the protocol for the encryption of provisioning material and audit records are derived through a pseudorandom function chain conformant with NIST SP 800-108 (Recommendation for Key Derivation Using Pseudorandom Functions) and applied through AES-256 in Galois/Counter Mode (AES-256-GCM).

The chain is structured as follows. The output of the ML-KEM-768 decapsulation is passed through HKDF-SHA-512, with an explicit context binding identifying the protocol context (CTID provisioning, audit record encryption, or session key derivation), an explicit version identifier of the protocol, and a salt derived from the message identifier. The output of HKDF-SHA-512 is used as the input keying material for AES-256-GCM.

The PRF-to-AES chain serves three purposes. First, the chain isolates the symmetric key derived for one context (CTID provisioning) from the symmetric key derived for another context (audit record encryption) even where the underlying ML-KEM-768 shared secret is the same. Second, the chain provides an explicit binding to the protocol version, supporting forward-secure operation across protocol version changes. Third, the chain provides nonce-misuse resistance through the salt derivation pattern.

3.6 Three-hash content addressing

A content unit signed under the Trust Identity Protocol is addressed by a triple of hash values computed under three distinct hash functions. The triple, taken together, constitutes the content identifier referenced by the TIP-CONTENT record.

The three hash functions are:

SHA2-256. The first member of the triple is the SHA2-256 hash of the canonical normalization output described in Part V. SHA2-256 is selected for breadth of platform support and is the principal identifier in most operational paths.

SHA3-256. The second member of the triple is the SHA3-256 hash of the canonical normalization output. SHA3-256 is based on the Keccak sponge construction and provides cryptanalytic diversity against developments that may reduce the security of SHA2-256 without affecting SHA3-256.

BLAKE3. The third member of the triple is the BLAKE3 hash of the canonical normalization output, computed at the default output length of 256 bits. BLAKE3 is selected for performance and is the principal hash used for high-throughput operational paths including the verification of bulk-imported archives.

The three-hash addressing serves three purposes. First, it provides defense in depth against cryptanalytic developments against any single hash function. A second-preimage attack on SHA2-256 that does not extend to SHA3-256 or to BLAKE3 does not enable an adversary to substitute a content unit under the same content identifier. Second, it provides interoperability

with the published address spaces of other systems (SHA2-256 is the address used by IPFS and by most existing distributed object stores; SHA3-256 is the address used by certain Ethereum-derivative systems; BLAKE3 is the address used by emerging high-throughput content stores). Third, it provides a robust input to the canonical content identifier used in CNA-2.2 normalization.

3.7 Threat model and formal security claims

The Trust Identity Protocol is designed against the threat model below. The model identifies the adversaries against which the protocol’s cryptographic guarantees are designed to hold and the adversaries against which the protocol’s guarantees are conditional on assumptions documented herein.

3.7.1 Adversaries the protocol contains

A polynomially-bounded classical adversary. An adversary with classical computing resources cannot, under the protocol, produce a signature that verifies under the public key of a CTID holder, a Verification Provider, a Node Operator, or the AI Trust Council, without possession of the corresponding private key. The claim rests on the conjectured hardness of the Module Learning With Errors problem and the Module Short Integer Solution problem underlying ML-DSA-65, and on the conjectured pseudorandomness of SHA2-256, SHA3-256, BLAKE3, and SHAKE-256.

A quantum-bounded adversary against present primitives. An adversary with a fault-tolerant quantum computer of the size and quality presently demonstrated cannot, under the protocol, forge signatures or recover encapsulated session keys, provided the adversary cannot solve the Module Learning With Errors problem or the Module Short Integer Solution problem at the parameter set sizes used. ML-KEM-768 and ML-DSA-65 are believed to remain secure against such an adversary.

A federation-level adversary controlling a minority of Node Operators. An adversary controlling fewer than the consensus threshold of Node Operators cannot, under the protocol, cause an invalid TIP-CONTENT record to be accepted by the federation. The consistency property of the federated DAG is preserved through the signature, timestamp, and reference structure of DAG entries.

A Verification Provider operating in a single jurisdiction. An adversary obtaining the cooperation or the compelled service of a single Verification Provider may obtain the mapping between particular CTIDs and the real-world identities of their holders, but cannot, under the protocol, suspend the operation of the federation outside the cooperating Verification Provider’s jurisdiction, cannot revoke CTIDs issued by other Verification Providers, and cannot manipulate Trust Scores beyond the deterministic effect on Trust Scores of the cooperating provider’s accreditation status. The warrant canary and the audit obligation supply notice of such cooperation to the extent permitted by the cooperating jurisdiction’s law.

3.7.2 Adversaries the protocol does not contain

An adversary controlling the holder’s device. An adversary that has compromised the WebAuthn resident key authenticator on the holder’s device may, for the duration of the compromise,

produce signatures under the holder's CTID indistinguishable from signatures produced by the holder. The protocol mitigates this exposure through the revocation mechanism described in Part IV and through the Trust Score updates described in Part VI, but does not eliminate it.

An adversary obtaining a future cryptanalytic advance. A cryptanalytic advance materially reducing the security of the Module Learning With Errors problem, the Module Short Integer Solution problem, or the underlying hash functions may reduce the security of the protocol. The protocol's response to such an advance is the migration path described in Section 3.7.3.

A regulatory adversary acting on the entire federation. A coordinated regulatory action across all jurisdictions in which Verification Providers and Node Operators operate may suspend the protocol's operation. The protocol's response to such an action is institutional (the AI Trust Council Charter, the published jurisdiction declarations, the standards organization engagement) rather than cryptographic.

3.7.3 Cryptographic migration path

The protocol contains a defined migration path against the possibility of a cryptographic compromise of the primitives selected in this Part. The migration path comprises four stages.

Stage 1: monitoring. The AI Trust Council maintains a continuous review of the public cryptanalytic literature applicable to the primitives identified in this Part. The Council publishes an annual cryptographic posture report.

Stage 2: advisory. On identification of a material reduction in the security of any primitive, the Council issues an advisory identifying the affected primitive and recommending operational measures.

Stage 3: supplementary signing. On a determination that a primitive's security has fallen below a defined threshold, the Council requires the dual signing of subsequent TIP-CONTENT records under a successor primitive identified by the Council. The successor primitive is selected from the NIST post-quantum suite, the IETF post-quantum suite, or a standardized successor at the time of the determination.

Stage 4: re-signing of long-term archives. Publishers and journalism organizations electing to maintain long-term archives may opt into a re-signing service operated by the Council, under which existing TIP-CONTENT records are dual-signed under the successor primitive and the dual signatures are published to the federated DAG.

The migration path is structured so that the protocol can respond to a cryptanalytic event without disruption of the federated network and without invalidation of the historical record.

3.8 Conformance testing

The canonical TIP Protocol Specification publishes test vectors for ML-KEM-768, ML-DSA-65, SLH-DSA-SHA2-192s, HKDF-SHA-512, AES-256-GCM, SHA2-256, SHA3-256, and BLAKE3. The test vectors are conformance test vectors derived from the published test vectors of NIST, NIST CAVP, the IETF, and the BLAKE3 reference implementation. A conformant implementation of the protocol is required to pass the published test vectors.

The AI Trust Council maintains a Conformance Registry identifying implementations that have passed the test vectors and have submitted to the conformance review process. Inclusion in

the Conformance Registry is voluntary and is not a condition of TIPCL-1.0 licensing. Inclusion supplies a public attestation of conformance suitable for reliance by regulators, auditors, and counterparties.

Contact The AI Lab regarding Part III

Technical Inquiries: tip@theailab.org Conformance Inquiries: compliance@theailab.org
Standards Engagement: standards@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part IV: TIP-ID: Identity Layer

This Part describes the identity layer of the Trust Identity Protocol. The identity layer comprises the Cryptographic Trust Identity (CTID), the institutional role of the Verification Provider, the issuance and revocation lifecycle, the optional biometric binding implemented through the WebAuthn resident key authenticator, the portability of identities across jurisdictions, and the dispute and governance mechanisms applicable to identity matters.

4.1 Cryptographic Trust Identity (CTID) construction

A CTID is a pseudonymous identifier derived from a public cryptographic key and is the principal artifact of the identity layer.

4.1.1 Generation

A prospective CTID holder, using software conformant with the protocol, instructs a WebAuthn resident key authenticator on the holder's device to generate a key pair. The authenticator generates the private key within its hardware security boundary, retains the private key under the boundary, and exports the public key to the calling software. The public key is then transmitted to a Verification Provider as part of the issuance request.

The protocol requires that the keypair be generated using ML-DSA-65 (FIPS 204) or, where the authenticator does not yet support post-quantum primitives, a classical scheme (Ed25519 or ECDSA P-256) augmented by an enveloping ML-DSA-65 signature produced by the Verification Provider. The transition path from classical authenticator support to native post-quantum authenticator support is described in Section 4.1.4.

4.1.2 Derivation of the CTID from the public key

The CTID is computed deterministically from the public key as follows:

```
CTID = BLAKE3( "TIP-CTID-v1" || version || vp_id || pub_key ) [first 30 bytes]  
formatted as 5 groups of 6 base32 characters separated by hyphens
```

The components of the input are: the protocol context tag "TIP-CTID-v1" preventing collision with hash outputs in other contexts; the protocol version identifier; the Verification Provider's identifier in the Verification Provider Registry; and the public key of the holder's WebAuthn credential. The output is truncated to thirty bytes, supplying 240 bits of preimage and collision

resistance, which exceeds the security strength of the underlying primitives. The truncated output is formatted as five groups of six base32 characters separated by hyphens, producing a representation suitable for human transcription, copy-paste, and display on user-facing surfaces. An example CTID has the form Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23.

4.1.3 Properties of the CTID

The CTID has the following properties.

Pseudonymous. The CTID is derived from the public key alone, augmented by the Verification Provider identifier and the protocol context. The CTID does not contain the holder's name, the holder's email address, the holder's date of birth, the holder's place of residence, or any other identifier directly linked to the holder's real-world identity. The mapping between the CTID and the holder's real-world identity is held by the Verification Provider and is disclosed only through the legal process of the Verification Provider's jurisdiction.

Stable across surfaces. The CTID does not change when the holder uses different platforms, surfaces, or devices to operate the CTID. The CTID is the same identifier whether displayed alongside a journalistic article, a social post, a published video, or a commerce listing.

Deterministically derivable. Any party in possession of the public key, the Verification Provider identifier, and the protocol version may compute the CTID. The Verification Provider does not maintain a separate identifier database; the CTID is computed from the inputs at the time of need.

Portable across jurisdictions. The CTID is recognized by readers, platforms, publishers, and regulators in jurisdictions other than the jurisdiction of the issuing Verification Provider, subject to the legal interoperability frameworks described in Part IX.

4.1.4 Migration to native post-quantum authenticator support

At the publication of this whitepaper, the principal hardware authenticators marketed under the FIDO2 specification (including the products of Yubico, Google, and Apple) do not support ML-DSA-65 natively. The protocol's transition path comprises three stages.

Stage 1: enveloped classical signature. A holder using a classical authenticator (Ed25519 or ECDSA P-256) submits the public key of the classical credential to the Verification Provider. The Verification Provider generates an ML-DSA-65 keypair on behalf of the holder, binds the ML-DSA-65 public key to the holder's classical public key in an enveloping signature, and supplies the ML-DSA-65 keypair to the holder's software for storage in encrypted form on the holder's device. The CTID is derived from the ML-DSA-65 public key.

Stage 2: dual-mode authenticator support. As FIDO2 authenticators add native ML-DSA-65 support (anticipated 2026 to 2028 based on the FIDO Alliance roadmap), holders may upgrade to authenticators producing ML-DSA-65 signatures natively. Existing CTIDs continue to operate; new CTIDs use the native authenticator.

Stage 3: native-only operation. When native ML-DSA-65 authenticator support becomes ubiquitous, the AI Trust Council retires the enveloped classical signature path for new CTID issuance. Existing CTIDs continue to operate. The retirement decision is published with at least eighteen months of notice.

4.2 The Verification Provider

A Verification Provider is an organization accredited by the AI Trust Council, on the recommendation of The AI Lab, to issue CTIDs.

4.2.1 Role

A Verification Provider performs three functions.

Identity verification. The Verification Provider verifies the identity of a prospective CTID holder against published criteria. The criteria are graded: a basic CTID may be issued on the basis of a verified email address; an enhanced CTID may be issued on the basis of a government-issued identity document together with a liveness check; a high-assurance CTID may be issued on the basis of in-person verification with documentary evidence and biometric binding. The criteria applicable to each grade are published in the Verification Provider's Accreditation Schedule.

CTID issuance. Following identity verification, the Verification Provider performs the issuance computation described in Section 4.1, signs the issuance certificate, and publishes the issuance event to the federated DAG.

Mapping maintenance and disclosure. The Verification Provider maintains the mapping between the CTID and the holder's real-world identity in a secure data store under the data protection law applicable to the Verification Provider's jurisdiction. Disclosure occurs only through the legal process of the jurisdiction, subject to the warrant canary published by the Verification Provider.

4.2.2 Accreditation

A prospective Verification Provider submits an application to the AI Trust Council under the procedure published at theailab.org/tip-verification-provider. The application includes:

1. Identification of the applicant entity, its jurisdiction of incorporation, and its principal place of business.
2. Identification of the natural persons responsible for the Verification Provider's operations.
3. Description of the applicant's identity verification practices at each grade.
4. Description of the applicant's information security program.
5. Description of the applicant's data protection program, including its lawful basis under applicable data protection law, its cross-border transfer mechanism, and its data subject rights response capacity.
6. Description of the applicant's incident response program.
7. Identification of the auditor proposed to conduct the annual audit.
8. Identification of the auditor's qualifications.
9. Identification of the legal process under the applicant's jurisdiction by which the mapping between a CTID and a holder's real-world identity may be compelled to disclosure.
10. A draft warrant canary statement.

The AI Trust Council reviews the application, may request supplementary information, and votes on accreditation under the supermajority threshold described in the Charter. Accreditation, on grant, is recorded in the Verification Provider Registry, is published, and is signed by the AI Trust Council.

4.2.3 Annual obligations

A Verification Provider is required to:

1. Pay the annual accreditation fee in the amount published at theailab.org/tip-verification-provider.
2. Procure an annual independent audit of identity verification, information security, and data protection practices, conducted by an auditor on the AI Trust Council's list of accepted auditors.
3. Publish a current warrant canary statement on the Verification Provider's website at an interval of not less than ninety days.
4. Publish a current jurisdiction declaration.
5. Notify the AI Trust Council of any change in the Verification Provider's jurisdiction declaration, in the natural persons responsible for operations, in the auditor, or in any material respect of the accreditation submission.
6. Notify the AI Trust Council of any significant incident affecting the integrity of identity verification, the security of the mapping, or the operation of the Verification Provider, in accordance with the incident notification timeline of the Verification Provider's jurisdiction (NIS2 timelines where applicable).
7. Comply with applicable data protection law in each jurisdiction in which the Verification Provider operates.

4.2.4 Suspension and revocation

The AI Trust Council may suspend or revoke the accreditation of a Verification Provider for cause. Causes include the failure to satisfy an annual obligation, a material adverse audit finding, a material adverse data protection finding by a supervisory authority, a material change in the Verification Provider's jurisdiction declaration affecting the integrity of the mapping, or a determination that the Verification Provider has engaged in conduct inconsistent with the Charter. Suspension and revocation events are signed by the AI Trust Council and published to the federated DAG. The procedure and the appeal rights are described in the Charter.

4.3 Pseudonymity, revocation, and succession

The CTID's pseudonymous character is structural to the protocol. The revocation lifecycle preserves the pseudonymous character while supplying the mechanisms necessary to operate the protocol against compromised keys, departed holders, and adversely adjudicated CTIDs.

4.3.1 Pseudonymity in operation

A CTID may be used by the holder across platforms, surfaces, and jurisdictions without disclosing the holder's real-world identity to the platforms, surfaces, or jurisdictions where the CTID is used. The platforms and surfaces see the CTID; the Verification Provider sees the mapping; no other party sees the mapping unless the Verification Provider discloses it.

4.3.2 Revocation by the holder

A holder may revoke the holder's CTID at any time, through a revocation request submitted to the issuing Verification Provider. The Verification Provider publishes the revocation event

to the federated DAG within the synchronization interval. Signatures produced by the CTID after the revocation timestamp are treated as suspect by readers and platforms.

A holder seeking to terminate the holder's presence in the protocol may instruct the issuing Verification Provider to delete the mapping between the CTID and the holder's real-world identity. Following deletion, the CTID's record on the federated DAG remains as an orphaned pseudonymous record. The holder's real-world identity is no longer derivable from any party's records subject to the deletion. The pseudonymization analysis under General Data Protection Regulation Recital 26, applied in Section 9.1.2, supports the position that the orphaned record is not personal data.

4.3.3 Revocation by the Verification Provider

The issuing Verification Provider may revoke a CTID on:

1. A determination that the keypair has been compromised.
2. Identification of a material misrepresentation in the identity verification by the holder.
3. A determination that the CTID has been used in a manner inconsistent with the Verification Provider's Accreditation Schedule.
4. A binding order from a court or supervisory authority of competent jurisdiction.
5. The conclusion of an adjudication process under Part VI in which revocation is the determined consequence.

4.3.4 Succession

A holder whose CTID is revoked and who seeks a successor CTID may apply for issuance under the same or a different Verification Provider. The successor CTID is a distinct identifier and is not derived from the revoked CTID. The issuing Verification Provider may, at the holder's request, publish a succession notice on the federated DAG identifying the revoked CTID and the successor CTID and signed by both the revoked CTID and the successor CTID. The succession notice supports continuity of the holder's reputation across the change of CTID.

The succession mechanism does not link the revoked CTID and the successor CTID for parties other than those receiving the succession notice. A holder seeking to terminate continuity may omit the succession notice; the successor CTID then operates without inherited reputation.

4.4 Biometric binding and the WebAuthn resident key flow

A CTID is bound to the holder's device through the WebAuthn resident key authenticator on the device. A CTID is bound to the holder, where the holder elects and the device supports it, through a biometric user verification gesture on the device.

4.4.1 The WebAuthn resident key authenticator

A WebAuthn resident key authenticator, sometimes called a "discoverable credential" authenticator, is an authenticator that stores the private key of a credential within the authenticator and that returns the credential identifier on user verification without prompting the calling software for the credential identifier. The principal commercial implementations of WebAuthn resident key authenticators at the publication of this whitepaper are the YubiKey 5 Series, the Google Titan Security Key, the Apple platform authenticator (Touch ID and Face ID acting as

a FIDO2 authenticator on macOS, iOS, and iPadOS), the Microsoft Windows Hello platform authenticator, and the Android platform authenticator. The protocol is conformant with the WebAuthn Level 3 specification published by the World Wide Web Consortium.

4.4.2 Biometric user verification

A holder may elect to enable biometric user verification at the time of CTID issuance. Biometric user verification, where enabled, requires the holder to present a biometric (a fingerprint, a face scan, or analogous) to the authenticator before the authenticator will produce a signature on the holder's behalf.

The principal architectural decision of the biometric binding is that the biometric template is generated, stored, and matched on the authenticator (and therefore on the holder's device). The biometric template does not leave the authenticator. The Verification Provider receives only (a) the public key of the WebAuthn credential, (b) the attestation that the authenticator has verified the holder, and (c) where the holder elects, an indication that biometric user verification was used.

This architecture has three regulatory consequences.

General Data Protection Regulation Article 9. The Verification Provider does not collect biometric data within the meaning of Article 9(1). The Verification Provider collects the public key and the attestation. The biometric data is collected, processed, and stored exclusively by the authenticator on the holder's device, where it is in the control of the holder.

Illinois BIPA and Texas CUBI. The Verification Provider does not collect biometric identifiers within the meaning of the Illinois Biometric Information Privacy Act or the Texas Capture or Use of Biometric Identifier Act. The corresponding obligations on collection, written informed consent, and disclosure that would otherwise apply to the Verification Provider are not triggered.

Privacy preservation across jurisdictions. A holder presenting a CTID with biometric user verification in a jurisdiction that imposes biometric data localization requirements is not exporting biometric data, because the biometric data remains on the holder's device. The corresponding cross-border transfer concerns are not implicated.

4.4.3 Backup, recovery, and device loss

A holder using a single authenticator faces a risk of CTID inaccessibility if the authenticator is lost, damaged, or otherwise becomes unavailable. The protocol supports three responses.

Multiple authenticators. A holder may, at the time of issuance or subsequently, register multiple authenticators (a primary and one or more secondary) under the same CTID. The Verification Provider associates each authenticator's public key with the CTID. Loss of one authenticator does not, in this configuration, terminate the CTID.

Verification Provider attended recovery. A holder who has lost all authenticators may, on satisfaction of the Verification Provider's recovery criteria, obtain the issuance of a successor CTID with a succession notice as described in Section 4.3.4. The recovery criteria are published in the Verification Provider's Accreditation Schedule.

Social recovery (optional, future). The protocol contemplates, as a future feature, the optional binding of a CTID to a set of trusted recovery contacts, the cooperation of a defined subset of whom would authorize recovery. The mechanism is not part of the operational protocol at the publication of this whitepaper; it is identified as a roadmap item in Part XI.

4.5 Identity portability across jurisdictions

A CTID issued by a Verification Provider in one jurisdiction is recognized by readers, platforms, publishers, and regulators in jurisdictions other than the jurisdiction of the issuing Verification Provider, subject to the legal interoperability frameworks described in Part IX.

4.5.1 Portability between protocol participants

The CTID is, as a technical artifact, the same identifier in every jurisdiction. The signature scheme, the federated DAG, and the Verification Provider Registry are common. A reader in any jurisdiction may verify a TIP-CONTENT record signed by a CTID issued in any other jurisdiction. The technical portability is not dependent on any cross-border arrangement.

4.5.2 Portability of the regulatory effect

The regulatory effect of a CTID, by contrast, depends on the jurisdiction in which the reader's regulator operates. A regulator in the European Union, evaluating a CTID issued by a Verification Provider in New Zealand, considers the CTID under the European Union's regulatory framework. The Verification Provider's jurisdiction declaration identifies the legal process by which the mapping may be compelled in the issuing jurisdiction; the regulator in the receiving jurisdiction may apply its own legal process to compel disclosure from a Verification Provider operating in the receiving jurisdiction.

4.5.3 Cross-border interoperability with the European Digital Identity Wallet

The AI Lab will, through the AI Trust Council, publish a CTID-to-EUDI-Wallet interoperability profile. The profile describes the mapping between a CTID issued by an accredited Verification Provider and a credential presented through a European Digital Identity Wallet under eIDAS 2.0. The profile is intended to permit a holder of a CTID to present an eIDAS-conformant credential bearing the CTID, recognized by relying parties in the European Union that accept EUDI Wallet credentials.

4.5.4 Cross-border interoperability with Aadhaar and DigiLocker

The AI Lab will, through the AI Trust Council, engage with the Unique Identification Authority of India and the Ministry of Electronics and Information Technology concerning interoperability between a CTID issued by a Verification Provider in India and credentials issued through Aadhaar and stored in DigiLocker, subject to applicable Indian law including the Digital Personal Data Protection Act 2023.

4.6 Identity governance and dispute

4.6.1 The CTID holder's standing

A CTID holder has standing to:

1. Use the CTID in accordance with the Accreditation Schedule of the issuing Verification Provider.
2. Revoke the CTID at any time.
3. Request deletion of the mapping between the CTID and the holder's real-world identity from the issuing Verification Provider, subject to applicable data protection law.
4. Apply for a successor CTID following revocation.
5. Initiate a dispute under Part VI concerning an alleged improper revocation, an alleged improper Blocking Item, or an alleged adverse Trust Score event.
6. Participate in the AI Trust Council Network Period reauthorization as a representative of CTID holders.

4.6.2 The Verification Provider's standing

A Verification Provider has standing to:

1. Issue, revoke, and adjust the grade of CTIDs under the Accreditation Schedule.
2. Receive notification of suspension and revocation events affecting the Verification Provider's accreditation.
3. Initiate an appeal of any suspension or revocation of its accreditation.
4. Participate in the AI Trust Council Network Period reauthorization as a representative of Verification Providers.

4.6.3 Disputes

Disputes concerning the issuance, revocation, or operation of a CTID are heard by bonded jurors under Part VI. The Verification Provider, the CTID holder, and a third party adversely affected by the CTID's use have standing under the rules described in that Part. Adjudication outcomes are published to the federated DAG and are signed by the bonded jurors and by the AI Trust Council.

Contact The AI Lab regarding Part IV

Verification Provider Accreditation: licensing@theailab.org , theailab.org/tip-verification-provider Technical Inquiries: tip@theailab.org General Counsel: legal@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part V: TIP-CONTENT: Provenance Layer

This Part describes the content provenance layer of the Trust Identity Protocol. The provenance layer comprises the Canonical Normalization Algorithm Version 2.2 (CNA-2.2), the dual operational mode (Publisher Mode and Creator Mode), the content versioning semantics, the Origin Code system, the signature payload format, and the reference packet specification. The provenance layer is the principal interface through which units of content acquire a verifiable cryptographic provenance under the protocol.

5.1 Overview of the Canonical Normalization Algorithm Version 2.2

The Canonical Normalization Algorithm Version 2.2 (“CNA-2.2”) is the deterministic procedure by which a unit of content is reduced to a canonical byte sequence and identified by the three-hash content addressing system described in Part III. The Algorithm is the foundation of the provenance layer because the verifiability of a signed content unit depends on a procedure that produces, given the same content unit, the same canonical byte sequence and the same content identifier irrespective of the platform, surface, transport mechanism, or rendering engine.

CNA-2.2 supersedes the prior CNA-1.0 normalization specified for the WordPress reference plugin and the prior CNA-2.0 and CNA-2.1 normalizations specified for the browser extension. The principal differences between CNA-2.2 and its predecessors are the addition of an extensible normalization framework permitting the introduction of new content modality variants (CNA-IMG, CNA-VID, CNA-AUDIO) under the same procedural envelope, the addition of an explicit content scope extraction step, and the strengthening of the platform-specific extraction step against changes in platform document structures.

CNA-2.2 operates on units of content categorized into three modalities at the publication of this whitepaper: text, structured documents (HTML pages, articles, blog posts, comments), and composite (a structured document together with embedded images, audio, or video). The image, audio, and video modalities are the subject of variants CNA-IMG, CNA-AUDIO, and CNA-VID respectively, which are specified independently and which produce hash outputs compatible with the three-hash addressing of Part III. The variants are referenced in this Part where relevant but are not specified here in detail.

5.2 The nine canonical normalization steps

CNA-2.2 comprises nine ordered steps. Each step is deterministic. The output of step N is the input of step N+1. The output of step 9 is the canonical byte sequence from which the three-hash content identifier is computed.

5.2.1 Step 1: Platform identification

The Algorithm identifies the platform on which the content unit was encountered or originated. Platform identification is performed by matching the document’s metadata, the document’s URL, and the document’s structural signatures against a published platform registry. The platform registry is maintained by The AI Lab and is updated by the AI Trust Council on the introduction of support for new platforms. The canonical list of supported platforms and the per-platform content-type mappings are published as part of the TIP browser extension source distribution, in the files `src/platform-categories.js` (category mapping) and `src/platform-content-types.js` (content-type mapping), and the same list is mirrored on `vp.theailab.org` for the Verification Provider and creator-registration surfaces.

At the publication of this whitepaper, the platform registry covers the following categories and platforms:

Category	Platforms
Microblog	X (formerly Twitter), Truth Social, Mastodon, Weibo, Threads, Reddit, Tumblr, Bluesky
Visual	Instagram, Facebook, Pinterest
Video	YouTube, TikTok
Audio	Podcast, Vimeo, SoundCloud, Spotify
News	Reuters, Associated Press (AP News), BOOM, The New York Times, The Wall Street Journal, BBC, generic news outlets
Articles	LinkedIn, Substack, Medium, WordPress, Blogger, Ghost, Dev.to, Hashnode, SlideShare, Scribd, generic blog platforms
Messaging	WeChat, Discord, Slack, Telegram, WhatsApp, Element
Other	Catch-all category for arbitrary HTML pages

Each platform exposes an ordered list of content types (for example, on X: tweet, tweet with image, tweet with video, and thread; on LinkedIn: post, photo, carousel, video, document, and article) defined in the `PLATFORM_CONTENT_TYPES` mapping in the extension source. The canonical hashing formula for each content type is defined in `src/tip-types.js`. The categories and per-platform type lists are maintained as a single source of truth across the browser extension, the WordPress reference plugin, the mobile web application at vp.theailab.org, and this whitepaper.

In addition to the per-platform content-type lists, the protocol exposes a universal content-type taxonomy presented to creators at the registration surface. The universal taxonomy maps onto the per-platform lists above, onto the CNA-2.2 modality variants described in Section 5.2, and onto the bitwise canonicalization rules defined in `src/tip-types.js`. At the publication of this whitepaper, the universal taxonomy comprises the following types:

Content type	Description
Blog post	Personal or professional blog entry
News article	Headline, URL, and byline
Newsletter issue	Issue of a newsletter on Substack, Beehiiv, ConvertKit, Ghost, or comparable platform
Essay	Long-form personal or opinion essay
Review	Review of a product, book, restaurant, film, or comparable subject
Tutorial or How-to	Step-by-step guide, instructional content, or recipe
Press release	Official announcement
Article	Canonical URL plus body content
Post	Short-form update
Text post	Written post or status
Comment or Reply	Signed reply to an article or post
Image	Infographic, illustration, chart, or screenshot
Photo	Image with caption
Video	Video URL with caption
Livestream	Live broadcast or recording
Audio	Audio file with title and show notes

Content type	Description
Document	Document file with title and description

The universal taxonomy is maintained as a single source of truth across the creator registration surfaces, the TIP browser extension, the WordPress reference plugin, the mobile web application at vp.theailab.org, and this whitepaper. New universal content types are added by the AI Trust Council in accordance with the Charter. The introduction of a new type triggers the publication of an addendum to the canonical bitwise canonicalization rules in `src/tip-types.js` and a corresponding amendment to the platform registry under Section 5.2.1.

The platform identifier is incorporated into the canonical byte sequence as a four-character platform code and is used by subsequent steps to select platform-specific extraction patterns.

5.2.2 Step 2: Content scope extraction

The Algorithm extracts the content scope from the document. The content scope is the portion of the document that constitutes the content unit being signed, excluding navigation, advertising, related-content recommendations, comment threads not part of the content unit, and other surrounding material. Content scope extraction is performed using platform-specific selectors maintained in the platform registry. For platforms supporting structured semantic markup (Schema.org, OpenGraph, JSON-LD article metadata), the markup is used as the primary source of scope identification. For platforms not supplying structured markup, the content scope is identified by a combination of CSS selector and document structure heuristics.

The content scope extraction step is the principal step at which the protocol distinguishes between the content the signer intends to sign and the surrounding material the signer does not intend to sign. The deterministic extraction is essential to verifiability: a reader performing verification recovers the same content scope from the same document.

5.2.3 Step 3: Structural canonicalization

The extracted content scope is canonicalized at the structural level. The canonicalization comprises (a) the resolution of relative URLs to absolute URLs, (b) the normalization of HTML element names to lowercase, (c) the normalization of HTML attribute names to lowercase, (d) the sorting of HTML attributes within each element in lexicographic order, (e) the canonicalization of attribute values according to the published rules for each attribute, (f) the removal of HTML comments, (g) the removal of script elements and the inline scripts of HTML elements, (h) the removal of style attributes and style elements, (i) the removal of presentational HTML elements (`font`, `b`, `i`, `u` where used presentationally), (j) the preservation of structural HTML elements (`h1` through `h6`, `p`, `ul`, `ol`, `li`, `blockquote`, `pre`, `code`, `a`, `img`, `figure`, `figcaption`, `table`, `thead`, `tbody`, `tr`, `th`, `td`, `em`, `strong`, `cite`), and (k) the canonicalization of whitespace within text nodes to single spaces with no leading or trailing whitespace.

5.2.4 Step 4: Character normalization

The extracted text content is normalized at the character level. The normalization comprises (a) the application of Unicode Normalization Form C (NFC) as specified by the Unicode Consortium in Unicode Standard Annex 15, (b) the replacement of “smart” punctuation marks with

their canonical ASCII equivalents where the character semantics are unchanged (curly quotes become straight quotes, ellipsis character becomes three periods, the figure-dash character is preserved), (c) the preservation of non-Latin script characters in their native form, (d) the preservation of mathematical and scientific characters in their native Unicode form, (e) the removal of zero-width and invisible control characters (U+200B, U+200C, U+200D, U+FEFF), and (f) the consistent treatment of combining characters and surrogate pairs.

The character normalization is designed to be resilient to the platform-specific substitutions that frequently occur when content is copied between editors and platforms, while preserving the semantic distinctions on which the content's meaning depends.

5.2.5 Step 5: Reference normalization

References within the content scope (hyperlinks, image references, video references, audio references) are normalized. URLs are decomposed into scheme, authority, path, query, and fragment components and are reassembled in canonical form. Query parameters are sorted in lexicographic order, with platform-specific tracking parameters removed in accordance with the published tracking parameter registry (utm_source, utm_medium, fbclid, gclid, mc_eid, and others identified by the AI Trust Council on review). Fragments are preserved where the content scope identifies the fragment as semantically meaningful, and are removed otherwise.

The reference normalization is designed to produce a stable identifier for the content scope's external references that is insensitive to the platform's tracking pattern but sensitive to the substantive content of the reference.

5.2.6 Step 6: Embedded media identification

Embedded media (images, audio, video) within the content scope are identified by their three-hash content identifiers computed under the applicable variant of CNA (CNA-IMG, CNA-AUDIO, CNA-VID). The three-hash content identifiers are substituted for the platform-specific URLs of the embedded media in the canonical byte sequence. The substitution permits the signed content unit to be verified against the embedded media independently of where the media is hosted, and permits an embedded medium to be re-hosted on a different surface without invalidating the signature.

5.2.7 Step 7: Metadata extraction

Document-level metadata is extracted. The metadata includes (a) the title of the content unit, (b) the publication date asserted by the publisher or the platform, (c) the language of the content unit identified by IETF BCP 47 language tags, (d) the author identifier (in Publisher Mode, the publisher's CTID and the byline of the natural author within the publication; in Creator Mode, the creator's CTID), (e) the canonical URL of the content unit as asserted by the publisher or the platform, (f) the Origin Code, and (g) the version number of the content unit as described in Section 5.5.

5.2.8 Step 8: Serialization

The output of steps 1 through 7 is serialized into a canonical byte sequence. The serialization uses a deterministic JSON encoding conformant with RFC 8785 (JSON Canonicalization

Scheme). Object keys are sorted in lexicographic order. Numbers are encoded in the canonical form specified by RFC 8785. String values are encoded using the canonical UTF-8 encoding. The serialized output is a deterministic byte sequence.

5.2.9 Step 9: Three-hash addressing

The canonical byte sequence produced by step 8 is the input to the three-hash content addressing system described in Part III. SHA2-256, SHA3-256, and BLAKE3 are computed over the canonical byte sequence. The three hashes, taken together, constitute the content identifier of the content unit. The content identifier is the principal artifact recorded in the TIP-CONTENT record.

5.3 Publisher Mode

Publisher Mode is the operational mode in which an editorial publishing organization signs content using a CTID held by the organization.

5.3.1 Architectural distinction from Creator Mode

Publisher Mode is distinguished from Creator Mode by the locus of the signing key. In Publisher Mode, the signing key is held by the organization, typically in a hardware security module operated by the organization's infrastructure team, and the act of signing is performed by an automated process invoked by the editorial workflow. In Creator Mode, the signing key is held by a natural person on a personal device, and the act of signing is performed by an explicit user verification gesture by the natural person.

The architectural distinction is significant for three reasons. First, the responsibility for the content's accuracy attaches at the level of the institution in Publisher Mode and at the level of the natural person in Creator Mode. Second, the recovery profile differs: a publishing organization with a compromised signing key is identifiable as an institution and can be assisted by the issuing Verification Provider through institutional channels; a natural person with a compromised authenticator follows the personal recovery procedure described in Section 4.4.3. Third, the regulatory treatment differs: Publisher Mode signatures attach to the publication's editorial responsibility for the content; Creator Mode signatures attach to the natural person.

5.3.2 Editorial workflow integration

Publisher Mode is integrated with the editorial workflow of the publisher. On the publisher's editorial system marking a content unit as ready for publication, the editorial system invokes the Publisher Mode signing service. The signing service performs CNA-2.2 normalization on the content unit, produces the three-hash content identifier, signs the identifier using the publisher's CTID private key, and publishes the resulting TIP-CONTENT record to the federated DAG. The editorial system records the TIP-CONTENT record's DAG location alongside the publisher's record of the content unit.

The Publisher Mode signing service is implemented by the publisher using the reference signing library distributed by The AI Lab, by a Commercial License vendor providing a Publisher Mode service, or by the publisher's own implementation conformant to the canonical TIP Protocol Specification. The reference signing library is distributed under TIPCL-1.0 and is available

at the publisher’s TIPCL-1.0 tier.

5.3.3 Byline attribution

Publisher Mode supports byline attribution. The Publisher Mode signature attaches to the publisher’s CTID. Where the publication identifies natural authors as the bylined creators of the content unit, the publication may, in the metadata of the TIP-CONTENT record, identify the bylined natural authors by their Creator Mode CTIDs. The publication’s Publisher Mode signature attaches the publication’s institutional commitment to the content; the bylined natural authors’ Creator Mode CTIDs identify the natural persons within the publication who claim authorship.

This dual attribution is the operational pattern for journalism: the publication’s institutional commitment to the content is the principal signal, and the bylined natural author is identified for purposes of authorship credit and accountability without requiring the natural author to operate the signing infrastructure.

5.4 Creator Mode

Creator Mode is the operational mode in which a natural person signs content using a CTID held by the natural person on a personal device.

5.4.1 Browser extension flow

The principal Creator Mode flow at the publication of this whitepaper is the browser extension flow. The browser extension, available for Chrome, Firefox, Arc, and Safari, observes the user’s interaction with platforms supporting Creator Mode (the platforms identified in the platform registry as Creator Mode targets). On the user invoking the extension on a content unit (typically a post the user is about to publish), the extension performs CNA-2.2 normalization on the content unit in the browser, presents the content identifier to the user, and invokes the user’s WebAuthn resident key authenticator to produce a signature. The signed TIP-CONTENT record is then published to the federated DAG, and the platform-specific publishing action proceeds.

The user verification gesture in Creator Mode is explicit. The browser extension does not produce signatures without the user’s authenticator confirming the user’s presence and (where biometric user verification is enabled) the user’s biometric identity. The extension does not maintain a signing session in which subsequent signatures are produced without user verification.

5.4.2 Mobile flow

The Creator Mode mobile flow is provided through the mobile web application. The mobile web application is a Progressive Web Application conformant with the W3C Web App Manifest specification, installable on iOS, iPadOS, Android, and platform-equivalent operating systems. The mobile flow uses the device’s platform authenticator (Touch ID, Face ID, Android biometric authenticator) for user verification and follows the same signing procedure as the browser extension flow.

5.4.3 Embedded creator flow

The protocol contemplates the embedding of a Creator Mode flow directly within platforms that elect to integrate. A platform integrating the Creator Mode flow embeds the protocol's signing library and supplies the user with a one-tap signing experience native to the platform's publishing interface. Embedded creator flow integration is licensed under TIPCL-1.0 at the platform's applicable Commercial Tier. The protocol publishes a reference embedded integration for the Mastodon and Bluesky platforms.

5.5 Content versioning and CONTENT_UPDATED semantics

A content unit may be modified after publication. The Trust Identity Protocol supports content versioning through the CONTENT_UPDATED event and the associated rules.

5.5.1 Versioning model

Each TIP-CONTENT record identifies a version number. The first publication of a content unit is recorded as version 1. A subsequent modification is recorded as version N, where N is the next integer after the highest existing version of the same content unit.

A modification of a content unit may result in:

1. A version increment for non-substantive modifications (the correction of typographical errors, the addition of an editor's note, the correction of factual misstatements with disclosure of the correction). The CONTENT_UPDATED event records the version increment, identifies the modification, and is signed by the original CTID. The Trust Score is not reduced.
2. A version increment for substantive modifications (substantial revisions, the addition or removal of material). The CONTENT_UPDATED event records the version increment, identifies the modification, and is signed by the original CTID. The Trust Score may be reduced by the Behavioral sub-score where the modification is silent (not disclosed).
3. A new content unit (a new CNA-2.2 content identifier) for modifications so substantial that the modified content is not the same content unit. The new content unit is recorded under a new version 1; the original content unit may carry an UPDATED_BY pointer to the new content unit.

The distinction between (a), (b), and (c) is made by the signer at the time of the CONTENT_UPDATED event. The reader may inspect the signed CONTENT_UPDATED events and may apply the reader's own threshold for materiality. The protocol does not adjudicate the materiality of the modification at the moment of update; adjudication is initiated by a disputant under the Part VI procedure.

5.5.2 The CONTENT_UPDATED event

The CONTENT_UPDATED event is a signed message containing (a) the content identifier of the prior version, (b) the content identifier of the new version, (c) the version number, (d) the timestamp, (e) the modification category ((a), (b), or (c) under Section 5.5.1), (f) a human-readable note describing the modification, and (g) the signature of the CTID. The CONTENT_UPDATED event is published to the federated DAG.

5.5.3 Reader-facing versioning

Reader-facing surfaces (the Global Seal of Trust, the browser extension panel, the <tip-badge> web component) display the version number of the content unit currently being read and supply a link to the version history of the content unit. The version history allows the reader to inspect the modifications made to the content unit since its original publication.

5.6 Origin Codes

The Trust Identity Protocol assigns to each TIP-CONTENT record an Origin Code identifying the provenance category of the content unit. The Origin Code is a two-character mnemonic recorded in the metadata of the TIP-CONTENT record and presented to the reader through the reader-facing badge.

5.6.1 OH: Original Human

Origin Code OH is assigned to a content unit produced by a natural human person, without the substantive participation of an artificial intelligence model in the production of the content. Routine assistance from spellchecking, grammar suggestion, and other tools that do not constitute the substantive production of content does not preclude the assignment of OH.

The OH origin code is the principal origin code for journalistic content produced by named natural journalists, for academic work, for personal communication, and for any content produced by a natural person without substantive AI participation.

5.6.2 AA: AI-Assisted

Origin Code AA is assigned to a content unit produced by a natural human person with the substantive assistance of an artificial intelligence model, in a manner in which the natural human person remains the principal author and exercises editorial judgment over the final content. Examples include a journalist using an AI model to summarize a long document, a novelist using an AI model to suggest variations of a passage, a researcher using an AI model to draft a literature review subsequently substantively edited by the researcher.

The AA origin code identifies content whose authorship is hybrid and supports the disclosure interest of the reader. The Trust Score does not penalize the AA origin code as such; the score reflects the behavioral history of the signer.

5.6.3 AG: AI-Generated

Origin Code AG is assigned to a content unit produced by an artificial intelligence model, including content produced without human authorial intervention and content produced under prompts so general that the model's output is the principal authorial contribution. The AG origin code identifies content that is, in substance, generated by a model.

The AG origin code is the principal origin code for content to which Article 50(2) of the EU AI Act applies and to which analogous disclosure obligations in other jurisdictions apply. A platform displaying a content unit with the AG origin code is supplied, by the protocol, with the machine-readable marking necessary to satisfy the platform's Article 50 obligations.

5.6.4 MX: Mixed

Origin Code MX is assigned to a content unit comprising substantively different origin categories within the same content unit. A documentary video comprising a journalist's narration over AI-generated B-roll, an essay comprising natural-person original text and AI-generated illustrations, or an article comprising a natural-person editorial frame around an AI-generated dialogue are examples of MX content.

The MX origin code identifies content in which the disclosure of the origin requires granularity beyond the single content unit. The reader-facing badge, when displaying the MX origin code, supplies the reader with a link to the granular Origin Code map of the content unit, identifying which portions are OH, AA, or AG.

5.6.5 Selection of the Origin Code

The Origin Code is selected by the signer at the time of signing. The signer's selection reflects the signer's representation of the origin of the content unit. A false Origin Code selection (assigning OH to a content unit that is, in substance, AG) is a misrepresentation actionable under the Trust Score's Adjudicated sub-score and may, on adjudication, result in the suspension or revocation of the CTID. The Origin Code's reliability rests on the institutional commitment of the signer and the reputational consequences of misrepresentation, not on a model-based detection of synthetic content. The protocol does not displace synthetic content detection technology; the protocol supplies a signed representation by the responsible party concerning the origin.

5.7 Signature payload format and reference packet

The TIP-CONTENT record is structured as a signed payload conformant with a defined wire format. The format is normative and is the basis on which verifiers, archivers, and platforms interoperate.

5.7.1 Payload structure

The TIP-CONTENT record comprises three sections.

The header section. The header identifies the protocol version, the CNA variant used (CNA-2.2, CNA-IMG, CNA-VID, CNA-AUDIO), the signature scheme (ML-DSA-65, optional SLH-DSA-SHA2-192s), the CTID of the signer, and the Origin Code.

The body section. The body contains the three-hash content identifier, the canonical URL, the platform identifier, the metadata (title, publication date, language, byline, version number), and (where the version number is greater than one) the content identifier of the prior version and the modification category.

The signature section. The signature is the ML-DSA-65 signature of the concatenation of the header section and the body section, computed by the signer using the private key associated with the CTID. Where the SLH-DSA-SHA2-192s signature is also applied, the signature section contains both signatures.

5.7.2 Wire format

The wire format is a deterministic JSON serialization conformant with RFC 8785, with binary fields (content identifiers, public keys, signatures) base64-encoded. The wire format is the format published to the federated DAG and the format consumed by readers, platforms, and archivers.

5.7.3 Reference packet

The “reference packet” is the canonical exemplar of a TIP-CONTENT record, used in documentation and in the protocol’s test vector suite. A reference packet for each Origin Code, each CNA variant, and each operational mode (Publisher Mode and Creator Mode) is published in Appendix D and in the canonical Specification’s test vector directory. A conformant implementation reproduces the reference packet content identifier given the reference packet input.

5.8 Content Classification, Transient Media Retention, and Classifier Improvement

When content is submitted for registration, the protocol records only a cryptographic fingerprint of the work on the federated DAG. To determine an origin classification, that is, a probability that the media was generated by artificial intelligence, alongside the holder's declared Origin Code, the underlying media must be analyzed by the origin classifier. Because a single submission may comprise multiple files, including several images, videos, or audio recordings, and larger or more numerous files require more processing time, submitted media is enqueued for asynchronous processing rather than analyzed synchronously.

The raw media file is retained solely for the duration of classification and in no event for longer than twenty-one (21) days, after which it is permanently and irreversibly deleted. Following classification, the protocol retains only the classification output, namely the AI-generated probability and the associated Origin Code declaration, together with the content fingerprint. The raw file is never written to the DAG, published, sold, shared, or retained in any content library. This bounded retention window exists to permit accurate classification of large and multi-file submissions and is a data-minimization measure consistent with Article 5(1)(c) and Article 25 of Regulation (EU) 2016/679.

The origin classifier is not static. It is improved over time by training on classification signals derived from processed content and on the outcomes of provenance disputes adjudicated under the procedures described in Part VI. This training does not use the raw media files, which have already been deleted under the retention rule above; it uses the resulting classification signals and adjudicated outcomes. The purpose is to increase the accuracy of the classifier's probability estimates for all participants, reducing both false positives against genuine human authorship and false negatives that would allow synthetic content to pass as human. Consistent with Part VI, the classifier is advisory: no automated classification, by itself, imposes an adverse determination on a holder without the human review and adjudication safeguards specified in Part VI.

Contact The AI Lab regarding Part V

Technical Inquiries: tip@theailab.org Conformance Inquiries: compliance@theailab.org
Standards Engagement: standards@theailab.org

Part VI: TIP-TRUST: Reputation Layer

This Part describes the reputation layer of the Trust Identity Protocol. The reputation layer comprises the Trust Score, its four constituent sub-scores, the rules by which the sub-scores are aggregated, the Blocking Items that suspend or revoke the operational utility of a CTID, the bonded jury adjudication procedure, the appeals process, the reader-facing display modes, and the Global Seal of Trust. The reputation layer translates the cryptographically verified identity and content provenance of the lower layers into a numeric signal a reader, a platform, or a regulator can act on.

6.1 The Trust Score

The Trust Score is the principal reputation artifact of the protocol. The Trust Score associated with a CTID is a numeric value between zero and one thousand inclusive, computed deterministically from four sub-scores under the aggregation rules described in Section 6.2.

The Trust Score is not a classification of the natural person, organization, or other entity behind the CTID. The Trust Score is a measure of the operational history of the CTID under the protocol: the cryptographic state of its credentials, the behavioral pattern of its use within the protocol, the adjudicated outcomes of disputes it has been a party to, and the network position of the CTID within the federated ecosystem. The Trust Score is, in regulatory terms, a content authenticity and operator reliability signal. It is not a social scoring system within the meaning of Article 5(1)(c) of Regulation (EU) 2024/1689, for the reasons set out in Part X Section 10.8.

6.2 The four sub-scores

The Trust Score is computed from four sub-scores. Each sub-score is itself a numeric value between zero and one thousand. The aggregation is a weighted average defined in Section 6.2.5.

6.2.1 Cryptographic sub-score

The Cryptographic sub-score reflects the cryptographic posture of the CTID. The inputs are:

1. The signature scheme used by the CTID at issuance and currently. A CTID issued and operating under ML-DSA-65 receives the full Cryptographic sub-score weight. A CTID operating under a classical scheme augmented by the enveloping signature pattern described in Part IV Section 4.1.4 receives a reduced Cryptographic sub-score weight that increases as the CTID migrates to a native post-quantum authenticator.
2. The integrity of the authenticator binding. A CTID whose authenticator attestation has been verified and remains valid receives the full Cryptographic sub-score weight. A CTID whose attestation has expired or whose authenticator has issued a key compromise notice receives a materially reduced Cryptographic sub-score.
3. The recency of the signing key in use. A CTID whose signing key is the original key associated with the CTID receives the full Cryptographic sub-score weight. A CTID whose signing key has been rotated under the procedure published by the issuing Verification

Provider receives the full Cryptographic sub-score weight; the rotation event is recorded on the federated DAG.

4. The integrity of any optional SLH-DSA fallback signatures recorded on long-term archival records.

The Cryptographic sub-score reflects the technical assurance attaching to the CTID's signing capability and is largely controlled by the CTID holder's choice of authenticator and the CTID holder's compliance with the issuing Verification Provider's procedures.

6.2.2 Behavioral sub-score

The Behavioral sub-score reflects the behavioral pattern of the CTID's use within the protocol. The inputs are:

1. The volume of TIP-CONTENT records published under the CTID over the observation window, with established CTIDs receiving a higher Behavioral sub-score than freshly issued CTIDs.
2. The Origin Code distribution of the published TIP-CONTENT records, with CTIDs publishing a distribution consistent with their declared role receiving a higher Behavioral sub-score than CTIDs publishing a distribution inconsistent with their declared role.
3. The CONTENT_UPDATED event distribution, with CTIDs disclosing modifications through CONTENT_UPDATED events receiving a higher Behavioral sub-score than CTIDs whose disclosed modifications are persistently silent or whose modifications are systematically miscategorized.
4. Anomaly indicators including but not limited to publication rates inconsistent with human or institutional operation, signature timestamp patterns inconsistent with the declared jurisdiction of the CTID, and platform identifier distributions inconsistent with the CTID's declared use.

The Behavioral sub-score is the principal vector by which Sybil and automation adversaries are reflected in the Trust Score.

6.2.3 Adjudicated sub-score

The Adjudicated sub-score reflects the outcomes of adjudication proceedings under Section 6.4 affecting the CTID. The inputs are:

1. Final determinations by bonded juries adverse to the CTID. Adverse determinations reduce the Adjudicated sub-score in proportion to the severity of the determination (the categories of severity are defined in the Charter and in the operative adjudication procedure).
2. Final determinations favorable to the CTID following a dispute initiated by a third party. Favorable determinations preserve the Adjudicated sub-score and, where the dispute was facially without merit, may modestly increase the Adjudicated sub-score of the CTID and reduce the Adjudicated sub-score of the initiating party.
3. Settlements, withdrawals, and dismissals of disputes, which are recorded but which do not by themselves alter the Adjudicated sub-score.

The Adjudicated sub-score is the principal vector by which the substantive evaluation of the CTID's conduct (as distinct from the cryptographic and behavioral statistical signals) enters the Trust Score.

6.2.4 Network sub-score

The Network sub-score reflects the network position of the CTID within the federated ecosystem. The inputs are:

1. The accreditation status of the issuing Verification Provider, with CTIDs issued by Verification Providers in good standing receiving the full Network sub-score weight, and CTIDs issued by Verification Providers under suspension receiving a reduced Network sub-score.
2. The graded character of the CTID. A high-assurance CTID (issued on in-person verification with documentary evidence) receives a higher Network sub-score than an enhanced CTID, which receives a higher Network sub-score than a basic CTID.
3. The jurisdiction declaration of the issuing Verification Provider, with the Network sub-score reflecting the publicly assessed strength of the legal protections and the legal process available in the jurisdiction. The jurisdiction-effect input is small and is published; its purpose is to give effect to the rule-of-law indicia material to the operational reliability of the CTID, not to disadvantage CTIDs from any particular jurisdiction.

The Network sub-score is the principal vector by which institutional context enters the Trust Score.

6.2.5 Aggregation

The Trust Score is computed from the four sub-scores by weighted average. The default weights at the publication of this whitepaper are: Cryptographic 0.20; Behavioral 0.30; Adjudicated 0.30; Network 0.20. The weights are subject to amendment by the AI Trust Council under the supermajority threshold described in the Charter. Any weight amendment is published with at least one hundred and twenty days of notice.

The aggregated Trust Score is rounded to the nearest integer. The Trust Score is recomputed on every event materially affecting any sub-score and at least daily.

6.3 Trust Score tiers

For reader-facing display, the Trust Score is associated with five normative tiers. The tiers at the publication of this whitepaper, as defined in Section 26 of the canonical TIP Protocol Specification v5.0, are:

Score range	Tier	Icon	Color
800 to 1000	HIGHLY_TRUSTED	check	Green (#1A8A5C)
600 to 799	TRUSTED	check	Blue (#2563A8)
400 to 599	REVIEW_ADVISED	exclamation	Amber (#A88B15)
200 to 399	LOW_TRUST	cross	Orange (#C07318)
0 to 199	NOT_TRUSTED	cross	Red (#C53030)

The tier is the principal element of the reader-facing badge. The numeric Trust Score is available on tap or hover; the default visibility mode is TIER_ONLY (Section 27 of the canonical Specification), under which only the tier label is visible to third parties and the numeric score is hidden unless the holder elects the FULL_PUBLIC mode. The tier names are normative; conforming implementations MUST NOT invent new tier names. The tier color is paired with text and icon so that the tier is perceptible without color (for visually impaired readers and readers with red-green color blindness).

Values shown are governed by the AI Trust Council and may evolve. The score-range boundaries, tier names, icons, colors, penalty points, juror eligibility thresholds, Blocking Item triggers, and related numeric parameters are amendable by the AI Trust Council in accordance with the Charter. The values published here represent the parameters in effect at the publication of this whitepaper. Future amendments are published at theailab.org/tip-protocol and reflected in subsequent versions of the canonical Specification and of this whitepaper.

6.4 Blocking Items B1 through B6

A Blocking Item is a defined condition that, when present, materially constrains or suspends the operational utility of a CTID. Blocking Items operate alongside the Trust Score; a CTID may have a high Trust Score and an active Blocking Item, in which case the reader-facing presentation reflects the Blocking Item.

The Blocking Items at the publication of this whitepaper are as follows.

B1: Pending Adverse Adjudication. An open adjudication proceeding alleging conduct that, if found, would result in suspension or revocation. The CTID's operational utility is preserved during the proceeding; the reader-facing presentation displays a Pending Adverse Adjudication marker, with a link to the proceeding's docket where the docket is public.

B2: Verification Provider Under Suspension. The issuing Verification Provider has been suspended by the AI Trust Council. CTIDs issued by the suspended Verification Provider operate with a B2 marker. The B2 marker reflects an institutional, not personal, concern; the CTID holder's individual conduct is not the basis for B2.

B3: Cryptographic Compromise. A determination by the issuing Verification Provider or by the AI Trust Council that the CTID's signing capability has been compromised. The CTID is suspended pending revocation. Signatures produced after the B3 timestamp are treated as suspect by the protocol.

B4: Court or Regulatory Order. A binding order from a court or regulatory authority of competent jurisdiction directed at the CTID. The B4 marker identifies the issuing authority and (subject to the constraints of the order) the substantive basis for the order. The CTID's operational utility is constrained in accordance with the terms of the order.

B5: Final Adverse Adjudication. A final determination by a bonded jury adverse to the CTID under a procedure described in this Section. The B5 marker is published and is durable; the marker may be lifted only by a subsequent successful appeal or by the operation of a remediation procedure described in the Charter.

B6: Material Misrepresentation in Verification. A determination by the issuing Verification Provider that the identity verification supporting the CTID's issuance was materially misrep-

resented. The CTID is suspended pending revocation. Successor issuance, where permitted, follows the procedure of Section 4.3.4 of Part IV.

Blocking Items B3 and B6 are revocation pathways: the operational result is revocation of the CTID, with the corresponding loss of operational utility. Blocking Items B1, B2, B4, and B5 are constraints on operational utility short of revocation, except where the underlying basis (a court order, a final adjudication) operates as a revocation.

6.5 Bonded jury adjudication

The principal mechanism by which a substantive dispute concerning a CTID is resolved is bonded jury adjudication. The procedure is designed to supply a human determination, on the merits, with a structural commitment by the determining jurors to the integrity of the determination.

6.5.1 The bonded juror

A bonded juror is a natural person who has:

1. Obtained a high-assurance CTID from an accredited Verification Provider.
2. Submitted to the bonded juror qualification procedure published by the AI Trust Council, including a background review and a topic-area screening.
3. Posted a bond, denominated in the manner published by the AI Trust Council, forfeit on a determination by the Charter procedure that the bonded juror has acted in bad faith.
4. Accepted the bonded juror code of conduct, including the conflict disclosure and recusal requirements applicable to Council Members.

The bonded juror panel is drawn from a population of bonded jurors selected for representativeness across jurisdictions, languages, and subject-matter expertise. A panel for a given dispute comprises three or five bonded jurors, depending on the severity category of the dispute, selected randomly from the qualified panel of jurors who have not declared a conflict.

6.5.2 Procedure

A dispute is initiated by a complainant submitting a complaint to the AI Trust Council. The complaint identifies the respondent CTID, identifies the substantive basis for the complaint, and supplies the evidence the complainant relies on. The complaint is published, redacted only for personal data of the complainant where the complainant has elected to remain pseudonymous and for confidential commercial information of third parties.

The respondent is notified and is afforded an opportunity to respond. The respondent's response is published on the same basis.

A bonded jury panel is convened. The panel reviews the complaint, the response, and any evidence submitted. The panel may request additional information from either party, may request technical analysis from the AI Lab or from a third party retained for the purpose, and may, in the case of complex matters, hold a hearing at which the parties are heard in writing or in audio.

The panel issues a final determination. The determination identifies the issues, the evidence, the panel's findings of fact, and the panel's disposition. The disposition is one of: complaint

dismissed; complaint upheld with no operational effect on the CTID; complaint upheld with a defined operational effect on the CTID (Trust Score reduction; B1 to B6 marker); complaint upheld with revocation of the CTID. The determination is published. The determination is signed by the bonded jurors and is recorded on the federated DAG.

6.5.3 AI-assisted pre-classification

The bonded jury panel may, in its discretion, use AI-assisted pre-classification tools to triage incoming complaints, to identify the substantive issues, and to identify analogous prior determinations. The use of AI-assisted pre-classification is a procedural aid; it does not displace the panel's responsibility for the determination. The use is recorded in the determination.

The human-in-the-loop structure satisfies the human oversight requirement of Article 14 of the EU AI Act in the event that any AI-assisted component of the adjudication path is, on a future regulatory determination, classified as a high-risk AI system.

6.6 Appeals

A party adversely affected by a final determination of a bonded jury panel may appeal to the AI Trust Council on the grounds set out in the Charter, including manifest error of fact, manifest error of law, procedural irregularity material to the determination, or the discovery of new evidence not reasonably available at the time of the original determination. The appeal is heard by a panel of three Members of the Council, none of whom were involved in the original determination. The appeal panel may affirm, modify, or reverse the original determination, or may remit the matter to a new bonded jury panel.

The appeal is the final administrative remedy within the protocol. Parties retain such judicial remedies as may be available to them under the law of the relevant jurisdiction; the protocol does not contract out of judicial remedies.

6.7 Reader-facing display

6.7.1 The Trust Score badge

The protocol publishes a reader-facing badge displayed alongside the content unit. The badge is implemented as the `<tip-badge>` web component, as the browser extension panel, and as the embedded creator integration on supporting platforms. The badge displays, in its default mode:

1. The CTID of the signer, in the formatted hyphenated representation.
2. The Trust Score tier (HIGHLY_TRUSTED, TRUSTED, REVIEW_ADVISED, LOW_TRUST, NOT_TRUSTED) as defined in Section 6.3.
3. The Origin Code (OH, AA, AG, MX).
4. The Global Seal of Trust where applicable (Section 6.8).
5. On tap or hover, the numeric Trust Score, the version number, and a link to the version history and to the jurisdiction of the issuing Verification Provider.

6.7.2 Display modes

The badge supports three display modes corresponding to the platform and reader context:

Standard mode. The full badge as described in Section 6.7.1, suitable for desktop and mobile web surfaces.

Compact mode. A reduced badge displaying the Trust Score class and the Origin Code only, with full information available on tap. Suitable for surfaces with constrained vertical or horizontal real estate.

Minimal mode. A single-character indicator (the Global Seal of Trust character, the Verified-class character, or the Blocked-class character) suitable for inline display within text or alongside short-form content.

Platforms and publishers select the display mode appropriate to their surface. The reference rendering of each mode is published in the canonical Specification.

6.7.3 Accessibility

The reader-facing badge is implemented with full accessibility support: the badge supplies appropriate ARIA labels and roles, the badge is keyboard-navigable, the badge is operable with screen readers, the badge's color usage is compliant with WCAG 2.2 Level AA contrast requirements, and the badge supports the user agent's prefers-reduced-motion and prefers-color-scheme settings. The accessibility implementation is normative and is part of the conformance requirements of TIPCL-1.0.

6.8 The Global Seal of Trust

The Global Seal of Trust is the principal reader-facing indication that a unit of content has been signed under the Trust Identity Protocol by a CTID with a Trust Score in the Verified class, with no active Blocking Items, and with an Origin Code consistent with the content's declared character. The Global Seal of Trust is visible to the reader as a distinct visual mark integrated into the badge.

The Global Seal of Trust is a trademark of The AI Lab Intelligence Unobscured, Inc. (USPTO Serial No. 99607461, pending). Use of the Global Seal of Trust mark in any surface conformant with the protocol is permitted under the Trademark Usage Policy described in Part X Section 10.5. Use outside the protocol's specifications is not authorized.

The reader who sees the Global Seal of Trust beside a unit of content is, by the structure of the protocol, presented with the following layered representation:

1. That the content has been cryptographically signed by a CTID whose private key is held in a user-verified WebAuthn resident key authenticator.
2. That the CTID has been issued by a Verification Provider accredited by the AI Trust Council and currently in good standing.
3. That the CTID has accumulated, through its operational history, a Trust Score in the Verified class under the four-sub-score system.
4. That the CTID is not under any Blocking Item.
5. That the Origin Code declared by the signer is recorded.
6. That the federated DAG contains the published TIP-CONTENT record and the published version history.

The Global Seal of Trust is not a guarantee of factual accuracy. The protocol does not adjudicate the factual accuracy of content at the moment of signing. The Global Seal of Trust is a guarantee that the protocol's identity, provenance, and reputation infrastructure has, on the evidence available to the federation, supplied no signal inconsistent with the publication of the content under the conditions represented by the badge.

Contact The AI Lab regarding Part VI

Adjudication and Disputes: council@theailab.org , theailab.org/ai-trust-council Technical
Inquiries: tip@theailab.org General Counsel: legal@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part VII: Federated DAG and Network Topology

This Part describes the network on which the Trust Identity Protocol operates. The network is a federation of Node Operators, each operating one or more nodes that together maintain the public, append-only directed acyclic graph on which TIP-CONTENT records, CTID issuance and revocation events, Trust Score recalculations, Blocking Item activations, adjudication outcomes, and Verification Provider accreditation events are recorded. This Part addresses the architecture of the federation, the categories of node, the three named operational periods (Pre-Genesis Period, Founding Period, Network Period), the synchronization protocol, the consistency and partition tolerance properties, the service level commitments, the geographic distribution of the federation, and the capacity model.

7.1 The federated directed acyclic graph

The Trust Identity Protocol records the operational events of the protocol on a federated directed acyclic graph (the "DAG"). The DAG is a distributed, append-only data structure replicated across the Node Operators in the federation. The DAG is not a blockchain in the conventional sense: the DAG does not require proof-of-work, does not require proof-of-stake, does not assign a token-denominated reward to Node Operators for the act of recording entries, and does not depend on Byzantine fault-tolerant consensus across an unbounded population. The DAG is a federated, credentialed, append-only record structure analogous to the certificate transparency logs operated by the public-key infrastructure community, extended for the additional categories of event recorded by the protocol.

7.1.1 Entry structure

Each entry in the DAG comprises (a) a payload corresponding to a defined protocol event, (b) a timestamp recorded by the producing Node Operator and verified against the federation's time consensus, (c) one or more references to immediately preceding entries forming the directed acyclic structure, (d) a signature by the producing Node Operator binding the payload, the timestamp, and the references, and (e) a unique entry identifier derived from the three-hash addressing of the payload.

7.1.2 Append-only property

Entries are added to the DAG and are not removed. The append-only property is structural to the protocol. The non-repudiation, auditability, and regulatory-alignment consequences of the append-only property are addressed in Part II Section 2.3. The pseudonymization analysis that reconciles the append-only property with the right to erasure under data protection law is addressed in Part IX Section 9.1.2 and in Part IV.

7.1.3 Public verifiability

Any party may, using software conformant with the protocol, retrieve a DAG entry by its entry identifier, verify the signature of the producing Node Operator, verify the references to preceding entries, and recompute the entry identifier from the payload. The federation does not depend on the discretion of any single Node Operator for the public verifiability of any entry.

7.2 Node Operator roles

The federation comprises three categories of node. A single Node Operator may operate nodes in any or all categories.

7.2.1 Light Node

A Light Node maintains a partial replica of the DAG. The Light Node retains the entry headers and the cryptographic commitments necessary to verify the integrity of entries on the DAG without retaining the full payloads of all entries. The Light Node retains this header and commitment data on a rolling one-year rotation, holding each entry for twelve months from the date it is observed and pruning it as newer entries are written, so the node continuously maintains a trailing twelve-month window of verifiable DAG history. The Light Node is suitable for readers, platforms, publishers, and integrators who require the ability to verify TIP-CONTENT records against the DAG without operating the full DAG storage and bandwidth burden.

A Light Node has the technical requirements set out in the Technical Requirements Specification (Document 02 of the Founding Node Operator Package). The Light Node's principal requirement is a moderate-throughput network connection (one gigabit per second symmetric is the published specification) and modest storage and computational capacity sufficient to maintain entry headers and the cryptographic commitments.

7.2.2 Full Node

A Full Node maintains a full replica of the DAG. The Full Node retains the full payloads of all entries within its observation horizon, supplies historical entries to other nodes on request, participates in the federation's synchronization protocol, and produces new entries on behalf of CTID holders and Verification Providers signing through the Full Node.

A Full Node has the technical requirements set out in the Technical Requirements Specification. The Full Node's principal requirements include a high-throughput network connection (ten gigabits per second dedicated is the published specification for the Cloud deployment profile), the storage capacity sufficient to retain the full DAG together with the operational margin for projected growth (the published specification provides three-year and ten-year sizing), and the computational capacity sufficient to verify incoming entries at the federation's published

throughput target. The Full Node’s specification includes one or more graphics processing units (the published minimum is one graphics processing unit with sixteen gigabytes of video memory; the published recommended specification is two graphics processing units with forty-eight gigabytes of video memory or more) for the optional AI-assisted dispute pre-classification described in Part VI Section 6.5.3.

7.2.3 Verification Provider node

A Verification Provider node is a node operated by a Verification Provider in support of the Verification Provider’s identity issuance and mapping maintenance functions. The Verification Provider node interacts with the Full Node infrastructure of the federation to publish CTID issuance events, suspension and revocation events, and accreditation status updates. The Verification Provider node’s specification depends on the Verification Provider’s operational scale and on the assurance grade of CTIDs the Verification Provider is accredited to issue.

7.3 Operational periods

The Trust Identity Protocol operates in three named periods. The first is the **Pre-Genesis Period**, commencing on June 1, 2026 and concluding on the Genesis Date, during which (a) the canonical TIP Protocol Specification is published, (b) this whitepaper is published, (c) the AI Trust Council Charter is ratified and the Council convenes, (d) the Founding Node Operator Agreements are executed and Founding Node Operators stand up Full Node deployments in pre-Genesis operating mode, and (e) the federated DAG does not accept TIP-CONTENT records, CTID issuance events, or other protocol events as production records. The second is the **Founding Period**, commencing on the Genesis Date and continuing through a date to be determined by The AI Lab in consultation with the AI Trust Council. The third is the **Network Period**, commencing on the conclusion of the Founding Period.

The **Genesis Date** is the date in June 2026 on which the federated network commences live operation, being the date determined by the sole director of The AI Lab on the satisfaction of the Operational Readiness Conditions and published by The AI Lab not less than three (3) business days in advance on theailab.org.

7.3.1 The Founding Node Operators

During the Pre-Genesis Period and the Founding Period that follows, the network is operated by a credentialed set of Founding Node Operators, each party to a Founding Node Operator Agreement with The AI Lab. The Founding Node Operators contracted as of the publication date of this whitepaper, each standing up a Full Node deployment in pre-Genesis operating mode and continuing into production operation from the Genesis Date, are:

Founding Node Operator	Jurisdiction	Principal node category
THE PRESCIENT PACHYDERM LTD	United Kingdom	Full Node
AZLogics Private Limited	India	Full Node
Apex Modular Solutions LLC	United States	Full Node
Timpi International Ltd	New Zealand	Full Node

Founding Node Operator	Jurisdiction	Principal node category
Lonestar Data Holdings Inc.	United States	Full Node
The AI Lab Intelligence Unobscured, Inc.	United States (theailab.org)	Full Node

Each Founding Node Operator has committed to the operation of a Full Node, the service level commitments described in Section 7.6, the warrant canary and the jurisdiction declaration described in Part X, and the conformance with the data protection and cybersecurity obligations of its jurisdiction. The Founding Node Operator Agreements incorporate by reference the Technical Requirements Specification, the Service Level Agreement, and TIPCL-1.0.

7.3.2 The character of the Founding Period

The Founding Period commences on the Genesis Date. It is operationally distinguished from the Network Period that follows in five respects.

Smaller federation. The Founding Period operates with a smaller and credentialed set of Node Operators, all of whom are party to written agreements with The AI Lab. The Network Period contemplates the expansion of the federation through a published accreditation procedure analogous to the Verification Provider accreditation procedure.

Operational coordination. The Founding Period contemplates a higher degree of operational coordination among Node Operators, including coordinated upgrade windows, coordinated incident response, and a single-channel escalation path through The AI Lab. The Network Period contemplates a decentralized operational pattern with peer-to-peer coordination protocols.

Protocol amendment cadence. The Founding Period contemplates a higher cadence of protocol amendments as the federation matures. The Network Period contemplates an amendment cadence stabilized at the level customary for international standards (a major revision approximately every twenty-four to thirty-six months).

Economic model. The Founding Period operates under an economic model in which Founding Node Operators contribute infrastructure under TIPCL-1.0 Free Tier terms (in the case of journalism organizations, educational institutions, and government), under Commercial License terms (in the case of commercial Node Operators), or under the Founding Node Operator Agreement's commitments. The Network Period contemplates a successor economic model under development by The AI Lab and the AI Trust Council. No representation is made in this whitepaper concerning the Network Period economic terms.

Governance posture. The Founding Period operates under the AI Trust Council in its Founding Member composition. The Network Period operates under the Council in its Network Period composition, expanded to include representatives of Node Operators, Verification Providers, publishers, civil society organizations, and academic institutions, as described in Part X.

7.3.3 Transition to the Network Period

The transition from the Founding Period to the Network Period is initiated by a supermajority ratification of the AI Trust Council on the satisfaction of the conditions identified in the Charter. The conditions include the establishment of the Network Period accreditation procedure for Node Operators, the establishment of the Network Period economic model, the expansion of the Council to its Network Period composition, the engagement with at least one standards organization for the formal standardization of the protocol, and the satisfaction of operational stability indicators for the Founding Period network.

7.4 Synchronization, consistency, and partition tolerance

7.4.1 Synchronization protocol

The federation operates a synchronization protocol by which entries produced by one Full Node are propagated to other Full Nodes. The synchronization protocol is a gossip protocol over an authenticated transport layer. A Full Node periodically (the published synchronization interval is one second under nominal operating conditions) announces to its peers the identifiers of entries it has produced or received since the prior announcement. Peers request the payloads of entries they do not already hold. The transport layer is TLS 1.3 with the hybrid post-quantum key agreement described in Part III.

7.4.2 Consistency model

The federation supplies an eventually consistent view of the DAG. An entry produced by a Full Node is propagated to other Full Nodes within a published synchronization horizon (the published target is five seconds at the ninety-fifth percentile for entries propagated to all Full Nodes in the Founding Period network). Within the synchronization horizon, different Full Nodes may have different views of the DAG; outside the synchronization horizon, the views converge.

The eventual consistency model is appropriate to the protocol's operational profile. A CTID holder signing a TIP-CONTENT record and a reader verifying the record are typically separated in time by an interval substantially greater than the synchronization horizon. Where a reader requires the highest-assurance verification of a recent record, the reader's software may query multiple Full Nodes and may withhold the Verified-class display until the record has propagated to a quorum of Full Nodes.

7.4.3 Partition tolerance

The federation tolerates the partition of one or more Full Nodes from the remainder of the federation. A partitioned Full Node continues to accept new entries from local CTID holders and Verification Providers, queues outgoing entries for propagation, and continues to serve historical entries to local readers. On the restoration of connectivity, the queued outgoing entries are propagated and the synchronization protocol resumes. The protocol does not require the agreement of a quorum of Full Nodes for the production of new entries; an entry produced by a partitioned Full Node is valid on its production and is recognized as valid by the federation on synchronization.

7.4.4 Resolution of structural ambiguity

The protocol's design avoids the structural ambiguity that would otherwise arise from concurrent production of entries by multiple Full Nodes. Each entry's identifier is derived deterministically from its payload, its timestamp, and its references; concurrent entries are accepted as siblings of the directed acyclic structure rather than as competing versions of the same entry. A reader retrieving an entry retrieves the same entry irrespective of which Full Node serves the retrieval.

7.4.5 Time consensus

The federation relies on a time consensus protocol to bound the divergence of timestamps recorded by Full Nodes. The protocol uses Network Time Protocol with authenticated NTP servers operated by national metrology institutions (the National Institute of Standards and Technology in the United States, the Bureau International des Poids et Mesures internationally, the National Physical Laboratory in the United Kingdom, and the Measurement Standards Laboratory of New Zealand). Each Full Node continuously monitors its time divergence from the consensus and publishes a divergence indicator. Entries produced by a Full Node whose divergence exceeds a published threshold are treated as suspect by the federation pending the Node Operator's resolution of the divergence.

7.5 Service level commitments

7.5.1 Service level targets

The Founding Node Operator Agreement specifies the service level targets a Founding Node Operator commits to maintain. The published service level targets at the publication of this whitepaper are:

Indicator	Target
Node availability (monthly)	99.9 per cent
Entry production latency (95th percentile)	Less than 250 milliseconds
Entry propagation latency (95th percentile, full federation)	Less than 5 seconds
Synchronization completeness (24-hour window)	100 per cent of produced entries propagated to all peer Full Nodes
Incident response, severity 1 acknowledgment	Within 30 minutes
Incident response, severity 1 mitigation	Within 4 hours
Incident response, severity 2 acknowledgment	Within 2 hours
Incident response, severity 2 mitigation	Within 24 hours
Scheduled maintenance windows	Published 30 days in advance
Reporting cadence	Monthly availability and incident report

7.5.2 Service level remedies

Service level shortfalls are addressed by the Service Level Agreement (Document 03 of the Founding Node Operator Package), which sets out the conditions under which a shortfall is excused (force majeure, scheduled maintenance within published windows, restoration following dependence failures upstream of the Node Operator) and the conditions under which a shortfall results in service credits, intensified monitoring, remediation plan submission, or, in extreme cases, suspension of Founding Node Operator status.

7.5.3 Incident classification

Incidents are classified into severity 1 (outage or compromise affecting the production of new entries, the propagation of entries to peers, or the integrity of historical entries), severity 2 (degradation affecting the timeliness of entry production or propagation without affecting availability or integrity), and severity 3 (operational matters not affecting availability, timeliness, or integrity). The classification is normative and is applied uniformly across the federation.

7.6 Geographic distribution of the Founding Period network

The Founding Period network, standing up in pre-Genesis operating mode during the Pre-Genesis Period and continuing into production operation from the Genesis Date, is distributed across four jurisdictions: the United Kingdom (THE PRESCIENT PACHYDERM LTD), India (AZLogics Private Limited), the United States (Apex Modular Solutions LLC, Lonestar Data Holdings Inc., and The AI Lab Intelligence Unobscured, Inc.), and New Zealand (Timpi International Ltd, deploying from Christchurch).

The four-jurisdiction distribution is selected to (a) avoid the concentration of operational and legal risk in any single jurisdiction, (b) supply a federation present in each of the principal regulatory zones with which the protocol is designed to interoperate (the European Union, through the United Kingdom node as a legally and geographically proximate operator with the European Union, recognising that the United Kingdom is not a Member State; the Indo-Pacific, through the India and New Zealand nodes; and North America, through the two United States nodes), (c) supply geographic latency reduction for readers in the principal demand regions, and (d) supply jurisdictional diversity for legal-process and warrant-canary purposes.

The Network Period network is intended to expand to additional jurisdictions on the conditions set out in the Charter, including representation in the European Union (the Founding Period network does not include a Node Operator in a Member State of the European Union, an absence that the Network Period transition is intended to remedy), East Asia, the Middle East, Africa, and South America.

7.7 Capacity model and scaling profile

The Founding Period network is sized for an operational throughput of one hundred entries per second sustained, with a peak capacity of one thousand entries per second. The sustained throughput supports the publication of approximately eight million entries per day, sufficient for the Founding Period reader, publisher, and creator population identified in the operational planning of The AI Lab.

The Network Period scaling profile contemplates the expansion of the federation to a throughput target of ten thousand entries per second sustained, sufficient for adoption by major social platforms, publishing platforms, and news organizations. The scaling is expected to be horizontal: additional Full Nodes are added to the federation as load and resilience requirements dictate. The protocol's design avoids the per-entry coordination overhead that would otherwise constrain horizontal scaling.

7.7.1 Storage profile

The storage profile of a Full Node is dominated by the long-term retention of TIP-CONTENT record payloads. The published estimate is approximately five kilobytes per TIP-CONTENT record averaged across origin codes and content modalities. At the Founding Period sustained throughput, the annual storage growth per Full Node is approximately one hundred and fifty terabytes. The Technical Requirements Specification published for Full Nodes provides for three-year and ten-year storage sizing.

7.7.2 Bandwidth profile

The bandwidth profile of a Full Node is dominated by the synchronization protocol and by reader queries. The published estimate is approximately one gigabit per second sustained and ten gigabits per second peak. The Technical Requirements Specification provides for the bandwidth sizing.

7.7.3 Computational profile

The computational profile of a Full Node is dominated by signature verification (the federation verifies every incoming entry's signature) and, where the Node Operator participates in adjudication, by AI-assisted pre-classification of disputes (described in Part VI Section 6.5.3). The published estimate is approximately twenty processor cores sustained for signature verification and the graphics-processing-unit specification described in Section 7.2.2 for adjudication support.

7.8 Industry support and program memberships

The Trust Identity Protocol's reference implementation is built and operated on commercial cloud, accelerator, and AI-model infrastructure provided by a small number of major industry programs. The AI Lab Intelligence Unobscured, Inc. maintains member status in the following programs:

- NVIDIA Inception: graphics processing unit credits, AI hardware access, technical resources, and the startup-development support furnished by NVIDIA Corporation through its accelerator program.
- AWS Activate: cloud infrastructure credits, architectural consultation, and developer-tools support furnished by Amazon Web Services, Inc.
- Microsoft for Startups: cloud infrastructure credits, Azure services, and developer-tools support furnished by Microsoft Corporation.
- Claude for Startups: AI model access and developer-tools support furnished by Anthropic PBC.

These memberships reflect the commercial relationships through which the protocol’s reference implementation, its development environments, and its operational infrastructure are funded and supported. They do NOT constitute:

1. Endorsement of the Trust Identity Protocol by NVIDIA Corporation, Amazon Web Services, Inc., Amazon.com, Inc., Microsoft Corporation, or Anthropic PBC.
2. Sponsorship of the AI Trust Council by any of the named companies.
3. Operational backing of the federated network or of any Founding Node.
4. Execution of a Founding Node Operator Agreement, or undertaking of any service-level commitment, by any of the named companies.

The Founding Node Operator network identified in Part VII Section 7.3.1 is constituted by the seven legal entities that have signed the Founding Node Operator Agreement and the three additional Founding Node Operators in accession during the Pre-Genesis Period. The named program memberships above are vendor and developer-program relationships, separate from and additional to, that operator network.

The AI Lab acknowledges the practical contribution that these program memberships make to the protocol’s accelerated path to production readiness and to the affordability of building public-infrastructure software at the scale described in this whitepaper.

Contact The AI Lab regarding Part VII

Founding Node Operator Applications: nodes@theailab.org , theailab.org/founding-node
Technical Inquiries: tip@theailab.org Operations: operations@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part VIII: Reference Implementation and Integration

This Part describes the reference implementations of the Trust Identity Protocol published by The AI Lab. The reference implementations comprise the REST API surface through which Verification Providers, Node Operators, publishers, platforms, and readers interact with the federation; the browser extension distributed for Chrome, Firefox, Arc, and Safari; the `<tip-badge>` web component for publisher integration; the WordPress reference plugin for publishers operating on the WordPress platform; the mobile web application; and the software development kits planned for general-purpose integration in commonly deployed application stacks. The reference implementations are distributed under the TIP Protocol Code License Version 1.0 (TIPCL-1.0) as described in Part X and Appendix F.

8.1 The REST API surface

The Trust Identity Protocol publishes a REST API surface comprising nineteen endpoints organized into five groups. The endpoints are described in summary in this Section. The byte-level specification of request and response schemas, error codes, authentication, rate limiting, and conformance test vectors is set out in the canonical TIP Protocol Specification.

8.1.1 Endpoint inventory

The endpoint inventory below summarizes the principal REST endpoints of a TIP node. The canonical and complete list, including request and response schemas, error codes, and conformance test vectors, is maintained in the TIP Protocol Specification v5.0 at github.com/theailaborg/tip-protocol. Endpoints are grouped here by functional area.

Group	Endpoint	Operation
Identity	POST /v1/identity/register	Register a TIP-ID (Verification Provider)
Identity	GET /v1/identity/{tip_id}	Retrieve TIP-ID metadata
Identity	GET /v1/identity/{tip_id}/history	Retrieve event history of a TIP-ID
Identity	GET /v1/identity/{tip_id}/profile	Retrieve TIP-ID public profile
Identity	POST /v1/identity/{tip_id}/profile	Update TIP-ID public profile
Identity	GET /v1/identity/{tip_id}/score	Retrieve the current Trust Score
Identity	GET /v1/identity/{tip_id}/seal	Render the AI Trust ID Seal as SVG
Identity	GET /v1/identity/lookup	Look up TIP-ID by query
Identity	GET /v1/identity/by-dedup-hash/{dedup_hash}	Key-recovery pre-flight check
Identity	POST /v1/identity/{tip_id}/keys/recover	Submit a key recovery transaction
Verification	POST /v1/verify/session	Open a verification session
Verification	GET /v1/verify/{id}/webauthn-challenge	WebAuthn challenge
Verification	POST /v1/verify/{id}/gov-id	Submit government ID step
Verification	POST /v1/verify/{id}/biometric	Submit biometric step
Verification	POST /v1/verify/{id}/liveness	Submit liveness step
Verification	POST /v1/verify/{id}/social	Submit social verification step
Verification	POST /v1/verify/{id}/submit	Submit completed verification session
Verification	GET /v1/verify/{id}/status	Poll session status
Verification	GET /v1/verify/{id}/resume	Resume an interrupted session
Provenance	POST /v1/content/register	Register a TIP-CONTENT record
Provenance	GET /v1/content/{ctid}	Retrieve a TIP-CONTENT record
Provenance	POST /v1/content/{ctid}/update-origin	Record a CONTENT_UPDATED event
Provenance	POST /v1/content/{ctid}/dispute	File a dispute against a content record
Provenance	GET /v1/disputes/{dispute_id}	Retrieve dispute state
Domain	POST /v1/domain/register	Register a publisher domain

Group	Endpoint	Operation
Domain	POST /v1/domain/{domain}/verify	Verify a publisher domain
Reviews	GET /v1/reviews	List reviews open to the requester
Reviews	POST /v1/reviews/{review_id}/confirm	Confirm a review
Reviews	POST /v1/reviews/{review_id}/dismiss	Dismiss a review
VP Registry	GET /v1/vps	List accredited Verification Providers
VP Registry	GET /v1/vp/{vp_id}	Retrieve VP details
VP Registry	POST /v1/vp/register	Submit a VP accreditation application
Network and DAG	GET /v1/node/info	Node identity and configuration
Network and DAG	GET /v1/node/peers	Node peer list
Network and DAG	GET /v1/dag/stats	DAG metrics
Network and DAG	GET /v1/dedup/merkle-root	Current dedup registry Merkle root
Network and DAG	POST /v1/dedup/check	ZK uniqueness check
Network and DAG	GET /v1/state-root	Federation state root
Network and DAG	GET /v1/revocations	Current revocations

8.1.2 Authentication

Authentication patterns vary by endpoint. Endpoints invoked by Verification Providers (the issuance, revocation, and succession endpoints) are authenticated by ML-DSA-65 signatures of the request payload using the Verification Provider's private key. Endpoints invoked by CTID holders (the content publication and dispute filing endpoints) are authenticated by ML-DSA-65 signatures of the request payload using the CTID holder's private key. Read endpoints (the retrieval endpoints in the Identity, Provenance, Trust, and Network groups) are unauthenticated to a default rate limit and admit signed authentication for an elevated rate limit.

8.1.3 Rate limiting

Rate limits are published in the canonical Specification and are reviewed by the AI Trust Council on an annual basis. The published rate limits are calibrated to support the operational profile of Verification Providers, Node Operators, publishers, platforms, and reader-side verification at the Founding Period scale. Commercial Licensees may obtain elevated rate limits under the terms of their Commercial License.

8.1.4 Error model

Errors are returned in a uniform structure conformant with RFC 9457 (Problem Details for HTTP APIs). The published error code registry is normative; conformant implementations return error codes from the registry only.

8.1.5 Versioning

The API surface is versioned at the path level. Version v1 is the version published at the publication of this whitepaper. Backward-incompatible changes to the API surface require ratification by the AI Trust Council under the supermajority threshold of the Charter, are published at a successor path version, and are accompanied by a deprecation schedule for the predecessor version of not less than eighteen months.

8.2 Browser extension

The Trust Identity Protocol browser extension is the principal Creator Mode integration at the publication of this whitepaper.

8.2.1 Supported browsers

The browser extension is published and distributed for the following browsers. Direct distribution URLs are provided where the extension is already live in the corresponding store.

Browser	Distribution channel	Manifest version	Distribution URL
Google Chrome	Chrome Web Store	Manifest V3	chromewebstore.google.com/detail/hkobdcb
Mozilla Firefox	Mozilla Add-ons (AMO)	Manifest V2 with Manifest V3 transition	addons.mozilla.org/en-US/firefox/addon/tip-know-what-s-real/
Brave	Chrome Web Store (Chromium-based)	Manifest V3	(installed via Chrome Web Store entry above)
DuckDuckGo Browser	Chrome Web Store (Chromium-based)	Manifest V3	(installed via Chrome Web Store entry above)
The Browser Company Arc	Chrome Web Store (Chromium-based)	Manifest V3	(installed via Chrome Web Store entry above)
Apple Safari	Apple App Store (Safari Web Extension)	Safari Web Extension API	(in submission)

8.2.2 Architecture

The extension comprises a background service worker responsible for managing the connection to the federation, a content script injected into pages of the platforms registered in the platform registry, a popup user interface invoked on user action, and an options page for configuration. The extension's content scripts apply CNA-2.2 normalization to content units within the page in

the browser, surface the user-actionable signing flow, and invoke the user's WebAuthn resident key authenticator through the Web Authentication API.

8.2.3 Publisher Mode and Creator Mode

The extension supports both Publisher Mode and Creator Mode. Publisher Mode operation requires the extension to be installed on the publisher's infrastructure and is appropriate to integrations in which the publisher's editorial system uses the browser as a deployment surface. Creator Mode is the principal Mode for individual creators using the extension on their personal browsers.

8.2.4 Privacy posture

The extension does not transmit content unit text to any party other than the user's chosen Full Node, at the user's explicit instruction. The extension does not record telemetry of the user's browsing, the user's content unit text, or the user's identifying information. The extension's operational logs are retained locally and are not transmitted. The extension's privacy posture is reviewed annually as part of the AI Trust Council's transparency report.

8.2.5 Localization

The extension is localized into the languages identified in the canonical Specification's locale message catalog. At the publication of this whitepaper, the extension is localized into English, French, German, Spanish, Portuguese, Italian, Dutch, Polish, Japanese, Korean, Simplified Chinese, Traditional Chinese, Arabic, Hindi, and te reo Maori. Additional locales are added through the AI Trust Council's localization workstream.

8.3 The <tip-badge> web component

The <tip-badge> is a web component implementing the reader-facing badge described in Part VI Section 6.7. The component is the principal Publisher Mode integration for web publishers and platforms.

8.3.1 Specification

The component is specified conformantly with the W3C Custom Elements V1 specification and is a standard HTML element usable in any compliant browser. The component is published as an importable JavaScript module, as a single-file embedded build, and as a server-rendered HTML representation for environments that do not execute JavaScript on the reader's device.

8.3.2 Use by publishers

A publisher integrates the <tip-badge> by placing the element alongside the content unit in the publisher's templates and supplying the content unit's CTID, content identifier, and version identifier as attributes. The component fetches the current Trust Score, the active Blocking Items, and the version history from the federation through the public read endpoints. The component renders the badge in the display mode specified by the publisher.

8.3.3 Customization

The component supports the customization of typographic and visual attributes through CSS Custom Properties published in the component's specification. The component does not permit the customization of the substantive content of the badge (the Trust Score tier, the Origin Code, the Global Seal of Trust, the Blocking Item indicators), to preserve the reader's reliance on the consistency of the substantive signal across publishers and platforms.

8.3.4 Accessibility

The component implements full WCAG 2.2 Level AA accessibility, including appropriate ARIA labels and roles, keyboard navigation, screen reader operability, and conformant color contrast. The component supports the user agent's prefers-reduced-motion and prefers-color-scheme settings.

8.4 WordPress reference plugin

The WordPress reference plugin is the principal Publisher Mode integration for publishers operating on the WordPress platform. The plugin implements the Canonical Normalization Algorithm Version 1.0 (CNA-1.0), the Publisher Mode signing workflow integrated with the WordPress editor, the CONTENT_UPDATED event flow integrated with the WordPress revisions system, and the <tip-badge> web component embedded in the publisher's theme.

8.4.1 CNA-1.0

CNA-1.0 is the canonical normalization variant applicable to WordPress content. CNA-1.0 is a six-step procedure described in the canonical Specification and operates on WordPress post objects (including the post title, post content, post excerpt, post meta, and post taxonomies). CNA-1.0 produces a content identifier compatible with the three-hash addressing of Part III. The plugin supplies the CNA-1.0 reference implementation in PHP.

8.4.2 Plugin architecture

The plugin operates within the WordPress Plugin API. On a post being published or updated, the plugin invokes the Publisher Mode signing service (a remote service operated by the publisher's chosen vendor or an on-premise service operated by the publisher) and records the resulting TIP-CONTENT record's DAG location in WordPress post meta. The plugin's settings page exposes the configuration of the signing service endpoint, the CTID and credential for the publisher, the Origin Code defaulting policy, and the byline mapping between WordPress user accounts and Creator Mode CTIDs of named natural authors.

8.4.3 Distribution

The plugin is published on the WordPress Plugin Directory at wordpress.org/plugins/tip-protocol/ and distributed under TIPCL-1.0. Publishers in the Free Tier under TIPCL-1.0 may install and use the plugin without fee. Publishers in the Commercial Tier obtain the plugin under the terms of their Commercial License. The plugin is live and installable at the publication date of this whitepaper.

8.5 Mobile web application

The mobile web application is a Progressive Web Application conformant with the W3C Web App Manifest specification. The application supplies a Creator Mode integration for users operating on mobile devices.

8.5.1 Architecture

The application operates as a client-side Progressive Web Application, with a service worker for offline capability and for the queueing of signing operations during periods of intermittent connectivity. The application uses the device's platform authenticator (Touch ID, Face ID, the Android biometric authenticator, or analogous) for user verification, through the Web Authentication API.

8.5.2 Installation

The application is live and installable at **vp.theailab.org** (<https://vp.theailab.org/>). It runs on iOS, iPadOS, Android, and analogous mobile operating systems through the device browser's installation flow. The application does not require distribution through a mobile application store as a native application, which avoids the dependence of the protocol on the policies of mobile platform operators.

8.5.3 Functionality

The application supplies the principal Creator Mode operations: the publication of TIP-CONTENT records for content composed within the application; the publication of TIP-CONTENT records for content composed outside the application and imported through the share extension; the publication of CONTENT_UPDATED events; the retrieval of the user's current Trust Score and active Blocking Items; the filing of disputes; and the management of the user's CTIDs.

8.6 Software development kits

The AI Lab publishes software development kits for the integration of the Trust Identity Protocol into commonly deployed application stacks. The publication schedule of the software development kits is set out in Part XI.

8.6.1 Published kits

Kit	Target	Status
TIP SDK for JavaScript and TypeScript	Node.js, Deno, Bun, browsers	Published v2.0
TIP SDK for Python	CPython 3.10 and above	Published v2.0

8.6.2 Planned kits

All remaining language kits are scheduled for publication by the end of calendar year 2026:

Kit	Target	Planned publication
TIP SDK for Rust	Stable Rust 1.85 and above	Q3 2026
TIP SDK for Go	Go 1.23 and above	Q3 2026
TIP SDK for Java	Java 21 LTS and above	Q4 2026
TIP SDK for .NET	.NET 9.0 and above	Q4 2026
TIP SDK for Swift	Swift 6 and above	Q4 2026
TIP SDK for Kotlin	Kotlin 2 and above	Q4 2026
TIP SDK for PHP	PHP 8.3 and above	Q4 2026
TIP SDK for Ruby	Ruby 3.3 and above	Q4 2026

8.6.3 Conformance

Software development kits published by The AI Lab are required by their license terms to pass the conformance test vector suite for the protocol. Software development kits developed by third parties may, voluntarily and on submission to the AI Trust Council, obtain inclusion in the Conformance Registry described in Part III Section 3.8.

8.7 Distribution and update channels

8.7.1 Browser extension distribution

The browser extension is distributed exclusively through the official distribution channels of each browser (the Chrome Web Store, Mozilla Add-ons, the Apple App Store, and analogous). The distribution channel is identified in the publisher metadata of the extension package and is verifiable by readers. Distribution through unofficial channels is not authorized.

8.7.2 Update cadence

The browser extension, the <tip-badge> component, the WordPress plugin, the mobile web application, and the software development kits follow a published update cadence. Substantive updates require ratification by the AI Trust Council under the supermajority threshold of the Charter, except for security updates which may be released by The AI Lab on an emergency basis with prompt Council notification.

8.7.3 Software bill of materials

Each published implementation supplies a Software Bill of Materials conformant with the United States Cybersecurity and Infrastructure Security Agency guidance and with the Cyber Resilience Act requirements identified in Part IX. The Software Bill of Materials identifies each open source component, its version, its license, and its known security advisories.

Contact The AI Lab regarding Part VIII

Technical Inquiries: tip@theailab.org Integration Support: integrations@theailab.org Security Reports: security@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part IX: Regulatory Alignment

This Part identifies the principal statutes, regulations, and international instruments within whose regulatory landscape the Trust Identity Protocol is designed to operate, and describes the technical and institutional features of the protocol that support compliance with those instruments by the parties to whom they apply. Statements in this Part are statements of design intent and of the technical controls implemented in the canonical specification. Such statements are not certifications of compliance, opinions of counsel, or attestations by a regulator or supervisory authority. A licensee is responsible for assessing the application of any instrument referenced in this Part to the licensee's particular use of the protocol and for obtaining qualified legal advice. The provisions of Appendix J apply to this Part.

9.1 European Union

9.1.1 Artificial Intelligence Act, Regulation (EU) 2024/1689

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (the "AI Act") entered into force on 1 August 2024. The Act applies in phases: the prohibited practices in Article 5 and the AI literacy obligations in Article 4 apply from 2 February 2025; the general-purpose AI model rules in Chapter V and the governance and penalties chapters apply from 2 August 2025; the principal high-risk obligations under Annex III apply from 2 August 2026; and the high-risk obligations for systems embedded in products covered by Annex I product-safety legislation apply from 2 August 2027.

Classification. The Trust Identity Protocol is not an "AI system" within the meaning of Article 3(1) of the AI Act. The protocol's core operations consist of cryptographic key generation, deterministic content normalization, digital signature production and verification, and the recording of signed events on an append-only directed acyclic graph. These operations do not constitute a machine-based system designed to operate with varying levels of autonomy producing outputs that influence physical or virtual environments. The protocol is therefore not subject to the obligations of providers, importers, distributors, or deployers of AI systems. Three operational components carrying AI elements are addressed in Section 10.8 and are not high-risk under Annex III.

Article 5(1)(c) social scoring. The Trust Score described in Part VI is not a social scoring system within the meaning of Article 5(1)(c). The Trust Score evaluates content authenticity and operator behavior within the protocol context. It does not classify natural persons by their social behavior, personality, or personal characteristics across contexts unrelated to the data's original collection. The Trust Identity Protocol Charter expressly prohibits the use of the Trust Score by Verification Providers and Node Operators to score natural persons in social contexts unrelated to the protocol.

Article 50 transparency obligations. Article 50(2) requires providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video, or text content to ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated. Article 50(4) requires deployers of an AI system that generates or manipulates image, audio, or video content constituting a deep fake to disclose that the content has been artificially generated or manipulated. The Trust Identity Protocol is

designed to supply the machine-readable marking required by Article 50(2) through the CNA-2.2 content fingerprint and the Origin Code system (codes OH, AA, AG, MX, defined in Part V). The Global Seal of Trust and the reader-facing badge supply the human-readable disclosure contemplated by Article 50(4). Providers and deployers using TIP to satisfy their Article 50 obligations should consult Section 9.1.1.bis of the AI Office implementing guidance issued from time to time.

Article 14 human oversight. Any optional AI-assisted dispute pre-classification component of the protocol is subject to the human-in-the-loop requirement structurally enforced by the bonded juror final-determination rule in Part VI.

Article 4 AI literacy. The AI Lab maintains a written AI literacy program for its officers, employees, and consultants. Founding Node Operator Agreements and Verification Provider accreditation agreements include an analogous obligation for the counterparty's personnel.

Article 95 voluntary codes of conduct. The AI Trust Council is constituted to be the kind of multi-stakeholder body contemplated by Article 95. The AI Lab and the Council will, at an appropriate time, seek recognition of the protocol's governance regime as a voluntary code of conduct under Article 95.

Article 56 codes of practice for general-purpose AI. Providers of general-purpose AI models electing to comply with the General-Purpose AI Code of Practice published by the European Commission on 10 July 2025 may use the Trust Identity Protocol as the provenance backbone for the marking and disclosure measures contemplated by the Code.

EU Authorized Representative under Article 22. The Trust Identity Protocol, as published in this whitepaper, does not include any component subject to the obligations of an authorized representative under Article 22. If, on a future application of the AI Act, designation of an authorized representative becomes required, The AI Lab will designate such a representative and will publish the designation on theailab.org.

9.1.2 General Data Protection Regulation, Regulation (EU) 2016/679

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (the "GDPR") applies to the processing of personal data of natural persons in the European Union, including processing by controllers and processors not established in the Union where the processing relates to the offering of goods or services to data subjects in the Union or the monitoring of their behavior within the Union (Article 3).

Lawful basis. Verification Providers process personal data on the lawful basis of consent of the data subject under Article 6(1)(a), supplemented where applicable by performance of contract under Article 6(1)(b) and legitimate interests under Article 6(1)(f). Verification Providers operating under public authority may rely on Article 6(1)(e).

Special category data. The optional biometric binding of a CTID involves the processing of biometric data within the meaning of Article 9(1). Verification Providers process such data on the basis of explicit consent under Article 9(2)(a). The architectural design of TIP-ID materially mitigates the scope of biometric data processing: the biometric template is generated, stored, and matched on the user's device (a WebAuthn resident key authenticator); the biometric template does not leave the device; the Verification Provider receives only the public key of the WebAuthn credential and the attestation that the device verified the user. The Verification Provider

does not store the biometric template.

Article 17 right to erasure and the append-only DAG. The structural tension between the right to erasure and the append-only nature of the federated DAG is addressed by a pseudonymization approach consistent with Recital 26 of the GDPR and with the guidance of the European Data Protection Board on pseudonymization. The CTID stored on the DAG is a pseudonymous identifier derived from a public cryptographic key, not directly identifiable personal data. Upon a verified erasure request, the Verification Provider that issued the CTID deletes the linkage between the CTID and the data subject's real-world identity from the Verification Provider's records. The CTID hash remains on the DAG as an orphaned pseudonymous record with no link to any identifiable natural person. The Node Operator is not required to delete DAG records, because the pseudonymized hash alone, in the absence of the Verification Provider-held linkage, does not constitute personal data within the meaning of Article 4(1).

Article 22 solely automated decision-making. The Trust Score is not a solely automated decision producing legal effects concerning natural persons or similarly significantly affecting them within the meaning of Article 22(1). Where a Trust Score change results from a dispute adjudication, the final determination is made by bonded human jurors, satisfying the meaningful human element. Where a platform uses the Trust Score to inform automated decisions about content visibility, the platform is the controller of that automated decision and is responsible for meeting Article 22 in its own implementation.

Article 35 data protection impact assessment. The AI Lab has conducted a data protection impact assessment for the reference architecture of the Trust Identity Protocol. Verification Providers are required by the accreditation agreement to conduct an Article 35 impact assessment for their own processing operations. A model impact assessment is published at theailab.org/dpia.

Chapter V cross-border transfers. Transfers of personal data from the European Union to Verification Providers established outside the European Union are made on the basis of (a) an adequacy decision adopted by the European Commission where applicable, (b) standard contractual clauses adopted by the European Commission, or (c) binding corporate rules approved by the competent supervisory authority. Verification Providers operating in jurisdictions with sectoral or general adequacy decisions are identified in the Verification Provider Registry.

9.1.3 Digital Services Act, Regulation (EU) 2022/2065

Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 (the "Digital Services Act") imposes obligations on providers of intermediary services, including hosting services and online platforms, with augmented obligations for very large online platforms and very large online search engines under Chapter III, Section 5.

Article 30 trusted flaggers. Accredited Verification Providers may apply to Digital Services Coordinators for designation as trusted flaggers under Article 22, where the Verification Provider's expertise relates to the detection of inauthentic accounts or manipulated content. The Trust Identity Protocol supplies a technical mechanism by which a trusted flagger may communicate to a platform that an account or unit of content has been verified, suspended, or revoked.

Article 35 systemic risk mitigation by very large online platforms. Providers of very large on-

line platforms and very large online search engines are required by Article 34 to assess systemic risks, including risks arising from the dissemination of illegal content, fundamental rights infringements, civic discourse and electoral processes, and gender-based violence and minors. Article 35 requires the implementation of mitigation measures. The Trust Identity Protocol supplies a technical mechanism by which a very large online platform may incorporate authenticated identity and content provenance into its risk mitigation measures.

Article 39 advertising transparency. Where a unit of content is an advertisement, the CNA-2.2 fingerprint and the TIP-CONTENT record permit the platform to associate the advertisement with the verified identity of the advertiser, supporting the transparency obligations of Article 39.

9.1.4 eIDAS and the European Digital Identity Wallet

Regulation (EU) 910/2014 of the European Parliament and of the Council of 23 July 2014, as amended by Regulation (EU) 2024/1183 of 11 April 2024 (“eIDAS 2.0”), establishes the framework for electronic identification and trust services for electronic transactions in the internal market and introduces the European Digital Identity Wallet (the “EUDI Wallet”).

Interoperability with the EUDI Wallet. The AI Lab will, through the AI Trust Council, publish a CTID-to-EUDI-Wallet interoperability profile describing how a CTID issued by a Verification Provider may be presented in, and verified against, an EUDI Wallet credential. The profile will align with the technical specifications published by the European Commission for the EUDI Wallet Architecture and Reference Framework.

Qualified electronic signatures. Verification Providers electing to issue qualified electronic signatures or qualified electronic attestations of attributes under eIDAS 2.0 may use ML-DSA-65 (FIPS 204) within signature suites recognised by the European Telecommunications Standards Institute under ETSI TS 119 312 as such suites are amended to incorporate post-quantum primitives. The selection of ML-DSA-65 as the primary signature scheme of the Trust Identity Protocol is aligned with the trajectory of the European Cybersecurity Certification Framework toward post-quantum readiness.

9.1.5 NIS2 Directive, Directive (EU) 2022/2555

Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 (the “NIS2 Directive”) establishes measures for a high common level of cybersecurity across the Union, including obligations on essential and important entities to implement cybersecurity risk-management measures and to notify significant incidents.

Founding Node Operators and Verification Providers established in the European Union and meeting the size thresholds and sectoral criteria of NIS2 may qualify as essential or important entities under Article 3. Such entities are required by Article 23 to notify significant incidents to the national computer security incident response team or the competent authority through (a) an early warning within 24 hours, (b) an incident notification within 72 hours, and (c) a final report within one month. The Founding Node Operator Agreement and the Verification Provider accreditation agreement include incident notification clauses aligned with these NIS2 timelines.

9.1.6 Cyber Resilience Act, Regulation (EU) 2024/2847

Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 (the “Cyber Resilience Act”) imposes essential cybersecurity requirements on products with digital elements placed on the European Union market. The principal obligations apply 36 months after entry into force.

Reference implementations of the Trust Identity Protocol published by The AI Lab are designed to meet the essential cybersecurity requirements set out in Annex I of the Cyber Resilience Act, including secure-by-default configuration, protection against unauthorised access, vulnerability handling processes, the provision of a software bill of materials, and a five-year support period from the date of placing on the market. Manufacturers of products with digital elements incorporating reference implementations of the Trust Identity Protocol are responsible for the conformity assessment of their own products.

9.1.7 Council of Europe Framework Convention on AI

The Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature on 5 September 2024, establishes principles for the activities within the lifecycle of AI systems consistent with human rights, the functioning of democracy, and the rule of law. The Trust Identity Protocol is designed to support each of the operative principles of the Framework Convention, including human dignity and individual autonomy, equality and non-discrimination, respect for privacy and personal data protection, transparency and oversight, accountability and responsibility, reliability, and safe innovation.

9.2 United States

9.2.1 Federal Trade Commission Act Section 5

Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, prohibits unfair or deceptive acts or practices in or affecting commerce. The Federal Trade Commission has, in recent guidance and enforcement actions, applied Section 5 to misrepresentations concerning the provenance, authenticity, and origin of digital content.

The Trust Identity Protocol is designed to supply a substantiated, machine-verifiable claim concerning content provenance and operator identity. Verification Providers issuing CTIDs and Node Operators recording CNA-2.2 fingerprints are required by the accreditation agreement and the Node Operator Agreement to ensure that the representations they make about the content and identities to which the protocol attaches are substantiated and not misleading. The semantics of the Trust Score, the meaning of the Global Seal of Trust, and the meaning of each Origin Code are documented normatively in this whitepaper and in the canonical TIP Protocol Specification.

The Trust Identity Protocol does not modify the substantive duties of platforms or publishers under Section 5. A platform or publisher using TIP to present a Trust Score or a Global Seal of Trust to a user is responsible for the accuracy and non-deceptive presentation of such information to that user in the platform’s own user interface.

9.2.2 NIST AI Risk Management Framework

The National Institute of Standards and Technology AI Risk Management Framework Version 1.0 (NIST AI 100-1, January 2023) and its companion Playbook articulate the GOVERN-MAP-MEASURE-MANAGE functions and an extensive catalog of categories and subcategories applicable to the lifecycle of AI systems and to the institutional context in which they operate. The compliance crosswalk in Appendix E maps the controls of the Trust Identity Protocol to the categories and subcategories of the NIST AI RMF.

9.2.3 Executive Order 14110 and successor instruments

Executive Order 14110 of 30 October 2023 on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, and the instruments issued under it (including the Office of Management and Budget Memorandum M-24-10), establish the federal posture toward AI provenance, including the obligation of federal agencies to use, where feasible, content provenance and authentication tools. The Trust Identity Protocol is positioned to be a content provenance and authentication tool for the purposes of those instruments.

9.2.4 State law mosaic

State law in the United States imposes obligations on biometric data collection, AI content disclosure, and consumer privacy. Verification Providers and Node Operators operating in the United States are responsible for assessing the application of these statutes to their own activities. The following are non-exhaustive examples relevant to the Trust Identity Protocol.

California. The California AI Transparency Act (Senate Bill 942, 2024) and the California AI Content Provenance legislation (Assembly Bill 853, 2024) require certain large online platforms and providers of generative AI systems to apply latent disclosures and provenance metadata to AI-generated content and to make a provenance reading tool available. The CNA-2.2 fingerprint, the Origin Code system, and the TIP-CONTENT record supply provenance metadata in machine-readable form suitable for use by platforms subject to these statutes.

Illinois, Texas, Washington. The Illinois Biometric Information Privacy Act (740 ILCS 14), the Texas Capture or Use of Biometric Identifier Act (Tex. Bus. & Com. Code § 503.001), and the Washington statute on biometric identifiers (RCW 19.375) impose obligations on the collection, storage, and use of biometric identifiers. The architectural decision in TIP-ID to retain biometric templates on the user's device, transmitting only the public key of a WebAuthn credential to the Verification Provider, materially reduces the scope of biometric processing by Verification Providers under these statutes.

Colorado, New York City. Statutes including the Colorado AI Act (Senate Bill 24-205) and New York City Local Law 144 of 2021 (Automated Employment Decision Tools) impose obligations on developers and deployers of certain AI systems. The Trust Identity Protocol is not a high-risk AI system within the meaning of these statutes; Verification Providers and Node Operators are responsible for assessing the application of these statutes to their own activities.

9.3 United Kingdom

9.3.1 Online Safety Act 2023

The Online Safety Act 2023 imposes duties on providers of user-to-user services and of search services in respect of illegal content, content harmful to children, and content involving particular forms of harm. Section 12 (and the accompanying age-assurance codes of practice issued by Ofcom) requires providers of services likely to be accessed by children to implement age-assurance measures.

A CTID issued by a Verification Provider may carry an optional age proof attribute, supplied without disclosing the underlying birth date or identity document, in conformance with privacy-preserving age assurance principles. The protocol does not require any provider of on-line services to use age proof attributes from a CTID, but supplies the technical mechanism by which such providers may, if they elect, meet their Section 12 duties using a privacy-preserving mechanism.

9.3.2 United Kingdom General Data Protection Regulation and Information Commissioner's Office guidance

The United Kingdom General Data Protection Regulation, read together with the Data Protection Act 2018, applies to the processing of personal data of data subjects in the United Kingdom. The architectural commitments described in Section 9.1.2 (lawful basis, special category data, right to erasure, automated decision-making, impact assessment, cross-border transfers) apply analogously to processing within the scope of the United Kingdom General Data Protection Regulation, subject to the differences in the United Kingdom regime including the United Kingdom adequacy framework for cross-border transfers.

9.4 New Zealand, Privacy Act 2020

The Privacy Act 2020 (Act No. 31 of 2020) and the Information Privacy Principles (IPPs) under that Act apply to the collection, use, disclosure, and storage of personal information by agencies in New Zealand. IPP 12 (cross-border disclosure) imposes obligations on the disclosure of personal information to a foreign person or entity. The Part 6, Subpart 1 notifiable privacy breach rules require notification of certain privacy breaches to the Office of the Privacy Commissioner and to affected individuals.

The Trust Identity Protocol architecture is consistent with the IPPs. Where a Verification Provider is established in New Zealand or processes personal information of natural persons in New Zealand, the Verification Provider is required by the accreditation agreement to comply with the Privacy Act 2020, to designate a Privacy Officer where required, and to notify privacy breaches in accordance with Part 6, Subpart 1.

9.5 India, Digital Personal Data Protection Act 2023

The Digital Personal Data Protection Act 2023 (Act No. 22 of 2023) establishes the framework for the processing of digital personal data in India. The Act requires consent of the data principal for processing, subject to enumerated exceptions, and imposes obligations on data fiduciaries including reasonable security safeguards, notification of personal data breaches, and the appointment of a Data Protection Officer for significant data fiduciaries.

The Trust Identity Protocol is consistent with the consent-based framework of the Act and is designed to interoperate with the Data Empowerment and Protection Architecture consent manager pattern through the optional issuance of consent receipts attached to CNA-2.2 fingerprints. Verification Providers operating in India are required by the accreditation agreement to comply with the Act and with the rules issued under it.

9.6 OECD AI Principles

The OECD Recommendation of the Council on Artificial Intelligence, adopted on 22 May 2019 and updated on 3 May 2024, sets out five values-based principles for the responsible stewardship of trustworthy AI: inclusive growth, sustainable development and well-being; respect for the rule of law, human rights and democratic values, including fairness and privacy; transparency and explainability; robustness, security and safety; and accountability. The Trust Identity Protocol is designed to support each of these principles. The compliance crosswalk in Appendix E maps the controls of the protocol to each principle.

9.7 Standards organization engagement

The AI Lab, through the AI Trust Council, will pursue engagement with the following standards organizations concerning the development of international standards relevant to the Trust Identity Protocol.

Standards organization	Workstream of interest
ISO/IEC JTC 1/SC 27 (International Organization for Standardization, Joint Technical Committee 1, Subcommittee 27)	Information security, cybersecurity, and privacy protection
World Wide Web Consortium (W3C)	Verifiable Credentials, Decentralized Identifiers, Content Authenticity
Internet Engineering Task Force (IETF)	Cryptography (CFRG), web authentication, transport security
European Telecommunications Standards Institute (ETSI)	Electronic Signatures and Infrastructures, Quantum-Safe Cryptography
National Institute of Standards and Technology (NIST)	Post-quantum cryptography standardization, AI Risk Management Framework, Cybersecurity Framework
International Press Telecommunications Council (IPTC)	Content metadata standards for journalism
Coalition for Content Provenance and Authenticity (C2PA)	Interoperability with C2PA manifests where the publisher elects to publish both

9.8 Export control posture

Cryptographic implementations within the Trust Identity Protocol are designed to qualify, where source code is made publicly available, for the publicly available source code exclusion under the Wassenaar Arrangement General Software Note and the mass market exception under United States Export Administration Regulations § 740.17. The AI Lab files Commerce Control List classifications and notifications as required by the United States Bureau of

Industry and Security and complies with the export control regimes of the jurisdictions in which it operates.

9.9 Antitrust posture

The TIPCL-1.0 fee schedule set out in Part X is uniform across all licensees within each tier and income group, is published in advance, is not negotiated on a per-licensee basis, and is not conditioned on any restriction unrelated to the practice of the Trust Identity Protocol. The patent license granted under TIPCL-1.0 Section 8 is royalty-free. The Apache License 2.0 conversion provision establishes a fixed horizon at which the protocol becomes available under a permissive open source license. These features are designed to comply with the fair, reasonable, and non-discriminatory principles applicable to licensing of standardized technologies, and to mitigate antitrust risk associated with the possible emergence of the protocol as a de facto standard.

Contact The AI Lab regarding Regulatory Alignment

General Counsel: legal@theailab.org AI Trust Council: council@theailab.org Licensing: licensing@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part X: Governance and Licensing

This Part describes the legal and institutional architecture under which the Trust Identity Protocol is published, evolved, licensed, and enforced. The technical layers described in Parts III through VIII operate within the framework set out below. Reference to this Part is incorporated into every Founding Node Operator Agreement, every Verification Provider accreditation, and every Commercial License issued by The AI Lab.

10.1 The AI Lab Intelligence Unobscured, Inc.

The Trust Identity Protocol™ is published, maintained, and stewarded by **The AI Lab Intelligence Unobscured, Inc.**, a corporation incorporated under the General Corporation Law of the State of Delaware on January 28, 2026, with United States Employer Identification Number 41-3998789. The corporation's principal executive offices are located in Wilmington, Delaware, United States. The corporation is registered to do business in the State of New Jersey as a foreign corporation.

The capital structure of The AI Lab, as restated by the Amended and Restated Certificate of Incorporation filed with the Delaware Secretary of State on May 12, 2026 (Delaware Service Request No. 20262490401), comprises 10,000,000 shares of Class A Common Stock and 10,000,000 shares of Class B Common Stock, with a voting ratio of ten to one. The Founder and Chairman, Dinesh Mendhe, holds the Class A shares. The capital structure includes sunset provisions on the Class A voting preference described in the Bylaws and the Voting Agreement.

The 2026 Stock Plan, adopted by the sole director on May 12, 2026, authorizes the issuance of up to 750,000 shares of Class B Common Stock to employees, advisors, and consultants.

Equity participation in The AI Lab is restricted, subject to a right of first refusal in favor of the corporation, and conditioned on vesting provisions and Section 83(b) election protocols described in the Stock Plan.

The corporation files annual franchise tax reports with the Delaware Division of Corporations and federal corporate income tax returns on Form 1120 with the United States Internal Revenue Service. The corporation maintains a Minute Book containing the records of corporate action required by the Delaware General Corporation Law.

The institutional separation between (a) the corporation that publishes and maintains the protocol and (b) the multi-stakeholder body that ratifies protocol changes is established in Section 10.2 and is structurally enforced by the AI Trust Council Charter.

10.2 The AI Trust Council

10.2.1 Establishment and independence

The AI Trust Council (the “Council”) is the independent multi-stakeholder body established by The AI Lab to ratify the governance, evolution, and enforcement of the Trust Identity Protocol. The Council is constituted under a written Charter ratified by the sole director of The AI Lab and published at theailab.org/ai-trust-council. The Charter is incorporated by reference into this Part. The Council was conceived, designed, and convened by Dinesh Mendhe, who is also the creator and founder of the Trust Identity Protocol, the AI Trust Registry, the AI Trust ID, and the Human Trust ID system to which the Council’s governance applies. **The AI Trust Council was convened on the third day of May, 2026.**

The Council is structured to be the kind of multi-stakeholder body contemplated by Article 95 of Regulation (EU) 2024/1689 (the “EU AI Act”) concerning voluntary codes of conduct, and is structured to satisfy the independence indicia that the AI Office and analogous regulatory bodies in other jurisdictions are expected to apply when evaluating governance regimes for content provenance and verifiable identity infrastructure.

The Council’s independence is structurally enforced by the Charter through the provisions summarized in Section 10.2.4.

10.2.2 Founding Members

The Council’s Founding Members, who serve from the date of the Council’s first plenary session through the conclusion of the Founding Period unless replaced earlier in accordance with the Charter, are:

Seat	Member	Notes
Founding Chair	Joshua Baron	Independent Member, principal responsibility for Council proceedings

Seat	Member	Notes
Founder Seat	Dinesh Mendhe	Founder and creator of the Trust Identity Protocol, the AI Trust Council, the AI Trust Registry, the AI Trust ID, and the Human Trust ID system; sole inventor named on the underlying United States provisional patent applications; Member with declared interest as Founder of The AI Lab, subject to a defined recusal scope (see below)
Founding Member	Ross Thorpe	Independent Member
Founding Member	Issa Nesheiwat	Independent Member
Ex Officio Observer	Chief Executive Officer of The AI Lab Intelligence Unobscured, Inc. (currently Dr. Sofia Martinez Gonzalez), serving by virtue of office	Observer capacity; attends to provide operational information; may be excluded from executive session under the procedure in the Charter

The Founder Seat holder shall recuse from deliberations on (a) enforcement actions against The AI Lab, (b) Commercial License fee schedule amendments, and (c) any matter in which the Founder Seat holder has a direct financial interest.

The Founding Chair and the two independent Founding Members (Baron, Thorpe, Nesheiwat) are independent of The AI Lab. The Founder Seat carries a declared interest and a defined recusal scope. The Ex Officio Observer attends in an observer capacity. This composition is designed to satisfy the independence threshold applied by regulators evaluating multi-stakeholder governance bodies.

10.2.3 Decision rights

The Council ratifies the following categories of action:

1. Amendments to the canonical TIP Protocol Specification, including version increments.
2. Amendments to the Canonical Normalization Algorithm, including the introduction of new normalization variants (CNA-IMG, CNA-VID, and successor variants).
3. Accreditation criteria for Verification Providers, and decisions to grant, suspend, or revoke individual Verification Provider accreditations.
4. Suspension or revocation of Founding Node Operator status for cause.
5. Modifications to the Trust Score sub-score weights and to the threshold values for Blocking Items B1 through B6.
6. Standards organization engagement positions (ISO/IEC JTC 1/SC 27, W3C, IETF, ETSI, NIST).
7. Statements made by The AI Lab to regulatory authorities concerning the conformance of the protocol with applicable law.

8. The transition from the Founding Period to the Network Period described in Part VII.
9. Charter amendments.

Council ratifications under (a) through (h) are required for the action to take effect. Charter amendments under (i) require both supermajority Council ratification and sole director consent of The AI Lab.

The AI Lab shall not depart from a duly ratified Council decision except by written instrument signed by the sole director, recorded in the Minute Book of the corporation, and published on theailab.org/ai-trust-council within fourteen days. Any such departure shall identify the Council decision departed from and shall state the reasons.

10.2.4 Independence covenants

The Charter contains, and the Council operates under, the following independence covenants:

1. **Quorum:** matters are taken up only when the Founding Members and the Founder Seat holder are present or have submitted a written position in accordance with the Charter.
2. **Decision threshold:** matters pass by three-of-five constituency supermajority for protocol amendments, Verification Provider accreditation suspensions, Charter amendments, and statements to regulatory authorities. Routine matters proceed under the consensus procedure described in the Charter.
3. **Conflict of interest:** mandatory written disclosure on appointment and annually thereafter; mandatory recusal from any matter in which the Member has a direct or indirect financial or personal interest.
4. **Right of dissent:** any Member may publish a written dissent attached to the Council's record of decision.
5. **Transparency:** minutes of plenary sessions are published within thirty days, redacted only for personal data and for confidential commercial information of third parties.
6. **Annual transparency report:** the Council publishes an annual transparency report in January of each year, summarizing the prior year's decisions, dissents, and accreditation actions.
7. **No instruction from The AI Lab:** no officer or director of The AI Lab may instruct, direct, or financially condition the conduct of a Member in connection with Council deliberations. The Founder Seat holder is bound by this covenant when acting in the Founder Seat capacity.
8. **Removal:** Members may be removed only for cause (defined narrowly in the Charter as incapacity, conviction of a crime of moral turpitude, or material breach of the Charter), by a three-of-five constituency supermajority of the remaining Members. The AI Lab does not retain the power to remove a Member.
9. **Term:** three-year staggered terms with a two-term limit, except the Founder Seat which is held during the Founder's life or until resignation.
10. **Succession:** the remaining Members select their successors by the consensus procedure described in the Charter, subject to a public call for candidates and a published selection record.
11. **Compensation:** Members may be compensated by The AI Lab for time spent on Council duties at a rate published annually; compensation does not give The AI Lab the right to instruct.

10.2.5 AI Trust Advisory Board

The Council is supported by an AI Trust Advisory Board (the “Advisory Board”). The Advisory Board is a consultative body composed of subject-matter experts appointed by the Council Chair, on the recommendation of any Member or of The AI Lab, and confirmed by the consensus procedure described in the Charter. Advisors hold no seat on the Council, do not participate in Council votes, and do not direct protocol matters. The Advisory Board’s role is to contribute domain expertise to specific workstreams (including but not limited to cryptography, journalism, public policy, education, regulation, civil society, information security, and content provenance) at the request of the Council or of any Member.

Advisors serve at the pleasure of the Council Chair for a renewable term of two years and may be removed by the Council Chair without cause or by the consensus procedure described in the Charter for cause defined under Section 10.2.4(8). Advisors are bound by the same conflict disclosure obligations as Members and shall recuse from any workstream in which they hold a direct or indirect financial or personal interest.

The constitution of the Advisory Board at the date of publication of this whitepaper is:

Advisor	Capacity	Affiliation
John DiVuolo, PMP, ITIL, CISSP	Information security and risk management	Director of Information Security, Rutgers University
Dale Whittaker	Higher education, philanthropy, and equity-focused artificial intelligence	Advisor and Principal Officer, US Education Research and Development, Gates Foundation
Nate Angell	Open knowledge, community, and adoption	Founder, Nudgital; formerly Director of Communications and Community, Creative Commons; formerly Director of Marketing, Hypothesis

The Advisory Board is expected to expand during the Pre-Genesis Period and the Founding Period as additional workstreams are identified and as additional expertise is required for the orderly governance of the protocol.

10.2.6 AI Trust Ambassadors

The Council further recognizes a public-advocacy body designated the AI Trust Ambassadors (the “Ambassadors”). Ambassadors are individuals who carry the mission of the Trust Identity Protocol into the communities, regions, industries, and disciplines they serve. Ambassadors are not Members of the Council, do not steward the protocol, do not advise specific workstreams in the manner of the Advisory Board, and do not hold seats. The Ambassador function is outward-facing public advocacy in conversations with creators, institutions, regulators, and platforms.

Ambassadors are nominated by any Council Member, Advisor, or the sole director of The AI Lab, reviewed by the Council Chair, and confirmed by the consensus procedure described in

the Charter. Ambassadors serve renewable two-year terms and may be removed by the Council Chair without cause. Ambassadors are bound by a code of conduct published by the Council and may not bind the Council, the Advisory Board, or The AI Lab without prior written authorization.

The Ambassador roster is established during the Pre-Genesis Period and the early Founding Period and is published, with annual updates, at theailab.org/ai-trust-council.

10.2.7 Activation and Network Period transition

The Charter takes effect on the publication date of this whitepaper. The Council holds its first plenary session within sixty days. At the conclusion of the Founding Period, the Council reauthorizes itself under a renewed Charter expanding membership to include representatives of Node Operators, Verification Providers, publishers, civil society organizations, and academic institutions, with the Founder Seat preserved and the ex officio CEO observer seat preserved.

10.3 TIPCL-1.0 License Summary

The TIP Protocol Code License, Version 1.0 (“TIPCL-1.0”), is the license under which implementations of the technical framework described in this whitepaper are governed. TIPCL-1.0 is published in its entirety at theailab.org/tip-license and is incorporated by reference as Appendix F. The summary in this Section is provided for the convenience of the reader and does not modify the operative text of TIPCL-1.0.

10.3.1 Free Tier eligibility

A license under TIPCL-1.0 is granted on a no-fee basis to the following categories of person and entity (the “Free Tier”):

1. Individual natural persons with annual gross revenue below one hundred thousand United States dollars (US\$100,000).
2. Small businesses, however organized, with annual gross revenue below one hundred thousand United States dollars (US\$100,000).
3. Nonprofit organizations, non-governmental organizations, and registered charities, irrespective of revenue.
4. Educational institutions, irrespective of revenue.
5. Government entities at any level, irrespective of revenue.
6. Journalism organizations, for editorial use only, irrespective of revenue.
7. Research and development or testing within the per-organization user count and duration ceilings published from time to time by The AI Lab at theailab.org/tip-license.

The Free Tier is the canonical no-fee threshold. References in any prior publication, draft, or third-party summary to a Free Tier ceiling other than US\$100,000, including any reference to a US\$500,000 figure, are superseded by this Section and by the operative text of TIPCL-1.0.

10.3.2 Commercial License tier schedule

Persons and entities not eligible for the Free Tier require a Commercial License from The AI Lab. The annual fee schedule for Commercial Licenses is as set out below (the “Commercial Tier Schedule”):

Tier	Annual revenue band	Annual fee (USD)
Micro	US\$100,000 to US\$250,000	US\$500
Seed	US\$250,000 to US\$500,000	US\$1,100
Starter	US\$500,000 to US\$5,000,000	US\$2,750
Growth	US\$5,000,000 to US\$25,000,000	US\$8,250
Business	US\$25,000,000 to US\$100,000,000	US\$27,500
Enterprise	US\$100,000,000 to US\$500,000,000	US\$71,500
Corporate	US\$500,000,000 to US\$2,000,000,000	US\$165,000
Strategic	US\$2,000,000,000 to US\$10,000,000,000	US\$385,000
Global	US\$10,000,000,000 and above	US\$550,000

The Commercial Tier Schedule is the canonical fee schedule for Commercial Licenses. The Commercial Tier Schedule states the fee applicable to a licensee whose country of principal place of business is a high-income economy. For a licensee with annual gross revenue below US\$25,000,000, the fee for its tier is multiplied by 1.00 for a high-income economy, 0.50 for upper-middle, 0.25 for lower-middle, and 0.12 for low income, subject to a minimum of US\$125. At US\$25,000,000 of annual gross revenue and above, a single fee applies worldwide. The fees are uniform across all licensees within each tier and income group, are published in advance, are not negotiated on a per-licensee basis, and are not conditioned on any restriction unrelated to the practice of the Trust Identity Protocol. This structure is designed to comply with fair, reasonable, and non-discriminatory principles applicable to licensing of standardized technologies.

Tier nomenclature used in prior publications, drafts, or third-party summaries, including the nomenclature “Sentinel,” “Guardian,” “Silver,” “Gold,” and “Platinum,” is retired and is no longer used. The retired names are documented here solely to inform readers of legacy materials that no longer reflect The AI Lab’s licensing schedule.

10.3.3 Grace period and renewal

A licensee whose annual gross revenue crosses a tier band during a license year is granted a grace period until the end of the then-current license year to obtain a Commercial License at the appropriate tier. Renewal terms are described in the operative text of TIPCL-1.0.

10.4 Patent licensing under TIPCL-1.0 Section 8

The AI Lab has filed five United States provisional patent applications covering the inventions implemented in the Trust Identity Protocol. Dinesh Mendhe is the sole inventor named on all five United States provisional patent applications, and the applications are assigned to The AI Lab Intelligence Unobscured, Inc. The five provisional applications, organized in Claim Groups A through BB, were filed as follows: (1) the foundational TIP Protocol v1.0 application filed March 11, 2026 (Application Number 64/003,066, Claim Groups A through E); (2) the TIP Protocol v2.0 refinements application filed March 14, 2026 (Application Number 64/005,947, Claim Groups F through J); (3) the TIP Protocol v2.0 Content-Layer application filed April 7, 2026 (Application Number 64/031,648, Attorney Docket AILAB-2026-PROV-03, Claim Groups K through P); (4) the TIP Protocol v2.0 Identity-Layer application filed May 5, 2026 (Application Number 64/058,152, Attorney Docket AILAB-2026-PROV-04, Claim Groups Q through

X, comprising eight inventions); and (5) the TIP Protocol v2.0 Community Verification-Layer application filed May 15, 2026 (Application Number 64/067,319, Claim Groups Y through BB, comprising four inventions). The earliest priority date in the portfolio is March 11, 2026. Collectively, the applications cover the foundational three-layer protocol architecture (TIP-ID, TIP-CONTENT, TIP-TRUST), the four-layer biometric verification stack, the mandatory origin declaration system (OH, AA, AG, MX Origin Codes), the deterministic trust scoring engine, the Canonical Normalization Algorithm (CNA-1 and CNA-2), the dual-mode (Publisher Mode and Creator Mode) verification flow, the content versioning semantics, the content scope extraction mechanism, the multi-layer verification delivery architecture, the content-type extensible normalization framework, the typed identity taxonomy with relational attribution modes, the multi-officer organization governance with platform-hosted content attestation, the centrally-configurable protocol policy framework, the org-scoped contributor registry with multi-author authorship array (CNA-2.2), the multi-model consensus classification with weighted accuracy scoring, the Classification Provider Registry with DAG-anchored accreditation, the two-phase protocol review with private self-correction window and three-step reader verification, and the staking-based community trust verification with reviewer badge progression, together with related inventions. The corresponding non-provisional applications will be filed within the deadlines required by the United States Patent and Trademark Office. International priority will be claimed under Article 4 of the Paris Convention for the Protection of Industrial Property and the Patent Cooperation Treaty.

TIPCL-1.0 Section 8 grants to every licensee in good standing a royalty-free, non-exclusive, non-transferable license under the issued and pending patent claims of The AI Lab, limited to the practice of the Trust Identity Protocol as described in the canonical TIP Protocol Specification and in this whitepaper. The Section 8 patent license is conditioned on (a) compliance with the operative terms of TIPCL-1.0, (b) attribution of The AI Lab as required by TIPCL-1.0 and by CC BY 4.0 for the underlying specification, and (c) the defensive termination provision of Section 8.

The defensive termination provision states that the Section 8 patent license terminates automatically with respect to a licensee that initiates patent infringement litigation against The AI Lab or against any other licensee in good standing alleging infringement by the practice of the Trust Identity Protocol. The provision is structured to discourage patent assertion against the protocol community without conferring a right to assert infringement against parties outside the community.

10.5 Trademark policy

The AI Lab maintains trademark applications before the United States Patent and Trademark Office for the marks identified in the following table.

Mark	USPTO Serial No.	Status
The AI Lab	99597929	Pending
Trust Identity Protocol	99603145	Pending
Global Seal of Trust	99607461	Pending
AI Trust Council	99749088	Pending

The marks are used in commerce by The AI Lab. Use of the marks by licensees, Founding Node Operators, Verification Providers, and other authorized parties is governed by a Trademark Usage Policy published at theailab.org/trademark, available on request from legal@theailab.org. The Trademark Usage Policy describes (a) the permitted forms of nominative fair use, (b) the conditions under which a licensee may describe an implementation as “TIP-compliant” or “Built on the Trust Identity Protocol,” (c) prohibited uses, and (d) the procedure for trademark complaints.

Authorized parties are required to display the marks with appropriate symbols (™ for marks for which registration is pending, ® for marks for which registration has been granted) on first use in each document, screen, or surface, and to include in such document, screen, or surface a notice that the marks are owned by The AI Lab.

Nominative fair use of the marks is expressly permitted for the following purposes without prior authorization: (i) reference to the protocol or to The AI Lab by regulators, supervisory authorities, courts, and legislative bodies in the conduct of their official functions; (ii) reference by academic researchers in peer-reviewed publications; (iii) reference by journalists in editorial coverage; and (iv) reference by standards-setting organizations in published standards. Use under this paragraph does not require a license fee or a license agreement, provided that the use is accurate, does not imply endorsement by The AI Lab, and does not modify the marks.

10.6 Apache 2.0 conversion provision

TIPCL-1.0 contains a self-executing conversion provision that, on January 1, 2031, automatically converts the license under which the canonical TIP Protocol Specification is published from CC BY 4.0 to the Apache License, Version 2.0, and converts the license under which the reference implementations are published from TIPCL-1.0 to the Apache License, Version 2.0.

The conversion provision is structured to provide a fixed horizon at which the protocol and its reference implementations become available under a permissive open source license, while preserving for the Founding Period and the initial Network Period the licensing architecture necessary to (a) compensate The AI Lab for the research, development, and operational costs of bringing the protocol to operational status, (b) maintain quality and security through Commercial License obligations, and (c) ensure that the protocol does not fragment into incompatible variants before the standards organization engagement described in Section 9.7 has had time to consolidate the canonical specification.

The conversion does not affect (a) trademark rights of The AI Lab, which are not transferred by the conversion, (b) Commercial License fees accrued and unpaid as of the conversion date, or (c) the AI Trust Council’s role in ratifying protocol changes.

10.7 Verification Provider role, accreditation, and obligations

A Verification Provider is an organization accredited by the AI Trust Council, on the recommendation of The AI Lab, under Part IV to issue CTIDs to natural persons and to organizations. Verification Providers operate under written accreditation agreements with The AI Lab.

The public-facing commitments of a Verification Provider are:

1. An annual accreditation fee paid to The AI Lab in the amount published from time to time at theailab.org/tip-verification-provider.

2. An annual independent audit of the Verification Provider’s identity verification practices, conducted by an auditor acceptable to The AI Lab.
3. The maintenance of a warrant canary published on the Verification Provider’s website, attesting to the absence of compelled disclosure orders received during the reporting period to the extent permitted by the law of the Verification Provider’s jurisdiction.
4. A jurisdiction declaration identifying the jurisdiction or jurisdictions under which the Verification Provider issues CTIDs and the jurisdictions in which the Verification Provider maintains personal data of CTID holders.
5. Compliance with applicable data protection law in each jurisdiction in which the Verification Provider operates, including (without limitation) Regulation (EU) 2016/679 (the General Data Protection Regulation), the United Kingdom General Data Protection Regulation, the New Zealand Privacy Act 2020, the Indian Digital Personal Data Protection Act 2023, and applicable state and federal data protection law in the United States.
6. Where the Verification Provider is established or operates in the European Union and meets the threshold criteria, compliance with Directive (EU) 2022/2555 (the NIS2 Directive) as transposed in the relevant Member State, and registration as a “trust service provider” under Regulation (EU) 910/2014 as amended (eIDAS 2.0) where the Verification Provider elects to issue qualified electronic signatures or qualified electronic attestations of attributes.

Verification Providers do not issue currency, securities, or other instruments of value. Verification Providers issue identity attestations. The economic terms under which Verification Providers will operate beyond the public-facing commitments set out above are being designed by The AI Lab and the AI Trust Council and are not part of the public launch of the Trust Identity Protocol described in this whitepaper. No representation is made in this whitepaper concerning the future economic terms applicable to Verification Providers, and no licensee, partner, or third party is entitled to rely on any statement, draft, or third-party summary that purports to describe such future economic terms. Inquiries from prospective Verification Providers should be directed to licensing@theailab.org.

10.8 EU AI Act classification posture

The Trust Identity Protocol is not an “AI system” within the meaning of Article 3(1) of Regulation (EU) 2024/1689 (the “EU AI Act”). The protocol is a cryptographic identity and provenance framework. The core operations of the protocol consist of cryptographic key generation, deterministic content normalization, digital signature production and verification, and the recording of signed events on an append-only directed acyclic graph. These operations do not constitute machine-based systems designed to operate with varying levels of autonomy producing outputs that influence physical or virtual environments within the meaning of Article 3(1).

Three operational components of the protocol carry AI elements and are addressed below.

Biometric verification under TIP-ID. The optional biometric binding of a CTID to a natural person using a WebAuthn resident key authenticator is biometric **verification** (one-to-one comparison against a template held on the user’s device) and is not biometric **identification** (one-to-many comparison against a database). Annex III, paragraph 1, of the EU AI Act applies to biometric identification systems. Biometric verification systems are excluded from Annex III by virtue of the definition of “biometric identification” in Article 3(34) and the clarification in Recital 54. The TIP-ID biometric binding component is therefore not high-risk under the EU

AI Act.

AI-assisted dispute pre-classification. Reference implementations of the dispute adjudication path described in Part VI permit the optional use of machine-learning models for the preliminary classification of disputes prior to assignment to a bonded jury. The final determination is made by bonded human jurors in accordance with the rules set out in Part VI. The human-in-the-loop requirement is structural, not advisory, and supplies the human oversight required by Article 14 of the EU AI Act for any AI component that might otherwise be characterized as high-risk.

Trust Score computation. The Trust Score described in Part VI is the aggregated output of four sub-scores (Cryptographic, Behavioral, Adjudicated, Network) computed by deterministic rules. The Trust Score evaluates content authenticity and operator behavior within the protocol context. It does not classify natural persons by their social behavior, personality, or personal characteristics across contexts unrelated to the data's original collection. The Trust Score is therefore not a "social scoring" system within the meaning of Article 5(1)(c) of the EU AI Act. The AI Lab and the AI Trust Council will publish, with each protocol release, a notice expressly disclaiming social scoring use cases and identifying the protocol's intended uses.

The Trust Identity Protocol is structured to supply the machine-readable marking and disclosure mechanisms contemplated by Article 50 of the EU AI Act. CNA-2.2 produces deterministic content fingerprints suitable for inclusion as machine-readable marking under Article 50(2). The Origin Code system (codes OH, AA, AG, MX) supplies machine-readable identification of artificially generated or manipulated content. The Global Seal of Trust and the reader-facing badge supply the human-readable disclosure contemplated by Article 50(4).

The AI Trust Council is structured to be the kind of multi-stakeholder body contemplated by Article 95 of the EU AI Act. The AI Lab and the Council will, at an appropriate time after the publication of this whitepaper, engage with the European AI Office concerning recognition of the protocol's governance regime as a voluntary code of conduct under Article 95.

10.9 Cross-jurisdictional governance commitments

In furtherance of the cross-jurisdictional operation of the Trust Identity Protocol, The AI Lab makes the following commitments:

1. **AI literacy.** The AI Lab will maintain a written AI literacy program for its officers, employees, and consultants consistent with the requirements of Article 4 of the EU AI Act. Founding Node Operator Agreements and Verification Provider accreditation agreements include an AI literacy clause requiring the counterparty to maintain an analogous program for its personnel involved in the operation of the protocol.
2. **EU Authorized Representative.** If, on a future application of the EU AI Act to any component of the protocol, The AI Lab is required by Article 22 of the EU AI Act to designate an authorized representative in the European Union, The AI Lab will designate such a representative and will publish the designation on theailab.org.
3. **AI Office engagement.** The AI Lab will, through the AI Trust Council, engage with the European AI Office concerning (i) recognition of the AI Trust Council as a multi-stakeholder body under Article 95 of the EU AI Act, and (ii) the conformance of the protocol with Articles 50(2) and 50(4) of the EU AI Act.

4. **Standards organization engagement.** The AI Lab will, through the AI Trust Council, pursue engagement with the International Organization for Standardization Joint Technical Committee 1 Subcommittee 27 (ISO/IEC JTC 1/SC 27), the World Wide Web Consortium (W3C), the Internet Engineering Task Force (IETF), and the European Telecommunications Standards Institute (ETSI) concerning the development of international standards for content provenance and verifiable identity infrastructure.
5. **Council of Europe Framework Convention on AI.** The protocol is designed to support the human rights, democracy, and rule of law principles set out in the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225, opened for signature September 5, 2024).
6. **Conditional commitments.** The commitments in this Section 10.9 are made in good faith and are conditional on (i) the continued operation of the protocol, (ii) the absence of material adverse change in the regulatory landscape that would render any commitment infeasible or unlawful, and (iii) the receipt by The AI Lab of sufficient cooperation from the relevant regulators, standards organizations, and counterparties.

Contact The AI Lab

Licensing and Commercial Implementation: licensing@theailab.org, theailab.org/tip-license
Founding Node Operator Applications: nodes@theailab.org, theailab.org/founding-node
AI Trust Council Membership: council@theailab.org, theailab.org/ai-trust-council
Verification Provider Accreditation: licensing@theailab.org, theailab.org/tip-verification-provider
General Counsel: legal@theailab.org Press and Public Affairs: press@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Part XI: Roadmap and Forward-Looking Statements

This Part identifies the operational milestones achieved by The AI Lab and the Trust Identity Protocol as of the publication date of this whitepaper, the milestones planned for the twelve-month, twenty-four-month, and thirty-six-month horizons that follow, the principal material risks bearing on the achievement of those milestones, and the conditions on which the transition from the Founding Period to the Network Period and the corresponding reauthorization of the AI Trust Council are conditioned.

Statements in this Part concerning future events, plans, or operational targets are forward-looking statements within the meaning of Appendix J Section J.7. Such statements are based on the present expectations of The AI Lab and are subject to risks and uncertainties, many of which are outside the control of The AI Lab. Actual results may differ materially from those projected. The AI Lab undertakes no obligation to update or revise any forward-looking statement except as may be required by applicable law. The reader is referred to Appendix J for the full safe harbor.

11.1 Operational milestones to date

The following milestones have been achieved by The AI Lab and the Trust Identity Protocol as of the publication date of this whitepaper.

January 28, 2026. Incorporation of The AI Lab Intelligence Unobscured, Inc. in the State of Delaware.

February through April 2026. Drafting of the canonical TIP Protocol Specification through versions v1.0 (internal), v2.0, v3.0 (filed with the United States Copyright Office, Case 1-15116205291, pending), and v4.0.

May 3, 2026. Restructuring of The AI Lab corporate documents to align with the pre-Genesis operational posture.

May 12, 2026. Filing of the Amended and Restated Certificate of Incorporation establishing the dual-class capital structure (Class A and Class B Common Stock, ten-to-one voting ratio) (Delaware Service Request No. 20262490401). Adoption of the 2026 Stock Plan authorizing the issuance of up to 750,000 shares of Class B Common Stock.

June 1, 2026. Publication of the canonical TIP Protocol Specification v5.0 under Creative Commons Attribution 4.0 International License and filing with the United States Copyright Office (Case 1-15175755931, pending) as a derivative work referencing the prior v3.0 filing.

June 2026. Execution of Founding Node Operator Agreements with the seven signed Founding Node Operators identified in Part VII Section 7.3.1 (THE PRESCIENT PACHYDERM LTD, AZ-Logics Private Limited, Apex Modular Solutions LLC, Timpi International Ltd, Lonestar Data Holdings Inc., 6Simplex Software Solutions Pvt Ltd, and The AI Lab Intelligence Unobscured, Inc.), and accession in progress for three further Founding Node Operators (Marist University, Rutgers University, and The Core / BOOM FactCheck), upon countersignature of the Founding Node Operator Agreement. Each signed Founding Node Operator stands up a Full Node deployment in pre-Genesis operating mode.

June 2026. Ratification by the sole director of The AI Lab of the AI Trust Council Charter establishing the independent multi-stakeholder governance body described in Part X Section 10.2.

June 2026. Publication of this whitepaper. Commencement of the Pre-Genesis Period.

June 2026, Genesis Date. Adoption by the sole director of The AI Lab of the Operational Readiness Resolution and determination of the Genesis Date, published not less than three (3) business days in advance on theailab.org. The federated network commences live operation. The Founding Period commences.

11.1.1 Operational Readiness Conditions

The Genesis Date is determined by the sole director on the satisfaction of the following Operational Readiness Conditions, each verified in writing and recorded in the Minute Book of The AI Lab:

1. Execution of Founding Node Operator Agreements with at least three (3) Founding Node Operators in at least two (2) jurisdictions.
2. Stand-up by each Founding Node Operator of a Full Node deployment satisfying the Technical Requirements Specification, in pre-Genesis operating mode, with operational status verified by The AI Lab.
3. Successful completion of synchronization protocol testing between Founding Node Operators, with the published synchronization target indicators met in the pre-Genesis test

window.

4. Ratification of the AI Trust Council Charter by the sole director.
5. Convening of the inaugural plenary session of the AI Trust Council, including the Council's written confirmation to the sole director of its observation that the operational, governance, and licensing conditions for Genesis have been satisfied.
6. Publication of the canonical TIP Protocol Specification v5.0 under Creative Commons Attribution 4.0 International License.
7. Publication of this whitepaper.
8. Filing of the United States Copyright Office derivative application for this whitepaper.
9. Capture of an Internet Archive snapshot of the published whitepaper for independent third-party timestamping.
10. Approval by the sole director of the Genesis Date.

On the determination by the sole director that the Operational Readiness Conditions have been satisfied, the Genesis Date is published on theailab.org not less than three (3) business days in advance of the Genesis Date.

11.2 Twelve-month horizon

The following milestones are planned for the twelve months following the publication of this whitepaper.

First plenary session of the AI Trust Council. The Council holds its inaugural plenary session during the Pre-Genesis Period and not later than sixty days after the Genesis Date, ratifies its inaugural rules of procedure, confirms the satisfaction of the Operational Readiness Conditions as a prerequisite to the sole director's determination of the Genesis Date, accepts the inaugural set of Verification Provider applications, and adopts the conformance program for the protocol.

Accreditation of the inaugural set of Verification Providers. The Council reviews and acts on Verification Provider applications received during the application window opening on the publication of this whitepaper. The Council's accreditation pipeline targets the issuance of the inaugural Verification Provider accreditations within ninety days of the publication of this whitepaper.

Expansion of the federation to an EU Member State Node Operator. The Founding Period network does not currently include a Founding Node Operator established in a Member State of the European Union, an absence acknowledged in Part VII Section 7.6. The AI Lab and the AI Trust Council intend to address this absence within the twelve-month horizon through the accreditation of one or more additional Founding Node Operators in EU Member States.

Engagement with the European AI Office. The AI Lab and the AI Trust Council, on behalf of the protocol, intend to commence engagement with the European AI Office concerning (a) the recognition of the AI Trust Council as a multi-stakeholder body under Article 95 of the EU AI Act and (b) the conformance of the protocol with Articles 50(2) and 50(4) of the EU AI Act, in the manner contemplated by Part IX Section 9.1.1.

Publication of CTID-to-EUDI-Wallet interoperability profile. The AI Lab and the AI Trust Council intend to publish the CTID-to-EUDI-Wallet interoperability profile described in Part IV Section 4.5.3 within the twelve-month horizon.

Publication of software development kits. The Rust, Go, Java, and .NET SDKs identified in Part VIII Section 8.6.2 are scheduled for publication within the twelve-month horizon under the published schedule.

Engagement with at least one standards organization. The AI Lab and the AI Trust Council intend to commence formal engagement with at least one standards organization identified in Part IX Section 9.7 within the twelve-month horizon, with priority on ISO/IEC JTC 1/SC 27 and on the IETF.

First annual transparency report of the AI Trust Council. The Council publishes its inaugural annual transparency report in January 2027 covering the inaugural period of its operation.

11.3 Twenty-four-month horizon

The following milestones are planned for the twelve months following the conclusion of the twelve-month horizon described in Section 11.2.

Throughput expansion of the federation. The federation is expected to expand to support sustained throughput of one thousand entries per second by the conclusion of the twenty-four-month horizon, with the horizontal scaling profile described in Part VII Section 7.7.

Geographic expansion. The federation is expected to include Founding Node Operators in additional jurisdictions including the European Union, East Asia, and the Middle East by the conclusion of the twenty-four-month horizon.

Accreditation of the second set of Verification Providers. The Council is expected to accredit additional Verification Providers expanding the geographic coverage of identity verification to jurisdictions including the European Union, the United Arab Emirates, Singapore, Japan, Brazil, and South Africa.

Publication of CNA-IMG, CNA-VID, and CNA-AUDIO variants. The image, video, and audio variants of the Canonical Normalization Algorithm described in Part V Section 5.1 are scheduled for publication within the twenty-four-month horizon following Council review and conformance test vector production.

Native ML-DSA-65 authenticator deployment. The transition from the enveloped classical signature pattern described in Part IV Section 4.1.4 to native ML-DSA-65 authenticator support is expected to make material progress within the twenty-four-month horizon as the FIDO Alliance and authenticator manufacturers publish ML-DSA-65 authenticator support.

Submission to standards organizations. Formal submissions to at least one of ISO/IEC JTC 1/SC 27, the IETF, the W3C, and ETSI are expected within the twenty-four-month horizon.

Engagement with regulators in additional jurisdictions. Engagement with the United Kingdom Office of Communications, the New Zealand Office of the Privacy Commissioner, the Indian Ministry of Electronics and Information Technology, and the United States Federal Trade Commission concerning the operation of the protocol within those jurisdictions is expected to advance within the twenty-four-month horizon.

11.4 Thirty-six-month horizon

The following milestones are planned for the twelve months following the conclusion of the twenty-four-month horizon described in Section 11.3.

Transition from the Founding Period to the Network Period. The conditions for the transition described in Section 11.6 are expected to be satisfied within the thirty-six-month horizon. The transition is initiated by a supermajority ratification of the AI Trust Council and is accompanied by the reauthorization of the Council under the Network Period composition described in Part X Section 10.2.7.

Throughput expansion. The federation is expected to expand to support sustained throughput of ten thousand entries per second by the conclusion of the thirty-six-month horizon, with the scaling profile described in Part VII Section 7.7.

Native post-quantum operation. The transition to native ML-DSA-65 authenticator support is expected to be substantially complete within the thirty-six-month horizon, with the enveloped classical signature pattern retired for new CTID issuance subject to the Council's notice provisions described in Part IV Section 4.1.4.

Publication of additional SDKs. The Swift, Kotlin, PHP, and Ruby SDKs identified in Part VIII Section 8.6.2 are scheduled for publication within the thirty-six-month horizon.

International standards process advancement. The standards processes commenced in the twelve-month and twenty-four-month horizons are expected to advance toward publication of an international standard within the thirty-six-month horizon.

Engagement with the Council of Europe. Engagement with the Council of Europe concerning the alignment of the protocol with the Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225) is expected to advance within the thirty-six-month horizon.

11.5 Material risks

The achievement of the milestones described in Sections 11.2, 11.3, and 11.4 is subject to material risks. The principal material risks identified by The AI Lab as of the publication date of this whitepaper are as follows.

Regulatory risk. Changes in applicable law, regulation, or executive policy in any jurisdiction in which the protocol operates may materially affect the operational posture of the protocol, the obligations of Founding Node Operators and Verification Providers, and the conformance of the protocol with regulatory expectations. The protocol's design is intended to be resilient to ordinary regulatory evolution; extraordinary regulatory action against the protocol or against its institutional sponsors cannot be reliably forecast.

Counterparty risk. The protocol depends on the performance of Founding Node Operators, Verification Providers, and standards organization counterparties under written agreements. The failure of one or more counterparties to perform under its agreement with The AI Lab may materially affect the operational milestones.

Intellectual property risk. The protocol is the subject of pending applications before the United States Copyright Office and the United States Patent and Trademark Office. Adverse outcomes

in any pending application may materially affect the protocol's intellectual property posture. As of the publication date of this whitepaper, one trademark portfolio refusal has been issued (SR 1-15116637136, the seal portfolio, refused under 17 U.S.C. § 1052(e)) and is not the subject of an appeal. Other applications are pending.

Cryptographic risk. A material reduction in the security of any cryptographic primitive identified in Part III may materially affect the security posture of the protocol. The cryptographic migration path described in Part III Section 3.7.3 is the protocol's structured response. The migration path is designed to be initiable on a cryptanalytic event without disruption of the federated network.

Adversary materialization risk. An adversarial attack model not within the threat model described in Part III Section 3.7 may materialize during the operational period of the protocol. The protocol's response to such an event is a combination of the bonded jury adjudication procedure described in Part VI Section 6.5, the protocol amendment procedure under the Charter, and (where the attack exceeds the protocol's institutional capacity) the engagement of supervisory authorities and standards organizations.

Standards risk. The standards organization engagement described in Part IX Section 9.7 may not converge on a published international standard within the horizons described above, or may converge on a standard incompatible with the canonical TIP Protocol Specification. The protocol's response is the continued operation of the canonical Specification as the reference for the federation, with such accommodations of the resulting standard as the AI Trust Council ratifies.

Market risk. The level of public, institutional, and platform interest in content provenance and verifiable identity may not develop at the rate or to the level necessary to support the operational scaling described in Sections 11.2 through 11.4. The protocol's response is the disciplined cost structure of the federation, the published licensing schedule, and the engagement strategy through standards organizations and regulators.

Operational risk. The operational risks bearing on the federation include the availability and capacity of the underlying internet infrastructure, the availability of cryptographic primitive implementations across the operating system and runtime base, and the response of platforms to integration requests. These risks are addressed by the Founding Node Operator Agreements, the Service Level Agreement, and the integration partnership pipeline.

11.6 Conditions for AI Trust Council Network Period reauthorization

The transition from the Founding Period to the Network Period, and the corresponding reauthorization of the AI Trust Council under the Network Period composition, are conditioned on the satisfaction of each of the following conditions as ratified by supermajority vote of the Council:

1. The publication of the inaugural annual transparency report of the Council, and the publication of at least one subsequent annual transparency report.
2. The accreditation of at least three Verification Providers operating in at least three jurisdictions.
3. The presence of at least one Node Operator in each of the principal regulatory zones

identified in Part VII Section 7.6 (North America, the European Union, the Indo-Pacific).

4. The completion of at least one bonded jury adjudication proceeding to a published determination.
5. The publication of the CTID-to-EUDI-Wallet interoperability profile.
6. The publication of at least one of CNA-IMG, CNA-VID, or CNA-AUDIO.
7. The substantial completion of the transition to native ML-DSA-65 authenticator support, as determined by the Council.
8. The commencement of formal engagement with at least one standards organization identified in Part IX Section 9.7.
9. The establishment of the Network Period economic model for Verification Providers, as ratified by the Council.
10. The publication of the Network Period Charter, expanding voting membership of the Council to include representatives of Node Operators, Verification Providers, publishers, civil society organizations, and academic institutions.

The Council may, on supermajority vote, defer the transition by extending the Founding Period for such period as the Council determines.

11.7 Statement on the design horizon

The protocol is designed with the operational horizon of decades. The cryptographic primitives are selected for durability against the cryptographic developments of the next several decades. The append-only DAG is structured for retention of TIP-CONTENT records on an analogous horizon. The licensing structure converts to the permissive Apache License 2.0 on January 1, 2031. The AI Trust Council Charter contemplates reauthorization and Network Period operation on an indefinite horizon.

The selection of a long design horizon reflects the underlying social purpose. Content authenticity, verifiable identity, and reputational accountability are conditions of a society in which the public can rely on the digital record of its institutions, its journalism, its courts, and its democratic processes. The Trust Identity Protocol is offered as a piece of foundational infrastructure for that condition over the long term.

Contact The AI Lab regarding Part XI

Founding Node Operator Applications: nodes@theailab.org, theailab.org/founding-node AI Trust Council: council@theailab.org, theailab.org/ai-trust-council Standards Engagement: standards@theailab.org General Counsel: legal@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Appendix A: Glossary of Defined Terms

Each Capitalized Term used in this whitepaper has the meaning given to it in this Appendix or, where the term is defined in the substantive Part, in the substantive Part. In the event of

inconsistency between this Appendix and a definition in a substantive Part, the substantive Part controls.

Accreditation Schedule means the published schedule under which a Verification Provider operates, identifying the grades of CTID the Verification Provider may issue and the identity verification practices applicable to each grade.

Adjudicated Sub-Score has the meaning given in Part VI Section 6.2.3.

AI Act means Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence.

AI Office means the European Artificial Intelligence Office established under the AI Act.

The AI Lab means The AI Lab Intelligence Unobscured, Inc., a Delaware corporation with principal executive offices in Wilmington, Delaware, United States.

AI Trust Council or **Council** means the independent multi-stakeholder body established by The AI Lab to ratify the governance, evolution, and enforcement of the Trust Identity Protocol, operating under the Charter.

Apache License 2.0 Conversion Provision has the meaning given in Part X Section 10.6.

Authenticator means a WebAuthn resident key authenticator as described in Part IV Section 4.4.1.

Behavioral Sub-Score has the meaning given in Part VI Section 6.2.2.

Blocking Item means a defined condition (numbered B1 through B6 in Part VI Section 6.4) that materially constrains or suspends the operational utility of a CTID.

Bonded Juror has the meaning given in Part VI Section 6.5.1.

CC BY 4.0 means the Creative Commons Attribution 4.0 International Public License published by Creative Commons Corporation.

Charter means the AI Trust Council Charter ratified by the sole director of The AI Lab and published at theailab.org/ai-trust-council.

Class A Common Stock means the Class A Common Stock of The AI Lab, par value as specified in the Amended and Restated Certificate of Incorporation, carrying ten votes per share.

Class B Common Stock means the Class B Common Stock of The AI Lab, par value as specified in the Amended and Restated Certificate of Incorporation, carrying one vote per share.

CNA-1.0 means the Canonical Normalization Algorithm Version 1.0, applicable to WordPress content as described in Part VIII Section 8.4.1.

CNA-2.2 means the Canonical Normalization Algorithm Version 2.2, the principal normalization specified in Part V.

CNA-IMG, CNA-VID, CNA-AUDIO mean the image, video, and audio variants of the Canonical Normalization Algorithm referenced in Part V Section 5.1.

Commercial License means a license issued under TIPCL-1.0 to a Person not eligible for the Free Tier, at the annual fee identified in the Commercial Tier Schedule.

Commercial Tier Schedule means the nine-tier schedule of Commercial License annual fees set out in Part X Section 10.3.2.

CONTENT_UPDATED means the event published to the federated DAG identifying the modification of a content unit, as described in Part V Section 5.5.2.

Council means the AI Trust Council.

Council of Europe Framework Convention means the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature on 5 September 2024.

Creator Mode has the meaning given in Part V Section 5.4.

Cryptographic Sub-Score has the meaning given in Part VI Section 6.2.1.

CTID means a Cryptographic Trust Identity, the pseudonymous identifier issued by a Verification Provider under Part IV.

Cyber Resilience Act means Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024.

DAG means the federated, append-only directed acyclic graph maintained by the federation of Node Operators on which protocol events are recorded.

Data Empowerment and Protection Architecture or **DEPA** means the consent manager pattern developed in India to operationalize consent-based data sharing.

Digital Services Act means Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022.

DPDPA means the Digital Personal Data Protection Act 2023 of India (Act No. 22 of 2023).

eIDAS 2.0 means Regulation (EU) 910/2014 as amended by Regulation (EU) 2024/1183.

Entry means a record in the federated DAG.

EU AI Act has the same meaning as AI Act.

EUDI Wallet means the European Digital Identity Wallet contemplated by eIDAS 2.0.

Federated Network means the network of Node Operators that together maintain the DAG.

FIPS 203 means Federal Information Processing Standard Publication 203 (ML-KEM), published by NIST in August 2024.

FIPS 204 means Federal Information Processing Standard Publication 204 (ML-DSA), published by NIST in August 2024.

FIPS 205 means Federal Information Processing Standard Publication 205 (SLH-DSA), published by NIST in August 2024.

Founder Seat has the meaning given in Part X Section 10.2.2.

Founding Member means a Founding Member of the AI Trust Council as identified in Part X Section 10.2.2 and in the Acknowledgments.

Founding Node Operator means a Node Operator party to a Founding Node Operator Agreement with The AI Lab.

Founding Period means the operational period commencing on the Genesis Date and concluding on a date determined by The AI Lab in consultation with the Council.

Genesis Date means the date in June 2026 on which the Trust Identity Protocol federated network commences live operation, being the date determined by the sole director of The AI Lab Intelligence Unobscured, Inc. on the satisfaction of the Operational Readiness Conditions and published by The AI Lab not less than three (3) business days in advance on theailab.org.

Free Tier has the meaning given in Part X Section 10.3.1.

Full Node has the meaning given in Part VII Section 7.2.2.

GDPR means Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016.

Global Seal of Trust has the meaning given in Part VI Section 6.8.

HKDF-SHA-512 means the HMAC-based Extract-and-Expand Key Derivation Function with SHA-512 as the underlying hash function, specified by RFC 5869.

Independent Member means a Member of the Council other than the Founder Seat holder.

Light Node has the meaning given in Part VII Section 7.2.1.

ML-DSA-65 means the Module-Lattice-Based Digital Signature Algorithm at parameter set 65, specified by FIPS 204.

ML-KEM-768 means the Module-Lattice-Based Key-Encapsulation Mechanism at parameter set 768, specified by FIPS 203.

Network Period means the operational period commencing on the conclusion of the Founding Period.

Operational Readiness Conditions means the conditions identified in Part XI Section 11.1.1 of this whitepaper on the satisfaction of which the Genesis Date may be set by the sole director.

Network Sub-Score has the meaning given in Part VI Section 6.2.4.

NIS2 Directive means Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022.

NIST means the United States National Institute of Standards and Technology.

Node Operator means an organization operating a node of the Federated Network under a Node Operator Agreement.

Online Safety Act means the Online Safety Act 2023 of the United Kingdom.

Origin Code means the four-letter classification (OH, AA, AG, MX) defined in Part V Section 5.6.

Person means a natural person or an organization, however constituted.

Pre-Genesis Period means the period commencing on June 1, 2026 and concluding on the Genesis Date, during which the federated network is being stood up by the Founding Node Operators in pre-Genesis operating mode but does not accept production protocol events.

Privacy Act 2020 means the Privacy Act 2020 of New Zealand (Act No. 31 of 2020).

Publisher Mode has the meaning given in Part V Section 5.3.

SLH-DSA means the Stateless Hash-Based Digital Signature Algorithm, specified by FIPS 205. SLH-DSA-SHA2-192s is the parameter set identified for the protocol’s optional long-term archival signature use case.

Specification means the canonical TIP Protocol Specification published by The AI Lab.

TIP or **Trust Identity Protocol** means the technical framework described in this whitepaper and in the Specification.

TIP-CONTENT means the content provenance layer of the Trust Identity Protocol defined in Part V.

TIP-CONTENT Record means the signed message structured as described in Part V Section 5.7.

TIPCL-1.0 means the TIP Protocol Code License, Version 1.0, as published by The AI Lab and incorporated by reference as Appendix F.

TIP-ID means the identity layer of the Trust Identity Protocol defined in Part IV.

TIP-TRUST means the reputation layer of the Trust Identity Protocol defined in Part VI.

Trademark Usage Policy means the policy described in Part X Section 10.5.

Trust Score has the meaning given in Part VI Section 6.1.

Trust Score Tier means one of the five normative reader-facing tiers (HIGHLY_TRUSTED, TRUSTED, REVIEW_ADVISED, LOW_TRUST, NOT_TRUSTED) defined in Part VI Section 6.3, with score-range boundaries, icons, and colors as set forth therein.

United States Copyright Office or **USCO** means the United States Copyright Office.

United States Patent and Trademark Office or **USPTO** means the United States Patent and Trademark Office.

Verification Provider or **VP** means an organization accredited by the Council, on the recommendation of The AI Lab, to issue CTIDs.

Verification Provider Registry means the public registry of accredited Verification Providers maintained by The AI Lab.

WebAuthn means the World Wide Web Consortium Web Authentication specification, Level 3 (or such successor as the Council recognizes).

Appendix B: Acronym Index

This Appendix indexes the acronyms used in this whitepaper. Each acronym is defined on first use in the relevant Part.

Acronym	Expansion	Reference
AA	AI-Assisted (Origin Code)	Part V Section 5.6.2
AES-256-GCM	Advanced Encryption Standard, 256-bit key, Galois/Counter Mode	Part III Section 3.5

Acronym	Expansion	Reference
AG	AI-Generated (Origin Code)	Part V Section 5.6.3
AICOP	Australian Information Commissioner Online Privacy (reference)	Part IX
AMO	Mozilla Add-ons distribution channel	Part VIII Section 8.2.1
ARIA	Accessible Rich Internet Applications	Part VI Section 6.7.3
BIPA	Biometric Information Privacy Act (Illinois)	Part IX Section 9.2.4
C2PA	Coalition for Content Provenance and Authenticity	Part I Section 1.3
CC BY 4.0	Creative Commons Attribution 4.0 International Public License	Glossary
CCSF 2.0	NIST Cybersecurity Framework Version 2.0	Appendix E
CETS	Council of Europe Treaty Series	Part IX Section 9.1.7
CFRG	Crypto Forum Research Group (IETF)	Part IX Section 9.7
CNA	Canonical Normalization Algorithm	Part V
CRA	Cyber Resilience Act	Part IX Section 9.1.6
CTID	Cryptographic Trust Identity	Part IV
CUBI	Capture or Use of Biometric Identifier Act (Texas)	Part IX Section 9.2.4
DAG	Directed Acyclic Graph	Part VII
DEPA	Data Empowerment and Protection Architecture	Part IX Section 9.5
DPDPA	Digital Personal Data Protection Act 2023 (India)	Part IX Section 9.5
DPIA	Data Protection Impact Assessment	Part IX Section 9.1.2
DSA	Digital Services Act	Part IX Section 9.1.3
ECDSA	Elliptic Curve Digital Signature Algorithm	Part III
EAR	United States Export Administration Regulations	Part IX Section 9.8
Ed25519	Edwards-curve Digital Signature Algorithm using Curve25519	Part IV
EDPB	European Data Protection Board	Part IX Section 9.1.2
EIN	Employer Identification Number	Part X Section 10.1
eIDAS	electronic IDentification, Authentication and trust Services	Part IX Section 9.1.4
ENISA	European Union Agency for Cybersecurity	Part III Section 3.1
ETSI	European Telecommunications Standards Institute	Part IX Section 9.7
EUDI	European Digital Identity	Part IV Section 4.5.3
FIDO2	Fast IDentity Online 2	Part IV Section 4.4.1
FIPS	Federal Information Processing Standards	Part III
FRIA	Fundamental Rights Impact Assessment	Part IX
FTC	Federal Trade Commission	Part IX Section 9.2.1
GDPR	General Data Protection Regulation	Part IX Section 9.1.2
GPAI	General-Purpose AI	Part IX Section 9.1.1
HKDF	HMAC-based Extract-and-Expand Key Derivation Function	Part III Section 3.5
HMAC	Hash-based Message Authentication Code	Part III
HTML	HyperText Markup Language	Part V
ICO	Information Commissioner's Office (UK)	Part IX

Acronym	Expansion	Reference
IETF	Internet Engineering Task Force	Part IX Section 9.7
IPP	Information Privacy Principle (NZ)	Part IX Section 9.4
IPTC	International Press Telecommunications Council	Part IX Section 9.7
ISO	International Organization for Standardization	Part IX Section 9.7
JTC	Joint Technical Committee (ISO/IEC)	Part IX Section 9.7
LCIA	London Court of International Arbitration	Part X
MeitY	Ministry of Electronics and Information Technology (India)	Part XI
ML-DSA	Module-Lattice-Based Digital Signature Algorithm	Part III Section 3.3
ML-KEM	Module-Lattice-Based Key-Encapsulation Mechanism	Part III Section 3.2
MX	Mixed (Origin Code)	Part V Section 5.6.4
NIS2	Network and Information Security Directive 2	Part IX Section 9.1.5
NIST	National Institute of Standards and Technology (US)	Part III
NIST AI RMF	NIST Artificial Intelligence Risk Management Framework	Part IX Section 9.2.2
NIST CSF	NIST Cybersecurity Framework	Appendix E
NTP	Network Time Protocol	Part VII Section 7.4.5
NZIAC	New Zealand International Arbitration Centre	Part X
OECD	Organisation for Economic Co-operation and Development	Part IX Section 9.6
OH	Original Human (Origin Code)	Part V Section 5.6.1
OPC	Office of the Privacy Commissioner (NZ)	Part IX Section 9.4
OSA	Online Safety Act 2023 (UK)	Part IX Section 9.3.1
PHP	PHP: Hypertext Preprocessor	Part VIII
PKCS	Public-Key Cryptography Standards	Part II Section 2.7
PQ	Post-Quantum	Part III
PRF	Pseudorandom Function	Part III Section 3.5
PWA	Progressive Web Application	Part VIII Section 8.5
QES	Qualified Electronic Signature	Part IX Section 9.1.4
RFC	Request for Comments (IETF)	Part V
RSA	Rivest-Shamir-Adleman	Part III
SBOM	Software Bill of Materials	Part VIII Section 8.7.3
SCC	Standard Contractual Clauses (EU)	Part IX Section 9.1.2
SDK	Software Development Kit	Part VIII Section 8.6
SHA2-256	Secure Hash Algorithm 2 with 256-bit output	Part III Section 3.6
SHA3-256	Secure Hash Algorithm 3 with 256-bit output	Part III Section 3.6
SHAKE-256	Secure Hash Algorithm KECCAK Extendable-Output Function	Part III
SLH-DSA	Stateless Hash-Based Digital Signature Algorithm	Part III Section 3.4
TIP	Trust Identity Protocol	Throughout
TIPCL	TIP Protocol Code License	Part X Section 10.3
TLS	Transport Layer Security	Part III Section 3.2

Acronym	Expansion	Reference
TUDPA	Texas Data Privacy and Security Act	Part IX
USCO	United States Copyright Office	Part X
USPTO	United States Patent and Trademark Office	Part X Section 10.5
UUID	Universally Unique Identifier	Part VIII
VLOP	Very Large Online Platform (DSA)	Part IX Section 9.1.3
VP	Verification Provider	Part IV Section 4.2
W3C	World Wide Web Consortium	Part IX Section 9.7
WCAG	Web Content Accessibility Guidelines	Part VI Section 6.7.3
WIPO	World Intellectual Property Organization	Part IX

Appendix C: CNA-2.2 Worked Example

This Appendix supplies a worked example of the Canonical Normalization Algorithm Version 2.2 applied to a representative content unit. The example is included for reference and is reproducible against the conformance test vectors published in the canonical TIP Protocol Specification.

C.1 Source content unit

The source content unit is a journalistic article published on a publisher conformant to the WordPress reference plugin, with the following salient characteristics:

Field	Value
Source URL	<code>https://example-publication.com/2026/06/02/article-slug?utm_source=newsletter</code>
Platform identifier	<code>wpref</code> (WordPress with reference plugin)
Publication time	<code>2026-06-02T14:30:00Z</code>
Author byline	<code>"Sample Journalist"</code>
Author Creator Mode CTID	<code>Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23</code>
Publisher Publisher Mode CTID	<code>R8MNOP-QRSTUV-WXYZ23-ABCDEF-GHIJKL</code>
Title (raw)	<code>"What the new law says about content provenance"</code>
Body (raw, abbreviated for the example)	Five paragraphs of text including one embedded image and one
external hyperlink	
Origin Code	<code>0H</code> (Original Human)
Version number	<code>1</code> (initial publication)

C.2 Step-by-step normalization

Step 1: Platform identification

The Algorithm matches the source URL pattern, the document's HTTP response headers, and the document's structural signatures against the platform registry. The matched platform code is `wpref`. The platform code is recorded as the first field of the canonical byte sequence.

Step 2: Content scope extraction

The Algorithm applies the wpref platform's content scope selectors:

Element	Source	Action
Article header (<article> element)	Document body	Retained
Navigation menu (<nav> element)	Document header	Removed
Sidebar (<aside> element)	Document body	Removed
Comments section	Document footer	Removed
Related articles	Document footer	Removed
Advertisement units	Various	Removed
Article body	<article> element	Retained

The content scope is the <article> element, comprising the article title, the article byline, the article publication date, the article body, and the embedded media within the article body.

Step 3: Structural canonicalization

The retained <article> element is canonicalized at the structural level:

- All HTML element names are normalized to lowercase.
- Attribute names are normalized to lowercase and sorted lexicographically within each element.
- Attribute values are canonicalized per the published rules (boolean attributes are normalized; URL attribute values are canonicalized in Step 5; class attribute values are sorted).
- HTML comments are removed.
- Inline scripts (<script> elements) and inline style elements are removed.
- Presentational HTML elements are removed; structural HTML elements are retained.
- Whitespace within text nodes is normalized to single spaces.

The output is a structurally canonicalized HTML fragment.

Step 4: Character normalization

The text content within the canonicalized HTML fragment is normalized at the character level:

- Unicode Normalization Form C is applied throughout.
- Curly punctuation marks are replaced with their straight ASCII equivalents.
- Zero-width and invisible control characters are removed.
- Non-Latin script characters are preserved in their native form.

Step 5: Reference normalization

References within the content scope are normalized:

Reference type	Source	Action
External hyperlink	https://example-target.com/some-page?utm_medium=link&utm_source=publisher&utm_medium,utm_source=referral	Tracking parameters (utm_medium, utm_source)

utm_source) removed; substantive parameters (id) preserved; canonical URL: <https://example-target.com/some-page?id=42> | | Embedded image | `` | Image identified by CNA-IMG variant; the URL is replaced by the three-hash content identifier of the image in the canonical byte sequence; the alt text is preserved as text | | Canonical URL of the article | https://example-publication.com/2026/06/02/article-slug?utm_source=newsletter | Tracking parameters removed; canonical URL: <https://example-publication.com/2026/06/02/article-slug> |

Step 6: Embedded media identification

The embedded image is processed by the CNA-IMG variant, producing a three-hash content identifier of the form:

```
{
  "sha2_256": "[SHA2-256 of the canonicalized image bytes]",
  "sha3_256": "[SHA3-256 of the canonicalized image bytes]",
  "blake3": "[BLAKE3 of the canonicalized image bytes]"
}
```

The three-hash identifier is substituted for the image URL in the canonical byte sequence.

Step 7: Metadata extraction

Document-level metadata is extracted:

Metadata field	Value
Title	"What the new law says about content provenance"
Publication date	2026-06-02T14:30:00Z
Language	en (per IETF BCP 47)
Author CTID	Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23
Publisher CTID	R8MNOP-QRSTUV-WXYZ23-ABCDEF-GHIJKL
Canonical URL	https://example-publication.com/2026/06/02/article-slug
Origin Code	0H
Version number	1

Step 8: Serialization

The output of Steps 1 through 7 is serialized into a canonical byte sequence using RFC 8785 JSON Canonicalization Scheme. The serialized output (abbreviated):

```
{
  "cna_variant": "CNA-2.2",
  "content_scope": "...[structurally and character canonicalized HTML]...",
  "embedded_media": [{"sha2_256": "...", "sha3_256": "...", "blake3": "..."}],
  "metadata": {
    "author_ctid": "Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23",
    "canonical_url": "https://example-publication.com/2026/06/02/article-slug",
    "language": "en",

```

```

    "origin_code": "0H",
    "platform_id": "wpref",
    "publication_date": "2026-06-02T14:30:00Z",
    "publisher_ctid": "R8MN0P-QRSTUV-WXYZ23-ABCDEF-GHIJKL",
    "title": "What the new law says about content provenance",
    "version": 1
  },
  "protocol_version": "TIP-1.0",
  "references": ["https://example-target.com/some-page?id=42"]
}

```

The serialization is byte-deterministic: object keys are sorted lexicographically, numbers are encoded in canonical form, and string values are encoded in canonical UTF-8.

Step 9: Three-hash addressing

The serialized canonical byte sequence is the input to the three-hash addressing system:

```

{
  "sha2_256": "8f3c2a1b...",
  "sha3_256": "7e4d3b2a...",
  "blake3": "9a6b5c4d..."
}

```

The three-hash content identifier is the principal artifact recorded in the TIP-CONTENT record.

C.3 Verification at the reader

A reader receiving the canonical URL of the article retrieves the article, applies the same nine-step CNA-2.2 procedure, computes the three-hash content identifier, retrieves the TIP-CONTENT record from the federation by content identifier, verifies the publisher's signature on the TIP-CONTENT record using the publisher's public key from the Verification Provider Registry, and presents the Trust Score and the Global Seal of Trust associated with the publisher's CTID through the reader-facing badge described in Part VI Section 6.7.

C.4 Verifiability across surfaces

A reader retrieving the same article through a different surface (a syndication republisher, a search engine cache, a social platform's preview) applies the same CNA-2.2 procedure. The deterministic Steps 1 through 9 produce the same content identifier, recovering the same TIP-CONTENT record from the federation. The verifiability is independent of the surface on which the content is encountered.

C.5 Failure mode, modified content

If the content is modified (a substantive edit) without the signer publishing a CONTENT_UPDATED event, the reader's CNA-2.2 procedure produces a content identifier that does not match the TIP-CONTENT record on the federation. The reader-facing badge displays a verification failure indicator. The reader is notified that the content is not bound to the publisher's signature.

C.6 Failure mode, revoked CTID

If the publisher's or author's CTID is revoked (a B3, B5, or B6 Blocking Item activation), the reader-facing badge displays the applicable Blocking Item, even where the content identifier matches. The reader is notified that the CTID under which the content was signed is no longer in the Verified class.

Appendix D: End-to-End Identity Issuance, Signing, and Verification Example

This Appendix supplies a worked example of the end-to-end identity issuance, content signing, and reader-side verification flow of the Trust Identity Protocol. The example illustrates the interaction of the layers described in Parts III, IV, V, and VI. Test vectors corresponding to this example are published in the canonical TIP Protocol Specification.

D.1 Scenario

A natural person, identified by the pseudonym "Sample Journalist," obtains a Creator Mode CTID from a Verification Provider in New Zealand, signs a journalistic article, publishes the signed article on a personal publishing site, and a reader in the European Union verifies the signed article and its associated Trust Score.

D.2 Step 1: Identity issuance

Sample Journalist navigates to the public application page of the New Zealand Verification Provider. The application page is operated by the Verification Provider under the Accreditation Schedule published in the Verification Provider Registry.

Sample Journalist undertakes the identity verification appropriate to the requested CTID grade. For an enhanced CTID, the verification comprises: (a) the submission of a government-issued identity document; (b) a liveness check performed by the Verification Provider's identity verification software through the device's camera; (c) the verification of an email address and a mobile telephone number; (d) the acceptance of the Verification Provider's data protection notice.

On successful verification, the Verification Provider instructs Sample Journalist's browser to invoke the device's WebAuthn resident key authenticator. The authenticator generates a keypair, retaining the private key within the authenticator's hardware security boundary, and returns the public key. The Verification Provider receives the public key and the attestation that the authenticator has verified Sample Journalist.

The Verification Provider computes the CTID:

```
CTID = BLAKE3("TIP-CTID-v1" || version || vp_id || pub_key) [first 30 bytes]
      = "Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23"
      (illustrative; the actual CTID depends on the actual public key)
```

The Verification Provider publishes the CTID issuance event to the federated DAG and signs the event with the Verification Provider's ML-DSA-65 private key. The event includes the CTID, the public key of Sample Journalist's authenticator, the issuance timestamp, the CTID grade

(enhanced), and the Verification Provider's identifier in the Verification Provider Registry. The Verification Provider also records the mapping between the CTID and Sample Journalist's real-world identity in the Verification Provider's secure data store under the New Zealand Privacy Act 2020.

Sample Journalist is presented with the CTID and is instructed to record it.

D.3 Step 2: Initial Trust Score computation

The federated DAG records the CTID issuance event. Each Full Node receiving the event computes the initial Trust Score:

Sub-score	Value	Notes
Cryptographic	1000	ML-DSA-65 native authenticator
Behavioral	500	New CTID, no history
Adjudicated	1000	No adverse adjudications
Network	850	High-assurance Verification Provider in good standing; enhanced
CTID grade		
Aggregate (weighted)	825	$(1000 \times 0.2) + (500 \times 0.3) + (1000 \times 0.3) + (850 \times 0.2) = 825$

The aggregate Trust Score of 825 places the CTID in the Verified class (700 to 1000).

D.4 Step 3: Content signing

Sample Journalist authors a journalistic article on a personal publishing site. The article is composed using the publishing site's editor. On the article being ready for publication, Sample Journalist invokes the Trust Identity Protocol browser extension on the editor page.

The browser extension applies the CNA-2.2 normalization to the article (the nine-step procedure described in Part V Section 5.2 and worked through in Appendix C). The output is the three-hash content identifier:

```
Content identifier: {
  "sha2_256": "8f3c2a1b...",
  "sha3_256": "7e4d3b2a...",
  "blake3": "9a6b5c4d..."
}
```

The browser extension presents the content identifier to Sample Journalist and asks for confirmation. Sample Journalist confirms.

The browser extension invokes Sample Journalist's WebAuthn authenticator. The authenticator prompts Sample Journalist for user verification (a biometric gesture in this example, since Sample Journalist enabled biometric user verification at issuance). The authenticator verifies the biometric on the device and produces an ML-DSA-65 signature of the TIP-CONTENT record header plus body.

The signed TIP-CONTENT record contains the header (protocol version, CNA variant, signature scheme, CTID, Origin Code 0H), the body (three-hash content identifier, canonical URL, platform identifier, metadata), and the signature.

The browser extension publishes the signed TIP-CONTENT record to a Full Node selected from the Verification Provider Registry. The publishing Full Node accepts the record, verifies the signature, propagates the record to peer Full Nodes through the synchronization protocol described in Part VII Section 7.4, and the publishing operation completes.

D.5 Step 4: Reader-side verification

A reader in the European Union encounters the article through a search engine result. The reader's browser, with the Trust Identity Protocol browser extension installed, retrieves the article.

The browser extension applies CNA-2.2 to the article content and computes the three-hash content identifier. The extension queries a Full Node (for the EU reader, the publication of an EU Member State Founding Node Operator within the twelve-month horizon described in Part XI Section 11.2 will reduce latency; for the present example, a United Kingdom Node Operator serves the EU reader) for the TIP-CONTENT record by content identifier.

The Full Node returns the TIP-CONTENT record. The browser extension verifies:

1. The signature on the TIP-CONTENT record using Sample Journalist's public key (retrieved from the CTID issuance event on the federated DAG).
2. The CTID issuance event signature using the Verification Provider's public key (retrieved from the Verification Provider Registry).
3. The Verification Provider's accreditation status (in good standing).
4. Sample Journalist's current Trust Score and active Blocking Items.

All verifications succeed. The current Trust Score remains in the Verified class. No Blocking Items are active. The Origin Code is 0H.

The browser extension displays the reader-facing badge:

```
[Global Seal of Trust]
CTID: Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23
Trust Score: Verified (825)
Origin: Original Human
Version: 1 (June 2, 2026)
Issuing Verification Provider: NZ-VP-001 (New Zealand)
```

D.6 Step 5: Subsequent modification

Three days after publication, Sample Journalist corrects a typographical error in the article. Sample Journalist invokes the browser extension on the corrected article. The extension computes a new content identifier, presents a CONTENT_UPDATED dialog identifying the modification category (Category (a), non-substantive correction per Part V Section 5.5.1), and prompts Sample Journalist for confirmation. Sample Journalist confirms.

The browser extension produces a signed CONTENT_UPDATED event, identifying the prior version's content identifier, the new version's content identifier, the version number (2), the

modification category, a human-readable note, the timestamp, and the signature. The event is published to the federated DAG.

A reader subsequently viewing the corrected article observes the badge with version number 2 and a link to the version history. The version history identifies the original publication and the CONTENT_UPDATED event with the modification category and the note. The Trust Score is unchanged.

D.7 Failure mode, authenticator loss

Sample Journalist loses the device on which the WebAuthn authenticator is configured. Sample Journalist initiates the Verification Provider's recovery procedure. The Verification Provider verifies Sample Journalist's identity by the recovery criteria published in the Accreditation Schedule. On successful recovery verification, the Verification Provider:

1. Revokes the CTID Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23 and publishes the revocation event to the federated DAG.
2. Issues a successor CTID S9MNOP-QRSTUV-WXYZ23-ABCDEF-GHIJKL to Sample Journalist using a new keypair generated on Sample Journalist's new device.
3. Publishes a succession notice signed by both the revoked CTID's stored backup key (held by Sample Journalist in the secure recovery medium published by the Verification Provider) and the successor CTID, identifying the relationship between the two CTIDs.

The reader subsequently encountering the article continues to verify it under the revoked CTID Q7ABCD-EFGHIJ-KLMNOP-QRSTUV-WXYZ23, which remains valid for verification of signatures produced prior to the revocation. The reader's badge presents the revocation event timestamp and notes that the signing CTID has been revoked but the signature was produced prior to the revocation. Sample Journalist's continued publishing under the successor CTID is presented under the successor CTID, with reputation continuity supplied by the succession notice.

D.8 Failure mode, adverse adjudication

A third party files a dispute alleging that an article signed by Sample Journalist's CTID contains a material misrepresentation. The complaint is filed under the bonded jury adjudication procedure described in Part VI Section 6.5.2. A bonded jury panel is convened, reviews the complaint and Sample Journalist's response, and (in this example) upholds the complaint with a Trust Score reduction. The Adjudicated sub-score of Sample Journalist's CTID is reduced from 1000 to a lower value reflecting the severity category of the adjudication. The aggregate Trust Score is recomputed. The Blocking Item B5 (Final Adverse Adjudication) is activated.

A reader subsequently viewing articles signed by Sample Journalist's CTID observes the reduced Trust Score and the B5 marker in the badge. The reader retains the ability to inspect the adjudication's published determination through the link supplied in the badge.

The successful adjudication does not invalidate the historical signatures of Sample Journalist's CTID; it qualifies the present and prospective reliance on the CTID through the reduced Trust Score and the Blocking Item marker.

Appendix E: Compliance Crosswalk

This Appendix maps the technical and institutional controls of the Trust Identity Protocol to the categories and subcategories of three widely adopted reference frameworks: ISO/IEC 27001:2022, the NIST Cybersecurity Framework Version 2.0 (NIST CSF 2.0), and the NIST AI Risk Management Framework Version 1.0 (NIST AI RMF). The Appendix also maps the protocol controls to the OECD AI Principles. The crosswalk is provided as a reference for licensees, Verification Providers, and Node Operators preparing their own conformance documentation. Statements in this Appendix are statements of design intent and of the technical controls implemented in the canonical specification; they are not certifications of compliance. The provisions of Appendix J apply.

E.1 ISO/IEC 27001:2022 Annex A controls

The following table identifies the principal Annex A controls of ISO/IEC 27001:2022 supported by the architecture of the Trust Identity Protocol.

Annex A control	TIP architectural support
A.5.1 Policies for information security	Charter of the AI Trust Council, TIPCL-1.0, Node Operator
Agreement, Verification Provider accreditation agreement	
A.5.7 Threat intelligence	DAG-based observability of suspension and revocation events
across the federated network	
A.5.15 Access control	CTID-based authentication using NIST FIPS 204 (ML-DSA-65)
digital signatures	
A.5.17 Authentication information	WebAuthn resident key authenticator with biometric binding;
private key never leaves the user's device	
A.5.23 Information security for use of cloud services	Verification Provider jurisdiction declaration; node operator
hosting requirements	
A.5.30 ICT readiness for business continuity	Federated DAG with multi-region Node Operator deployment
A.5.34 Privacy and protection of PII	Architectural commitment to pseudonymous DAG records; biometric
data	
device-local	
A.5.35 Independent review of information security	Annual Verification Provider audit by an auditor acceptable to
The	

Annex A control	TIP architectural support
AI Lab	
A.6.3 Information security awareness, education and training Provider accreditation agreement and Node Operator Agreement	Article 4 EU AI Act AI literacy commitments in Verification
A.8.5 Secure authentication	WebAuthn FIDO2 authenticator with platform attestation
A.8.9 Configuration management	Reference implementation configuration baselines published by The
AI Lab	
A.8.10 Information deletion	Pseudonymization-based right to erasure pattern under GDPR
Article 17	
A.8.11 Data masking	CTID is a pseudonymous identifier derived from a public key; no
real-world identifier on the DAG	
A.8.12 Data leakage prevention	Architectural decision that biometric templates do not leave
the user's device	
A.8.14 Redundancy of information processing facilities	Federated DAG with multi-region Node Operator deployment
A.8.15 Logging	Append-only DAG record of all signed events
A.8.16 Monitoring activities	Trust Score and Blocking Item monitoring; warrant canary
attestation	
A.8.20 Networks security	TLS 1.3 with hybrid post-quantum key agreement for transport
between protocol participants	
A.8.24 Use of cryptography	NIST FIPS 203, FIPS 204, FIPS 205 post-quantum primitives; AES
key protection chain	
A.8.26 Application security requirements	Reference implementation security requirements published by The
AI Lab	
A.8.28 Secure coding	Reference implementations follow secure coding standards
published	

Annex A control	TIP architectural support
by The AI Lab	
A.8.31 Separation of development, test and production environments	Reference deployment configuration
A.8.32 Change management	AI Trust Council ratification of protocol amendments
A.8.34 Protection of information systems during audit testing	Read-only audit access to DAG; segregation of audit identifiers

E.2 NIST Cybersecurity Framework 2.0 functions

The following table maps the NIST CSF 2.0 functions and categories to TIP controls.

NIST CSF 2.0 function and category	TIP architectural support
GOVERN (GV), Organizational Context	The AI Lab corporate governance; AI Trust Council Charter
GOVERN (GV), Risk Management Strategy	Material risks register in Part XI; Verification Provider and
Node	
Operator risk assessments	
GOVERN (GV), Roles, Responsibilities, and Authorities	Defined roles for Verification Provider, Node Operator, AI
Trust Council	
GOVERN (GV), Policy	TIPCL-1.0; Verification Provider accreditation agreement; Node
Operator Agreement	
IDENTIFY (ID), Asset Management	DAG record of registered Verification Providers and Node
Operators	
IDENTIFY (ID), Risk Assessment	Continuous monitoring of Trust Score and Blocking Items
PROTECT (PR), Identity Management, Authentication, and Access	
Control	CTID; ML-DSA-65 signatures; WebAuthn resident key
PROTECT (PR), Awareness and Training	AI literacy clauses in counterparty agreements
PROTECT (PR), Data Security	Post-quantum cryptographic primitives; pseudonymization
PROTECT (PR), Information Protection Processes and Procedures	Reference implementation security requirements
PROTECT (PR), Maintenance	Vulnerability handling process under Cyber Resilience Act

NIST CSF 2.0 function and category	TIP architectural support
compliance	
PROTECT (PR), Protective Technology	Append-only DAG; warrant canary; bonded juror requirement
DETECT (DE), Anomalies and Events	Trust Score updates; Blocking Items
DETECT (DE), Security Continuous Monitoring	Annual Verification Provider audit
DETECT (DE), Detection Processes	Dispute reporting and adjudication pathway
RESPOND (RS), Response Planning	NIS2-aligned incident response timelines in counterparty
agreements	
RESPOND (RS), Communications	Notification obligations under counterparty agreements
RESPOND (RS), Analysis	Bonded juror adjudication of disputes
RESPOND (RS), Mitigation	Suspension and revocation of CTIDs; Verification Provider
accreditation suspension	
RECOVER (RC), Recovery Planning	Multi-region federated DAG
RECOVER (RC), Improvements	AI Trust Council ratification of post-incident protocol
amendments	
RECOVER (RC), Communications	Transparency report published annually by the AI Trust Council

E.3 NIST AI Risk Management Framework

The NIST AI RMF organizes guidance into four functions: GOVERN, MAP, MEASURE, MAN-AGE. The following table maps the principal subcategories to TIP controls.

NIST AI RMF subcategory	TIP architectural support
GOVERN 1.1 Legal and regulatory requirements understood and managed	Part IX and Appendix E of this whitepaper
GOVERN 1.2 The characteristics of trustworthy AI are integrated into	
organizational policies	TIPCL-1.0 and AI Trust Council Charter
GOVERN 1.4 The risk management process and its outcomes are	
established through transparent policies	TIPCL-1.0 published; AI Trust Council transparency report
GOVERN 1.6 Mechanisms are in place to inventory AI systems	DAG record of Verification Providers and Node Operators
GOVERN 4.1 Organizational policies and practices are in place	

NIST AI RMF subcategory	TIP architectural support
to foster a critical thinking culture	AI literacy clauses
MAP 1.1 Intended purposes, prospective settings, etc. understood and documented	Part I and Part XI of this whitepaper Reference implementation test vectors; CNA-2.2 deterministic
MAP 2.3 Scientific integrity and TEVV considerations specification	
MAP 3.4 Processes for operator and practitioner proficiency	AI literacy clauses; Verification Provider audit
MAP 4.1 Approaches for mapping AI technology and legal risks of its components	Patent license under TIPCL-1.0 Section 8; defensive termination
provision	
MEASURE 1.1 Approaches and metrics for measurement of AI risks	Trust Score sub-scores; Blocking Items
MEASURE 2.7 AI system security and resilience analysis	Post-quantum primitives; cryptographic adversarial robustness
MEASURE 2.8 Risks associated with transparency and accountability	Audit trail through append-only DAG
MEASURE 3.2 Risk tracking approaches	Trust Score evolution recorded on DAG
MANAGE 1.2 Treatment of documented AI risks	Suspension, revocation, dispute pathways
MANAGE 2.2 Mechanisms for sustaining the value of deployed AI systems	Apache 2.0 conversion provision; AI Trust Council Charter
MANAGE 2.4 Mechanisms for superseding, disengaging, or deactivating AI systems	Verification Provider accreditation revocation
MANAGE 4.1 Post-deployment AI system monitoring plans transparency report	Annual Verification Provider audit; AI Trust Council

E.4 OECD AI Principles

OECD Principle	TIP architectural support
Inclusive growth, sustainable development and well-being	Free Tier eligibility for individuals at any income level, small businesses under US\$100,000 revenue,

nonprofits, educational institutions, government, journalism (Part X.10.3.1) | | Respect for the rule of law, human rights and democratic values, including fairness and privacy | Architectural commitment to pseudonymity; biometric template never leaves device; Article 5(1)(c) social scoring disclaimer; Council of Europe Framework Convention alignment | | Transparency and explainability | Normative definition of Trust Score sub-scores and Origin Codes; published CNA-2.2 specification; AI Trust Council transparency report | | Robustness, security and safety | NIST FIPS 203/204/205 post-quantum primitives; warrant canary; bonded juror requirement | | Accountability | TIPCL-1.0 attribution requirement; Verification Provider audit; AI Trust Council ratification |

Contact The AI Lab regarding the Compliance Crosswalk

General Counsel: legal@theailab.org Conformance Inquiries: compliance@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Appendix F: TIPCL-1.0 License Summary

This Appendix provides a structured summary of the TIP Protocol Code License, Version 1.0 (“TIPCL-1.0”). The operative text of TIPCL-1.0 is published at theailab.org/tip-license. In the event of any inconsistency between this summary and the operative text, the operative text controls.

F.1 Scope of license

TIPCL-1.0 governs the use, reproduction, modification, and distribution of (a) software implementations of the technical framework described in this whitepaper and in the canonical TIP Protocol Specification, (b) integration interfaces and software development kits published by The AI Lab, and (c) reference data, schemas, and test vectors published by The AI Lab in connection with the foregoing.

TIPCL-1.0 does not govern the canonical TIP Protocol Specification itself, which is published under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0).

F.2 Eligibility tiers

Tier category	Eligibility	Annual fee
Free Tier, individual	Individual natural person, any income level	No fee
Free Tier, small business	Business with annual gross revenue below US\$100,000	No fee

Tier category	Eligibility	Annual fee
Free Tier, nonprofit	Nonprofit, NGO, or registered charity, any size	No fee
Free Tier, educational	Educational institution, any size	No fee
Free Tier, government	Government entity, any level, any size	No fee
Free Tier, journalism	Journalism organization, editorial use only, any size	No fee
Free Tier, R&D	Research, development, or testing within published per-organization user and duration ceilings	No fee
Commercial, Micro	Annual gross revenue US\$100,000 to US\$250,000	US\$500
Commercial, Seed	Annual gross revenue US\$250,000 to US\$500,000	US\$1,100
Commercial, Starter	Annual gross revenue US\$500,000 to US\$5,000,000	US\$2,750
Commercial, Growth	Annual gross revenue US\$5,000,000 to US\$25,000,000	US\$8,250
Commercial, Business	Annual gross revenue US\$25,000,000 to US\$100,000,000	US\$27,500
Commercial, Enterprise	Annual gross revenue US\$100,000,000 to US\$500,000,000	US\$71,500
Commercial, Corporate	Annual gross revenue US\$500,000,000 to US\$2,000,000,000	US\$165,000
Commercial, Strategic	Annual gross revenue US\$2,000,000,000 to US\$10,000,000,000	US\$385,000
Commercial, Global	Annual gross revenue US\$10,000,000,000 and above	US\$550,000

F.3 Required conditions of use

A licensee is required to:

1. Include a copy of TIPCL-1.0, or a hyperlink to theailab.org/tip-license, in any redistribution.
2. Preserve the copyright, trademark, and attribution notices of The AI Lab in source code, documentation, and user-facing surfaces.
3. Provide attribution to The AI Lab in any publication that describes, references, or implements the Trust Identity Protocol, in the form set out in Section F.7.
4. Generate a NOTICE file in conformance with Section F.4 and include the NOTICE file in any redistribution.
5. Refrain from describing an implementation as “TIP-compliant,” “TIP-certified,” or by any equivalent phrase that implies endorsement by The AI Lab, unless the licensee has obtained a written compliance attestation from The AI Lab.
6. Refrain from using the trademarks of The AI Lab except in accordance with the Trademark Usage Policy published at theailab.org/trademark.
7. Where the licensee is in the Commercial Tier, pay the applicable annual fee in accordance with the operative text of TIPCL-1.0.
8. Where the licensee processes personal data of natural persons in connection with the Trust Identity Protocol, comply with applicable data protection law, including the regulations and statutes listed in Part IX of this whitepaper.

F.4 NOTICE file requirement

A licensee is required to maintain in the root directory of any source code redistribution, and to include in any user-facing documentation, a NOTICE file containing:

This product incorporates the Trust Identity Protocol (TIP), a technical framework published by The AI Lab Intelligence Unobscured, Inc., Wilmington, Delaware, United States.

The TIP Protocol Specification is licensed under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0). This implementation is licensed under the TIP Protocol Code License, Version 1.0 (TIPCL-1.0), available at <https://theailab.org/tip-license>.

The marks "The AI Lab," "Trust Identity Protocol," "TIP," "Global Seal of Trust," and "AI Trust Council" are trademarks of The AI Lab Intelligence Unobscured, Inc.

For licensing inquiries, contact licensing@theailab.org.

F.5 Patent license

TIPCL-1.0 Section 8 grants to every licensee in good standing a royalty-free, non-exclusive, non-transferable license under the issued and pending patent claims of The AI Lab that are necessarily infringed by the practice of the Trust Identity Protocol in accordance with the canonical specification. The patent license is conditioned on (a) compliance with the operative terms of TIPCL-1.0, (b) attribution, and (c) the defensive termination provision.

The defensive termination provision states that the Section 8 patent license terminates automatically with respect to a licensee that initiates patent infringement litigation against The AI Lab or against any other licensee in good standing alleging infringement by the practice of the Trust Identity Protocol.

F.6 Trademark license

TIPCL-1.0 Section 9 confirms that the license granted under TIPCL-1.0 does not grant any right to use the trademarks, service marks, trade names, or logos of The AI Lab, except as permitted by nominative fair use under applicable trademark law and except as expressly permitted by the Trademark Usage Policy.

F.7 Required attribution form

In any publication, presentation, advertisement, marketing material, technical documentation, public website, or other surface that describes, references, implements, or relies on the Trust Identity Protocol, the licensee shall include attribution in substantially the following form:

“Implements the Trust Identity Protocol™ (TIP), a technical framework published by The AI Lab Intelligence Unobscured, Inc. The TIP Protocol Specification is available at theailab.org under CC BY 4.0. Trust Identity Protocol™, TIP™, The AI Lab™,

Global Seal of Trust™, and AI Trust Council™ are trademarks of The AI Lab Intelligence Unobscured, Inc.”

F.8 Apache 2.0 conversion provision

TIPCL-1.0 Section 12 provides that on January 1, 2031, the canonical TIP Protocol Specification and the reference implementations published by The AI Lab become available under the Apache License, Version 2.0. The conversion does not affect (a) trademark rights of The AI Lab, (b) Commercial License fees accrued and unpaid as of the conversion date, or (c) the AI Trust Council’s role in ratifying protocol changes.

F.9 Termination

A license under TIPCL-1.0 terminates automatically on:

1. A material breach of the operative terms of TIPCL-1.0 that is not cured within thirty (30) days of written notice from The AI Lab.
2. The licensee’s initiation of patent infringement litigation as described in Section F.5.
3. The licensee’s commencement of a voluntary case under any chapter of the United States Bankruptcy Code, or the commencement against the licensee of an involuntary case that is not dismissed within sixty (60) days, in either case as such provisions are applied in the jurisdiction of the licensee’s primary place of business.

On termination, the licensee shall cease all use of the Trust Identity Protocol implementations covered by TIPCL-1.0, except that the licensee may retain copies for archival and legal compliance purposes.

F.10 Governing law and dispute resolution

The operative text of TIPCL-1.0 specifies the law of the State of Delaware as the governing law of the license, with disputes resolved by binding arbitration administered by the American Arbitration Association under its Commercial Arbitration Rules, seated in Wilmington, Delaware. Licensees outside the United States may elect, in writing at the time of license formation, an alternative seat and administering institution from a schedule published by The AI Lab.

F.11 Disclaimer

TIPCL-1.0 Section 10 contains a disclaimer of warranties and a limitation of liability in the form customary for open source licenses, qualified by applicable consumer protection law in the licensee’s jurisdiction. The disclaimer is set out in capital letters in the operative text of TIPCL-1.0 and is not reproduced here.

Contact The AI Lab regarding TIPCL-1.0

Licensing inquiries: licensing@theailab.org , theailab.org/tip-license General Counsel: legal@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Appendix G: References

This Appendix identifies the principal published instruments, standards, statutes, regulations, and academic and technical works on which this whitepaper relies. References are organized by category. The reference to a specific provision of any instrument identified herein incorporates by reference the most current published version of the instrument as of the publication date of this whitepaper.

G.1 Cryptographic standards

National Institute of Standards and Technology, Federal Information Processing Standard Publication 203, *Module-Lattice-Based Key-Encapsulation Mechanism Standard* (ML-KEM), Gaithersburg, MD, August 2024.

National Institute of Standards and Technology, Federal Information Processing Standard Publication 204, *Module-Lattice-Based Digital Signature Standard* (ML-DSA), Gaithersburg, MD, August 2024.

National Institute of Standards and Technology, Federal Information Processing Standard Publication 205, *Stateless Hash-Based Digital Signature Standard* (SLH-DSA), Gaithersburg, MD, August 2024.

National Institute of Standards and Technology, Special Publication 800-108 Revision 1, *Recommendation for Key Derivation Using Pseudorandom Functions*, Gaithersburg, MD, 2022.

Internet Engineering Task Force, Request for Comments 5869, *HMAC-based Extract-and-Expand Key Derivation Function* (HKDF), May 2010.

Internet Engineering Task Force, Request for Comments 8785, *JSON Canonicalization Scheme* (JCS), June 2020.

Internet Engineering Task Force, Request for Comments 9457, *Problem Details for HTTP APIs*, July 2023.

World Wide Web Consortium, *Web Authentication: An API for accessing Public Key Credentials Level 3*, W3C Recommendation.

National Institute of Standards and Technology, *Cybersecurity Framework Version 2.0*, NIST CSWP 29, February 2024.

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework*, NIST AI 100-1, January 2023.

International Organization for Standardization and International Electrotechnical Commission, *ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection, Information security management systems, Requirements*, October 2022.

G.2 European Union legislation

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (the Artificial Intelligence Act).

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (the General Data Protection Regulation).

Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services (the Digital Services Act).

Regulation (EU) 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market, as amended by Regulation (EU) 2024/1183 of 11 April 2024 (eIDAS 2.0).

Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (the NIS2 Directive).

Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements (the Cyber Resilience Act).

Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data (the Data Act).

G.3 United States legislation, regulation, and executive instruments

15 U.S.C. § 45 (Federal Trade Commission Act, Section 5).

15 C.F.R. § 740.17 (Export Administration Regulations, mass market exception).

37 C.F.R. § 202.1 (United States Copyright Office Regulations).

17 U.S.C. § 103 (Subject matter of copyright: Compilations and derivative works).

17 U.S.C. § 1052(e) (Trademark Act of 1946, Section 2(e) refusal grounds).

Executive Order 14110 of October 30, 2023 on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.

Office of Management and Budget Memorandum M-24-10 of March 28, 2024 on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence.

G.4 United States state legislation

California Senate Bill 942 of 2024 (AI Transparency Act).

California Assembly Bill 853 of 2024 (Content Provenance).

Colorado Senate Bill 24-205 (Colorado AI Act).

Illinois Biometric Information Privacy Act, 740 ILCS 14.

Texas Capture or Use of Biometric Identifier Act, Tex. Bus. & Com. Code § 503.001.

Washington Revised Code § 19.375 (biometric identifiers).

G.5 United Kingdom legislation

Online Safety Act 2023, 2023 c. 50.

Data Protection Act 2018, 2018 c. 12.

G.6 New Zealand legislation

Privacy Act 2020, Act No. 31 of 2020.

G.7 India legislation

Digital Personal Data Protection Act 2023, Act No. 22 of 2023.

G.8 Council of Europe instruments

Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, CETS No. 225, opened for signature 5 September 2024.

G.9 OECD instruments

OECD Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449, adopted 22 May 2019, updated 3 May 2024.

G.10 Standards organizations and industry bodies

Coalition for Content Provenance and Authenticity, *Content Credentials Technical Specification*, current published version.

FIDO Alliance, *Client to Authenticator Protocol (CTAP) Specification*, current published version.

Coalition for Content Provenance and Authenticity (C2PA), *Manifest Specification*, current published version.

International Press Telecommunications Council, *IPTC Photo Metadata Standard*, current published version.

World Wide Web Consortium, *Verifiable Credentials Data Model 2.0*, W3C Recommendation.

World Wide Web Consortium, *Decentralized Identifiers (DIDs) v1.0*, W3C Recommendation.

G.11 The AI Lab published works

The AI Lab Intelligence Unobscured, Inc., *TIP Protocol Specification, Version 5.0*, Wilmington, Delaware, June 2026 (United States Copyright Office Application No. 1-15175755931, pending; derivative of Application No. 1-15116205291, pending).

The AI Lab Intelligence Unobscured, Inc., *TIP Protocol Code License, Version 1.0 (TIPCL-1.0)*, current published version at theailab.org/tip-license.

The AI Lab Intelligence Unobscured, Inc., *AI Trust Council Charter, Version 1.0*, current published version at theailab.org/ai-trust-council.

The AI Lab Intelligence Unobscured, Inc., *Founding Node Operator Agreement template, Version 2.0*, current published version.

The AI Lab Intelligence Unobscured, Inc., *Technical Requirements Specification for Founding Node Operators, Version 2.0*, current published version.

The AI Lab Intelligence Unobscured, Inc., *Service Level Agreement template for Founding Node Operators, Version 2.0*, current published version.

The AI Lab Intelligence Unobscured, Inc., *Trademark Usage Policy, Version 1.0*, current published version at theailab.org/trademark.

Appendix H: Contact Directory

This Appendix consolidates the published contact addresses for The AI Lab Intelligence Unobscured, Inc. in connection with the Trust Identity Protocol. Inquiries directed to the addresses below are received and routed by The AI Lab during ordinary business hours in the Eastern Time Zone of the United States.

H.1 Licensing and commercial implementation

Purpose	Address
Commercial License inquiries (any tier)	licensing@theailab.org
Free Tier eligibility questions	licensing@theailab.org
Trademark Usage Policy	legal@theailab.org
Web	theailab.org/tip-license

H.2 Network participation

Purpose	Address
Founding Node Operator applications	nodes@theailab.org
Web	theailab.org/founding-node
Verification Provider accreditation inquiries	licensing@theailab.org
Web	theailab.org/tip-verification-provider

H.3 Governance

Purpose	Address
AI Trust Council inquiries	council@theailab.org
Membership applications (Network Period)	council@theailab.org
Web	theailab.org/ai-trust-council
Dispute filings	council@theailab.org

H.4 Technical engagement

Purpose	Address
Technical inquiries	tip@theailab.org
Integration support	integrations@theailab.org
Conformance inquiries	compliance@theailab.org

Purpose	Address
Security reports	security@theailab.org
Standards engagement	standards@theailab.org
Operations	operations@theailab.org
Web	theailab.org/tip-protocol

H.5 Legal

Purpose	Address
General Counsel	legal@theailab.org
Trademark and copyright inquiries	legal@theailab.org
Regulatory and supervisory authority engagement	legal@theailab.org

H.6 Press and public affairs

Purpose	Address
Press inquiries	press@theailab.org
Public affairs and policy inquiries	press@theailab.org

H.7 Postal address

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware United States.

Specific street addresses, courier delivery instructions, and registered agent service-of-process addresses are available on written request to legal@theailab.org.

Appendix I: Suggested Citation

The suggested citation block below is reproduced from the inside front cover for the convenience of the reader.

I.1 Plain text

Mendhe, D. (2026). *TIP Protocol Whitepaper, Version 1.1: Trust Identity Protocol for the Verifiable Internet*. The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware. United States Copyright Office Application No. 1-15175755931. Available at <https://theailab.org/whitepaper>.

I.2 IEEE

D. Mendhe, "TIP Protocol Whitepaper, Version 1.1," The AI Lab Intelligence Unobscured, Inc., Wilmington, DE, 2026.

I.3 ACM

Dinesh Mendhe. 2026. *TIP Protocol Whitepaper, Version 1.1*. The AI Lab Intelligence Unobscured, Inc., Wilmington, DE.

I.4 APA

Mendhe, D. (2026). *TIP Protocol Whitepaper, Version 1.1: Trust Identity Protocol for the Verifiable Internet*. The AI Lab Intelligence Unobscured, Inc.

I.5 Chicago

Mendhe, Dinesh. 2026. *TIP Protocol Whitepaper, Version 1.1*. Wilmington, DE: The AI Lab Intelligence Unobscured, Inc.

I.6 BibTeX

```
@techreport{mendhe2026tip,  
  author      = {Mendhe, Dinesh},  
  title       = {TIP Protocol Whitepaper, Version 1.1:  
                Trust Identity Protocol for the Verifiable Internet},  
  institution = {The AI Lab Intelligence Unobscured, Inc.},  
  address     = {Wilmington, Delaware},  
  year        = {2026},  
  type        = {Technical Report},  
  note        = {USCO Application No.\ 1-15175755931},  
  url         = {https://theailab.org/whitepaper}  
}
```

I.7 Citation of specific Parts and Sections

When citing a specific Part or Section of this whitepaper, the citation should identify the whitepaper version (v1.1), the Part number (Roman numeral), the Section number (Arabic numeral with decimal subdivisions where applicable), and the publication date. Example: Mendhe (2026), *TIP Protocol Whitepaper v1.1*, Part X Section 10.8 (August 23, 2026).

I.8 Citation of the canonical Specification

The canonical TIP Protocol Specification is cited separately from this whitepaper. The suggested citation for the canonical Specification is:

The AI Lab Intelligence Unobscured, Inc. (2026). *TIP Protocol Specification, Version 5.0*. Wilmington, Delaware. United States Copyright Office Application No. 1-15175755931, derivative of Application No. 1-15116205291, both pending. Available at https://github.com/theailaborg/tip-protocol/blob/main/spec/TIP_Protocol_Specification_v5_0.md, licensed under CC BY 4.0.

Appendix J: Notice and Disclaimer

This Appendix sets out the legal notices, disclaimers, and forward-looking statements safe harbor applicable to this whitepaper. The provisions of this Appendix apply to the entirety of the whitepaper, including the Foreword, the Executive Summary, the substantive Parts, the other Appendices, and any accompanying media.

J.1 Copyright notice

© 2026 The AI Lab Intelligence Unobscured, Inc. All rights reserved. This whitepaper is licensed under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0), available at creativecommons.org/licenses/by/4.0. The CC BY 4.0 license applies to the text, structure, and figures of this whitepaper, and does not extend to the trademarks of The AI Lab identified in Section J.4 or to the implementations of the Trust Identity Protocol, which are governed by TIPCL-1.0 as described in Part X and Appendix F.

This whitepaper is the subject of a copyright application filed with the United States Copyright Office referencing prior applications in the same series (United States Copyright Office Case Nos. 1-15175755931 and 1-15116205291, both pending). The filing dates of those applications establish priority irrespective of registration outcome.

J.2 No warranty

This whitepaper is provided on an “as is” basis without warranty of any kind, express or implied, including without limitation any warranty of merchantability, fitness for a particular purpose, non-infringement, accuracy, completeness, currency, or quiet enjoyment. The AI Lab does not represent or warrant that the information in this whitepaper is suitable for any purpose, and the reader is responsible for verifying the suitability of the information for the reader’s purposes before acting on it.

J.3 No legal advice and no fiduciary relationship

This whitepaper is a technical and policy reference document. It is not legal advice. It is not investment advice. It is not tax advice. It does not establish an attorney-client relationship, a fiduciary relationship, or any other professional relationship between The AI Lab and the reader. Readers seeking advice on the application of any statute, regulation, or other instrument referenced in this whitepaper should consult qualified counsel in the relevant jurisdiction.

J.4 Trademark notices

The marks “The AI Lab,” “Trust Identity Protocol,” “TIP,” “Global Seal of Trust,” and “AI Trust Council” are trademarks of The AI Lab Intelligence Unobscured, Inc., with applications pending before the United States Patent and Trademark Office under Serial Nos. 99597929, 99603145, 99607461, and 99749088. The marks are used with the TM symbol on first occurrence in each Part of this whitepaper. The reader is advised that other marks referenced in this whitepaper, including without limitation “NIST,” “FIPS,” “GDPR,” “ISO,” “W3C,” “IETF,” “C2PA,” and other standards and statutory references, are the property of their respective owners and are used here in their nominative, descriptive sense to identify standards, statutes, or organizations.

J.5 Patent notice

The Trust Identity Protocol is the subject of five United States provisional patent applications. Dinesh Mendhe, Founder and Chairman of The AI Lab, is the sole inventor named on each application. The applications are assigned to The AI Lab Intelligence Unobscured, Inc. The applications, organized in Claim Groups A through BB, cover the foundational three-layer proto-

col architecture, the v2.0 refinements, the content-layer normalization framework, the identity-layer typed taxonomy and multi-officer governance, and the community verification layer with multi-model consensus classification and staking-based trust verification. Corresponding non-provisional applications and foreign filings will be pursued in accordance with the applicable filing deadlines and treaty arrangements. No license to any patent is granted by this whitepaper. Patent licenses are granted exclusively under the terms of TIPCL-1.0 Section 8 as described in Part X and Appendix F. Implementation of the Trust Identity Protocol without a TIPCL-1.0 license, or by a person or entity whose TIPCL-1.0 license has terminated, may infringe the patents of The AI Lab.

J.6 Regulatory references and compliance posture

References in this whitepaper to statutes, regulations, executive orders, and other instruments of public law, including without limitation Regulation (EU) 2024/1689 (the EU Artificial Intelligence Act), Regulation (EU) 2016/679 (the General Data Protection Regulation), Regulation (EU) 2022/2065 (the Digital Services Act), Regulation (EU) 910/2014 as amended (eIDAS 2.0), the Federal Trade Commission Act, the NIST AI Risk Management Framework, Executive Order 14110, the Online Safety Act 2023 (United Kingdom), the Privacy Act 2020 (New Zealand), and the Digital Personal Data Protection Act 2023 (India), are provided for the convenience of the reader and to identify the regulatory landscape within which the Trust Identity Protocol is designed to operate.

Statements in this whitepaper that the protocol or its implementations are aligned with, conformable to, or designed to support compliance with any such instrument are statements of design intent and of the technical controls implemented in the canonical specification. Such statements are not (a) certifications of compliance, (b) opinions of counsel, (c) attestations by a regulator or supervisory authority, or (d) representations that any specific implementation of the protocol by any specific licensee will satisfy the licensee's compliance obligations. A licensee is responsible for assessing the application of any such instrument to the licensee's particular use of the protocol and for obtaining qualified legal advice.

J.7 Forward-looking statements safe harbor

This whitepaper contains forward-looking statements within the customary meaning of that term. Forward-looking statements include without limitation statements concerning (a) operational milestones, (b) the activation of the AI Trust Council, (c) the formation of standards organization workstreams, (d) the conversion of the licensing regime under TIPCL-1.0 Section 12, (e) the expansion of the Federated Network, (f) the introduction of new technical features, and (g) the response of regulators, publishers, and other third parties to the Trust Identity Protocol.

Forward-looking statements are based on the present expectations of The AI Lab and are subject to risks and uncertainties, many of which are outside the control of The AI Lab. Actual results may differ materially from those projected. Material risks include without limitation:

1. Changes in applicable law, regulation, or executive policy in any jurisdiction in which the protocol operates.
2. Failure of one or more Founding Node Operators or Verification Providers to perform under their agreements with The AI Lab.

3. Adverse outcomes in pending or future patent, trademark, or copyright proceedings before the United States Patent and Trademark Office, the United States Copyright Office, or analogous foreign authorities.
4. Adverse outcomes in pending or future civil or regulatory proceedings.
5. Cryptographic developments that materially affect the security assumptions of the post-quantum primitives identified in Part III.
6. Materialization of an adversarial attack model that the protocol does not, as designed, mitigate.
7. Failure of the standards organization engagement described in Section 9.7 to converge on a published international standard.
8. Reduction in market interest in content provenance and verifiable identity.

The AI Lab undertakes no obligation to update or revise any forward-looking statement, whether as a result of new information, future events, or otherwise, except as may be required by applicable law.

J.8 No reliance and no offer

No person is entitled to rely on this whitepaper as a basis for entering into a commercial relationship with The AI Lab or with any Founding Node Operator, Verification Provider, or licensee, except pursuant to a written agreement signed by an authorized officer of the counterparty. This whitepaper is not an offer to sell, a solicitation to buy, or a recommendation to invest in any security of The AI Lab or of any other person. Any commercial relationship with The AI Lab is governed exclusively by the written agreement entered into between the parties.

J.9 Limitation of liability

To the maximum extent permitted by applicable law, in no event shall The AI Lab, its directors, officers, employees, advisors, or agents be liable to the reader or to any third party for any indirect, incidental, consequential, special, exemplary, or punitive damages arising out of or in connection with the reader's use of this whitepaper, even if The AI Lab has been advised of the possibility of such damages. The aggregate liability of The AI Lab arising out of or in connection with this whitepaper shall not exceed one hundred United States dollars (US\$100). This limitation applies notwithstanding the failure of essential purpose of any limited remedy.

J.10 Governing law and dispute resolution

This Appendix and any non-contractual obligations arising out of or in connection with it are governed by the law of the State of Delaware, without regard to its conflict of laws principles. Any dispute, controversy, or claim arising out of or relating to this Appendix, including any question regarding its existence, validity, or termination, shall be referred to and finally resolved by binding arbitration administered by the American Arbitration Association under its Commercial Arbitration Rules, seated in Wilmington, Delaware. The arbitral tribunal shall consist of one arbitrator unless the amount in dispute exceeds US\$1,000,000, in which case the tribunal shall consist of three arbitrators. The language of the arbitration shall be English.

J.11 Severability

If any provision of this Appendix is held to be invalid, illegal, or unenforceable by a court or tribunal of competent jurisdiction, such provision shall be severed and the remaining provisions shall continue in full force and effect.

J.12 Entire understanding

This Appendix, together with TIPCL-1.0 and any written agreement signed between The AI Lab and the reader, constitutes the entire understanding between The AI Lab and the reader concerning the subject matter of this whitepaper, and supersedes all prior or contemporaneous oral or written communications, proposals, and representations.

Contact The AI Lab regarding this Appendix

General Counsel: legal@theailab.org

The AI Lab Intelligence Unobscured, Inc. Wilmington, Delaware, United States.

Appendix K: Document History

This Appendix records the version history of the TIP Protocol Whitepaper.

Version	Date	Changes	Ratifying authority
v1.0	June 2, 2026	Initial publication, written and edited by Dinesh Mendhe, Founder and Chairman of The AI Lab Intelligence Unobscured, Inc., sole inventor of the Trust Identity Protocol. Eleven Parts, twelve Appendices, Foreword by the Founder, Executive Summary.	Sole director of The AI Lab
v1.1	August 23, 2026	Corrections and amendments published at theailab.org/whitepaper/errata . Free Tier eligibility of an individual natural person corrected to any income level, with the US\$100,000 threshold applying to businesses only. TIPCL-1.0 commercial fee schedule amended to scale by the World Bank income group of the licensee's country of principal place of business below US\$25,000,000, subject to a US\$125 minimum, with a single worldwide rate at and above US\$25,000,000. Section 5.8 on content classification, transient media retention, and classifier improvement added. Appendix F.2 eligibility table repaired.	Sole director of The AI Lab

The v1.0 publication locks in: the TIPCL-1.0 nine-tier Commercial Tier Schedule with the canonical US\$100,000 free-tier threshold (no US\$500,000 ceiling); the AI Trust Council with a Found-

ing Chair, two independent Founding Members, a Founder Seat, an Ex Officio Observer, and Joining Members in accession; identification of seven signed Founding Node Operators and three Founding Node Operators in accession across four jurisdictions; the Genesis Date framework (commencement of live operation on a date in June 2026 published not less than three business days in advance on the satisfaction of the Operational Readiness Conditions); the Pre-Genesis Period commencing June 1, 2026; the EU AI Act classification posture (the protocol is not an AI system under Article 3(1) and is not a social scoring system under Article 5(1)(c)); positioning of the AI Trust Council under Article 95; designation of the EU Authorized Representative as a conditional commitment under Article 22; the nine-step CNA-2.2 normalization specification; the four sub-score Trust Score with weights 0.20 / 0.30 / 0.30 / 0.20; six Blocking Items B1 through B6; the bonded jury adjudication procedure; the Article 50(2) and 50(4) compliance infrastructure positioning; and the Apache License 2.0 conversion provision on January 1, 2031. The portfolio of five United States provisional patent applications by Dinesh Mendhe (Claim Groups A through BB) is recorded in Section 10.4 and Appendix J.5.

K.1 Cadence of subsequent versions

Subsequent versions of the whitepaper are published on the schedule set by the AI Trust Council. The Council's customary review cadence is annual, with extraordinary revisions on the occurrence of a material change in the regulatory landscape, a material change in the protocol's technical architecture, or a material change in the institutional posture of The AI Lab or the Council.

K.2 Amendment procedure

Substantive amendments to the whitepaper require ratification by the AI Trust Council under the supermajority threshold described in the Charter. Non-substantive amendments (the correction of typographical errors, the updating of dated references to standards-organization positions, the refresh of contact addresses) may be made by The AI Lab on the responsibility of the General Counsel without separate Council ratification. Each published version of the whitepaper carries (a) the version number, (b) the publication date, (c) the ratifying authority, (d) the cryptographic hash of the canonical PDF rendering, and (e) the United States Copyright Office application number associated with the version.

K.3 Withdrawal

A published version of the whitepaper may be withdrawn by ratification of the AI Trust Council on the determination that the version contains an error materially adverse to a relying party. Withdrawal is published with the same prominence as the original publication and identifies the basis for withdrawal. The withdrawal of a version does not affect the operational status of the protocol or the rights of licensees acquired prior to the withdrawal.

K.4 Archival

Each published version of the whitepaper is archived in (a) the Minute Book of The AI Lab Intelligence Unobscured, Inc., (b) the public repository at theailab.org/whitepaper, (c) the United States Copyright Office deposit accompanying the version's copyright application, and (d) the

Internet Archive snapshot captured on the publication date for independent third-party timestamping.

Appendix L: The Architect and Author

This Appendix records, for the historical record, the architect and author of the Trust Identity Protocol and of this whitepaper.

L.1 Identification

The architect and author is **Dinesh Mendhe**, Founder and Chairman of The AI Lab Intelligence Unobscured, Inc.

L.2 Works of authorship and invention

Dinesh Mendhe is the creator and founder of the following works, each of which is identified by its canonical name and the year of its creation:

- The **Trust Identity Protocol** (2026): the open, post-quantum cryptographic standard for verifying human identity and AI content provenance on the public internet, comprising the TIP-ID identity layer, the TIP-CONTENT provenance layer, and the TIP-TRUST reputation layer, as documented in this whitepaper and in the canonical TIP Protocol Specification.
- The **AI Trust Council** (2026): the independent multi-stakeholder governance body established to ratify the governance, evolution, and enforcement of the Trust Identity Protocol, modeled on the consensus traditions of ICANN, the IETF, and the W3C.
- The **AI Trust Registry** (2026): the public, append-only, cryptographically anchored registry of TIP identities and content provenance declarations, maintained by the federation of Founding Node Operators and Verification Providers.
- The **AI Trust ID** (2026): the umbrella identity framework under which TIP issues and verifies the identity of a verified natural person. The AI Trust ID framework comprises two named components: the **Cryptographic Trust Identity (CTID)**, which is the globally unique, post-quantum-signed cryptographic identifier bound to the holder; and the **TIP-ID**, which is the protocol-layer identifier used to address the holder in TIP transactions and in the AI Trust Registry. Both CTID and TIP-ID are sub-components of the AI Trust ID, not synonyms for it.
- The **Human Trust ID** (2026): the externally rendered, human-readable verification credential that asserts the bearer's identity has been verified by an accredited Verification Provider under the AI Trust Council's accreditation framework.
- The **Canonical Normalization Algorithm** at every published version (CNA-1, CNA-1.0, CNA-2, CNA-2.1, CNA-2.2): the deterministic content-canonicalization procedure underlying the three-hash content addressing scheme of TIP-CONTENT.
- The **Origin Code system** (OH, AA, AG, MX): the four-value taxonomy by which the origin of a content unit is declared at publish time.

He is the sole inventor named on the five United States provisional patent applications underlying the Trust Identity Protocol, organized in Claim Groups A through BB and recorded in Section 10.4 and Appendix J.5.

L.3 Founder’s commitment

The works named in L.2 were conceived as public infrastructure for the post-AI internet. The reference implementation converts irrevocably to the Apache License 2.0 on January 1, 2031, and the canonical specification is published in perpetuity under the Creative Commons Attribution 4.0 International Public License. These commitments are operative and cannot be revoked. The mission framing under which the works were created is “intelligence, unobscured”: the proposition that the integrity of human knowledge in the AI era depends on the cryptographic verifiability of who made online content and how, and on the institutional independence of the body that governs the verification standard.

L.4 Statement for the historical record

The Trust Identity Protocol, the AI Trust Council, the AI Trust Registry, the AI Trust ID, and the Human Trust ID were created during the period 2024 to 2026. They were published, in their canonical form, under the architectural and authorial responsibility of Dinesh Mendhe. The institutional and operational framework through which the protocol enters production at Genesis Date, as defined in Part XI, was likewise designed by him. The historical record of the protocol’s creation, of the Council’s establishment, of the Founding Node Operator network, of the licensing framework, of the patent portfolio, and of the regulatory posture, is set out in this whitepaper, in the canonical TIP Protocol Specification, in TIPCL-1.0, in the AI Trust Council Charter, and in the corporate records of The AI Lab Intelligence Unobscured, Inc., which records are preserved in the Minute Book of the corporation and in the United States Copyright Office under the application numbers cited in the colophon and Appendix J.

The Trust Identity Protocol is intended to outlast its architect and to remain the public infrastructure through which the question **who made this content** can be answered, cryptographically and verifiably, for as long as the internet exists.