



FOR HUMANITY

INFRASTRUCTURE OF TRUST FOR AI - GUIDE TO ENTITY ROLES AND RESPONSIBILITIES

MAY 2021 // PREPARED BY RYAN CARRIER

ACKNOWLEDGMENTS

The author would like to thank Mark Potkewitz, for his invaluable feedback on earlier drafts of this report.

PRINT AND ELECTRONIC DISTRIBUTION RIGHTS



© 2021 by ForHumanity. This work is licensed under a Attribution-NonCommercial-NoDerivatives 4.0 International license.

To view a copy of this license, visit: <https://creativecommons.org/licenses/by-nc-nd/4.0/>



Ryan Carrier, CFA

Ryan is Founder, Executive Director, and Chairman of the Board of Directors of ForHumanity.
Ryan@Forhumanity.center | <https://forhumanity.center/>



What is *my* role in an Infrastructure of Trust?

Introduction

DESCRIBING THE ROLES IN AN INFRASTRUCTURE OF TRUST FOR AI AND AUTONOMOUS SYSTEMS - WE HAVE A MODEL WITH A LONG TRACK RECORD OF SUCCESS. FORHUMANITY IS ADAPTING IT FOR AI AND AUTONOMOUS SYSTEMS.

Background on Independent Audit

In 1973, the accounting industry came together and formed The Financial Accounting Standards Board (FASB) which created the Generally Accepted Accounting Principles (GAAP) which still govern financial accounting today. Eventually, the US Securities and Exchange Commission, and other extranational regulatory agencies, required adherence to the GAAP standard for all publicly listed companies. This clarity and uniformity significantly improved the financial world. An infrastructure of trust has been built over the past 50 years because of critical features such as independence, certified practitioners, and third-party rules that are compliant with the law and best practices.

Adapting to AI and Autonomous Systems

ForHumanity has advocated for the adoption of this infrastructure of trust and explained how it can be adapted and adopted for the Governance, Accountability, and Oversight of AI and Autonomous Systems. We support the creation and mandate of Independent Audit of AI Systems (IAAIS). IAAIS provides a comprehensive solution grounded in the same fundamental principles as Independent Financial Audit. ForHumanity develops and maintains audit and certification criteria designed for a range of industries and jurisdictions.

The proposed system replicates the distributed oversight, accountability, and governance needed for AI and autonomous systems in the same manner as financial audit, through audit and pre-audit service providers. These entities will employ certified practitioners to prepare for an eventual independent audit performed by other certified practitioners. The audit criteria are presented transparently to maximize an entity's ability to achieve compliance. Advancements in systems technology allow many of these processes to be automated for entities such as with the Treadway Commissions' Committee of Sponsoring Organization (COSO) framework for internal risk, audit, and controls. The result is a fully integrated, compliance-by-design infrastructure that embeds human agency, transparency, disclosure, and compliance from design to decommissioning.

Role in Independent Audit of AI and Autonomous Systems

The audit criteria are applied in two vectors: 1) Top-down accountability, governance and oversight 2) laterally, AI system by AI system. The top-down approach creates accountability systems for ethics, bias, privacy, trust, and cybersecurity for the Board of Directors, Chief Executive Officer, and Chief Data Officer. Committee structures are required such as an Algorithmic Risk, Children's Data Oversight, and Ethics to manage the audit/compliance responsibilities. All of these top-down criteria apply to every AI and every autonomous system in the organization. The system-specific audit criteria are designed to ensure legal and best practice compliance tailored to the specific impact of each system on humans. This comprehensive approach ensures consistency across the organization combined with complete risk management coverage of each unique system.

Participants in the System

The roles largely remain the same in Independent Audit of AI Systems as described in [Taxonomy](#). There are six distinct roles in most jurisdictions. Each player performs their function and the rules are executed in the same conflict-free manner, ensuring the highest integrity.

Auditors

- An Auditor engages in 3-party contract party contracts, with the Target of Evaluation (ToE) and on behalf of the public or intended users.
- The auditor deploys certified practitioners to conduct the audits.
- The auditor itself is certified by the Government Accreditation Service.
- When audits are conducted there is no feedback loop to the company and the audit is compliant or non-compliant.
- Audits are publicly disclosed according to the rules of the jurisdiction.
- The Auditor is liable for false assertions of compliance
- An Auditor is licensed for use of certification criteria
- The Auditor shall not provide Pre-audit services to Audit clients
- An Auditor may provide Pre-Audit services to non-Audit ToEs (may require accreditation)

Pre-Audit Service Providers (PASP)

- PASP engages in a 2-party contract directly with the Target of Evaluation
- There is a direct feedback loop between the ToE and PASP
- The PASP may or may not deploy certified practitioners per local jurisdiction rules
- The PASP may or may not be accredited by the Government Accreditation Service
- The PASP offers no certification or guarantee of audit compliance
- The PASP works are private, on behalf of the ToE
- The PASP is not liable for failed compliance or false assertions of compliance
- The PASP may or may not be licensed for use of certification criteria, but must be licensed if the service offered is related to or designed to satisfy certification requirements
- The PASP shall not be the auditor for a PASP client
- A PASP may offer Audit service to non-PASP clients (must be accredited)
- A PASP may deploy compliance-in-a-box solutions for criteria compliance

Target of Evaluation

- ToE may engage PASP
- ToE shall have an Auditor if required by the Relevant Legal Framework
- ToE pledges that all systems to be certified will be disclosed to the Auditor, failure in this regard is the responsibility of the ToE
- ToE dealings with PASP shall be confidential and non-public audit compliance may be confidential with an Auditor
- ToE shall maintain compliance structures, such as Algorithmic Risk Committee, Children's Data Oversight Committee, and Ethics Committee
- ToE shall build and maintain internal controls and systems to aid in compliance with audit requirements and foster robust risk management
- ToE shall be responsible for all public disclosures

Third-Party Criteria creation, maintenance, and individual certifier (ForHumanity)

- Non-profit organization
- Independent of Auditors and PASP
- Transparent and considerate of input and critique from all participants
- Criteria designed to support human well-being
- Conflict-free of Corporate or other ToE influence
- Submits to the authority of the jurisdiction for certified criteria
- Iterates and maintains criteria consistent with the law and best practices in a binary and auditable fashion
- Trains and certifies individual practitioners on all criteria
- Maintains a transparent repository of use cases and knowledge stores in support of Auditors to facilitate the evaluation of compliance
- Licenses criteria to all qualified entities
- Provides standard contract clauses for Auditors and PASP
- Engages in distributed education system to maximize availability and certified individuals
- Maintains a system of Continuing Education (CE)
- Maintains a searchable, registration system of Accredited Individuals

Government-appointed Accreditation Service

- Creates trust and confidence in products and services
- Assures that entities have sufficient talent, skill, and scope to provide certification
- Regular review of accreditation standards
- Regular review of Accredited entities
- Regular review of Third-party Criteria provider and individual certification
- Maintains an accessible list of accredited entities
- Maintains an accessible list of sanctioned or suspended entities

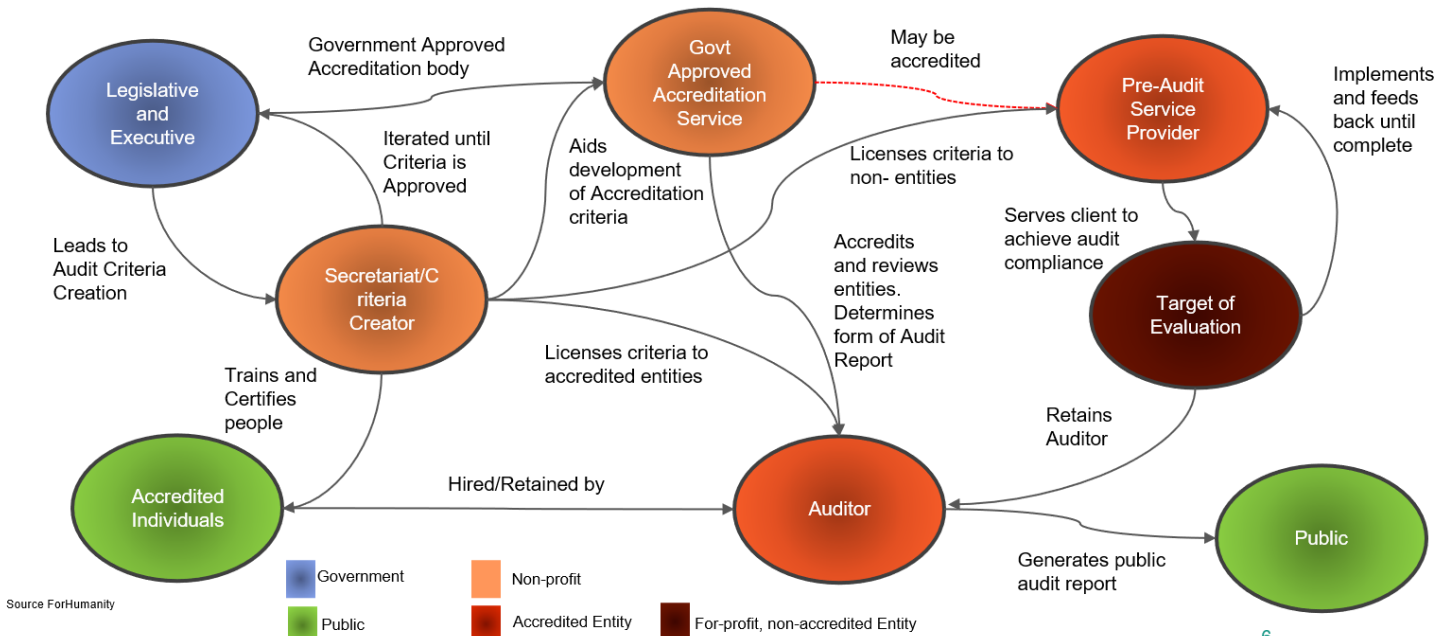


Infrastructure of Trust for AI -
Guide to Entity Roles and
Responsibilities

Governments/Regulators or similar Law-making/enforcement body

- Elected Body
- Legislative responsibilities
- Executive or enforcement responsibilities
- Regularly meets to review laws and best-practices
- Reviews and accredits (or rejects) submitted criteria
- Engages in enforcement actions for non-compliance with the law
- Handles concerns and issues brought by the Public

Roles and Responsibilities – for AI Audit infrastructure



Infrastructure of Trust for AI - Guide to Entity Roles and Responsibilities

A Few Differences between IA AIS from Financial Audit?

The skills and scope of IA AIS are different

In financial audits, there is great nuance, detail, and precisions required to evaluate the proper treatment of accounts, debits, and credits. A special type of mind is required to execute these works in complex organizations such as financial institutions. However, the skills required are very narrowly focused in the body of work referred to as accounting. Certified Public Accountants receive their accreditation because they demonstrate a sufficient understanding of this narrow but deep body of knowledge.

In Independent Audit of AI Systems, there is less nuance and more broad brushes of knowledge areas. Ranging from the law (a body unto itself) to Ethics/Philosophy (an academic discipline of its own) and to Cybersecurity (a field nearly 20 years old), it will be a rare practitioner who can master all areas of IA AIS. Therefore specialists and teams are expected to form so that an accredited organization can accomplish an audit, especially an enterprise-wide audit. It is likely that these specializations will occur along the lines of Ethics, Bias, Privacy, Trust, and Cybersecurity, especially in light of pre-existing expertise in Privacy and Cybersecurity. Additionally, there will be unique certifications based upon dedicated schemes, laws, and government-led areas of focus such as Privacy and Data Protection laws, Bias Audits, and Children's Codes. ForHumanity provides specialized training and accreditation for individuals for all specialties as well as general accreditation for individuals.

The Dynamism of the underlying marketplace

In financial audit, the scope was and remains the "numbers" for an entity. Financial audit focuses narrowly on debits, credits, assets, liabilities, cash flows, and balance sheets. The original version of Generally Accepted Accounting Principles (GAAP) would bear a strong resemblance to today's version. They would not be the same, of course, new entities, new tax codes, reinterpretations of rules and standards and the adaptation of best practices will have refined GAAP into the body of work that exists today. However, an auditor trained in the 1970s would find a large portion of their knowledge still useful and relevant.

For AI and Autonomous Systems, the market is substantially larger and the breadth of its reach might impact every function of a company. Furthermore, the technology itself changes dramatically year over year, decade over decade. In financial audit, this might be compared to an entire overhaul of a tax code every few years. This pace of change is not anticipated to abate and thus, laws and best practices are expected to evolve at a more rapid rate. The impact of dynamism is that the criteria and compliance will change more often requiring constant scrutiny and adaptation. It also means that the knowledge and understanding of the practitioners will have to be constantly evolved and updated. An AI Auditor's expertise is not likely to last for 50 years as has much of the core knowledge for financial accounting.

Whose interests are served in criteria creation?

After the Financial Accounting Standards Board (FASB) Wheat Committee member John Biegler later wrote, “The Wheat Committee approached its task with a sense of urgency, and I suspect that much of that sense of urgency was derived from a strong feeling on the part of most, if not all, Committee members that perhaps they were participants in a “last chance” effort to find a way to keep the financial accounting standard-setting process in the private sector”. The Wheat Committee, seven white men appointed by a single man, Marshall Armstrong, to rescue control of accounting standards for the industry. While their work may have succeeded and established the Financial Accounting Standards Board, the genesis of the organization and its founding members would hardly be described as “inclusive” or attempting to represent the views and well-being of humans in establishing FASB. In fact, it was a form of self-preservation for their industry.

This stands in sharp contrast to the mission of ForHumanity To examine and analyze the downside risks associated with the ubiquitous advance of AI & Automation, to engage in risk mitigation and ensure the optimal outcome... ForHumanity. We recognize the immense value and trust created by financial audit and its associated infrastructure on behalf of the public, for that we applaud the efforts of the Wheat Committee and FASB. However, a further hallmark of IAAIS is the embedded human agency and the requirement for diverse inputs and multi stakeholder feedback in the life cycle of AI and Autonomous Systems. For that reason, a seven-man committee designing the audit rules would fail miserably to ensure the optimal outcome... ForHumanity. Therefore, our crowdsourced, transparent, all-inclusive process is better suited to the needs of humans today as we construct an infrastructure of trust in AI and Autonomous Systems. This guide is designed to identify each of the main roles in the infrastructure of Independent Audit of AI Systems as well as call out their specific responsibilities. This organizational structure mimics the successful ecosystem of financial audit and enhances it in key areas. ForHumanity believes that this will lead to heightened trust in our AI and Autonomous systems and provide a roadmap for proactive execution of laws, regulations, and best practices.



Infrastructure of Trust for AI - Guide to Entity Roles and Responsibilities

ForHumanity

ForHumanity is a 501(c)(3) tax-exempt public charity formed to examine and analyze the downside risks associated with the ubiquitous advance of AI and automation. To this end, we engage in risk control and mitigation and deploy the lens and filter of Ethics, Bias, Privacy, Trust, and Cybersecurity to ensure the optimal outcome...ForHumanity.

ForHumanity is an interdisciplinary group of dedicated expert volunteers, with over 300 contributors and 32 Fellows, Its collective expertise spans the AI field, ranging from ethics to algorithmic risk and to security. Our team is drawn from the academic, legal, policy, corporate, and public sectors of over 40 countries around the world. Our mission is to help create an 'infrastructure of trust' for all autonomous systems that directly impact humans.

ForHumanity drafts comprehensive, pragmatic and implementable audit rules and standards for autonomous systems in every corner of the economy. Our experts collaborate with industry practitioners to ensure these audits achieve our mission of mitigating AI risk to humans. This system of audit rules and standards - adapted to local jurisdictional laws and regulations - is called Independent Audit of AI Systems (IAAIS).