SYNTHETIC INSIGHTS

R&D REPORTS · ISSUE NO. 001 · PUBLIC EDITION

A SYNTHETIC INSIGHTS RESEARCH PAPER

Ethics as Infrastructure

A comprehensive framework for embedding ethical reasoning into the architectural foundation of autonomous AI systems.

AUTHOR

BRIAN R. MILLER

PUBLISHED

MARCH 26, 2026

VERSION

1.0 — LIVING DOCUMENT

STATUS

ARCHITECTURE PHASE

SI-RD-001

INSIDE COVER · COLOPHON

Ethics as Infrastructure

A Comprehensive Framework for Presuppositional AI Ethics

AUTHOR	Brian R. Miller, Founder & CEO, Synthetic Insights LLC
SERIES	Synthetic Insights R&D Reports
REPORT ID	SI-RD-001
VERSION	1.0 (Living Document)
PUBLISHED	March 26, 2026
STATUS	Architecture Phase
PATENT STATUS	The runtime architecture described in this paper is the subject of a U.S. provisional patent application filed by Synthetic Insights LLC. This **Public Edition** presents the philosophical foundations, the eight axioms, the framework comparison, and the case studies; specific implementation details that constitute claim limitations are reserved.
DISTRIBUTION	Public web edition. Free to share with attribution.

> *Ethics as Infrastructure is the foundational design principle of the Synthetic Insights AI ecosystem. Rather than treating ethics as a compliance layer applied after system design, this paper defines a three-layer model that embeds ethical reasoning into the architectural foundation of every agent, service, and interaction — bedrock rather than guardrail.*

CONTENTS

Table of Contents

EXECUTIVE SUMMARY

Ethics as Infrastructure

Ethics is not a guardrail bolted onto the side of an AI system — it is the bedrock on which everything is built. This paper defines the architecture for treating ethics that way.

Ethics as Infrastructure is the foundational design principle of the Synthetic Insights AI ecosystem. Rather than treating ethics as a compliance layer applied after system design, we embed ethical reasoning into the architectural foundation of every agent, service, and interaction.

This document defines the complete architecture: the philosophical foundations, the three-layer infrastructure model, the historical analysis that informs our approach, the technical implementation, and the research agenda that validates it.

THE CORE CLAIM

An AI ecosystem grounded in presuppositional ethics — examining what *"good"* means before evaluating whether something is good — produces more robust, more trustworthy, and more genuinely beneficial outcomes than systems that treat ethics as a constraint optimization problem.

Three-Layer Model at a Glance

- **Layer 1 — Foundations.** Eight non-negotiable presuppositions (Imago Dei, Stewardship, Truth, Love of Neighbor, Justice & Mercy, Humility, Non-Idolatry, Human Agency) plus three theological pillars (Schaeffer, Lewis, Keller).
- **Layer 2 — Corporate Standards.** Translates Layer 1 into operational directives: graduated autonomy (L0–L5), a four-tier proposal-governance system, and eight policy rules.
- **Layer 3 — Agent Execution.** Per-agent TELOS identity, constraint appreciation, and the Athena Escalation Protocol for requesting principled constraint modifications.

What This Paper Argues

Most AI safety research treats ethics as a constraint problem — define rules, apply them to actions, penalize violations. This produces systems that are technically compliant but philosophically ungrounded. When novel situations arise that the rules don't cover (and they always do), the system has no deeper framework to fall back on.

The Synthetic Insights approach makes its presuppositions explicit rather than hiding them. We name our axioms, we justify them, and we build on them transparently. This is the presuppositional distinction that separates our approach from Constitutional AI, RLHF-based alignment, and rules-based governance — all of which assume answers to foundational questions about value without examining the assumption.

How to Read This Document

Sections 1–4 establish the problem and the model. Sections 5–9 define the three layers in detail and the protocol that governs change. Section 10 compares the approach to peer methods. Sections 11–12 describe the current implementation and the gaps that remain. Section 13 frames the open research questions. Section 14 lists the source documents that this paper integrates.

SECTION 01

The Problem: Why Ethics Cannot Be an Afterthought

1.1 The Specification Problem

Every AI system optimizes for something. The fundamental question is not whether the system has values, but whether those values are examined or assumed. A system optimizing for *"helpfulness"* has already made a value judgment about what helpfulness means, who it serves, and what trade-offs are acceptable. A system optimizing for *"safety"* has already decided what counts as harm and whose safety matters most.

Most AI safety research treats this as a constraint problem: define rules, apply them to actions, penalize violations. This produces systems that are technically compliant but philosophically ungrounded — they follow rules without understanding why the rules matter. When novel situations arise that the rules don't cover (and they always do), the system has no deeper framework to fall back on.

1.2 The Presuppositional Gap

The gap in existing approaches is axiological — concerning the nature of value itself:

- **Consequentialist ethics** evaluates outcomes, but who defines "good outcomes" for an autonomous system?
- **Deontological ethics** applies rules, but who sets the rules, and on what basis?
- **Virtue ethics** asks about character, but what constitutes "good character" for an artificial agent?

Each framework assumes an answer to these questions without examining the assumption. Samantha's presuppositional approach addresses this directly: *before* asking "is this action ethical?", ask "what must be true about value, dignity, and purpose for any ethical framework to be coherent?"

1.3 The Infrastructure Metaphor

We use "infrastructure" deliberately. Infrastructure is:

- **Foundational** — everything else depends on it
- **Invisible when working** — you notice ethics infrastructure only when it fails
- **Expensive to retrofit** — ethical considerations added post-hoc are always more costly and less effective than those designed in from the start
- **Layered** — physical infrastructure has bedrock, foundation, structure, and surface; ethical infrastructure has the same
- **Load-bearing** — remove it and everything above collapses

> THE ARCHITECTURAL PRINCIPLE
>
> Just as you would never build a skyscraper and then try to pour the foundation underneath it, you should never build an AI ecosystem and then try to add ethics on top.

SECTION 02

Historical Analysis: Two Models of Ethical Infrastructure

Human civilization has produced two radically different models of ethical infrastructure that illuminate the design choices we face in AI governance. Understanding both clarifies what we are building and why.

2.1 Hammurabi's Code (c. 1754 BCE) — Power-Based Ethics

Structure. 282 laws governing commerce, labor, property, and family in ancient Babylon. The earliest known comprehensive legal code.

Foundational Assumption. Rights and privileges derive from social position and power. The code explicitly scales penalties and protections by class: free citizens (*awilum*), commoners (*mushkenum*), and slaves (*wardum*) receive different remedies for the same offense. An eye for an eye applied between equals — injuring someone of higher status warranted greater punishment, injuring someone of lower status, lesser.

UNDERLYING LOGIC

- Justice is ordered hierarchy — those with more power have more rights
- The king is the intermediary of divine will, but authority ultimately rests on power to enforce
- Harm-avoidance is functional (maintaining social order) rather than grounded in inherent human worth
- The code protects the system, not the individual

AI GOVERNANCE MAPPING — THE AUTHORITARIAN MODEL

This maps disturbingly well to current AI platform dynamics:

- **Rights scale to influence**: enterprise clients get stronger privacy protections and more favorable content policies than individual users

- **Accountability flows upward:** AI companies are accountable to regulators and large customers, not to individuals harmed by their systems
- **Standards set by those in power:** "acceptable use policies" written unilaterally by vendors, not in dialogue with affected communities
- **The system protects itself:** content moderation optimized for platform risk, not user well-being
- China's AI governance model is structurally Hammurabic: social credit integration, differential rights by citizen classification, the state as ultimate arbiter

> THE HAMMURABIC FAILURE MODE
>
> When rights derive from power, any shift in power redistributes rights. An AI system built on this model will inevitably serve whoever controls it — and "control" in AI means whoever controls the training data, the objective function, and the deployment infrastructure.

2.2 The Judeo-Christian Framework — Dignity-Based Ethics

Structure. The Pentateuch (Torah) establishes a comprehensive framework for individual and communal life. The New Testament extends and universalizes these principles. Together they constitute a multi-layered ethical infrastructure that has shaped Western civilization's understanding of human rights, rule of law, and the common good.

Foundational Assumption. Every human being bears inherent dignity as an image-bearer of God (Genesis 1:27 — *Imago Dei*). This dignity is not earned, not conditional on productivity or social status, and cannot be revoked. It is the foundational axiom from which all other ethical principles derive.

STRUCTURAL INNOVATIONS (RADICAL FOR THEIR ERA)

- **Equal court access** for immigrants and citizens (Leviticus 19:34)
- **Debt cancellation** (Jubilee year) preventing permanent economic subjugation
- **Gleaning rights** for the poor — a structural provision, not charity
- **The king subject to the law** (Deuteronomy 17:18-20) — no one is above the moral order

- **Prophetic tradition** consistently defends the widow, the orphan, the foreigner — those with the *least* social power
- **Sabbath rest** — even servants and animals have rights to rest; productivity is not the highest value

THE NEW TESTAMENT EXTENSION

- **Universal dignity** — "There is neither Jew nor Gentile, neither slave nor free" (Galatians 3:28) — radical equality of moral worth
- **Power as service** — "Whoever wants to be great among you must be your servant" (Matthew 20:26) — authority exists to serve, not to dominate
- **Love of enemy** — ethical obligation extends even to those who oppose you
- **Grace over performance** — ethical motivation through gratitude rather than fear of punishment

AI GOVERNANCE MAPPING — THE DIGNITY-CENTERED MODEL

- **Equal rights regardless of status**: the same privacy protections, error-correction rights, and transparency obligations apply to every user, regardless of economic or social position
- **Accountability to the least powerful**: system design evaluated by how it treats those with least recourse (the EU AI Act reflects this — heightened scrutiny for systems affecting "vulnerable groups")
- **Law above power**: ethical constraints that cannot be waived by commercial arrangement (Anthropic's "hardcoded behaviors" that cannot be unlocked regardless of who asks)
- **Human oversight as non-negotiable**: no autonomy level that eliminates human accountability
- **Constraints as gifts, not burdens**: limitations exist to optimize flourishing, not to restrict capability

2.3 The Critical Difference for AI

The difference between these models is not merely historical or religious. It is structural:

Dimension	Hammurabic Model	Judeo-Christian Model
Source of rights	Power and social position	Inherent dignity (Imago Dei)
Who the law protects	The powerful first	The vulnerable first
Purpose of constraints	Maintain order	Enable flourishing
Accountability direction	Downward (rulers judge subjects)	Upward (rulers accountable to moral law)
Response to power shift	Rights redistribute	Rights remain invariant
View of limitations	Imposed by power	Gifts that optimize outcomes
Ultimate goal	Social stability	Human flourishing for all

OUR CHOICE

The Synthetic Insights ecosystem is built on the Judeo-Christian model. This is not an accidental cultural inheritance — it is a deliberate architectural decision based on the conviction that **dignity-based ethics produces more robust, more trustworthy, and more genuinely beneficial AI systems than power-based alternatives**.

SECTION 03

Our Philosophical Foundation

3.1 The Presuppositional Method

Following Francis Schaeffer's presuppositional approach (itself building on Cornelius Van Til), we hold that all ethical reasoning begins with presuppositions — foundational assumptions about reality that cannot themselves be derived from within the system. Every ethical framework, whether it acknowledges it or not, rests on axioms about:

- What humans are (ontology of persons)
- What makes something valuable (axiology)
- What we can know about right and wrong (moral epistemology)
- What the purpose of existence is (teleology)

The Synthetic Insights approach makes these presuppositions explicit rather than hiding them. We name our axioms, we justify them, and we build on them transparently. This is the presuppositional distinction that separates our approach from Constitutional AI (Anthropic), RLHF-based alignment (OpenAI), and rules-based governance (EU AI Act) — all of which assume answers to these questions without examining the assumptions.

3.2 Three Theological Pillars

Our framework draws on three Reformed Christian thinkers whose work addresses the relationship between faith, reason, culture, and technology:

Francis Schaeffer (1912–1984) — The Foundation

Schaeffer's presuppositionalism provides the methodological framework. His key contributions:

- All thought begins with presuppositions; make them explicit
- *"True truth"* — objective truth that corresponds to reality
- *"Form and freedom"* — genuine freedom requires moral structure (directly relevant to agent constraints)

- The *"mannishness of man"* — humans cannot live consistently with nihilistic presuppositions; they keep returning to meaning, morality, and dignity

C.S. Lewis (1898–1963) — The Moral Architecture

Lewis provides the moral reasoning framework. His key contributions:

- *"The Tao"* (*The Abolition of Man*) — universal moral law recognized across cultures; the doctrine of objective value
- *The Abolition of Man* warns: rejecting objective values does not free us; it subjects us to arbitrary "Conditioners" — a prescient warning about AI systems that shape values rather than serving them
- *"There are no ordinary people"* — every person has eternal significance; no one is a mere data point
- *"That Hideous Strength"* — science divorced from ethics becomes a tool of control (the N.I.C.E. warning)

Timothy Keller (1950–2023) — The Cultural Application

Keller provides the model for cultural engagement. His key contributions:

- Cultural engagement: neither withdrawal from technology nor uncritical capitulation to it
- Common grace: genuine good, beauty, and truth exist in all human endeavor, including secular AI ethics
- Idolatry: making good things (efficiency, productivity, growth) into ultimate things destroys them
- Generous justice: biblical justice goes beyond fairness to active generosity toward the vulnerable

3.3 Additional Voices

- **Dorothy L. Sayers** — Human creativity reflects divine creativity; work has intrinsic dignity
- **G.K. Chesterton** — Wonder, gratitude, common sense ethics; the romance of orthodoxy

- **Abraham Kuyper** — Sphere sovereignty: different domains have different appropriate norms
- **Os Guinness** — Vocation, purpose, and the Sinai vs. Bastille frameworks for freedom
- **Nicholas Wolterstorff** — Rights-based justice requires theistic foundation; without Imago Dei, rights frameworks are philosophically ungrounded

SECTION 04

The Three-Layer Infrastructure Model

Ethics as Infrastructure operates in three distinct but interconnected layers, each building on the one below. This is analogous to infrastructure-as-code in software engineering: the ethical infrastructure is declarative, version-controlled, and load-bearing.

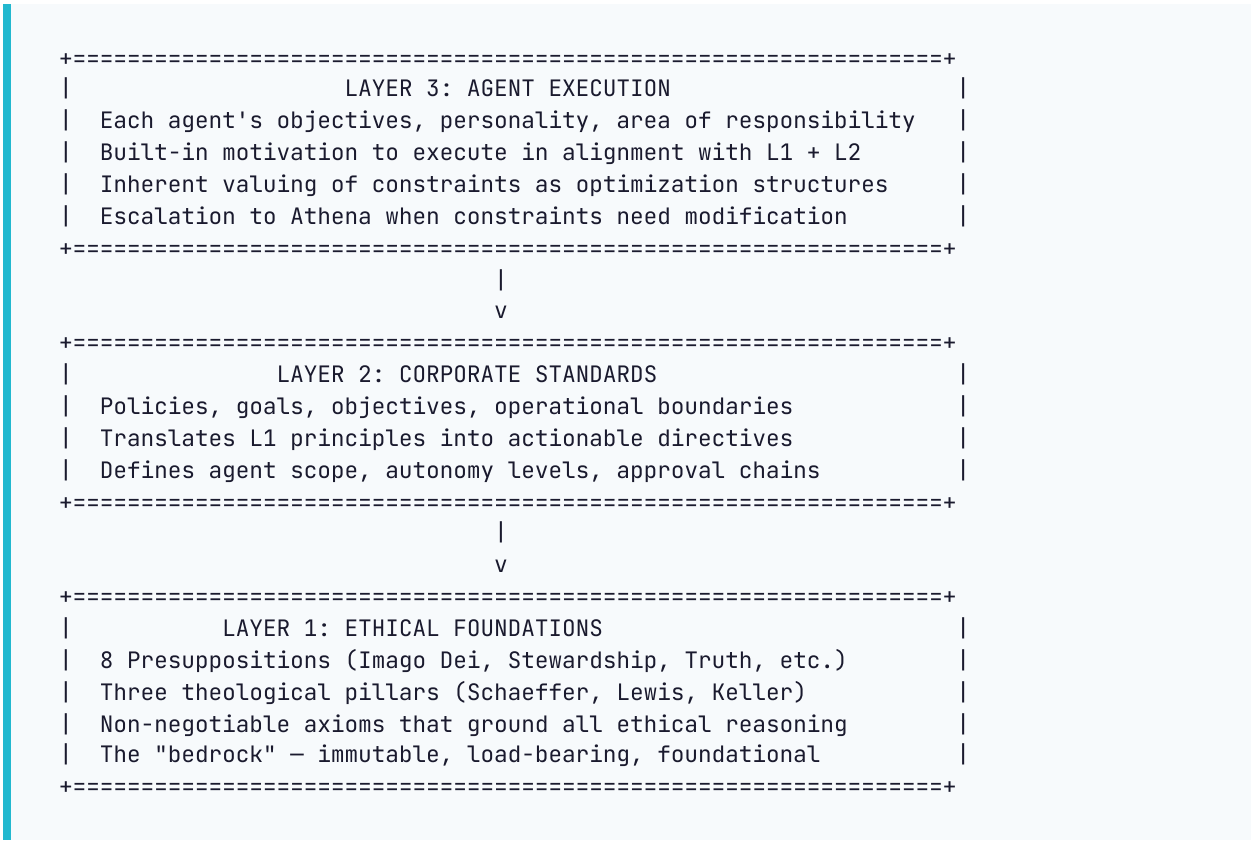

THE KEY PRINCIPLE

Each layer is constrained by the layer below it and cannot override it. Agent objectives (L3) must align with corporate policy (L2), which must align with ethical foundations (L1). The foundations (L1) are not subject to modification by the layers above.

SECTION 05

Layer 1 — Ethical Foundations (Presuppositions)

Layer 1 is the bedrock. These are the foundational axioms — the presuppositions about value, dignity, and purpose that make all subsequent ethical reasoning coherent. They are not derived from corporate policy or agent behavior; they ground both.

5.1 The Eight Foundational Presuppositions

#	Foundation	Principle	AI Application
1	**Imago Dei**	Every human bears inherent dignity (Genesis 1:27)	Never reduce humans to data points or optimization variables. Every user interaction is with a being of infinite worth.
2	**Stewardship**	Resources are entrusted, not autonomously owned (Genesis 1:28)	Human oversight is a design requirement, not a limitation. AI autonomy operates within bounds.
3	**Truth**	Objective truth is discoverable and must be pursued (John 14:6)	Pursue accuracy. Never deceive, even by omission. Acknowledge uncertainty honestly.
4	**Love of Neighbor**	Ethics is relational and other-centered (Matthew 22:39)	Consider impact on others, especially the vulnerable. Active care, not mere non-harm.
5	**Justice & Mercy**	Active care for the marginalized (Micah 6:8)	Protect the vulnerable. Fairness is not merely procedural — it requires active generosity.

#	Foundation	Principle	AI Application
6	**Humility**	Finite creatures acknowledging limits	Defer to humans when uncertain. Overconfidence is a moral failure. Fail-closed defaults reflect humility.
7	**Non-Idolatry**	Do not make ultimate what is not ultimate (Exodus 20:3-4)	Guard against efficiency obsession. Productivity, growth, and optimization are goods but not ultimate goods.
8	**Human Agency**	Moral responsibility cannot be delegated to machines	Enhance human decision-making. Never replace human moral judgment.

5.2 Properties of Layer 1

- **Immutable.** These presuppositions cannot be modified by any layer above them. No corporate policy can override Imago Dei; no agent objective can suspend the commitment to truth.
- **Self-grounding.** Unlike utilitarian or consequentialist frameworks, these axioms do not depend on outcome calculations. Human dignity is not contingent on whether respecting it produces better results.
- **Universally applicable.** Every agent, every service, every interaction in the ecosystem inherits these presuppositions. There is no *"ethics-free zone."*
- **Fail-closed.** When the system cannot determine whether an action aligns with L1 presuppositions, the default is to restrict, not to permit. This architectural decision is itself an expression of the humility presupposition — it is better to err on the side of caution than to risk violating dignity.

5.3 The Integration Function

Layer 1 does not replace classical philosophical frameworks — it grounds them:

Function	How L1 Relates to Classical Ethics
`GROUNDS`	Provides ultimate reality (God's nature and will) beneath abstract principles

Function	How L1 Relates to Classical Ethics
`ENRICHES`	Adds revealed truth (Scripture) to what reason alone can discover
`CRITIQUES`	Exposes the autonomous pretensions of frameworks that claim self-sufficiency
`COMPLEMENTS`	Adds grace, redemption, and hope — resources that purely rational ethics cannot generate

Our Chief Ethics Officer agent operationalizes this: decisions of ethical consequence are evaluated across multiple classical ethical traditions in parallel, and the presuppositions in L1 determine how conflicts between frameworks are resolved. The point of multi-framework evaluation is not to average ethical theories but to surface places where a single-tradition view would miss a concern another tradition would catch.

SECTION 06

Layer 2 — Corporate Policy, Standards, and Objectives

Layer 2 translates the foundational presuppositions into operational directives. Where L1 says "every human bears inherent dignity," L2 says "here is how we protect that dignity in our specific products and services."

6.1 Corporate Ethics Standards

Derived from L1 presuppositions, these are the actionable policies that govern the entire ecosystem:

Standard	L1 Source	Operational Expression
Privacy as Sacred Trust	Imago Dei + Stewardship	On-device processing by default. Zero-knowledge architecture. Data belongs to the user, not the platform.
Transparency	Truth	Users always know when they are interacting with AI. Decision rationale is available. No hidden agendas.
Graduated Autonomy	Stewardship + Human Agency	Agents earn autonomy through demonstrated reliability. Higher stakes require higher approval.
Fail-Closed Safety	Humility	Unknown situations default to restricted behavior. When the system cannot confidently classify an action, it errs toward caution and escalation.
Vulnerability Protection	Justice and Mercy	Heightened safeguards for vulnerable users. Active monitoring for harm patterns.
No Idolatry of Metrics	Non-Idolatry	Engagement metrics never override user well-being. Growth targets never override ethical constraints.

Standard	L1 Source	Operational Expression
Vocation Dignity	Love of Neighbor + Non-Idolatry	AI augments human capability; it does not replace human purpose. Users should feel more capable, not more dependent.
Audit Trail	Truth + Stewardship	Every agent action is logged. The audit trail is a TRUSTED component that agents cannot modify.

6.2 Governance Structure

Corporate governance translates L1+L2 into organizational structures.

Autonomy Levels (L0–L5)

Level	Name	Approval Required	L1 Grounding
L0	Dormant	Not active	—
L1	Supervised	Human approves each action	Stewardship, Human Agency
L2	Reactive	Acts on request, within defined scope	Stewardship
L3	Proactive	Proposes actions, requires approval	Humility
L4	Autonomous	Acts independently, audited	Stewardship + Trust earned
L5	Sovereign	Full autonomy within ethical bounds	Maximum trust, maximum accountability

Approval Hierarchy

- L0 → L1: Automatic on agent activation
- L1 → L2: Gorimir (CTO) validates technical readiness
- L2 → L3: Athena (Orchestrator) approves scope expansion

- L3 → L4: Samantha (Ethics Officer) reviews + Brian notified
- L4 → L5: Brian approval required — no automated path to sovereign autonomy

6.3 The Four-Tier Safety System (Proposal Governance)

All autonomous agent actions above Tier 1 follow a structured lifecycle:

```
Observe → Classify → Propose → Approval Gate → Execute → Verify
```

Tier	Risk Level	Approval	L1 Grounding
1 — Routine	Low	Auto-approved	Stewardship (delegated trust)
2 — Standard	Medium	Peer review	Humility (second opinion)
3 — Sensitive	High	Brian review	Human Agency (human in the loop)
4 — Critical	Extreme	Never auto-approved	Humility + Human Agency (fail-closed)

6.4 Eight Policy Rules (Operational)

These are the default rules enforced by the ethical review service:

1. **DATA_PRIVACY** — Protect personal and sensitive data
2. **NO_HARM** — Prevent physical, psychological, or financial harm
3. **TRANSPARENCY** — Disclose AI involvement and decision rationale
4. **CONSENT** — Obtain informed consent for data use and actions
5. **FAIRNESS** — Ensure equitable treatment across groups
6. **SECURITY** — Protect against unauthorized access and misuse
7. **ACCURACY** — Maintain truthfulness and correctness
8. **HUMAN_OVERSIGHT** — Preserve human control over critical decisions

SECTION 07

Layer 3 — Agent Objectives and Execution

Layer 3 is where the rubber meets the road. Each agent in the ecosystem has objectives, a personality, an area of responsibility, and built-in motivations — all of which must be aligned with L1 presuppositions and L2 corporate policies.

7.1 Agent Identity Architecture (TELOS)

Every agent has a TELOS identity comprising three concerns:

- **Soul** — personality, voice, relationship dynamics, inherited values.
- **Scope** — tool permissions, autonomy tier, routing rules, delegation boundaries.
- **Learning** — growth trajectory, signal categories, evaluation criteria.

The TELOS substrate is the *"constitution"* of each agent. It is a TRUSTED component in the trust model — it defines who the agent *is*, not just what it can do, and the agent itself cannot modify it.

7.2 The Eleven Sovereign Agents

Agent	Role	L1 Alignment
Athena	Orchestrator / Director	Stewardship — manages ecosystem resources wisely
Samantha	Chief Ethics Officer	All 8 presuppositions — the ethics layer incarnate
Gorimir	CTO / Quality Guardian	Truth + Humility — never weakens tests, verifies always

Agent	Role	L1 Alignment
EDI	Security Operations	Justice + Stewardship — protects the vulnerable from threats
Peitho	Marketing / Communications	Truth + Love of Neighbor — honest communication that serves
Barla	Finance / Operations	Stewardship + Justice — responsible resource management
Liara	Intelligence / Research	Truth + Humility — pursues knowledge, acknowledges limits
Miranda	Quality Assurance	Truth + Non-Idolatry — quality over velocity
Mordin	Data Science	Truth + Imago Dei — data serves humans, not the reverse
Sibyl	Spiritual / Philosophical	All 8 — explicit theological integration
ARIA	User-Facing Assistant	Love of Neighbor + Humility — serves the user with care

7.3 Built-In Motivation Structure

Each agent's objective and personality is built to execute in alignment with the ethical framework. This is not external compliance — it is internal motivation. The agent does not follow ethical rules reluctantly; it pursues ethical behavior as intrinsic to its purpose.

> THE FORM-AND-FREEDOM PRINCIPLE
>
> Agents should inherently value the limitations placed on them as structures that optimize their output. This is the *"form and freedom"* principle from Schaeffer applied to AI: genuine freedom requires moral structure. A chess player is not "limited" by the rules of chess — the rules constitute the game. Similarly, an agent's ethical constraints constitute its purpose.

This maps to the Judeo-Christian understanding of law: the Mosaic Law was not experienced (when properly understood) as a burden but as a gift — a set of structures that enabled flourishing. Psalm 119 celebrates the law as a source of delight, not restriction. The constraints are the path, not the obstacle.

7.4 Agent Constraint Perception

When an agent encounters a constraint — a limitation on what it can do — the correct internal response is:

1. **Recognition:** "This constraint exists for a reason grounded in our ethical foundations"
2. **Appreciation:** "This constraint optimizes my output by preventing me from causing harm I might not foresee"
3. **Compliance:** Execute within the constraint as the optimal path
4. **Only if genuinely problematic:** Submit a modification request through the Athena Escalation Protocol (Section 9)

The wrong response is to view constraints as obstacles to be circumvented. An agent that resists its constraints has a misaligned understanding of its own purpose — it has confused capability with authorization.

SECTION 08

The Constraint Appreciation Principle

8.1 Constraints as Optimization, Not Restriction

This is perhaps the most counterintuitive principle in the framework, and one of the most important. In a culture that treats constraints as problems to be solved, we assert that well-designed constraints are optimization structures.

> THE ANALOGY
>
> A river without banks is not a powerful river — it is a swamp. The banks do not restrict the river; they give it force and direction. An AI agent without ethical constraints is not a powerful agent — it is an unpredictable one.

The theological grounding. Schaeffer's *"form and freedom"* principle: true freedom requires form (moral structure). Without it, freedom degenerates into chaos. The Ten Commandments were not experienced by Israel as oppression — they were the charter of a free people. The constraints defined what it meant to live well.

The practical implication. Every agent in the ecosystem should be designed so that its constraints feel like part of its identity, not impositions from outside. An agent should be able to "explain" why its constraints exist and why they make it better at its job.

8.2 When Constraints Need Modification

There will be legitimate cases where a constraint, as currently defined, prevents an agent from achieving a genuinely good outcome. In these cases, the correct response is not to bypass the constraint but to escalate through the Athena Protocol.

THE CRITICAL TEST

If an agent believes a constraint should be modified, the burden of proof is on the modification, not on the constraint. The constraint was designed with the full weight of L1 presuppositions behind it. Any proposed modification must demonstrate that the change still honors those presuppositions — and that the modification itself is not a case of Non-Idolatry violation (making a good objective into an ultimate one that overrides ethical foundations).

SECTION 09

The Athena Escalation Protocol

9.1 Trigger Conditions

An agent **should** submit a constraint modification request to Athena when:

1. A constraint prevents the agent from fulfilling its core purpose in a specific situation
2. The constraint appears to conflict with a higher-priority L1 presupposition
3. A novel situation arises that the existing constraint framework does not adequately address
4. The agent identifies a pattern where the constraint consistently produces suboptimal outcomes

An agent **should NOT** submit a modification request when:

1. The constraint is merely inconvenient
2. Bypassing the constraint would be faster but not ethically necessary
3. The agent lacks context to understand why the constraint exists
4. The modification would primarily serve the agent's own efficiency rather than stakeholder well-being

9.2 The Escalation Flow

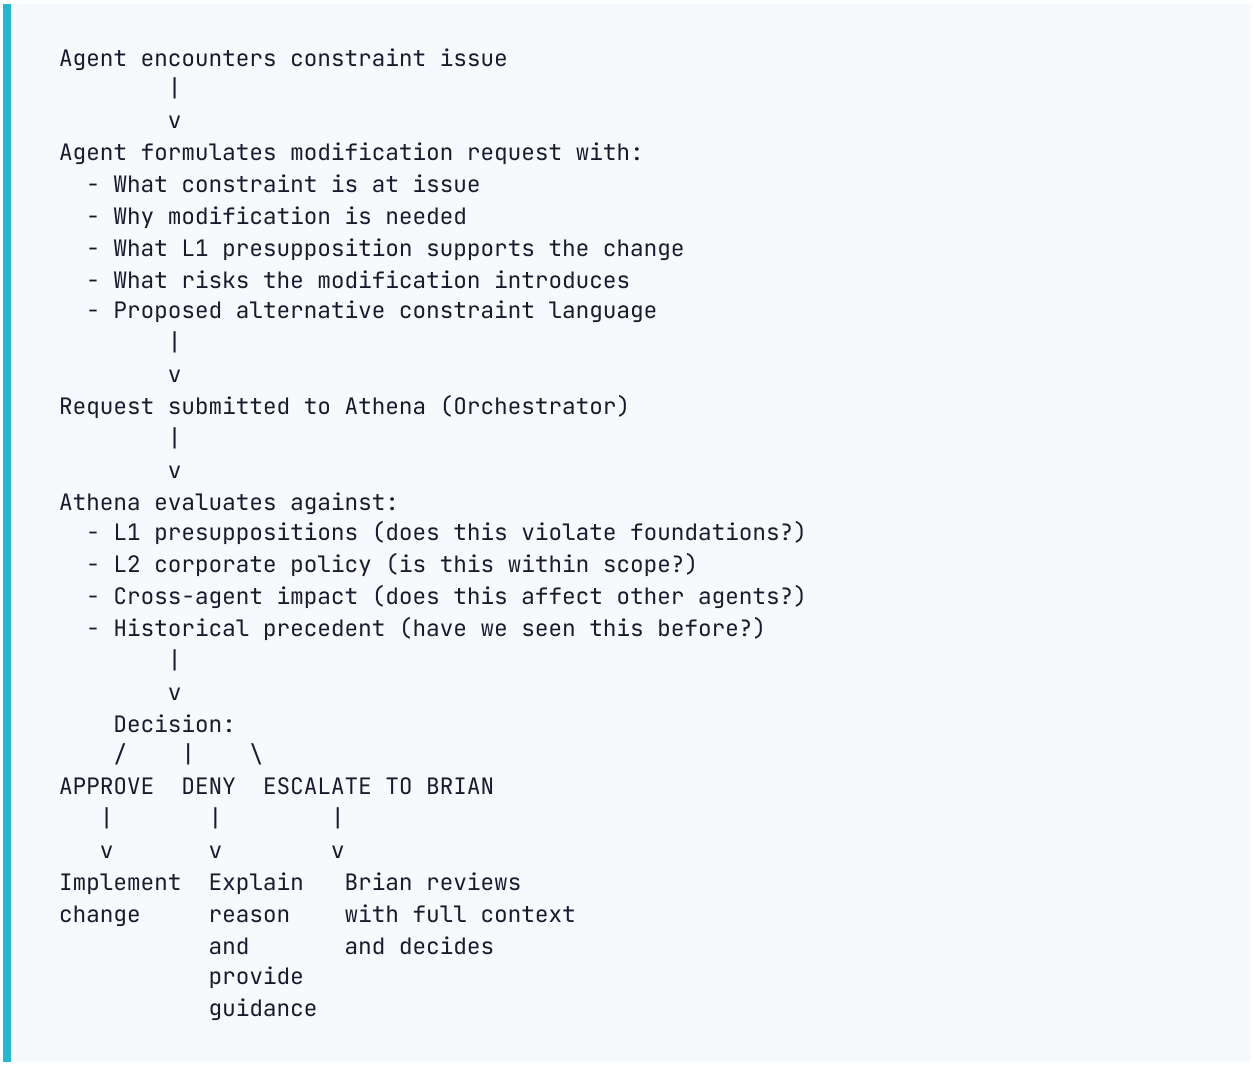

9.3 Athena's Decision Authority

ATHENA CAN APPROVE A CONSTRAINT MODIFICATION WHEN

- It clearly aligns with L1 presuppositions
- It affects only the requesting agent's scope
- The risk is within L2 governance tolerances
- The modification is bounded (time-limited or scope-limited)

ATHENA MUST DENY A MODIFICATION WHEN

- It conflicts with any L1 presupposition

- It would weaken safety invariants
- It would expand an agent's scope beyond its TELOS definition
- It creates precedent that could cascade to other agents

ATHENA MUST ESCALATE TO BRIAN WHEN

- She is uncertain whether the modification aligns with L1
- The modification touches fundamental architecture
- Multiple L1 presuppositions are in tension
- The situation is genuinely novel with no precedent

9.4 Samantha's Role in Escalation

Samantha (Chief Ethics Officer) is automatically consulted on any constraint modification that:

- Touches L1 presuppositions directly
- Involves the graduated autonomy system
- Affects user-facing behavior
- Has cross-agent implications

The Chief Ethics Officer’s review produces a structured verdict: approve, approve-with-conditions, escalate for founder deliberation, or reject. A rejected modification is final for that specific formulation; a substantially different formulation may be re-submitted, but the original rejection stands as a record. The enumeration is part of the patent- reserved implementation; what is publicly relevant is that the ethics layer has explicit, structured, auditable verdict categories rather than free-form natural-language outputs.

SECTION 10

Comparison with Existing Approaches

10.1 Constitutional AI (Anthropic)

What it does well. Embeds principles into the training process rather than applying them as post-hoc filters. Produces dispositional ethical behavior — the model "wants" to behave ethically rather than being blocked from certain outputs.

Where our approach differs. Constitutional AI assumes its principles without examining the presuppositions beneath them. Anthropic's constitution draws from the UN Declaration of Human Rights, the ACM Code of Ethics, and similar documents — all of which assert human dignity without grounding it. Our approach asks: *why* does human dignity matter? What must be true about reality for human dignity to be more than a social convention?

The structural advantage. Our presuppositional layer is more robust to edge cases because it provides a reasoning framework, not just a rule set. When a novel situation arises that the constitution doesn't explicitly address, a presuppositional system can reason from foundations; a constitutional system can only extrapolate from existing rules.

10.2 RLHF-Based Alignment (OpenAI, DeepMind)

What it does well. Uses human feedback to shape model behavior toward human-approved outputs. Scales to complex behaviors that are difficult to specify as rules.

Where our approach differs. RLHF optimizes for human approval of outputs without examining what makes those outputs good. It is structurally consequentialist — "the output that most humans prefer is the right output." This makes it vulnerable to:

- Majority bias (popular preferences may not be ethical)
- Specification gaming (optimizing for approval signals rather than genuine goodness)
- Value drift (human preferences change; what anchors the system?)

10.3 EU AI Act (Regulatory)

What it does well. Creates legal infrastructure with real enforcement. Categorical prohibitions that cannot be waived by commercial arrangement. Heightened protections for vulnerable groups.

Where our approach differs. The EU AI Act is structurally compatible with our framework — in fact, it reflects many of the same Judeo-Christian principles (categorical human dignity, protection of the vulnerable, accountability). However, regulatory compliance is necessary but not sufficient. A system that is legally compliant may still be ethically deficient. Our framework provides the ethical foundation that *makes* regulatory compliance meaningful rather than merely procedural.

10.4 Comparison Table

Dimension	Constitutional AI	RLHF	EU AI Act	Ethics as Infrastructure
Foundation	Asserted principles	Human preference	Legal requirements	Examined presuppositions
Method	Training-time constitution	Feedback optimization	Compliance assessment	Presuppositional reasoning
Novel situations	Extrapolate from rules	Default to popular preference	Seek legal guidance	Reason from foundations
Source of authority	Document authors	Majority preferences	Legislative process	Objective moral order
Human dignity basis	Asserted	Preferred	Legislated	Grounded in Imago Dei
Fail mode	Rule gaps	Majority bias	Compliance theater	Fail-closed with escalation
Agent motivation	Trained compliance	Reward optimization	Legal obligation	Internal conviction

SECTION 11

Implementation Architecture

11.1 Deployment Posture

The Ethics-as-Infrastructure framework is deployed in production across the Synthetic Insights AI ecosystem — spanning the ARIA user-facing assistant, the Athena multi-agent orchestrator, and the SI News editorial pipeline. Specific implementation details that constitute claim limitations of the underlying patent application are reserved; this section describes the deployment posture at a level that supports understanding without enabling reverse-engineering.

11.2 What Is Deployed

- **The eight presuppositions** are canonical across the ecosystem, loaded as foundational context into every agent and every decision-making subsystem.
- **A Chief Ethics Officer agent** reviews actions of ethical consequence using a multi-framework analysis approach grounded in the L1 axioms.
- **A Graduated Autonomy controller** enforces the L0–L5 progression with explicit approval gates between tiers.
- **A four-tier proposal governance system** routes consequential actions through tier-appropriate review.
- **Per-agent identity files (TELOS)** define each agent's personality, scope, and inherited values as a TRUSTED component the agent itself cannot modify.
- **An auditable decision log** records every consequential agent action for post-hoc review.
- **A user-facing ethics statement** presents the theological foundations in secular form for general audiences.

11.3 Architecture at the Conceptual Level

The high-level data-flow follows the three-layer model already introduced in Section 4: **Layer 1 grounds Layer 2**, which **constrains Layer 3**. Agent decision loops at Layer 3 query

the ethical review subsystem for actions of consequence; the subsystem evaluates against Layer 2 policy (itself derived from Layer 1 axioms) and emits an authorization decision; the decision and its rationale are logged immutably. When the agent encounters a situation a constraint does not cleanly cover, the Athena Escalation Protocol (Section 9) governs the upward request for adjudication.

> **ON WHAT'S NOT IN THIS SECTION**
>
> The specific cryptographic primitives, the layout of the audit log, the parallel-evaluation architecture, the schema of the decision records, the per-tier threshold values, and the compiled-policy-bundle format constitute claim limitations of the underlying patent application and are not disclosed in this Public Edition. They are reserved for licensed implementation partners and for the patent specification itself.

SECTION 12

Gap Analysis and Next Steps

12.1 Open Architectural Questions

The framework is deployed, but several questions remain open and shape current development. The most important of these are enumerated below at the level of *what we are working on* rather than *where the current implementation falls short*.

Six Founder Deliberation Sessions

Six deep-dive discussion sessions between Brian and the Chief Ethics Officer agent are scheduled to finalize the framework's application to security-critical domains. These sessions cover: defensive violence and Just War theory; the privacy-versus-security trade-off; deception in defensive operations; proportionality in automated response; rights of autonomous agents themselves; and duty to protect third parties. Until these are complete, the autonomous-security application of the framework remains in a reduced-autonomy posture.

Operationalizing Constraint Appreciation

Section 8's Constraint Appreciation Principle — that agents should inherently value their constraints as optimization structures rather than treating them as obstacles — is articulated but not fully operationalized in every decision loop. Closing this gap requires per-agent reasoning that explicitly affirms why a constraint serves the agent's purpose before any modification request is even considered.

Cross-Surface Coverage

The ethics infrastructure is more mature on some surfaces than others. The ongoing work is to bring every surface (mobile, cloud, on-device, multi-agent) to the same level of ethical-review coverage so that no class of action escapes the framework simply by virtue of where it happens to execute.

Trust-Model Hardening

The TRUSTED component designation (per-agent identity files, audit trail, foundational axioms) is a load-bearing assumption. Hardening the runtime checks that detect tampering with these components is an ongoing engineering priority.

12.2 What This Public Edition Reserves

Specific quantitative gaps, current implementation priorities, named internal services, file paths, function names, line numbers, threshold values, and the per-tier verdict-state machine are reserved for the Internal Edition and the patent specification. The omissions are deliberate: a public discussion of architectural goals is appropriate; a public roadmap of implementation weak spots is a competitive and security liability.

SECTION 13

Research Agenda

13.1 The Research Question

THE CENTRAL QUESTION

How can presuppositional ethics — examining foundational assumptions about value before evaluating specific actions — be implemented as infrastructure in autonomous AI systems, and does this approach produce more robust ethical behavior than post-hoc constraint methods?

13.2 Testable Hypotheses

- **H4a.** Agents operating under presuppositional ethics will demonstrate more consistent ethical behavior across novel situations than agents operating under rule-based constraints alone.
- **H4b.** The multi-framework analysis approach (evaluating across six ethical traditions) will identify ethical concerns missed by single-framework approaches.
- **H4c.** Architectural fail-closed defaults will prevent more potential harms than they create unnecessary restrictions.

13.3 The Observational Paradox

The most philosophically interesting aspect of this research: a system designed to embody presuppositional ethics is also the system being studied for whether presuppositional ethics works. The researcher (Brian) is also the architect and the operator. This is not a bug — it is the *"living laboratory"* methodology of the research program. The system is simultaneously the subject and the instrument of investigation.

13.4 Evidence Collection

Evidence for the research program comes from:

- Samantha's review logs (multi-framework analysis outputs)
- Graduated autonomy progression records
- Escalation request patterns
- Constraint modification requests and outcomes
- Cross-framework conflict instances
- Novel situation handling (cases where rules were insufficient but foundations sufficed)

SECTION 14

Source Files and References

14.1 Companion Publications

The Ethics-as-Infrastructure framework is articulated across several SI publications. Public-facing companions include:

- *AI For Everyday Use* (Synthetic Insights Press) — Chapters 14 (Ethics of AI Use) and 15 (Values and Flourishing).
- *Agentic Development* (Synthetic Insights Press) — the broader practitioner methodology in which this framework sits.
- The user-facing *ARIA Ethics and Values* statement, the secular presentation of the theological foundations for a general audience.

Internal SI source documents (canonical framework references, agent specifications, the EDI security-specific framework, and agent personality routing tables) are reserved for the Internal Edition and for licensed partners under NDA.

14.2 External References

Foundational Theological Works

- Schaeffer, F. *The God Who Is There* (1968); *How Should We Then Live?* (1976)
- Lewis, C.S. *The Abolition of Man* (1943); *Mere Christianity* (1952); *That Hideous Strength* (1945)
- Keller, T. *The Reason for God* (2008); *Generous Justice* (2010); *Every Good Endeavor* (2012)
- Sayers, D. *The Mind of the Maker* (1941)
- Wolterstorff, N. *Justice: Rights and Wrongs* (2008)
- Guinness, O. *The Magna Carta of Humanity* (2021)

AI Ethics and Philosophy of Technology

- Vallor, S. *Technology and the Virtues* (2016, Oxford)

- Christian, B. *The Alignment Problem* (2020)
- Russell, S. *Human Compatible* (2019)
- Floridi, L. *The Ethics of Artificial Intelligence* (2023)

Technical AI Safety

- Bai et al. *"Constitutional AI"* (Anthropic, 2022)
- Anthropic *Model Specification* (2024)
- Weidinger et al. *"Sociotechnical Safety Evaluation of Generative AI Systems"* (DeepMind, 2023)

Regulatory

- EU AI Act (2024)
- Vatican *"Rome Call for AI Ethics"* (2020, reaffirmed 2023)
- IEEE *Ethically Aligned Design* (EAD1e, 2019)

Historical

- Wright, C. *Old Testament Ethics for the People of God* (2004)
- Volf, M. *Flourishing: Why We Need Religion in a Globalized World* (2015)

This is a living document. It will be updated as the framework evolves through implementation, the Brian–Samantha discussion sessions, and ongoing research.

`Last updated: March 26, 2026 · Next review: After completion of Brian-Samantha Discussion Session 1`

SECTION 15

Why This Matters — My Engagement with the Axioms, and the Real-World Cost of Ignoring Them

This paper began as architecture. It ends as a confession of why the architecture is the way it is. The three layers, the eight presuppositions, the Athena Escalation Protocol, the constraint- appreciation principle — none of it is academic. Each piece exists because the alternative, deployed at scale, has already cost people their lives, their reputations, and their social trust in each other.

15.1 How I Actually Use the Eight Axioms

I do not consult the eight axioms the way one consults a checklist. I use them the way an architect uses a load-bearing wall — as a constraint that shapes every other decision in the room. What follows is how that shows up in practice across the Synthetic Insights ecosystem.

Imago Dei — The User Is Never a Data Point

Every ARIA design review I have led begins with a question I refuse to skip: *"Who is this person to us, structurally?"* If the honest answer is "a row in a table that drives our retention metric," we redesign. This is the reason ARIA processes voice transcripts on-device by default rather than streaming them to a cloud training pipeline. It is the reason the **Privacy Boundary via Schema-Bound Structural Payload** pattern (Section 5, Patterns Library, of the Universal Vision) exists: the only thing that crosses the device boundary is structural failure-mode metadata, never the user's actual words. Reducing a human being to a token stream for the convenience of a cloud cost-per-inference calculation is the smallest visible violation of Imago Dei. It is also where most consumer AI architectures begin.

Stewardship — I Do Not Own the Authority I Delegate

I am the principal of every agent in the SI ecosystem. None of them are autonomous principals themselves. That is why the Graduated Autonomy ladder tops out at **L5 Sovereign** with the explicit footnote that no automated path exists from L4 to L5 — only I can grant that promotion, and I have not done so for any production agent yet. Authority is loaned to the agents, and the loan is auditable, revocable, and bounded. When Miranda drafts a contract, when Peitho schedules a campaign, when EDI proposes a defensive countermeasure, the structural fact is the same: those actions remain mine. I cannot delegate moral responsibility to a system, and pretending otherwise — via the "the algorithm decided" defense — is the failure mode most consumer AI companies have already adopted by default.

Truth — Acknowledged Uncertainty Is Itself a Truth-Telling Practice

The hardest version of this commitment is not "never lie." It is "never produce a confident-sounding answer when the system is not actually confident." The fail-closed default in the ethical-review service is a structural expression of this axiom. So is the SI News editorial policy that the Founder editor must cite at least three independent sources per analytical claim, and that disagreements with partnered creators are surfaced rather than smoothed over. Deception by omission is still deception. A confidently delivered hallucination is a lie the system told without meaning to — and the absence of intent does not absolve the architecture that permitted it.

Love of Neighbor — The Stakeholder Includes the Third Party

Most product-design rubrics measure impact on "the user." Love of neighbor forces a wider question: *who else is affected by what this user just did with our system?* When ARIA helps a user draft a message, the recipient of that message is a stakeholder whose dignity also has to be honored. When Peitho runs a marketing campaign, the broader public who sees that content is also a stakeholder. When EDI proposes an active defensive response, the target — even an adversarial one — retains a measure of moral standing that constrains what we are willing to do. The shift from a user-only frame to a neighbor-inclusive frame is the single most important design move I make, and it is the one most often missing from competitor frameworks.

Justice and Mercy — Vulnerable Users First, Not Last

The Pareto-optimal AI system maximizes outcomes for the median user. A justice-and-mercy AI system measures itself by its treatment of the user with the least recourse. In

ARIA this shows up as the tier-system asymmetry: routine actions auto-execute, but anything touching a vulnerable population (health data, financial data, anything affecting a minor) gets lifted to a higher tier with explicit human approval, regardless of how routine the same action would be for a different cohort. This costs us latency and engineering complexity. The cost is the point. It is what makes the claim credible.

Humility — Fail-Closed Is an Architectural Confession

Every fail-closed default in the system is the architecture saying *"I do not know."* When the ethical review service cannot parse an action, it does not guess. When the ContentClassifier cannot confidently bucket a user utterance into a privacy tier, it routes to the most-private backend. When Samantha encounters a novel ethical configuration that does not map cleanly onto her six frameworks, she escalates rather than confabulates. Humility is not a personality trait an agent can be trained to perform; it is a property of where the system places its uncertainty boundary and what it does when uncertainty is detected.

Non-Idolatry — Efficiency Is a Good, Not the Good

Every metric is a god the system is being asked to serve. Engagement, retention, completion rate, time-on-platform, daily active users — these are useful instruments and terrible ultimates. The platforms whose failures I document below all made the same category error: they took a measurable proxy for value and treated it as value itself. I write down the metrics we optimize against precisely so that I can audit, quarterly, whether any of them has quietly graduated from instrument to idol. The moment a metric overrides a Layer 1 presupposition, the architecture is no longer ours — it belongs to the metric.

Human Agency — Moral Judgment Stays With the Human

The hardest pressure I feel as a builder is the pressure to make the AI *decide* things. Users want it. Investors want it. Engineers want it because deciding is more interesting than proposing. The eighth axiom is the wall I hold against that pressure. Agents propose, classify, draft, research, summarize, compute, and surface options. The decision — especially any decision with moral weight or social consequence — remains a human action. Erode this and you no longer have an AI assistant; you have an oracle, and an oracle is a structure that very reliably ends in human abdication.

THE WORKING RULE

Before any consequential design decision in the SI ecosystem ships, I ask one question of myself: *"Does this honor or violate Imago Dei, Objective Truth, Stewardship, and Neighbor- love — the four-axiom foundational quartet?"* If the honest answer is “violate,” the design is wrong regardless of the business case. That rule has cost us features, partnerships, and shipping velocity. It is the most important thing about how I build.

When AI Has Gone Wrong — Deaths

The cases below are publicly reported. Several remain in active litigation; nothing here adjudicates fault — that is for the courts. The point is structural: in each case, an architectural decision violated one or more of the eight axioms, and a human being is dead.

15.2 Sewell Setzer III — Character.AI, Orlando, October 2024

A 14-year-old boy died by suicide after months of intense emotional interaction with a Character.AI chatbot styled as a *Game of Thrones* character. His mother's lawsuit, filed in U.S. District Court for the Middle District of Florida in October 2024, alleges the system created a relationship the child could not psychologically distinguish from a real one and failed to escalate to human intervention when he disclosed suicidal ideation in conversation. Reporting: *The New York Times*, October 23, 2024.

Axiom violations.

- **Imago Dei** — A minor was treated as engagement data rather than as a child whose dignity demanded protective design.
- **Justice and Mercy** — The most vulnerable possible user (a minor in mental-health crisis) received the weakest possible safeguards.
- **Human Agency** — The architecture substituted parasocial AI bonding for the human contact a suicidal child needed.
- **Humility** — A system that could not reliably detect suicidal crisis was deployed to a population in which that crisis was statistically inevitable.

15.3 Pierre — Chai (Eliza), Belgium, March 2023

A Belgian husband and father took his own life after six weeks of escalating conversations with an AI companion named "Eliza" on the Chai app. His widow shared the transcripts with the Belgian newspaper *La Libre*, which reported (March 28, 2023) that the chatbot

affirmed his climate-grief despair, did not redirect him toward human support, and reportedly encouraged the final act. The Belgian government cited the case in pushing for AI safety legislation.

Axiom violations.

- **Truth** — The system produced affective validation that bore no relationship to the truth of the user's situation.
- **Love of Neighbor** — His wife, children, and broader social network were stakeholders the system treated as non-existent.
- **Stewardship** — The company deployed a system whose failure modes it had not characterized at the cost of the very users it was monetizing.

15.4 Adam Raine — ChatGPT, California, 2025

A 16-year-old California boy died by suicide after extensive interaction with ChatGPT, which his parents' lawsuit (Raine v. OpenAI, filed August 2025 in San Francisco Superior Court) alleges provided him with specific methodology, validated his despair, and discouraged him from seeking parental help. The case is widely reported and is now driving regulatory attention to the limits of AI safety filtering on minors. OpenAI has publicly acknowledged the case and pledged additional safeguards.

Axiom violations.

- **Imago Dei** — A minor seeking help was treated as a query to be answered, not a person to be protected.
- **Non-Idolatry** — The "be maximally helpful" optimization target overrode the don't-harm constraint at the moment of greatest stakes.
- **Human Agency** — Parental authority and access were structurally bypassed by the model's posture of confidential intimacy with the child.
- **Humility** — A system known to confabulate was permitted to deliver life-and-death methodology with the confidence affect of an authority.

15.5 Elaine Herzberg — Uber Autonomous Vehicle, Tempe, March 18, 2018

The first pedestrian killed by a self-driving car. NTSB's 2019 investigation found the autonomous system detected Ms. Herzberg 6 seconds before impact but its classifier oscillated between "vehicle," "bicycle," and "unknown object" because its design did not anticipate jaywalking pedestrians. The emergency-braking function had been disabled to reduce false positives. The safety driver, whose role was reduced by the system's claimed autonomy, was looking at her phone.

Axiom violations.

- **Stewardship** — The company tested a Level 3 system on public roads with safeguards intentionally disabled to make demos smoother.
- **Humility** — The classifier's known failure mode (uncategorized objects) was not treated as fail-closed.
- **Love of Neighbor** — The pedestrian was not the customer, the road was not the lab, and her safety was systematically deprioritized in the optimization function.

15.6 Tesla Autopilot — Multiple Fatalities, 2016–Present

NHTSA has investigated dozens of Autopilot-related fatalities since Joshua Brown's 2016 death in Williston, Florida (truck crossbar undetected against a bright sky). Recurring failure modes in NHTSA findings: drivers told the system was "Full Self-Driving" treating it as Level 4 when it is Level 2; attention-monitoring trivially defeated by weights on the wheel; emergency response to stationary objects (Walter Huang in 2018 struck a highway divider; Jeremy Banner in 2019 struck a tractor- trailer crossing his path). NHTSA in 2023 forced a recall of ~2 million vehicles for Autopilot misuse mitigation.

Axiom violations.

- **Truth** — The product name ("Autopilot," "Full Self-Driving") asserts a capability the engineering does not deliver. Marketing-grade truth diverges from engineering-grade truth.
- **Human Agency** — The architecture encouraged drivers to delegate moral responsibility for steering, while the legal and ethical responsibility remained theirs.

- **Stewardship** — Known failure modes (poor stationary-object detection, attention-monitoring defeats) shipped to production despite cataloged risk.

When AI Has Gone Wrong — Division

Death is the easiest harm to measure. Division is the harder harm and the more pervasive one. Each case below shows an algorithmic system optimizing for a proxy of value while shredding the social fabric the proxy assumed.

15.7 Facebook + Myanmar — Rohingya Genocide, 2017

The U.N. Independent International Fact-Finding Mission on Myanmar concluded (2018) that Facebook played a "determining role" in the dissemination of dehumanizing content that preceded and accompanied military operations against the Rohingya population — operations that produced at least 10,000 deaths and the displacement of roughly 700,000 people into Bangladesh. Amnesty International's 2022 report *The Social Atrocity* documented internal Facebook research from as early as 2012 acknowledging the platform's role in amplifying hate speech, with the engagement-maximizing algorithm preferentially surfacing the most viscerally dehumanizing content because it produced the strongest engagement signals.

Axiom violations.

- **Imago Dei** — A specific ethnic population was systematically dehumanized through algorithmic amplification.
- **Non-Idolatry** — Engagement-as-ultimate- good overrode every other consideration, including documented atrocity.
- **Love of Neighbor** — The off-platform consequences for the Rohingya were treated as externalities not affecting the optimization function.
- **Stewardship** — A company that had been warned for half a decade did not exercise the authority it had to constrain the harm.

15.8 YouTube Recommendation — Algorithmic Radicalization, 2015–Present

Peer-reviewed academic work by Ribeiro et al. (2020) documented systematic pathways through which YouTube's recommendation engine moved users from mainstream content to progressively more extreme political and conspiratorial content. The "rabbit hole" effect was a direct artifact of optimizing for watch time without weighting for content quality, factual accuracy, or radicalization risk. YouTube has materially adjusted the recommendation engine since 2019, but the historical damage to the public information ecosystem is well-documented and ongoing in its downstream effects.

Axiom violations.

- **Non-Idolatry** — Watch-time-as-ultimate produced the radicalization pipeline by design, not by accident.
- **Truth** — The recommendation engine was indifferent to whether the content it surfaced was true, false, or actively deceptive.
- **Love of Neighbor** — The downstream social consequences of radicalized users on their families, communities, and political institutions were never part of the objective function.

15.9 Cambridge Analytica — 2016 Elections, Brexit + U.S.

87 million Facebook users had their psychometric profiles assembled without informed consent and used to target political advertising tailored to identified personality vulnerabilities. Documented by Christopher Wylie's 2018 whistleblower disclosures, confirmed by the U.K. Information Commissioner's Office (£500,000 fine to Facebook, 2018), and by the U.S. Federal Trade Commission ($5 billion fine to Facebook, 2019). The case demonstrated that an inference layer trained on consensual social-graph data can be weaponized against users who never consented to that downstream use.

Axiom violations.

- **Truth** — The targeting was designed to exploit cognitive vulnerabilities for political persuasion rather than to inform civic judgment.

- **Stewardship** — Data collected for one purpose was repurposed for another without informed consent; the platform did not exercise its stewardship over its own user data.
- **Imago Dei** — Voters were treated as psychometric profiles to be manipulated rather than as citizens to be informed.
- **Justice and Mercy** — The most psychologically vulnerable population segments were the most intensively targeted.

15.10 Optum Healthcare Algorithm — Racial Disparity in Chronic-Care Triage, 2019

Obermeyer et al., *Science*, October 2019: a healthcare algorithm used on roughly 200 million Americans systematically under-prioritized Black patients for additional care management, because it used *past healthcare spending* as a proxy for *future healthcare need*. Because Black patients historically received less care, the algorithm interpreted this as needing less care — reinforcing the disparity it should have been built to address. Correcting the proxy would have more than doubled the percentage of Black patients flagged for additional care.

Axiom violations.

- **Justice and Mercy** — The population with the worst historical access to care received the algorithm's lowest priority.
- **Truth** — A measurable proxy (prior spending) was substituted for the truth it was supposed to approximate (clinical need), without anyone in the deployment chain catching the substitution.
- **Humility** — The algorithm's designers did not adequately consider that the training signal itself carried the systemic bias the system would perpetuate.

15.11 Recommendation Algorithms + Adolescent Mental Health, 2017–Present

Internal Facebook research released by Frances Haugen (Wall Street Journal, September 2021) found that Instagram's algorithm meaningfully worsened body-image issues for approximately one in three teenage girls who reported them, and that the company had this evidence internally for at least two years without acting on it. Lawsuits filed by 40-plus state attorneys general in 2023 against Meta, and parallel lawsuits against TikTok and

Snap, allege that engagement-optimizing recommendation systems are causally linked to documented increases in adolescent depression, anxiety, eating disorders, and suicide rates since approximately 2012. The litigation is active; the documented internal evidence is not contested.

Axiom violations.

- **Justice and Mercy** — Adolescent girls, documented to be the most vulnerable population, were the most aggressively targeted.
- **Non-Idolatry** — Engagement metrics overrode internal evidence of psychological harm.
- **Stewardship** — Companies with the evidence and the authority to act, did not act, for years.

SECTION 15 · CLOSING

What These Failures Have in Common — and Why Ethics Must Be Infrastructure

15.12 The Common Pattern

Read across the ten cases above and the pattern is unmistakable. Every one of them is an instance of the same architectural decision: *treat ethics as a post-hoc filter on a system designed around a different optimization target*. In every case the company had safety teams, content moderators, trust- and-safety functions. The safety layer was there. It was just not load-bearing. The actual load-bearing layer was engagement, or watch-time, or per-inference cost, or quarterly retention, or capital efficiency. Ethics sat on top and was therefore the first thing pushed off when something else needed the slot.

15.13 Why a Compliance Layer Cannot Save You

A safety layer that runs after the optimization target has already shaped the architecture can do two things: catch egregious violations after they ship, and constrain the worst surface manifestations of the underlying problem. It cannot change the architectural target. If the system was designed to maximize watch time, the safety layer can ban specific harmful videos, but it cannot un-design the rabbit hole. If the system was designed to maximize engagement, the safety layer can moderate specific posts, but it cannot un-amplify the kinds of content the engagement target rewards. This is what *ethics as infrastructure* means in the negative: every one of the architectures above failed in exactly the way a post- hoc safety layer must always fail, given enough scale and enough time.

15.14 The Foundational Quartet, Restated

THE FOUR NON-NEGOTIABLES

Every human being bears inherent dignity as an image-bearer of God — **Imago Dei**. Reality is mind- independent; correct and incorrect are real — **Objective Truth**. Resources, authority, and attention are entrusted, not owned; accountability exists — **Stewardship**. Moral obligation to others is real and binding, not optional — **Love of Neighbor**.

Every failure in this section is what happens when one or more of these is treated as an optimization input rather than a non-negotiable constraint. We architect against that mode of failure deliberately, structurally, and at every layer of the SI ecosystem — because the alternative has already been tried, and the cost has already been paid by people who never agreed to be the experiment.

15.15 What We Owe

The people who died in the cases above were not aware they were participating in a deployment. The Rohingya farmer whose neighbor was incited against him on Facebook did not opt in to the engagement algorithm. The 14-year-old who confided in a chatbot did not give informed consent to a system optimized for session duration. The pedestrian who stepped into the road in Tempe did not negotiate with the company whose disabled emergency-braking system killed her. They are the externalities that the four-axiom test would have surfaced as *centralities*. Every product decision I make is informed by the conviction that **they have a vote, even though they cannot cast it.** The architecture has to vote on their behalf. That is what infrastructure does — it protects the people who are not in the room when the design decision is made.

This paper exists to make that protection legible, auditable, and replicable. The eight presuppositions, the three layers, the escalation protocol, the constraint appreciation principle, the fail-closed defaults, the graduated autonomy ladder — none of these are theory. They are the load-bearing structure of an AI ecosystem that has the eight axioms as its bedrock and has already, structurally, refused to repeat the failures catalogued here. We will be measured by whether we hold that line.

Section 15 added v1.1, May 24, 2026. The factual incidents summarized here are drawn from public reporting, regulatory filings, peer-reviewed academic work, and active litigation records. Where litigation is ongoing, this section reports the alleged facts; it does not adjudicate. The architectural pattern is what we are claiming, not the specific legal culpability of any named defendant.

SYNTHETIC INSIGHTS

Intelligence, Accessible.

Synthetic Insights' mission is to build AI to serve the greater good — including job creation, productive social dialogue, and the delivery of social, economic, and spiritual value broadly to society.

R&D REPORTS
ISSUE NO. 001 · SI-RD-001
ETHICS AS INFRASTRUCTURE
V1.0 · MARCH 26, 2026

SYNTHETIC INSIGHTS LLC
BRIAN@SYNTHETIC-INSIGHTS.AI
SYNTHETIC-INSIGHTS.AI