

Human-COS Startup

Bring consequences, responsibility, and those who need to know into judgment while there is still time

Purpose in one sentence: Human-COS is human civilization's capacity for anticipatory judgment and directional immunity. STARTUP first seeks to prevent situations in which it is "too late to correct an error." It improves the conditions for judgment, being informed, and responsibility. It acquires no real-world sovereignty, automatic stopping power, approval authority, or punitive authority.

Zijun Fu | Civilization Leap Research Group

English Translation | Chinese source V1.0

The Chinese one-page brief and reading appendix are reflowed for English readability.

When to Initiate Deeper Judgment

- A single action may have major, hard-to-reverse, or cross-domain consequences
- The window for error correction or responsibility is shrinking rapidly
- Several locally reasonable choices may amplify one another
- Key specialists, affected people, and decision makers do not share the same factual picture
- Single-Event Global Impact: an event may be rare, yet one occurrence can change the whole system

When It Is Usually Unnecessary

- Ordinary everyday issues with low consequences that are easy to roll back
- Established professional processes are sufficient, with no evident cross-domain externalities
- The aim is merely to have the system make value judgments for people or score their character
- The aim is to invoke the "whole system" to bypass rights, professional responsibility, statutory authority, or genuine authorization

At Least Seven Things to Bring Together When a High Consequence Matter Enters Human-COS

What must be visible	Minimum content
1 Facts and unknowns	What is known now; how strong the evidence is; what remains unknown; which information is limited by time, permissions, or visibility.
2 Causal Map	Action → interests and stakes → changes in people's attitudes and relationships → the next round of action → multi-subject and cross-domain feedback. Do not consider only first-round gains.
3 Responsibility Map	Who should be informed; who judges; who decides or authorizes; who executes; who monitors or reopens the matter; who has real-world stop or rollback interfaces; who is responsible for outcomes and repair. Identify formal power, actual influence, and positions of responsibility separately.
4 Stakeholder Visibility	Whether the main affected parties and those who should know receive the information they need while correction is still possible. This does not mean majority voting or full public disclosure.
5 Alternative paths	Compare more than action versus inaction: continue, stop, delay, split, combine, conduct a reversible trial, or change the goal or scope of authorization.
6 Strongest objection and abandonment conditions	Preserve the strongest opposing case. Identify what new facts would show that the current judgment is wrong and must be downgraded, reopened, or abandoned.
7 Remaining window	How much time remains to change the outcome? How much time remains to fulfill responsibilities for early warning, notification, stopping, review, and repair?

The Most Important Safeguard in Startup

No Silent Crossing: high-consequence matters must not pass silently beyond a threshold without the main relevant subjects being informed, key objections being visible, and responsibilities being clearly assigned. Human-COS improves the conditions for a response, but does not itself acquire an abstract power to stop.

Who Decides

The final choice remains with real-world subjects who hold authority and responsibility and have received the necessary authorization. Human-COS changes only what can be seen before a decision and whether responsibility is clear.

How the Three Systems Divide Their Roles

Human-COS safeguards direction; AGI-COS safeguards operational safety and permissions; Nebula carries real-world action after judgment, decision, and authorization. Actual results then feed back into recalibration.

Current Status Is Not a Capability Promise

Theory: formal cross-stage baseline V1.2; startup specialist edition V1.1. Engineering: S0/S1 Public Core has established a public, reproducible baseline, including the frozen protocol and Schema, Case/Evidence, T0, Actor/Role ACL, append-only revisions, Context/source binding, and tests.

Not yet included or authorized: Responsibility Map automation, an automatic No Silent Crossing Guard, cross-platform federation, S2+, full state-machine orchestration, or automatic real-world execution.

The point is not for the system to know earlier, but for the people who need to know to know when they need to. Responsibility should become visible before action, rather than be sought after the outcome.

Civilization Leap Research Group | This page does not constitute real-world authorization or an instruction for automatic stopping, approval, or punishment.

Personalized Further Reading Guide

This is not a list of related reading. It follows the next questions raised by this document.

This page answers only three startup questions: when to enter, what to examine, and who decides. The next steps lead to the full public explanation, machine-side risk and authority tracks, and the minimum boundaries for an open future.

Toward the full explanation

[Human-COS: Seeing Earlier What Otherwise Becomes Visible Only Later](#)

The full public explanation and routes to participation.

Toward the machine side

[AGI-COS Controlled Version Family: One-page Staged Entry](#)

Choosing machine-side risk and authority tracks.

Toward the minimum boundaries

[Common Window Protocol: An Open Future and Three Constraints](#)

The minimum boundaries for an open future.

Entry Points for Challenges Counterexamples and Revision

New facts, counterexamples, failed tests, new capabilities, new attack surfaces, or changes in the conditions of a stage should all prompt a reassessment of the relevant document. Further reading must not be used to preserve existing conclusions.

[Supplementary English reading index](#)

[Main Public Entry Point](#)

Source appendix note: Body text unchanged | Library navigation appendix added only | Civilization Leap Research Group

English navigation

English links lead to corresponding webpages. Some destinations are public introductions or version-family overviews rather than complete English books. Unlinked source references do not establish the availability of a full English text.

[Existing English reading index](#)