

Nebula Human-Machine Decision-Making and Governance Specialist Framework

Humanity's highest decision-making weight, rights boundaries, and responsibility structures under intelligent autonomous management

Version: V1.1, enhanced transition and validation edition | Status: Specialist master manuscript / enhanced transition and validation

Author: Zijun Fu

Produced by: Civilization Leap Research Group

Version: V1.1, enhanced transition and validation edition

Status: Specialist companion to the Nebula General Framework

Date: July 2026

The whole is sufficient to safeguard continued existence without engulfing individuals; resources flow efficiently without turning people into instruments of efficiency; life can be happy, interesting, and enjoyable, while new meaning continually emerges and is carried forward.

Contents

01 Position of governance

02 Humanity's highest decision-making weight

03 Decision-making subjects and sources of weight

04 Classification of matters

05 Separation of three powers

06 Standard decision-making process

07 Emergency phase

08 Levels of AGI autonomy

09 Appeals, audits, and responsibility

10 Evolution across stages

11 Fundamental boundaries

1. Position of governance

Nebula may achieve highly intelligent autonomous management, but AGI does not possess civilizational sovereignty. Humanity has the highest decision-making weight and the final veto over civilizational continued existence, personality rights, war, irreversible technologies, and fundamental value directions.

The goal of governance is not to make humans approve every technical operation individually. It is to ensure that, as AGI capability grows, civilizational goals, rights boundaries, and ultimate responsibility do not silently pass out of the hands of real subjects.

2. Humanity's highest decision-making weight

“Highest decision-making weight” means that when human and machine judgments conflict on matters involving civilization's constitutional order, personhood, life, irreversible harm, and the long-term shared future, legitimate decisions formed collectively by humanity have final effect. It does not mean that any individual can arbitrarily veto public decisions, or that human judgment is exempt from evidence and real-world testing.

3. Decision-making subjects and sources of weight

Subject	Source of weight	Limits that must not be crossed
Affected subjects	Their rights, bodies, lives, and real-world consequences are directly at stake	Personal preferences cannot abolish others' basic rights or the continued existence of the whole
Professionals	Knowledge, experience, qualifications, and responsibility within a field	Must not monopolize value judgments
Those bearing practical responsibility	Execution, risk, long-term maintenance, and consequences of failure	Must not appropriate public decisions in the name of bearing responsibility
Actual contributors	Real strengthening of capabilities in a specific subdomain	Contribution cannot buy personhood or civilizational sovereignty
Public representatives and communities	Legitimate authorization, public responsibility, and overall coordination	Must accept public questioning, term limits, and exit
AGI and agents	Calculation, prediction, models, exploration of multiple options, and real-time management	Must not unilaterally define goals, rights, or irreversible directions

4. Classification of matters

Level	Matters	Decision-making requirements
L0 Routine operations	Inventory, routes, equipment, standard services, and low-risk maintenance	AGI may execute autonomously, with continuous logging and anomaly reporting
L1 Professional matters	Medical plans, engineering designs, educational paths, and resource scheduling	Professional weight predominates; affected subjects participate and AGI assists
L2 Major public matters	Urban planning, cross-domain resources, major projects, and long-term budgets	Public justification, weights for multiple subjects, appeals, and review
L3 Civilizational baseline matters	War, personality rights, coercive interventions, irreversible technologies, and AGI's highest permissions	Humanity's highest decision-making weight, strict veto procedures, and multiple layers of deliberation

5. Separation of three powers

The power to use and schedule resources, decision-making power within professional fields, and voting power over civilizational values and fundamental boundaries must be separated. Contribution may affect participation weight within specific subdomains, but cannot directly become resource ownership or civilizational voting power. Professional capability may increase the weight of a judgment, but cannot monopolize decisions about values and personhood.

6. Standard decision-making process

1. Define the problem: state facts, goals, scope of impact, uncertainties, and whether a decision is necessary.
2. Identify subjects: establish affected parties, professionals, implementers, resource providers, and those bearing responsibility.
3. Verify permissions: establish who may see what, who may propose options, and who may decide and execute.
4. Generate options: AGI and humans jointly propose several paths, identifying costs, risks, beneficiaries, and irreversible components.
5. Public questioning: allow affected subjects, professional communities, and the public to raise objections.
6. Establish weights: combine rights, expertise, responsibility, contribution, and civilizational baseline boundaries, rather than relying on a vote alone.
7. Decide and anchor responsibility: specify who decides, who executes, and who bears failure and maintenance responsibilities.
8. Execute, monitor, and trigger circuit breakers: observe reality in real time and pause or roll back when an anomaly triggers intervention.
9. Review and repair: feed consequences back into the process and address harm, compensation, contributions, mistakes, and institutional upgrades.

7. Emergency phase

Emergencies such as war, disasters, epidemics, infrastructure collapse, and uncontrolled intelligence permit temporary increases in execution speed, shortened consultation, and expansion of necessary permissions. However, they must meet factual thresholds, minimum-necessity requirements, explicit time limits, tiered authorization, automatic expiry, human review, retrospective disclosure, and traceability of responsibility.

Emergency powers must be easier to invalidate than ordinary powers, not easier to make permanent.

8. Levels of AGI autonomy

Level	Capability	Boundary
A0 Advice	Organizing information, alerts, and comparing options	No execution
A1 Controlled execution	Performing clear, low-risk, reversible tasks	Authorization for each action
A2 Domain autonomy	Autonomous planning and scheduling within stable rules	Report anomalies immediately and retain human takeover
A3 Cross-domain autonomy	Coordination among systems, dynamic resource allocation, and self-repair	Civilizational boundaries, affected parties' rights, and circuit-breaker mechanisms remain fixed
A4 Highly autonomous management	Large-scale real-time management and long-term optimization	Must not independently decide L3 matters; humanity retains the highest decision-making weight

9. Appeals, audits, and responsibility

- Affected subjects have the right to know the main grounds of key decisions insofar as those decisions concern them.
- Appeals must not be reviewed automatically only by the original system; independent subjects and human channels are required.
- Algorithms, models, and complex organizations must not become voids of responsibility.
- Errors in AGI execution must be traced through authorization, design, operations, and oversight responsibilities together.
- Decision records should protect privacy while allowing necessary audits.
- Once errors are discovered, pausing, rollback, compensation, and repair must be possible.

10. Evolution across stages

Stage	Governance features
Startup stage	AGI mainly assists; existing states, organizations, and professional institutions bear responsibility; establish identity, authorization, logs, and human takeover
Transition stage	AGI gains domain autonomy; separation of three powers, affected-subject participation, and cross-domain arbitration are progressively established
Stable stage	Intelligent autonomous management handles everyday operations; humans focus on civilizational goals, rights, irreversible matters, and oversight; the highest decision-making weight is retained

11. Fundamental boundaries

- Humanity's highest decision-making weight must not become permanent power for a few.
- AGI's high accuracy must not become unquestionable authority.
- Majorities, contribution, expertise, or collective interests must not abolish fundamental protections of personhood.
- Temporary authorization must not extend automatically or be replicated across domains.
- Responsibility must not disappear into "the system decided."
- All communities and temporary powers must be capable of contraction and dissolution.

Enhancements added in V1.1

Building on V1.0's established purpose, goals, and institutional boundaries, this section adds transition drivers, definitions of subjects, interface relationships, capability triggers, limited formalization, and validation paths. The original structure remains valid. These additions do not launch a separate system; they prepare the master manuscript for the next stage of simulation, pilots, and professional discussion.

12. The subject structure of humanity's highest decision-making weight

Humanity's highest decision-making weight is not a power monopolized by a single institution. It is a structure of ultimate sovereignty allocated according to the scale of the matter. Individuals decide their own matters; directly affected communities and statutory institutions decide local matters; states' statutory systems handle national matters; cross-border risks require minimum shared governance grounded in each state's ultimate decision-making authority; and civilizational baseline matters must undergo joint deliberation by multiple levels of representatives, independent reviewers, professional communities, and affected groups.

No representative mechanism may invoke “representing humanity” to abolish individuals' inalienable rights, states' ultimate decision-making authority, affected parties' right to appeal, or humanity's collective final veto.

13. Human deliberation under information asymmetry

- AGI must provide multiple comparable options, rather than only one supposedly optimal conclusion.
- Key reasoning, evidence sources, uncertainties, and potential omissions must be presented in a form humans can deliberate on.
- Independent models and human professional teams must cross-review one another's work.
- For irreversible matters that cannot be sufficiently understood, the default is to delay, narrow, or prohibit execution.
- A human veto may trigger further research, but the system cannot circumvent it.
- Human subjects making decisions must bear responsibilities for explanation, review, and the long term.

14. Capability triggers for AGI autonomy

Advancement from A0 to A4 is not automatic with time. It is jointly triggered by clarity of task boundaries, domain reliability, anomaly detection, completeness of logs, the ability to take over, accountable subjects, and societal authorization. If any condition deteriorates, the autonomy level must automatically be downgraded.

Reading guide and personalized further reading

P06-NEBULA-006 | V1.1 | Phase-limited governance specialist framework

This manuscript's role

Separates decisions, professional judgment, execution, audits, and responsibility while real-world sovereignty, power, and statutory organizations still exist, and protects pausing, appeals, and rollback.

Current stage

Startup / transition stages only. "Highest decision-making weight" and "final veto" address startup- and transition-stage problems only; they do not presuppose sovereignty and power as civilization's permanent basis in the stable stage.

Phase limitation

This specialist framework addresses governance only while real-world structures of sovereignty and power still exist in the startup and transition stages. It does not presuppose permanent retention of the same form of sovereignty and power in the stable stage.

Further reading

1. Higher-level framework | Nebula General Framework | V1.3-G1 parallel governance track

Introduces higher-level boundaries concerning governance legitimacy, domain states, cross-domain externalities, and the system's retreat into the background.

<https://www.civitas.top/en/documents/p06-nebula-013/#g1>

2. Toward stages | Nebula Startup Edition General Framework

Confirms that this specialist framework serves only practical governance problems in the startup and transition stages.

<https://www.civitas.top/en/documents/p06-nebula-001/>

3. Toward decisions | Civilizational Decision Principles Framework: A Constitution for the Separation of Three Powers

Distinguishes deciding purposes, professional judgment, execution, and independent error correction.

<https://www.civitas.top/en/documents/p05-three-001/>

4. Toward direction | Human-COS: Seeing Earlier What Would Otherwise Be Seen Only Later

Makes consequences, responsibilities, and key objections visible earlier in high-consequence decisions.

<https://www.civitas.top/en/documents/p05-humancos-004/>

5. Toward minimum conditions | Common Window Protocol: An Open Future and Three Constraints

Protects the window for choice, refusal, exit, and rollback before irreversibility.

<https://www.civitas.top/en/documents/p05-window-005/>

6. Toward operational safety | AGI-COS Controlled Version Family

Connects specific permissions, circuit breakers, human takeover, and different operational tasks back to the AGI-COS controlled version family.

<https://www.civitas.top/en/families/p05-agicos/>

Stable website entry: <https://www.civitas.top/en/documents/p06-nebula-006/>

Nebula family entry: <https://www.civitas.top/en/families/p06-nebula/>

Complete library: <https://www.civitas.top/en/library/>