

# Higher Semantics of Agent Swarms

## Contexts, Directed Executions, and Certified Reconfiguration

Anonymous

12 October 2026 - Version 0.1 - Draft for author review

### Abstract

An agent swarm is specified by a semantic context, typed agent families, interaction protocols, directed executions, and admissibility certificates. An infinity-topos provides dependent state families, higher identity, and effective descent; a weaker semantic base is sufficient when the precise finite-limit and descent requirements are supplied. We construct finite network and trace objects, prove coherent replacement and contextual descent, and give a compositional certificate rule for guarded asynchronous updates. Reconfiguration changes architectures through explicit transition witnesses rather than identification of state spaces. Counterexamples separate protocol compliance from state validity, coherent realization from local existence, and safety from progress. A finite coordinator-and-worker model checks exclusive leases and revision-sensitive completion. The contribution is an explicit semantic contract assembled from established categorical constructions, with its operational and representational limits stated.

### 1. Semantic contexts and specifications

Fix universes and let  $\mathcal{S}$  denote the infinity-category of spaces. All limits, mapping spaces, functors, and equivalences are homotopy coherent. An infinity-topos  $\mathcal{X}$  is an accessible left-exact localization of  $\text{Fun}(\mathcal{C}^{\text{op}}, \mathcal{S})$  for a small infinity-category  $\mathcal{C}$  [1]. Objects can retain higher identities; they are not restricted to sets of machine states.

**Definition 1.1 (Admissible semantic base).** A semantic base is an infinity-category  $\mathcal{C}$  with finite limits. For contextual descent, specify pullback-stable covers  $p : U \rightarrow T$  whose augmented Čech nerves  $U_{\bullet}$  satisfy

$$\mathcal{C}_{/T} \simeq \lim_{[k] \in \Delta} \mathcal{C}_{/U_k}. \quad (1)$$

The comparison uses pullback functors. A descent-compatible semantic base has these specified equivalences; not every finitely complete category does. Infinity-topoi satisfy (1) for effective epimorphisms by topos descent [1, Section 6.1]. Elementary higher topoi provide another setting with appropriate descent and dependent-type structure [2]. Only covers and constructions actually supported by the chosen base are used.

**Definition 1.2 (Predicate and realization).** A predicate on  $Z$  is a monomorphism  $m : V \hookrightarrow Z$ , equivalently a  $(-1)$ -truncated map. A certificate for  $z : T \rightarrow Z$  is a lift  $T \rightarrow V$  over  $Z$ ; its space is empty or contractible. A realization family  $R \rightarrow Z$  may instead have multiple inequivalent implementations or proof witnesses. A local section of that family is not automatically a global section.

Finite limits interpret conjunction and equality witnesses. Dependent products, universes, geometric realization, probability, and resource orders are additional structure when used. An infinity-topos supports the categorical constructions below; it does not by itself supply a scheduler, numerical model, communication channel, or correct implementation.

## 2. Agent families and semantic worlds

**Definition 2.1 (Swarm specification).** A swarm specification consists of a semantic base  $\mathcal{X}$ ; an architecture object  $A$ ; a role family  $r : R \rightarrow A$ ; a local state family  $S \rightarrow R$ ; a complete configuration object  $z : Z \rightarrow A$ ; typed ports and protocol data; a transition correspondence

$$Z \xleftarrow{s} E \xrightarrow{t} Z; \quad (2)$$

and a predicate  $V \hookrightarrow Z$ . Event labels, request identities, scheduler choices, queue contents, revisions, and environmental responses are retained in  $Z$  or  $E$  whenever they affect the specified behavior. A witness in  $E$  is a permitted directed step, not an identity path in  $Z$ .

When dependent products exist, a state component can be constructed as  $\Pi_r S \rightarrow A$ ; shared buffers and resources supply further factors. A fixed finite roster uses finite products. Transport of the dependent-product construction requires more than preservation of finite limits. Alternatively  $Z \rightarrow A$  is supplied directly with its interpretation. No equivalence between these alternatives is assumed.

**Definition 2.2 (Heterogeneous interpretation).** If agent  $i$  has its own semantic base  $\mathcal{C}_i$ , specify a common comparison base  $\mathcal{X}$  and a finite-limit-preserving functor  $F_i : \mathcal{C}_i \rightarrow \mathcal{X}$ , together with interpreted port maps and protocol agreements. For infinity-topoi such a functor can be the inverse image  $h_i^*$  of a geometric morphism  $h_i : \mathcal{X} \rightarrow \mathcal{C}_i$ . The direction matters: the inverse image carries the agent's specification into the comparison world. Arbitrary semantic worlds are not implicitly interchangeable.

Interactions occur after the required objects and maps have been interpreted in that common base. Modalities, resource annotations, or probabilistic enrichments must also be transported by specified structure maps if a claim involves them. Finite-limit preservation alone does not preserve every enrichment.

**Remark 2.3 (Higher and directed structure).** An operational Segal object  $C_\bullet : \Delta^{\text{op}} \rightarrow \mathcal{X}$  can additionally encode coherent composition, with  $C_0 = Z$  and  $C_n \simeq C_1 \times_Z \cdots \times_Z C_1$ . Primitive events map to its arrows. This structure is extra data [3]. A graph (2) already defines finite event paths, but is not a Segal object until identities, composition, and their coherence have been provided. Groupoid completion of operations would identify irreversible computation with reversible transport and generally changes the semantics.

### 3. Interaction networks as finite limits

For fixed finite agents  $i \in I$ , let  $B_i$  be local execution witnesses. They can describe one step or a fixed finite horizon; mixing these interpretations requires an explicit timing adapter. For wires  $e \in J$ , choose signal types  $P_e$  and component endpoint maps. Write  $C = \prod_i B_i$ ,  $Q = \prod_e P_e$ , and let  $o, u : C \rightarrow Q$  collect the two ends of each wire.

**Definition 3.1 (Wired behavior).** The execution object of the network is

$$N = C \times_{Q \times Q} Q, \quad C \xrightarrow{(o,u)} Q \times Q, \quad Q \xrightarrow{\Delta} Q \times Q. \quad (3)$$

A generalized element of  $N$  includes component executions and compatibility paths on every wire. Protocol constraints  $K \hookrightarrow C$  are imposed by  $N_K = N \times_C K$ . Shared-resource constraints belong here unless already present in the component contracts. External boundary maps turn  $N_K$  into an open behavior.

Retain source and target readouts in the same diagram. A supplied complete-configuration interpretation  $N_K \rightarrow Z \times Z$  then gives (2). Without this interpretation, the network is an open behavior, not yet a global state-transition model.

**Theorem 3.2 (Composition and coherent replacement).** Finite network constraints may be imposed in any order, with coherently equivalent results. A coherent equivalence of the complete component, port, and protocol diagram induces an equivalence of network execution objects, over fixed external boundaries.

**Proof.** For every test object  $T$ , maps  $T \rightarrow N_K$  are precisely coherent cones satisfying all displayed component, wire, and protocol conditions. Successive pullbacks impose the same cone data regardless of order. The mapping-space universal property identifies each construction with the same finite limit; its space of limit cones is contractible. An equivalence in the relevant diagram infinity-category is preserved by the limit functor. Keeping the boundary cone in that diagram gives the equivalence over the external ports.  $\square$

This is a finite-limit interpretation of established compositional systems methods [4,5]. It proves compatibility and invariance of the specified execution object, not nonemptiness, efficient execution, or physical realizability.

### 4. Finite traces and contextual descent

**Definition 4.1 (Finite event path).** For (2), define  $P_0 = Z$  and, for  $n \geq 1$ ,

$$P_n = \underbrace{E \times_Z \cdots \times_Z E}_{n \text{ factors}}, \quad (4)$$

where the target of one step is identified with the source of the next. Endpoints and labels are retained. Homotopy pullbacks include compatibility witnesses, not just equal endpoint values. There is no claim that the  $P_n$  alone form a simplicial object: an inner face would require composition of primitive events.

**Proposition 4.2 (Temporal concatenation).** For finite  $m, n$ ,

$$P_{m+n} \simeq P_m \times_Z P_n, \quad (5)$$

using final and initial endpoints. The identifications are coherent under further subdivision.

**Proof.** A cone specifying  $m + n$  composable events is the same as two cones for the first  $m$  and last  $n$  events, with identification of their shared state. This proves (5) on mapping spaces. Every parenthesization is a limit of the same path-incidence diagram, giving coherence. The case of a zero-length path uses the diagonal state interface.  $\square$

Temporal subdivision retains every transition edge. Agreement only at a set of observed vertices cannot reconstruct a missing event. This is the discrete finite-path analogue of temporal sheaf composition [4]; it introduces no continuity or real-time bound.

**Theorem 4.3 (Descent of swarm structure and traces).** Let (1) hold. A coherent Čech descent datum for a swarm diagram in the slices  $\mathcal{X}_{/U_k}$  descends to a diagram in  $\mathcal{X}_{/T}$ . Its monomorphisms and finite-limit constraints descend. For a global diagram  $D$  and every finite network or path construction  $L$ ,

$$p^*L(D) \simeq L(p^*D). \quad (6)$$

In an infinity-topos its resulting object  $W \rightarrow T$  also satisfies

$$W \simeq \operatorname{colim}_{[k] \in \Delta^{\text{op}}} (W \times_T U_k). \quad (7)$$

**Proof.** Apply  $\operatorname{Fun}(K, -)$  to (1), where  $K$  indexes the specified swarm diagram. Functor categories preserve limits, so objects, maps, and higher coherence descend together. Being a limiting cone is invariant under equivalence of infinity-categories. Monomorphy is characterized by an equivalence of the diagonal with its self-pullback, hence is also detected by descent. Pullback between slices preserves finite limits, which gives (6). For an infinity-topos  $|U_\bullet| \simeq T$  and colimits are universal; pulling that colimit back along  $W \rightarrow T$  gives (7). The space of a descended diagram with a specified identification to the datum is contractible, not its unrestricted automorphism space.  $\square$

Equation (7) is contextual Čech reconstruction. It is not a limit of local objects in the ambient category, nor a statement about all hypercovers without further hypotheses. Pairwise agreements without the required higher coherence are not the datum in this theorem.

## 5. Certificates for asynchronous updates

**Definition 5.1 (Certified correspondence).** With  $E|_V = E \times_{s,Z} V$ , a safety certificate is a map  $q : E|_V \rightarrow V$  and a specified identification

$$mq \simeq t \operatorname{pr}_E. \quad (8)$$

It asserts that every specified transition starting in  $V$  ends in  $V$ . Restricting  $E$  by a guard changes the allowed protocol and must be recorded as such.

**Theorem 5.2 (Safety of finite executions).** If (8) holds, paths whose initial state is in  $V$  have every state in  $V$ . Moreover

$$E \times_{Z \times Z} (V \times V) \simeq E|_V. \quad (9)$$

**Proof.** The source lift and  $q$  provide both endpoint lifts. Their spaces are propositions because  $m$  is monic, so forgetting the target lift has contractible fibers; this proves (9). Apply the same construction successively along the path-incidence diagram. Induction on the finite number of steps supplies all intermediate lifts, with the compatibility dictated by monomorphy.  $\square$

**Theorem 5.3 (Compositional guarded updates).** Suppose  $V = \bigcap_{i=1}^k V_i \hookrightarrow Z$  is a finite intersection. A guarded update  $u_j : G_j \rightarrow Z$ , with  $G_j \hookrightarrow Z$ , preserves  $V$  if, for every  $i$ , there is a lift

$$G_j \times_Z V \longrightarrow V_i \quad \text{of } u_j \text{ pr}_{G_j}. \quad (10)$$

Certificates  $G_j \times_Z V_i \rightarrow V_i$  suffice by restriction. A coherent frame equality supplies a lift when  $V_i$  is pulled back from an unchanged component readout. If event correspondences all admit these lifts, their coproduct correspondence preserves  $V$  whenever that coproduct exists and is stable under pullback.

**Proof.** The lifts in (10) have the same underlying target map to  $Z$ . The universal property of the intersection combines them into a lift to  $V$ , giving (8) for the guarded update span. Pullback-stable coproducts identify the restriction of a coproduct of events with the coproduct of their restrictions; the component lifts assemble there. Theorem 5.2 supplies finite-path preservation.  $\square$

Thus a local component proof needs certificates or frames for all affected global constraints. Two workers independently testing an available resource do not satisfy this rule for simultaneous allocation. Serializable commits, explicit ownership tokens, or a separate atomic arbitration relation are possible protocol choices, not consequences of (10).

**Proposition 5.4 (Locality of predicate satisfaction).** In an infinity-topos, if a generalized state satisfies  $V$  after an effective epimorphic cover of its domain, it satisfies  $V$  globally. For a fixed global transition correspondence, coherent local certificates of (8) descend; since their lift spaces are propositions, their existing overlap compatibility is unique.

**Proof.** For  $x : T \rightarrow Z$ , form  $W = T \times_Z V \hookrightarrow T$ . A local lift factors the effective epimorphism through  $W$ , so its image is all of  $T$  and  $W \simeq T$ . For (8), the desired lift is a section of the monomorphism  $(E|_V) \times_Z V \rightarrow E|_V$  along the target map. Apply the same argument, or descent (1), on a covering pullback. Monomorphy gives uniqueness and all coherence.  $\square$

## 6. Reconfiguration and semantic transport

A transition in (2) need not lie over the diagonal of  $A \times A$ . Its source and target architectures can differ. The complete state includes the role assignment and all outstanding obligations; retiring an agent, changing a model, and migrating a queue are explicitly typed transitions.

**Definition 6.1 (Certified reconfiguration).** A reconfiguration correspondence  $H \rightarrow Z \times Z$  is admissible when its restriction to valid sources has a valid-target lift as in (8). Migration witnesses record the treatment of pending requests, shared resources, and externally visible behavior. If behavior preservation is claimed, supply a map or equivalence of the relevant execution diagrams as well. Admissible migration is weaker than behavioral equivalence.

**Theorem 6.2 (Mixed execution safety).** Interleaving certified agent, communication, scheduler, and reconfiguration steps preserves  $V$  for every finite path, provided their combined correspondence satisfies the coproduct condition of Theorem 5.3.

**Proof.** Assemble the certificates by Theorem 5.3 and apply Theorem 5.2. Architecture changes are already coordinates of the complete state, so the induction does not require a constant roster.  $\square$

**Theorem 6.3 (Finite semantic transport).** A finite-limit-preserving functor  $F : \mathcal{X} \rightarrow \mathcal{Y}$  transports a supplied swarm diagram, its monic predicates, its certified correspondences, and every fixed finite network and path construction. In particular

$$F(N_K) \simeq N_{F(K)}, \quad F(P_n) \simeq P_n(FE), \quad (11)$$

with all component and endpoint objects interpreted by  $F$ . A supplied Segal object remains Segal under degree-wise application of  $F$ .

**Proof.** Each displayed construction uses finite limits. Monomorphy is the diagonal pullback condition, which  $F$  preserves. Apply  $F$  to the certificate square (8). The Segal conditions are finite pullback equivalences in each degree. No limit over all degrees is taken.  $\square$

Geometric inverse images supply such functors. The theorem interprets an already supplied  $Z$ ; it does not assert  $F(\Pi_r S) \simeq \Pi_{Fr} F(S)$ . Preservation of a combined coproduct, an unbounded population, or an infinite trace limit needs the corresponding additional comparison equivalence. Likewise descent-compatible interpretation requires preservation of the designated covers and their descent structure.

## 7. Obstructions and countermodels

**Proposition 7.1 (Local realizability does not select a global agent).** In  $\mathcal{X} = \text{Fun}(BC_2, \mathcal{S})$ , let  $L = \{a, b\}$  carry the swapping action. The map  $L \rightarrow 1$  is an effective epimorphism, and its pullback along itself has a diagonal section. Yet it has no global section.

**Proof.** It is objectwise surjective on components and hence effective epimorphic. A global section is a homotopy fixed point of the discrete swapping action. Such a point would be fixed by the swap, and neither point is. The local diagonal section therefore has no coherent descent to a global section.  $\square$

This models symmetric leader selection without symmetry-breaking data. It is distinct from Proposition 5.4:  $L \rightarrow 1$  is a realization family, not a monomorphism.

**Proposition 7.2 (Observations lose higher execution structure).** Let  $Z = 1$  and compare  $E = BC_2$  with  $E' = 1$ , both with constant endpoints. Their endpoint image relations agree, but their one-step execution objects are inequivalent.

**Proof.** Both image relations are  $1 \hookrightarrow 1 \times 1$ . However  $\pi_1(BC_2) = C_2$  and  $\pi_1(1) = 0$ . Thus equal observed endpoints do not satisfy Theorem 3.2’s diagram-equivalence hypothesis.  $\square$

The distinction extends to every higher degree: for a nonzero abelian group  $H$  and  $d \geq 2$ , choose  $E = K(H, d)$  with constant endpoints. Then  $P_n = K(H, d)^n$  and  $\pi_d(P_n) = H^n$  for  $n \geq 1$ . Any fixed finite truncation discards some such execution objects. The swarm specification therefore declares which higher witnesses it retains.

**Proposition 7.3 (Truncation need not commute with wiring).** In spaces,  $\tau_{\leq 0}$  does not preserve all finite homotopy pullbacks.

**Proof.** For nontrivial discrete  $G$ ,  $1 \times_{BG} 1 \simeq G$ . Its zero-truncation is  $G$ , whereas the pullback of the three zero-truncated objects is 1. Their cardinalities differ.  $\square$

Consequently retaining only sets of observation classes can alter a swarm’s compatibility witnesses. Higher semantics must specify which truncation is adequate and prove the needed comparison laws.

**Countermodel 7.4 (Well-typed deadlock).** A synchronous Boolean feedback contract  $x = 1 - x$  has an empty execution set. Typing every port does not make the network object inhabited. A progress theorem needs a realization, enabledness, and scheduling assumptions in addition to finite-limit composition.

**Countermodel 7.5 (Safety without progress).** A two-state system has an idle self-loop at  $w$  and a completion step  $w \rightarrow d$ . The predicate  $V = Z$  is preserved, but the infinite idle schedule never completes. Fairness or a well-founded rank is extra structure. Existence of a completing path is different from completion under every allowed scheduler.

**Countermodel 7.6 (Left exactness has a finite scope).** Let  $K = \{0\} \cup \{1/n : n \geq 1\}$  with its subspace topology and let  $B$  be the sheaf of locally constant Boolean functions. Define a section  $b_n$  that is zero only at  $1/n$ , and let  $c_n$  be constantly one. The two sections  $(b_n)$  and  $(c_n)$  of  $\prod_n B$  have different germs at zero: every neighborhood contains some  $1/n$ . Each individual pair  $b_n, c_n$  has the same germ there. The stalk functor therefore does not preserve this countable product, although it is a geometric inverse image and preserves finite limits. This prevents extending (11) to arbitrary infinite agent families by left exactness alone.

## 8. Finite coordinator-and-worker instance

The companion uses sets for a three-agent swarm: a coordinator and workers  $a, b$ . A configuration is  $(e, A, o, r)$ , where  $e \in \{0, 1, 2\}$  is the architecture revision,  $A \subseteq \{a, b\}$  is the active roster,  $o \in \{\emptyset, a, b\}$  is the exclusive resource owner, and  $r$  is the lease revision. The predicate is

$$o = \emptyset \Rightarrow r = -1, \quad o \neq \emptyset \Rightarrow (o \in A \wedge r = e). \quad (12)$$

The owner type enforces exclusivity. There are 24 valid states. The 17 attempted events are activation, retirement, and acquisition for either worker; completion for either worker with a reported revision in  $\{0, 1, 2, 3\}$ ; and reconfiguration to a revision in  $\{0, 1, 2\}$ .

Acquisition requires an active worker and a free resource. Completion requires matching ownership and revision. Retirement requires that the worker hold no lease. Reconfiguration requires a free resource and a strictly larger revision. Failed events retain the state and a rejection observation. These guards supply certificates for (12); the completion guard also specifies protocol compliance beyond that invariant.

The executable checker independently enumerates legal pre/post conditions and compares them with an imperative interpreter for all 408 state/event pairs: 120 accepted and 288 rejected. It checks 408 complete role-renaming comparisons, 125,280 initial-state/event-word pairs through length three, and 493,296 temporal subdivision cases. Every resulting state is valid. Negative controls detect simultaneous double allocation, retirement of an owner, revision changes with a live lease, and acceptance of a stale completion. The last defect can leave (12) true while violating the specified observation relation. Results and bounds are recorded separately; these finite checks do not formalize the general infinity-categorical proofs.

This instance interprets categorical structure in a finite ordinary model. It is not an implementation of an arbitrary infinity-topos, a distributed consensus protocol, or a verified deployment of language-model agents.

## 9. Attribution, scope, and continuation

The foundational ingredients are standard descent, finite-limit composition, dependent families, and Segal objects [1-5]. The proposed contribution is their assembly into an agent-swarm contract that explicitly retains scheduler effects, cross-agent constraints, higher execution witnesses, and architecture-changing certificates. The preceding higher-lens framework [6] supplies a compatible interface formulation; the present treatment makes directed correspondences and heterogeneous interpretations explicit.

Probabilistic kernels, learning objectives, physical resource accounting, fairness, faults, and unbounded populations need additional supplied structures and proofs. A semantics-preserving compiler must relate concrete state, observations, and effects to this specification; abstract existence is not that relation. The written proofs have not been checked by a proof assistant. The draft remains open to author and independent mathematical review.

**Automation disclosure.** OpenAI ChatGPT/Codex assisted with the framework, written proofs, reference checking, finite checker, and typesetting. The Anonymous byline is provisional for author review. No journal submission or independent review is asserted.

## References

1. Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009. Definition 6.1.0.4, Theorem 6.1.0.6, and Section 6.1.3. [Author’s electronic edition](#).
2. Nima Rasekh. *A Theory of Elementary Higher Toposes*. 2018. [arXiv:1805.03805](#).
3. Charles Rezk. *A model for the homotopy theory of homotopy theory*. Transactions of the American Mathematical Society 353 (2001), 973-1007. [arXiv:math/9811037](#).
4. Patrick Schultz, David I. Spivak, Christina Vasilakopoulou. *Dynamical Systems and Sheaves*. [arXiv:1609.08086](#).
5. Georgios Bakirtzis, Christina Vasilakopoulou, Cody H. Fleming. *Compositional Cyber-Physical Systems Modeling*. EPTCS 333 (2021), 125-138. [arXiv:2101.10484](#).
6. Anonymous. *Higher Lenses and Temporal Descent for Certified Adaptive Systems*. Higher Categorical Semantics, 2026. [Published note](#). The historical version 1.1 supplies the interface precedent; the present draft is self-contained.