

Four maps, one structure

Functor Coherence

ZP Companion | Version 1.11 | May 2026


This companion explains the ideas in plain language with diagrams and real-world examples. It is not the formal document — every claim here restates a result already proved in the corresponding technical document. Consult that document for the authoritative mathematics. This document assumes familiarity with the ZP-G Illustrated Companion.

What Is ZP-H Doing?

ZP-G built an abstract category C and proved that its structure — an initial object with no return morphisms — captures the essential shape of the Zero Paradox. But an abstract category on its own is just a skeleton. ZP-H puts flesh on the bones.

ZP-H constructs four explicit functors — structure-preserving maps — from the abstract category C into the four concrete mathematical frameworks of the Zero Paradox: the lattice algebra of ZP-A, the p-adic topology of ZP-B, the information theory of ZP-C, and the Hilbert space of ZP-D.

The result is a verification that all four frameworks are, in a precise sense, realizations of the same abstract structure. They are not four separate arguments for the same conclusion — they are four different windows looking at one thing. This agreement is coherence: all four frameworks are built on the same structural axioms (A1-A4, AX-B1) and the structural identification CC-1 ($S_0 = \perp$, derived in ZP-J), so the Binary Snap appearing in all of them reflects a shared foundation, not independent confirmation from unrelated starting points.



Four functors carry the abstract structure of C into four concrete mathematical frameworks

ZP-H constructs four functors from the abstract category C (center) into the four domain frameworks. Each functor F carries objects and morphisms faithfully, preserving the initial object and the forward-only structure.

The Structural Floor

Before the categorical machinery, it helps to understand one key property that \perp has in every ZP framework: the bottom element is not a limit point of the elements above it. There is a gap — a structural floor — that nothing above \perp can close.

The simplest example of this property is one most readers have already encountered: the collection of all subsets of a set. Take any set S — say $\{a, b, c\}$ — and collect every possible subset: $\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}$, and so on up to S itself. Order them by inclusion: $A \leq B$ means A is contained in B . The empty set \emptyset sits at the bottom, below everything.

Now ask: can a sequence of nonempty subsets get "closer and closer" to \emptyset ? No. Every nonempty subset has at least one element — its size is at least 1. There is no subset with size between 0 and 1, because you cannot have half an element. The gap between \emptyset and every nonempty subset is exactly 1, a discrete, combinatorial gap with no room to subdivide. In this discrete setting, no element is a limit point of elements above it — there is no way to subdivide the gap of one element. \emptyset has this property in its most extreme form: it sits below the entire lattice, so the minimum gap is between nothing and something. No sequence of nonempty sets closes that gap.

The same property appears in the ZP framework, but for deeper reasons. In ZP-B, the 2-adic valuation v_2 assigns 0 the value $+\infty$: $v_2(0) = +\infty$, meaning 0 is divisible by every power of 2. Every nonzero element has a finite v_2 . No sequence of nonzero elements can close this valuation gap: a nonzero element always has finite v_2 , while 0 has infinite v_2 , and no finite accumulation of finite values reaches infinity. In \mathbb{Q}_2 , sequences can get metrically close to 0, but the valuation gap cannot be closed. In ZP-A, \perp is the bottom of the lattice and the axioms forbid any non-trivial return. In ZP-C, the informational cost of approaching \perp diverges — infinite cost means no finite path reaches it. In ZP-D, the basis vector e_0 is orthogonal to every nonzero state — a right-angle separation that cannot be gradually closed.

The power set example shows that the structural floor property is not exotic: it appears in the most elementary object in set theory. What is non-trivial about the ZP framework is that the same property appears in four analytic settings — topology, algebra, information theory, Hilbert space — each with its own structural reason for why the bottom cannot be approached. These settings are not independent: they share the axioms (A1-A4, AX-B1) and the structural identification CC-1 ($S_0 = \perp$, derived in ZP-J) that produce this behavior in all four. ZP-H verifies that these four reasons are consistent.

What Are the Morphisms of C ?

ZP-G defined C abstractly — objects, morphisms, composition, identity — without specifying what the morphisms actually are. That abstraction was intentional: the theorems in ZP-G hold for any category with AX-G1 and AX-G2, regardless of the morphism content.

ZP-H fills in the detail (Definition D-H1): morphisms in C are state transitions. A morphism $f: A \rightarrow B$ represents a net addition of informational content — you can only move forward. Morphisms compose by sequential accumulation: going from A to B then B to C produces a state at least as large as B . This design

choice is what makes the four functors well-defined.

D-H1 is a design commitment, not a derived result. A different choice of morphism structure would produce different functors. ZP-H states this explicitly: the choice is the one that maps correctly to all four domain frameworks, and it is not laundered as a derivation.

The Four Functors

Each functor takes the abstract objects and morphisms of C and maps them into one of the four frameworks, verifying that the mapping preserves composition, identity, and the privileged role of the initial object.

FA (Lattice): The initial object 0 maps to \perp (bottom of the lattice). Each morphism $f: A \rightarrow B$ maps to the join operation $\perp \vee S = S$ that witnesses the transition. Composition corresponds to iterated joins. The forward-only structure of C maps to the monotone structure of ZP-A. FA is fully verified in Lean 4, sorry-free.

FB (p-Adic Topology): The initial object 0 maps to the element $0 \in \mathbb{Q}_2$. Each morphism maps to a discrete jump across a clopen boundary — formalized as antitone depth in \mathbb{Q}_2 BallDepth. Composition corresponds to sequential jumps. The irreversibility of ZP-B C3 (no path returns to 0 in \mathbb{Q}_2) is the topological realization of AX-G2. FB is a full Lean 4 functor (fb_functor, sorry-free) — not a proxy witness.

FC (Information Theory): The initial object 0 maps to the zero distribution $P = (1, 0)$. Each morphism maps to an informational transition with a non-negative cost measured in bits. The fundamental transition costs exactly 1 bit (ZP-C T1b). The informational singularity of ZP-G maps to the diverging surprisal of ZP-C T2. FC has a concrete ZPCategory categorical witness (NNRealZPCat, $\mathbb{R} \geq 0$ with \leq) grounded by T1b. The full abstract Lean functor for the information space codomain remains future work.

FD (Hilbert Space): The initial object 0 maps to the basis vector e_0 . Each morphism maps to an orthogonal extension — a step to a perpendicular basis vector. The Binary Snap becomes a right-angle turn in state space. The design commitment DP-1 (orthogonality represents clopen separation) is inherited here. FD has a concrete ZPCategory categorical witness (NNRealZPCat) grounded by T4. The full abstract Lean functor for the Hilbert space codomain remains future work.

Real-world analogy — Four instruments, one melody

Imagine the same musical phrase played on four different instruments: violin, trumpet, piano, and voice. Each sounds different, but any musician can hear that they are playing the same structure — the same intervals, the same rhythm. ZP-H shows that ZP-A through ZP-D are doing the same thing. They are four instruments playing one mathematical structure.

T-H1: Each Functor Preserves the Initial Object

For each of the four functors, the image of 0 is an initial object in the target framework — a universal source from which every other object has a unique morphism. The privileged status of 0 is not an artifact of the abstract category: it is preserved faithfully in every concrete realization.

Two Descriptions of One Obstruction

ZP-G said: there is no morphism from any non-initial object back to 0 (AX-G2). ZP-C said: the informational cost of approaching 0 is unbounded — the accumulated surprisal grows without bound. These look like two different statements about the same thing. Are they compatible?

T-H2 proves they are. The two characterizations are the same obstruction seen from different vantage points. In ZP-C, ask: what happens to surprisal as you approach 0 along an infinite path? The answer is divergence. In ZP-G, ask: can you reach 0 by a morphism from outside? The answer is no. If infinite cost is required to reach 0, then no finite morphism can accomplish it — which is exactly what "no morphism exists" means. Divergence and non-existence are two faces of the same fact.

T-H2: Singularity Compatibility


The categorical singularity (no return morphism to 0) and the information-theoretic singularity (diverging surprisal approaching 0) are two descriptions of the same structural fact. Under the functor FC, they correspond precisely.

The Binary Snap Under All Four Functors

The Binary Snap is the transition from 0 to the first non-initial object. In the abstract category C, it is the unique morphism from 0 to ε_0 — guaranteed to exist by AX-G1 and to be unique by the definition of an initial object. In ZP-E, this transition was derived as a theorem (T-SNAP), not assumed as an axiom.

T-H3 shows that all four functors agree on what the Snap is:

- In lattice algebra: the first join — $\perp \vee \varepsilon_0$ — producing the first non-bottom state.
- In p-adic topology: a discrete jump from 0 to ε_0 across a clopen boundary, irreversible by the topological structure of \mathbb{Q}_2 .
- In information theory: an informational transition costing exactly 1 bit — the minimum possible information cost for any state change.
- In Hilbert space: an orthogonal shift from basis vector e_0 to e_1 , with inner product $\langle e_0, e_1 \rangle = 0$.



All four frameworks describe the Binary Snap (T-SNAP) in their own language. The vertical line marks the moment of the Snap. Left: \perp . Right: the minimum nonzero state.

The agreement across four frameworks reflects coherence, not independent confirmation. All four share the same axioms — A1-A4 (lattice axioms), AX-B1 (binary existence) — and the structural identification CC-1 ($S_0 = \perp$, derived in ZP-J). The Binary Snap appears in all four because those foundations are built into each framework, not because four separate arguments from unrelated starting points happened to agree.

T-H3: The Binary Snap Under All Four Functors

The four functor images of the Binary Snap are mutually consistent: each framework establishes that the transition is irreversible, costs something (informational work, topological separation, orthogonal displacement), and is the minimal first step. This agreement is coherence across shared structural commitments (A1-A4, AX-B1, CC-1), not independent replication. T-SNAP is a derived theorem inherited here from ZP-E. The only additional design premise is DP-1 (ZP-D). CC-1 ($S_0 = \perp$) is derived in ZP-J, not a free commitment.

What this means: The Binary Snap is not a construct of any one framework. It is a structural fact that survives translation into four distinct mathematical languages. ZP-H is the document that verifies this translation is faithful.