Watching the Graph Grow.

30 frames in the operational lifecycle of a 506-node intent graph.

TOBIAS MARTENS · FOUNDER & CEO

Day 0. Empty graph.

506 latent intent nodes across 22 sub-groups. No traffic. No edges. No data.

TOTAL NODES

506

Every dot is an intent waiting to be activated. Every cell is a sub-group of human need.

A query enters the system.

First user types a 'who else?' question into the router.

QUERIES

1

USER QUERY

“who else delivers food past midnight?”

Intent classifier resolves to canonical node.

FOOD DELIVERY activates. The first node lights up.

ACTIVE NODES

1

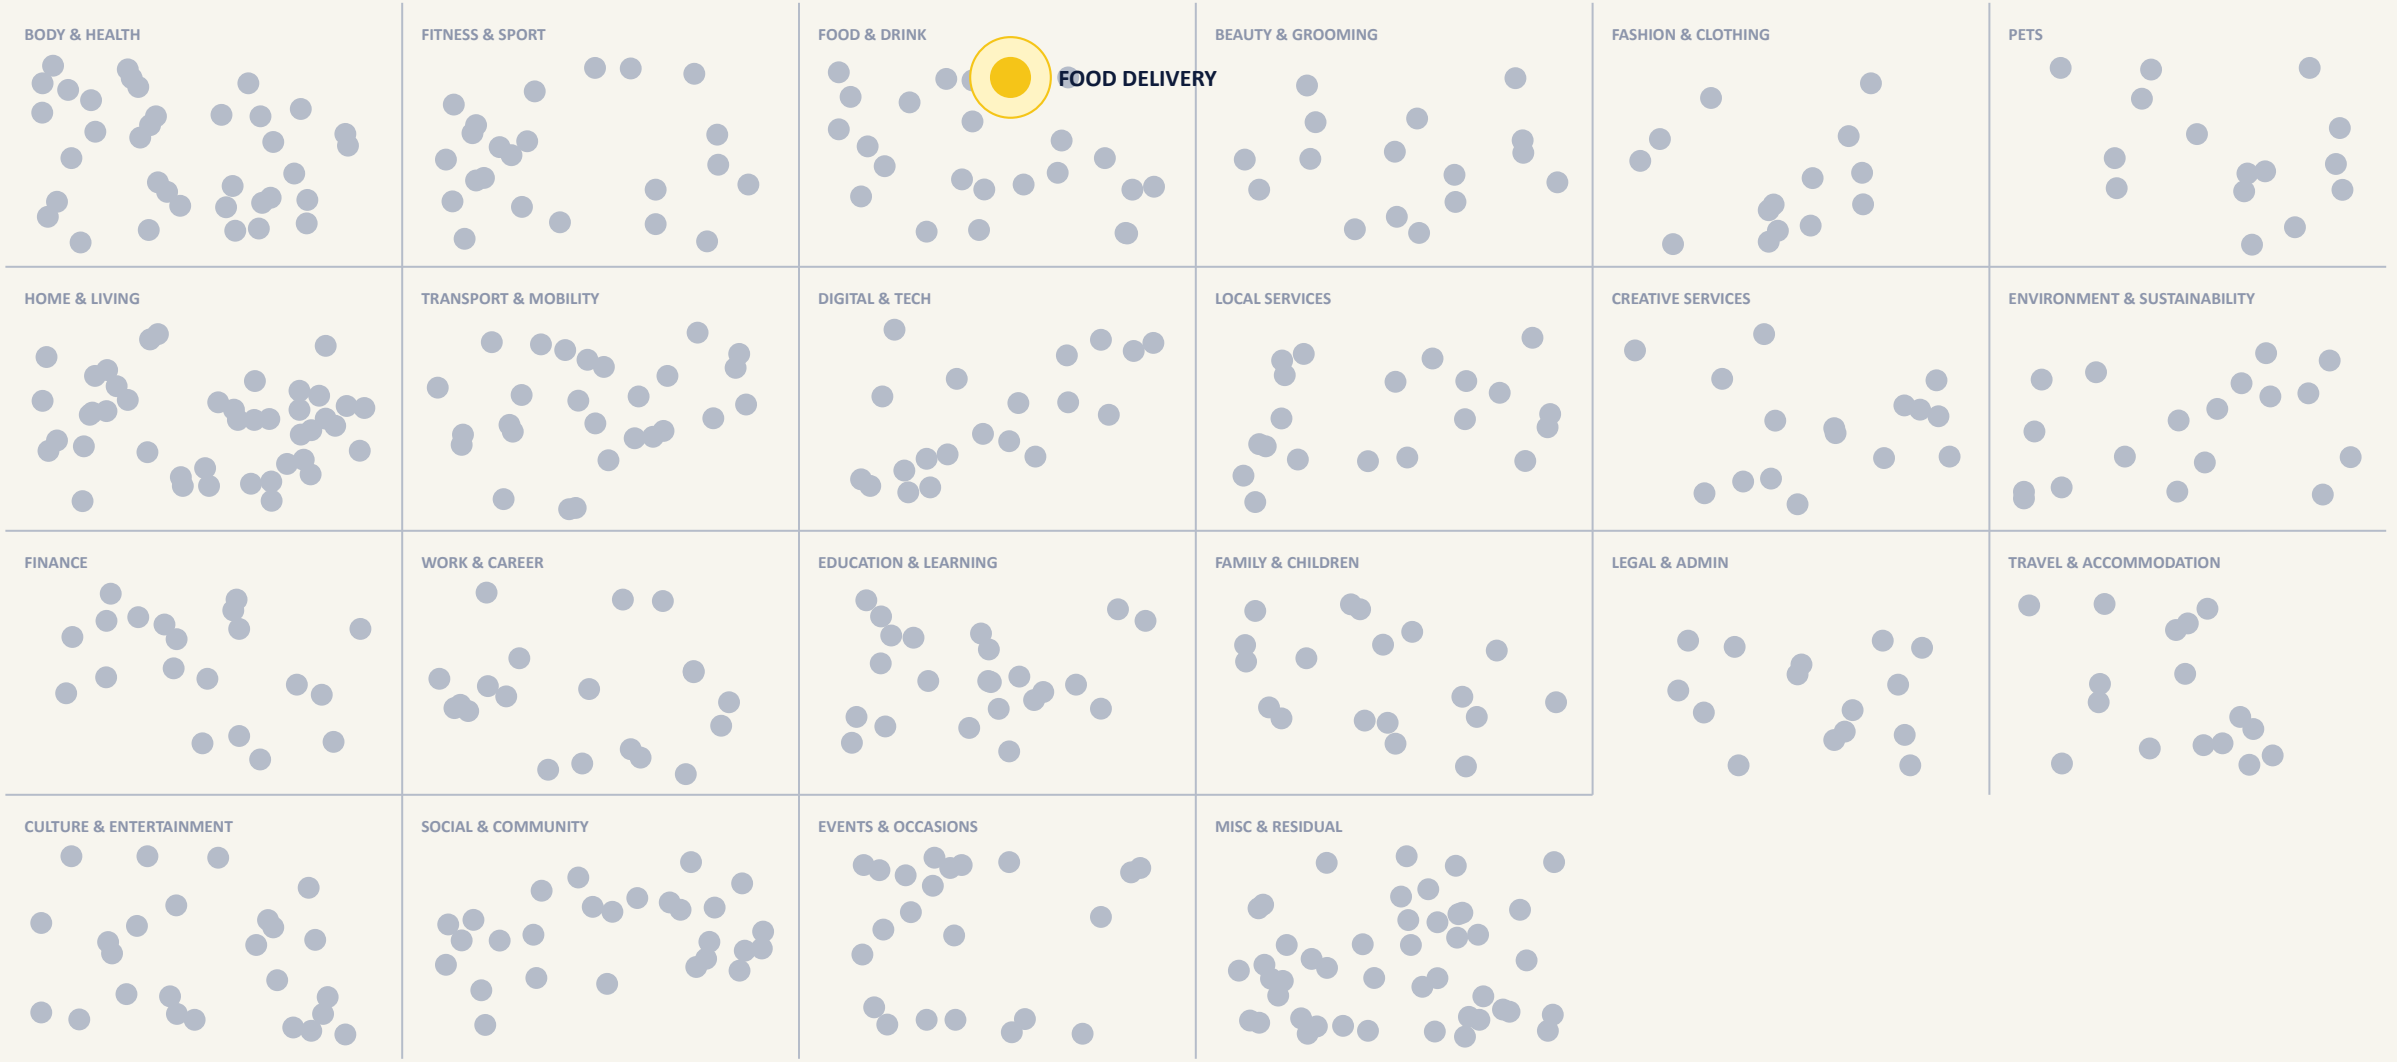

From this point on, every routing event sharpens what the graph knows.

First provider attaches.

FOOD DELIVERY → DoorDash. Edge weight = 1.

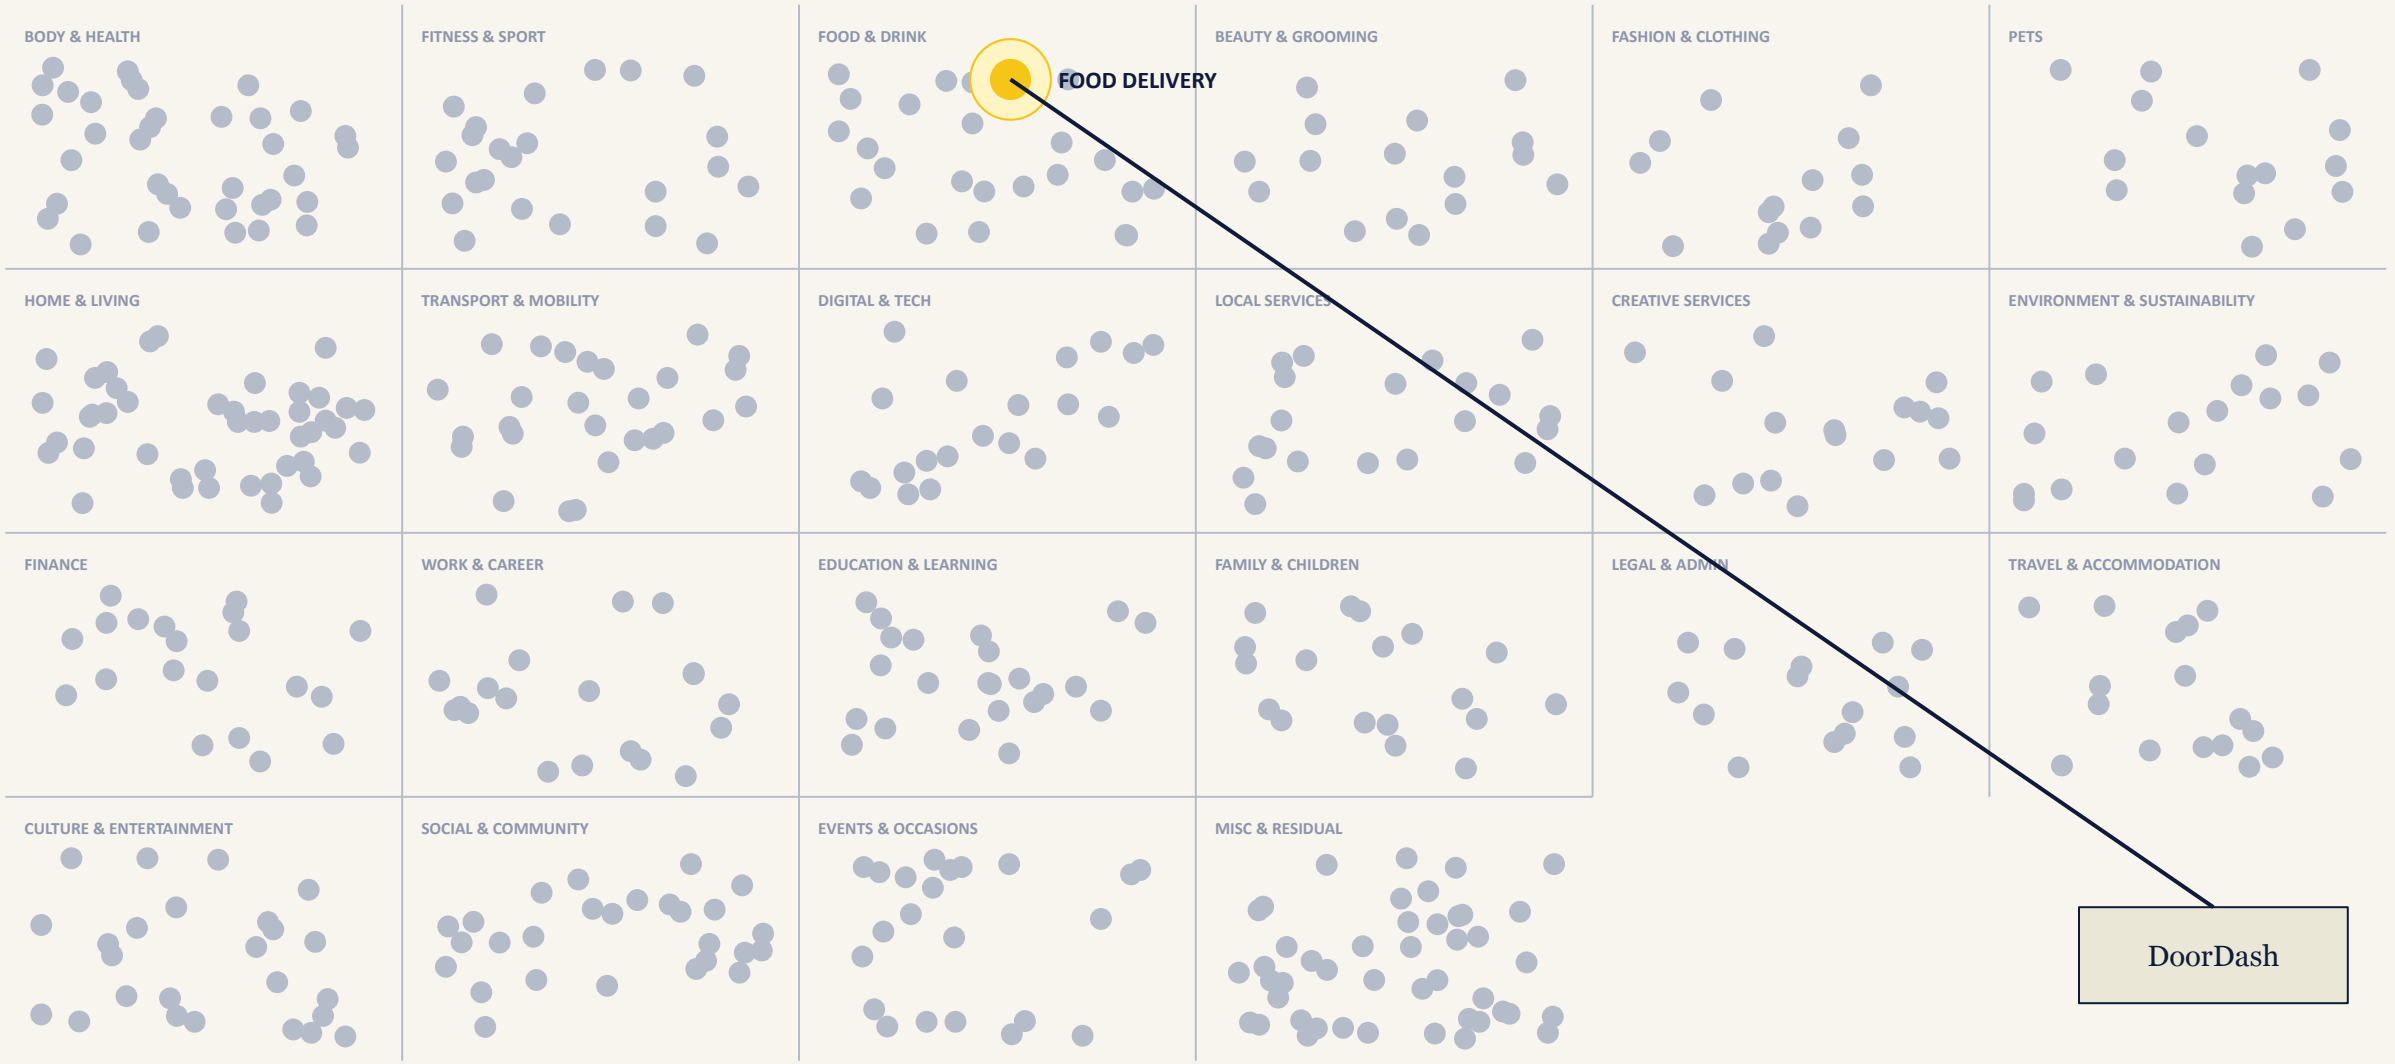

The edge is the unit of value. One observation linking intent to provider.

Constraint attaches.

User wanted 'past midnight' — HOURS becomes a routable attribute.

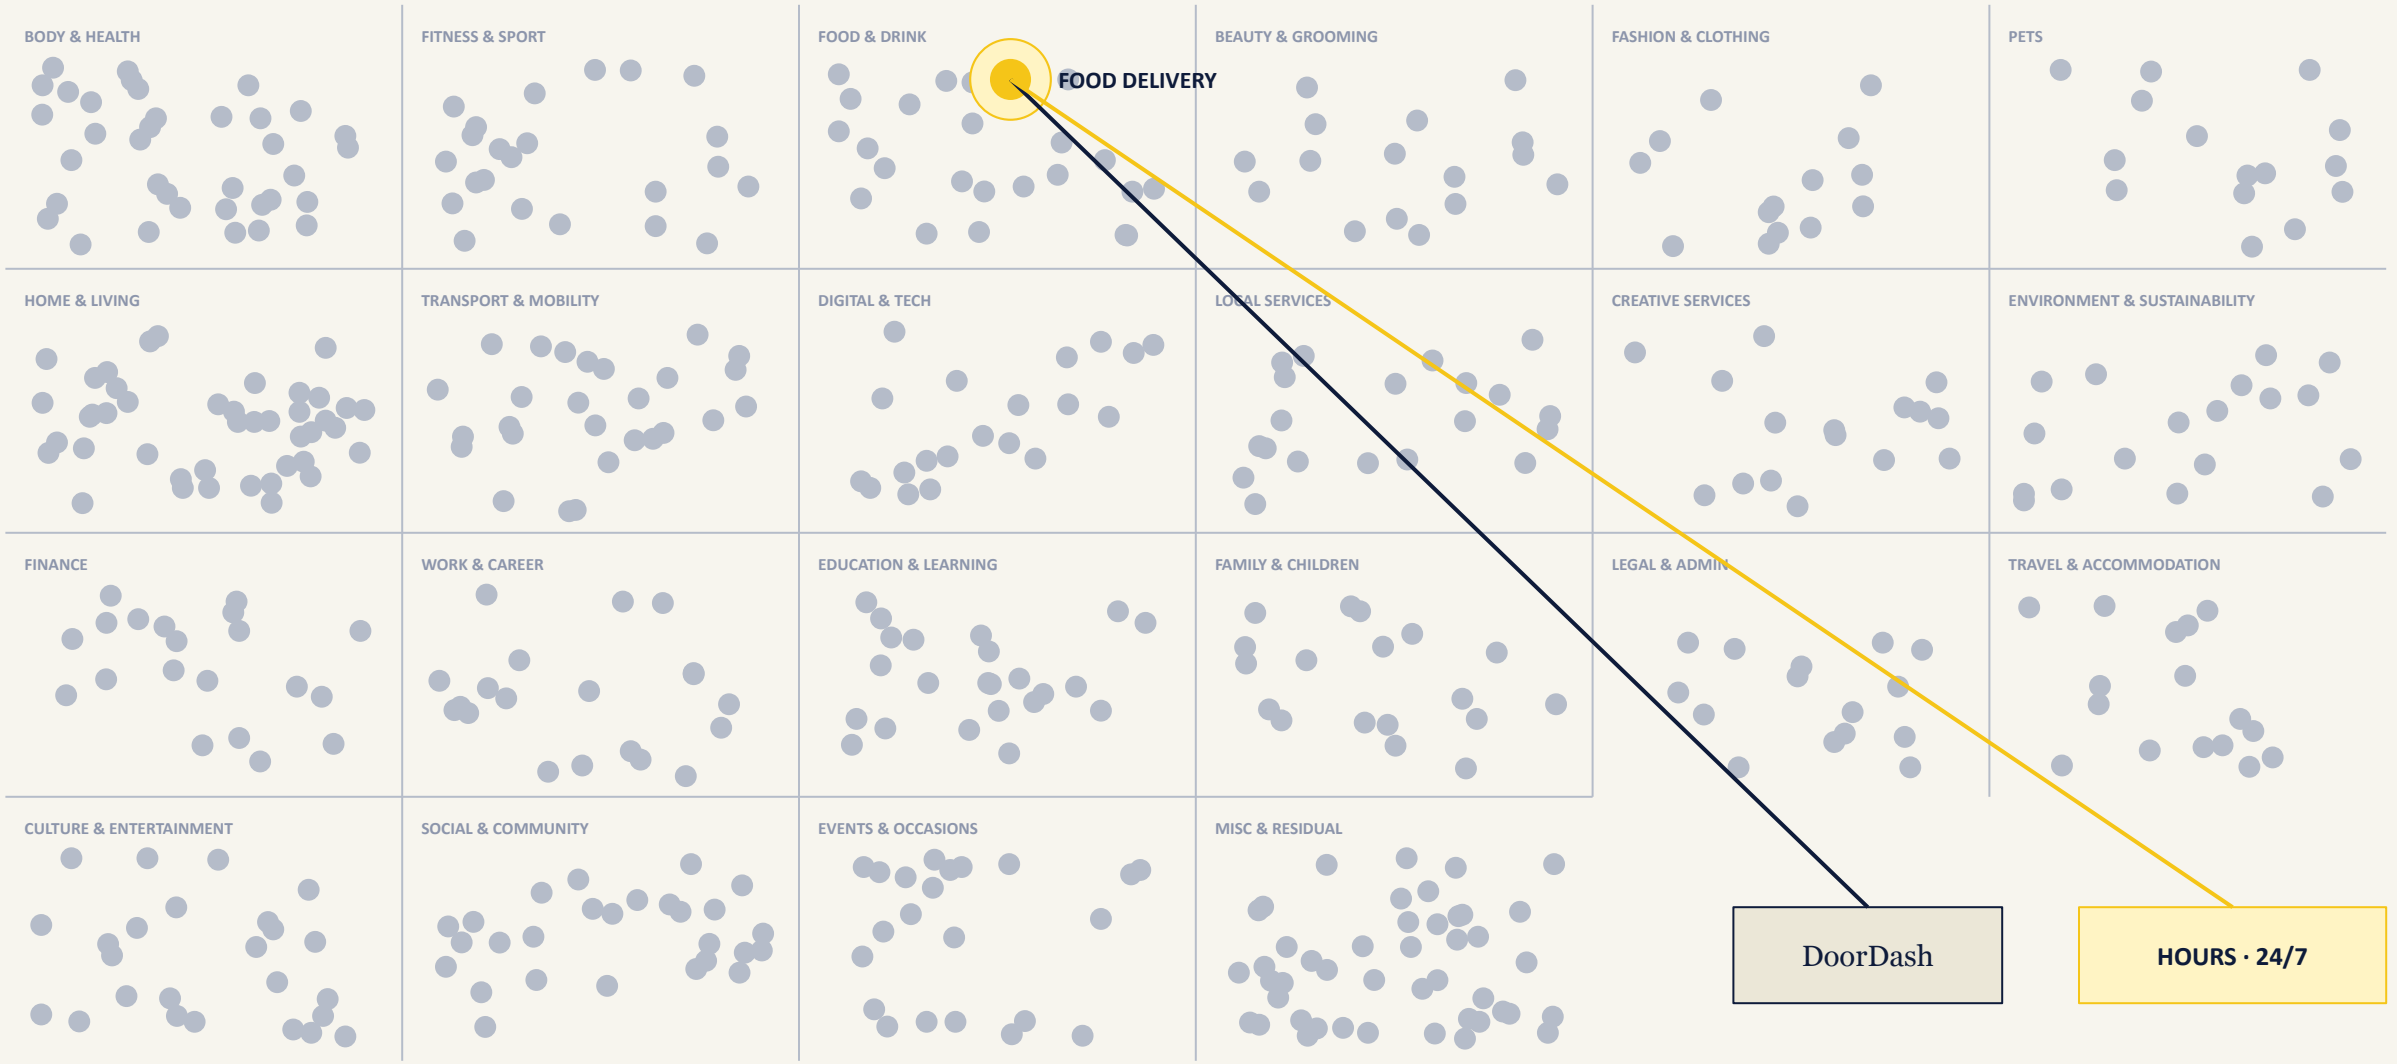

Constraints are first-class. The graph stores not just WHO but UNDER WHAT CONDITIONS.

Outcome observed. Edge thickens.

User clicks DoorDash. Booking confirmed. The edge weight increments. The graph remembers.

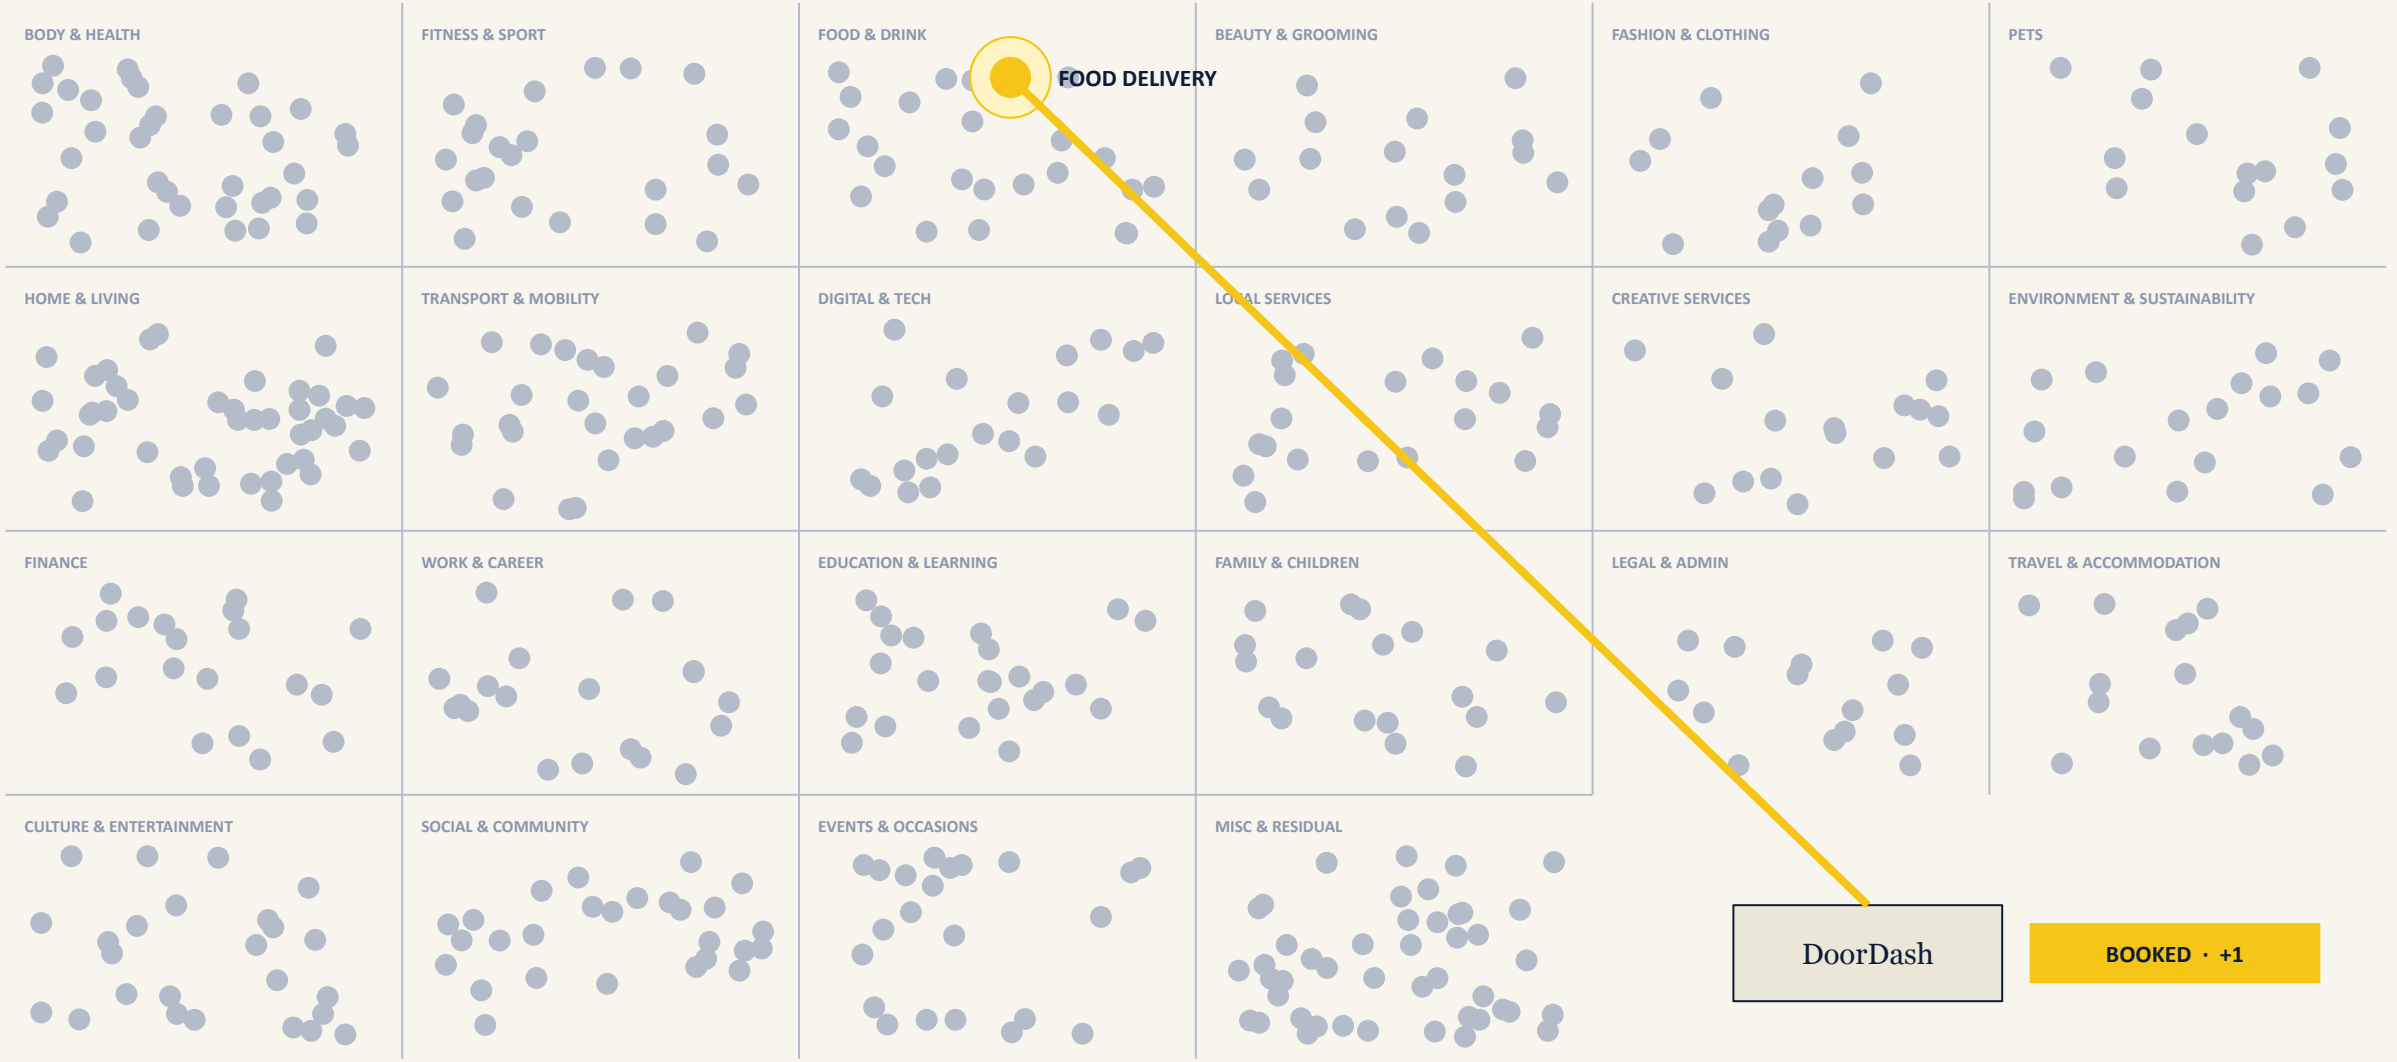

Every signal — click, dwell, return, booking — becomes an edge weight update.

Second query. New constraint emerges.

"who else delivers vegan food past midnight?" — DIET becomes routable.

EDGES

5

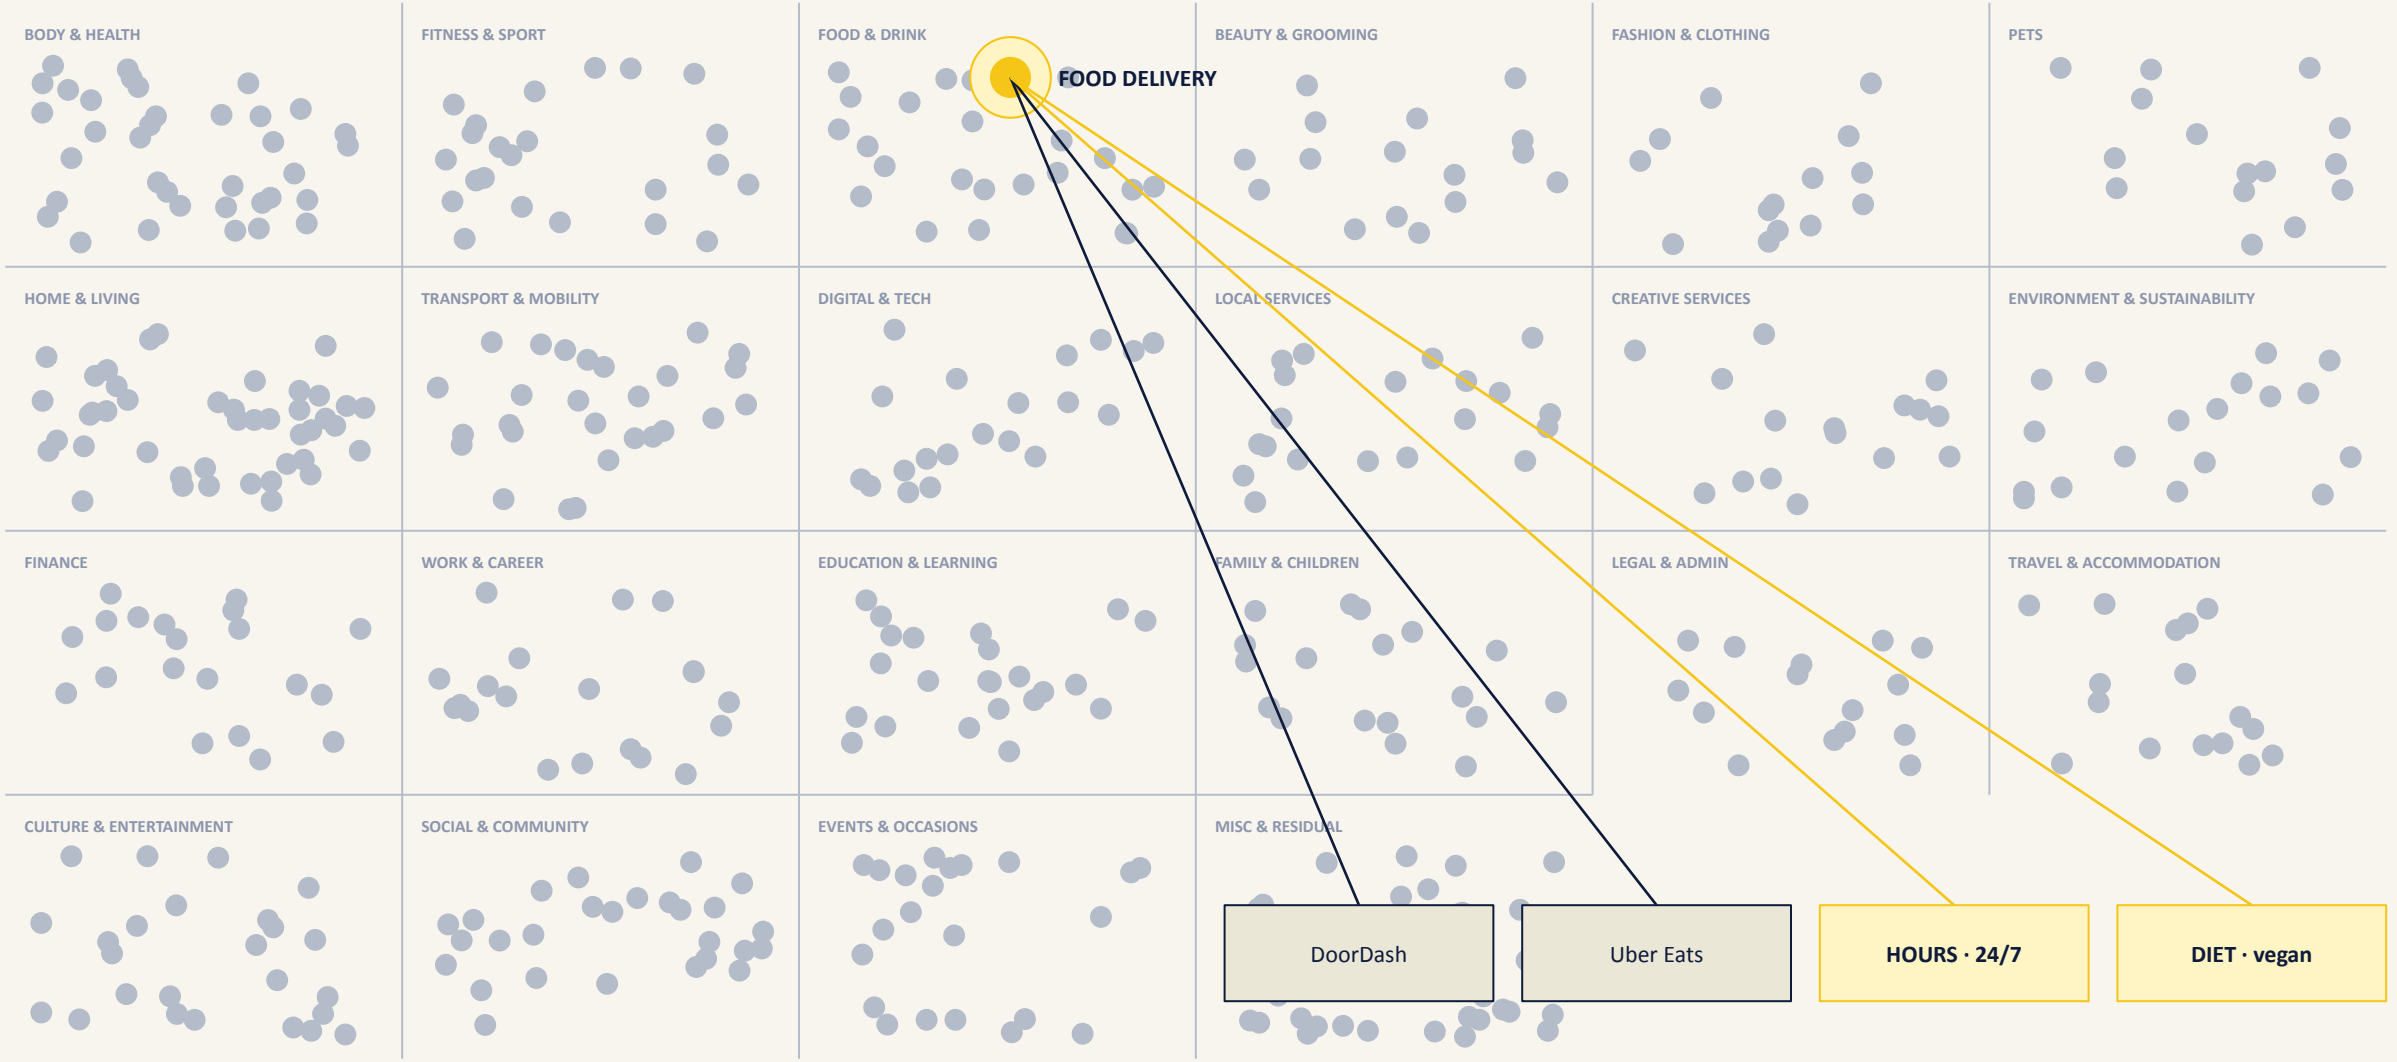

Two queries. The graph already knows providers AND attribute axes. The flywheel begins.

FOOD DELIVERY now has 6 edges.

After ~50 queries, the node is surrounded by constraint petals. A first sub-cluster.

EDGES

6

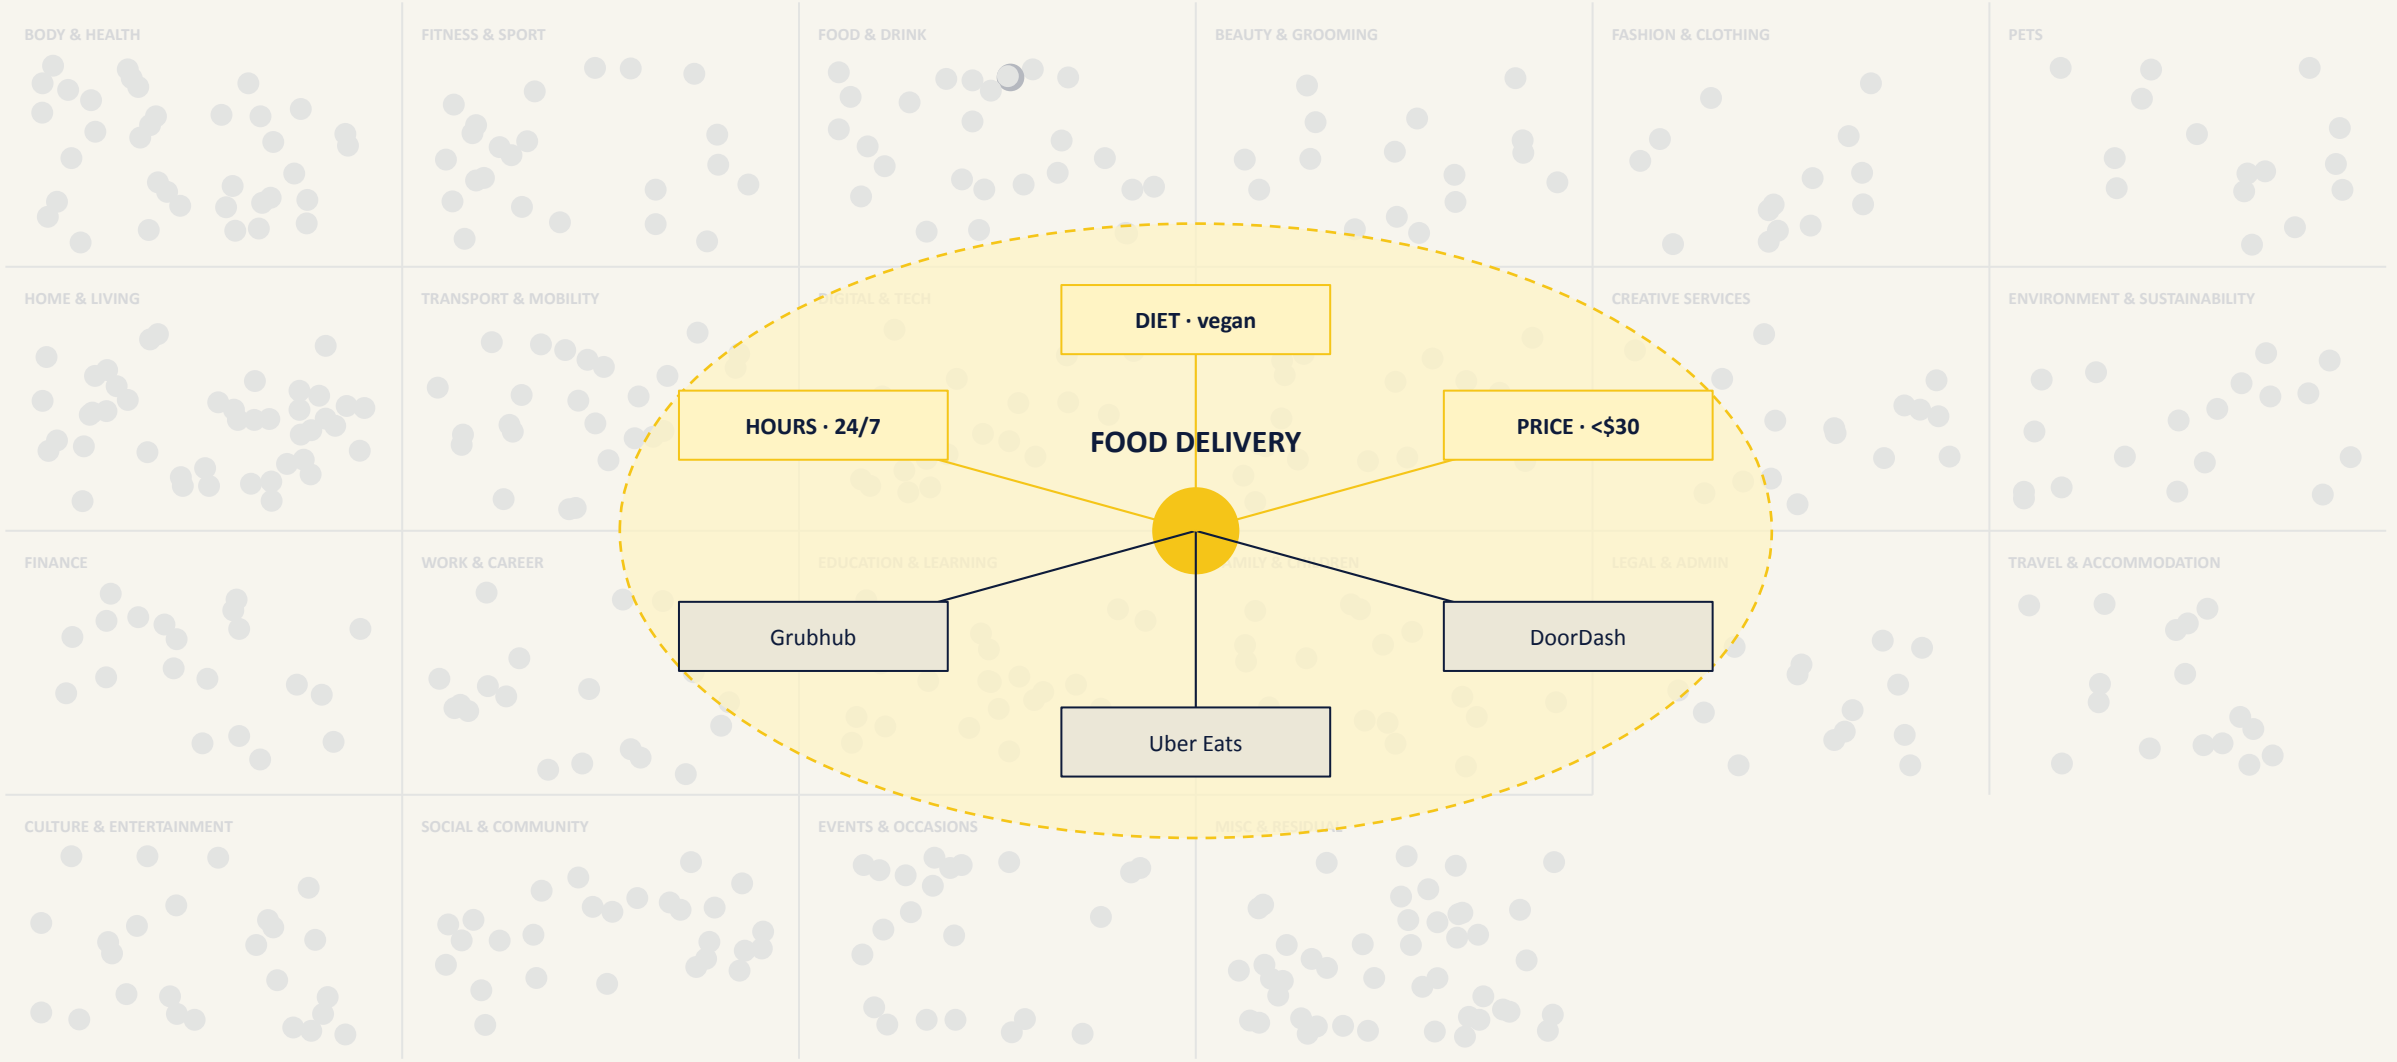

One node, six edges, two attribute types. This pattern will repeat 506 times.

Adjacent nodes activate.

Users keep asking food questions. RESTAURANT, GROCERY, MEAL KIT, BAKERY light up next.

ACTIVE NODES

12

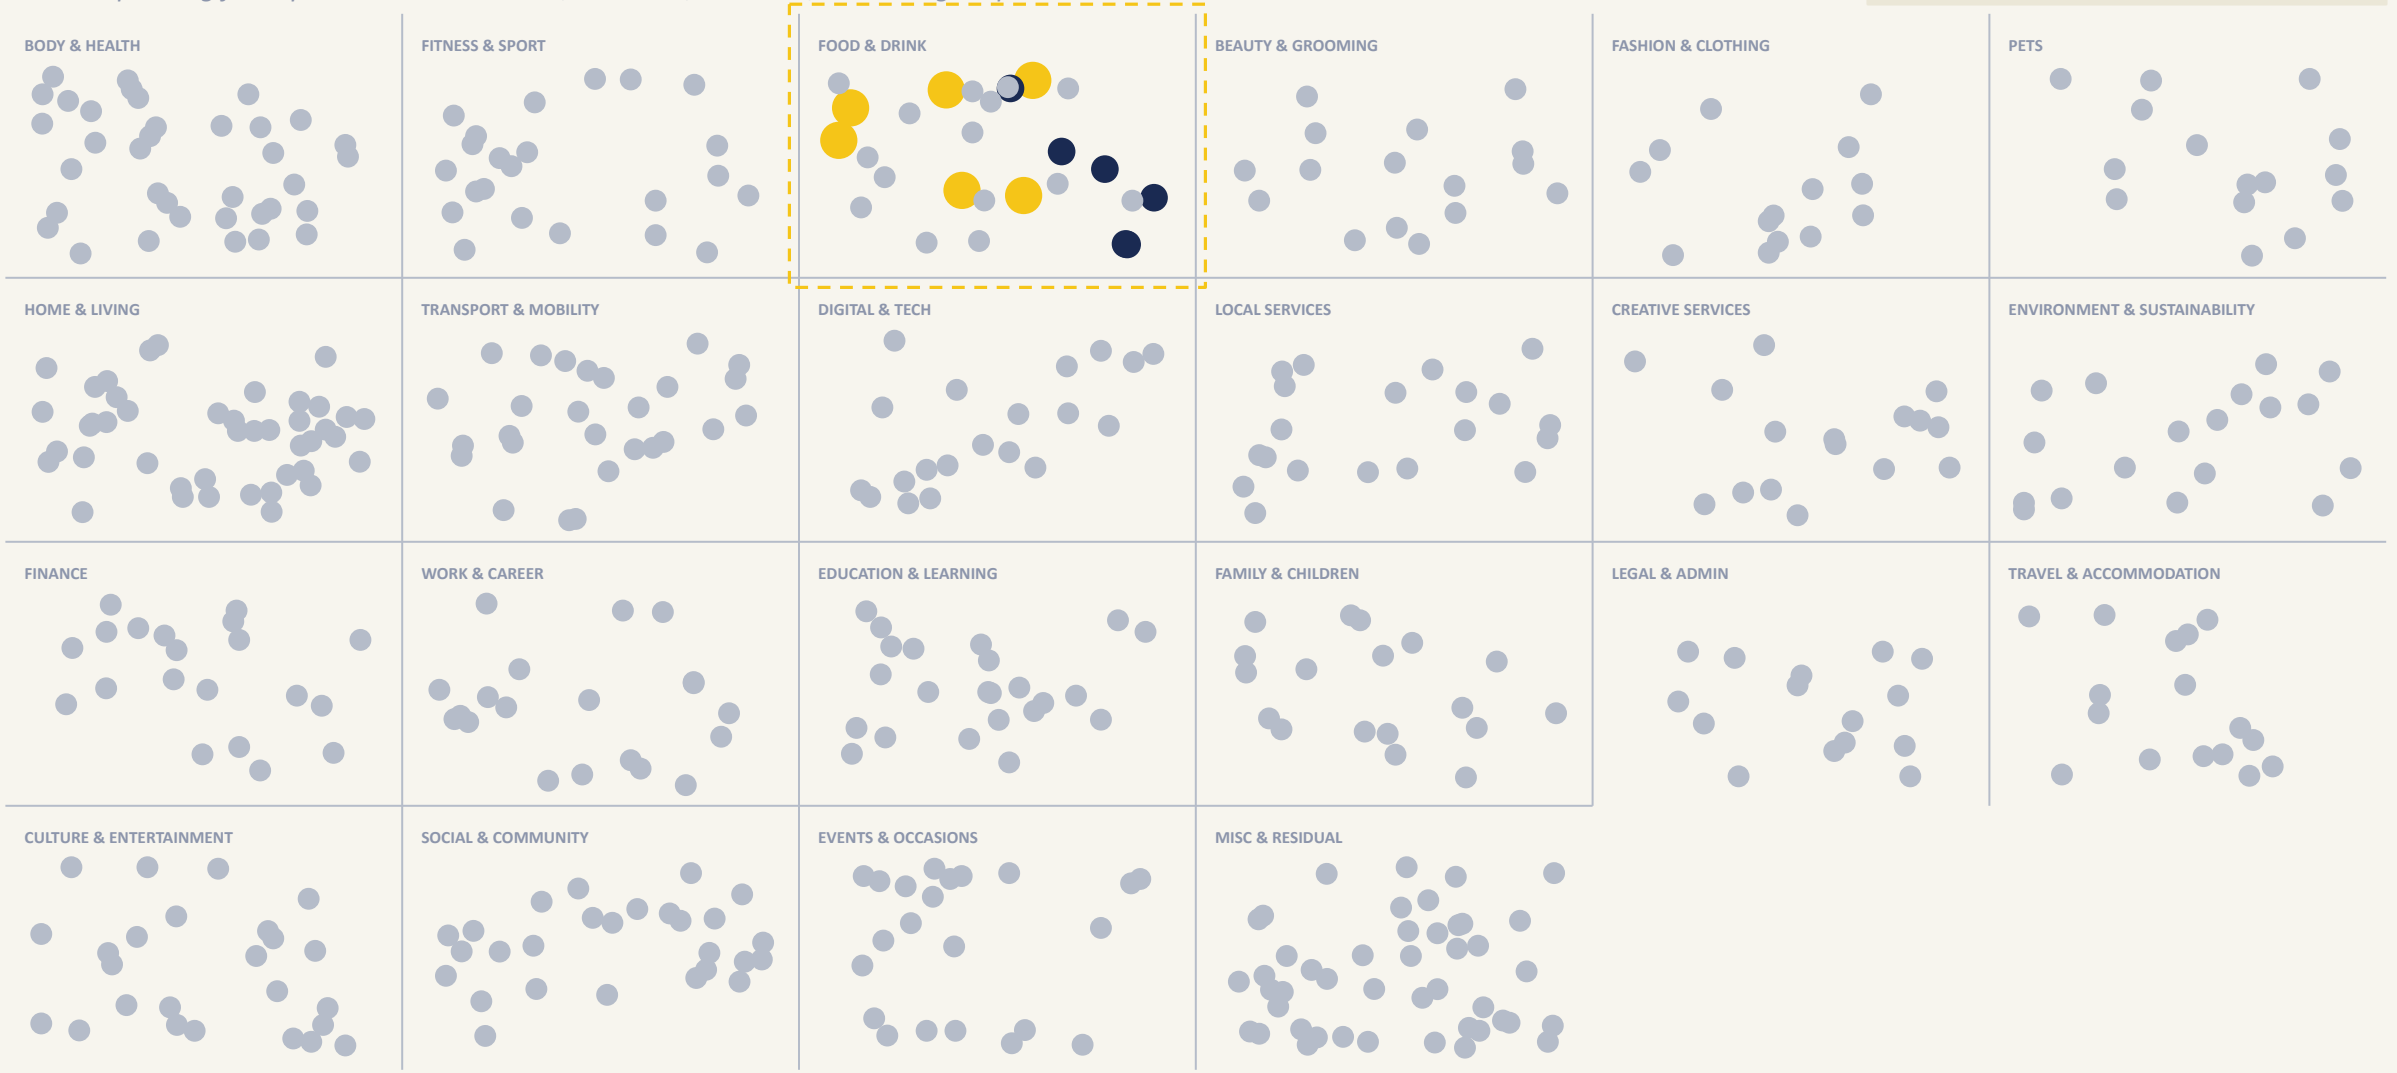

Activation spreads inside a sub-group. The cell starts to glow.

Cross-edges form.

The same providers appear across multiple food intents. Users move between them. Edges sprout.

EDGES IN CLUSTER

~80

Inside each cluster, providers and constraints are SHARED. Routing learns the structure for free.

FOOD & DRINK · 27 nodes routed.

First complete sub-group. Cluster has interior structure: dense centre, sparse edges, shared attributes.

CLUSTER #1

DONE

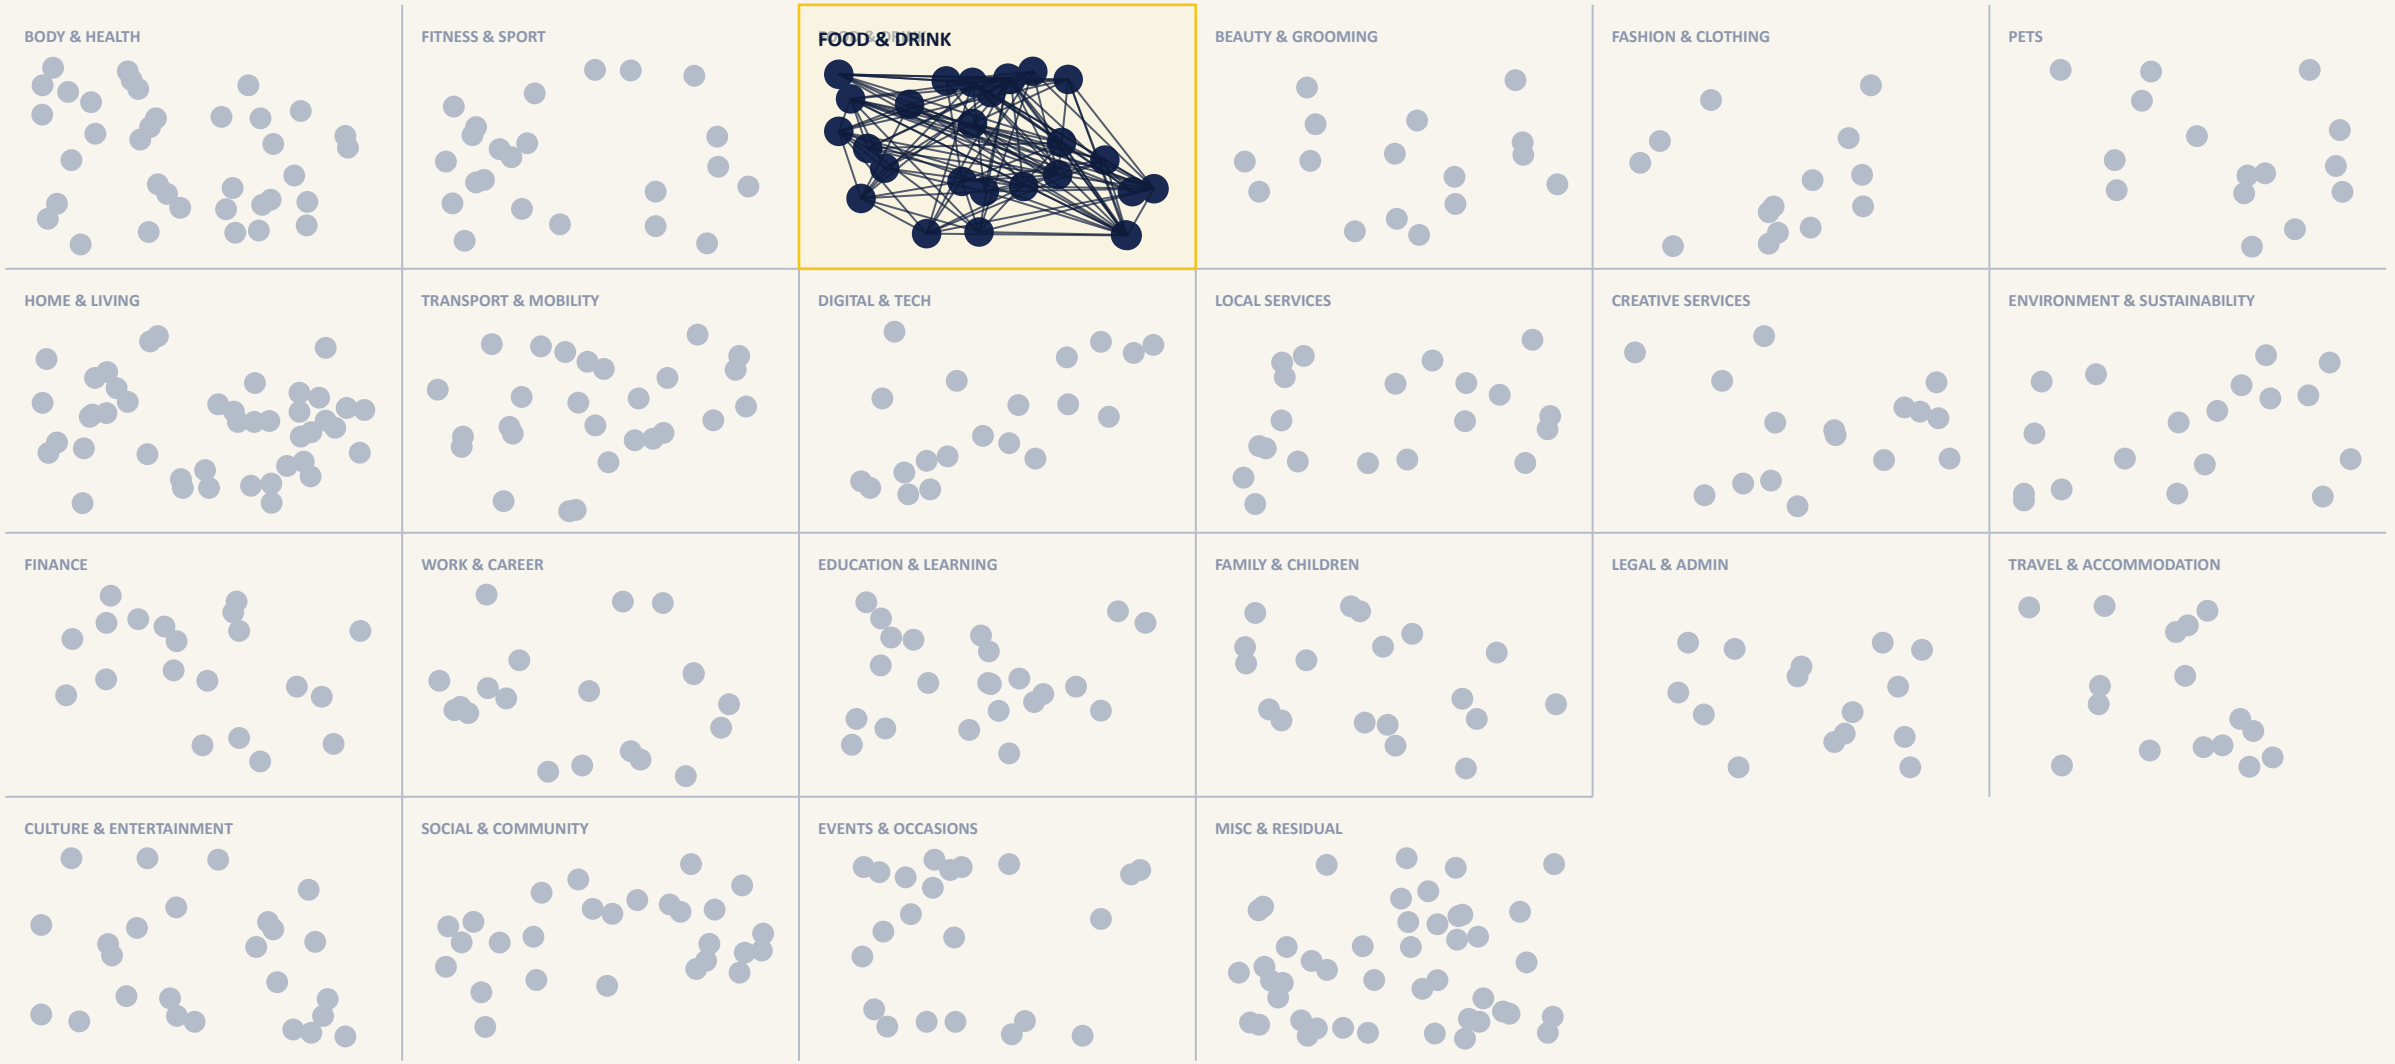

One sub-group routed. Twenty-one to go. The same pattern.

Second cluster awakens.

Transport intents start lighting up — different region of the graph, same flywheel.

ACTIVE NODES

54

Different category. Same routing primitive. Same cluster shape.

Constraints are SHARED across clusters.

OPEN NOW · LANGUAGE · BUDGET · INSURANCE — these don't belong to a single cluster. They span many.

CROSS-CLUSTER

DETECTED

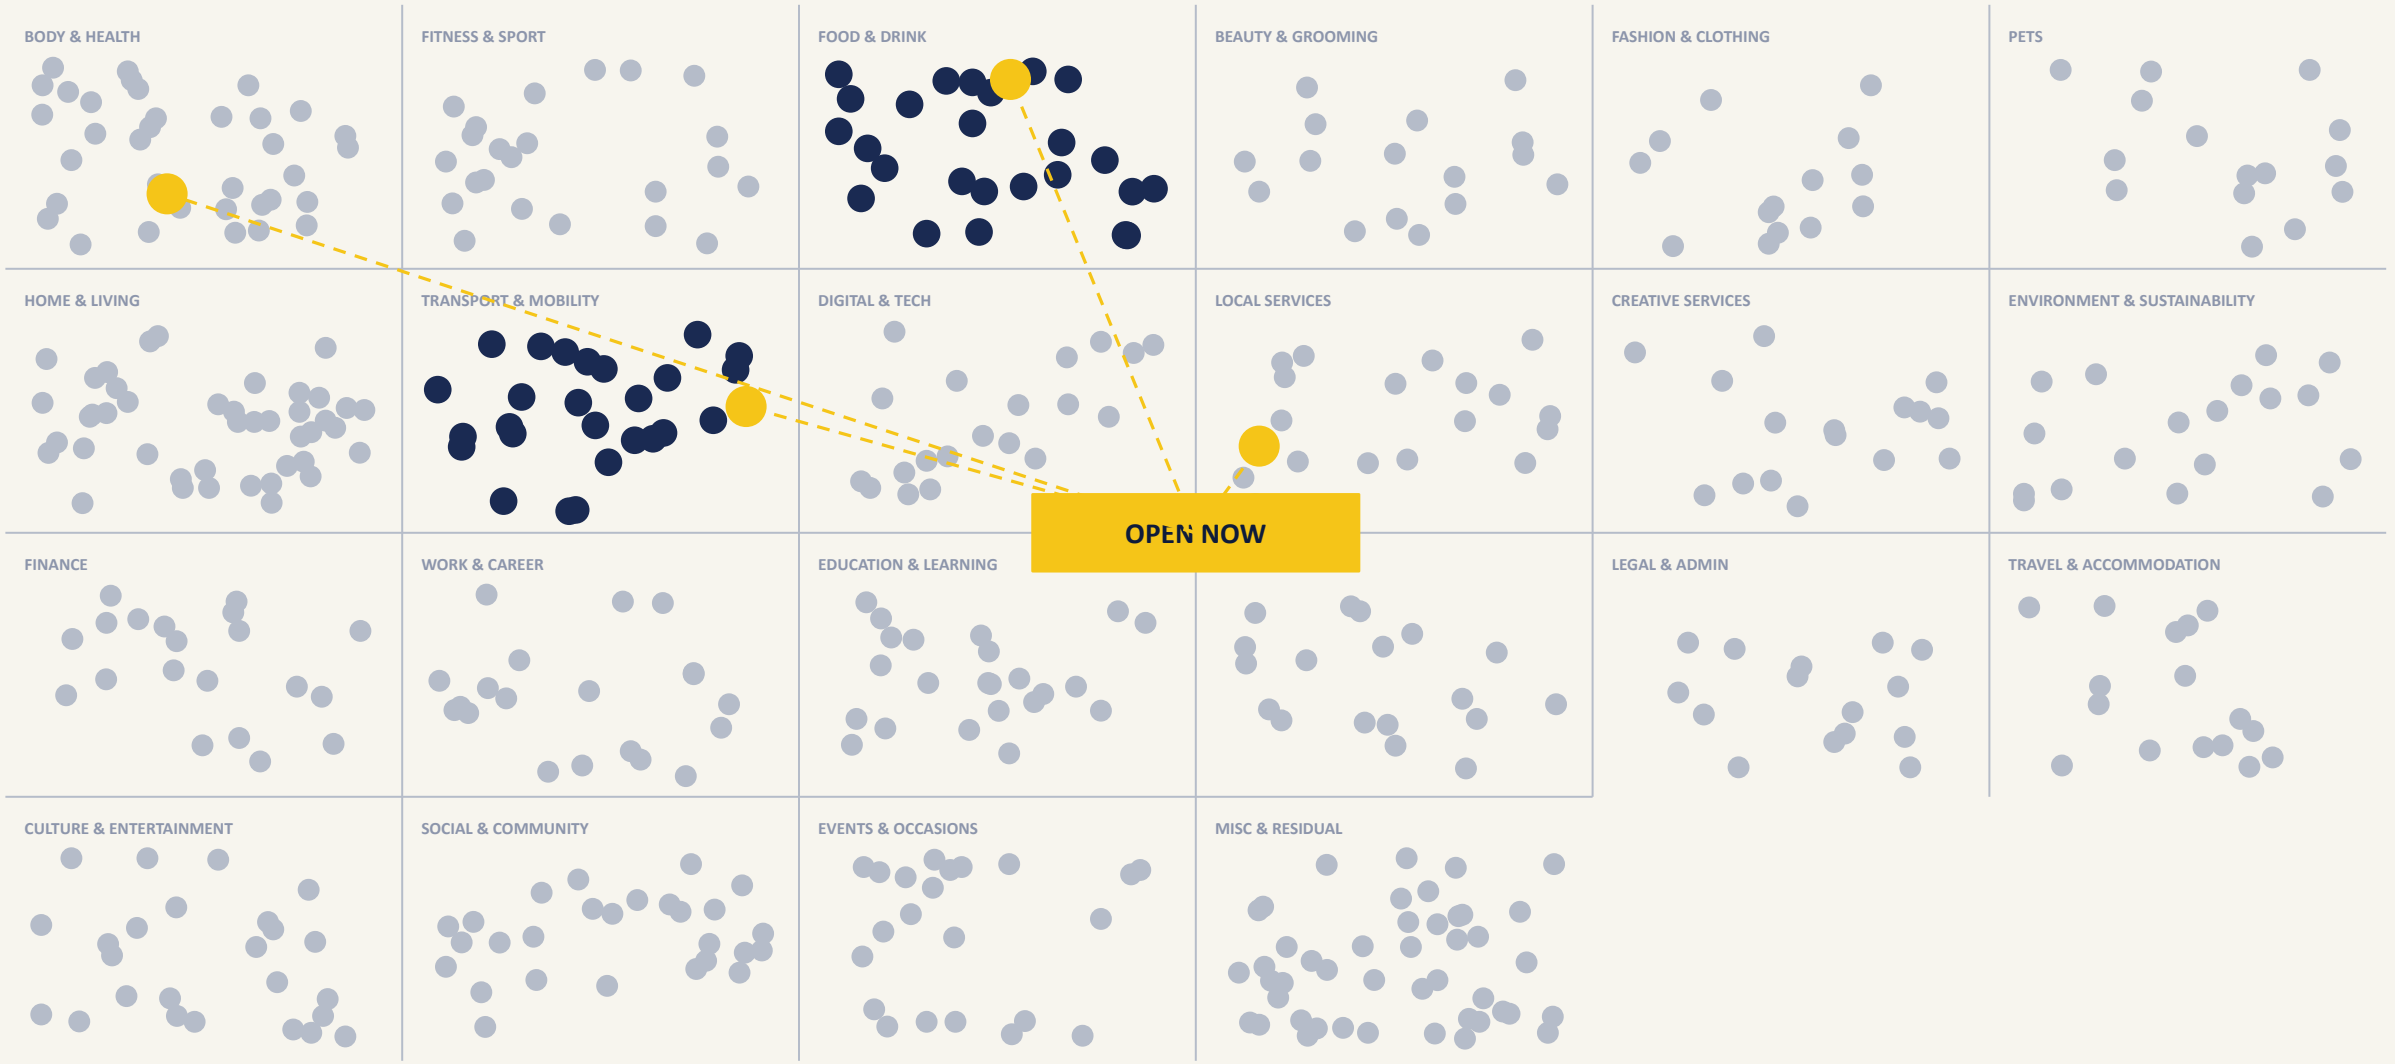

One constraint, many intents. The graph compresses what would otherwise be N separate verticals.

Routing reaches the long tail.

Categories no incumbent has aggregated start activating. LOCKSMITH. CHIROPRACTOR. BEEKEEPING.

FRAGMENTED ACTIVE

3

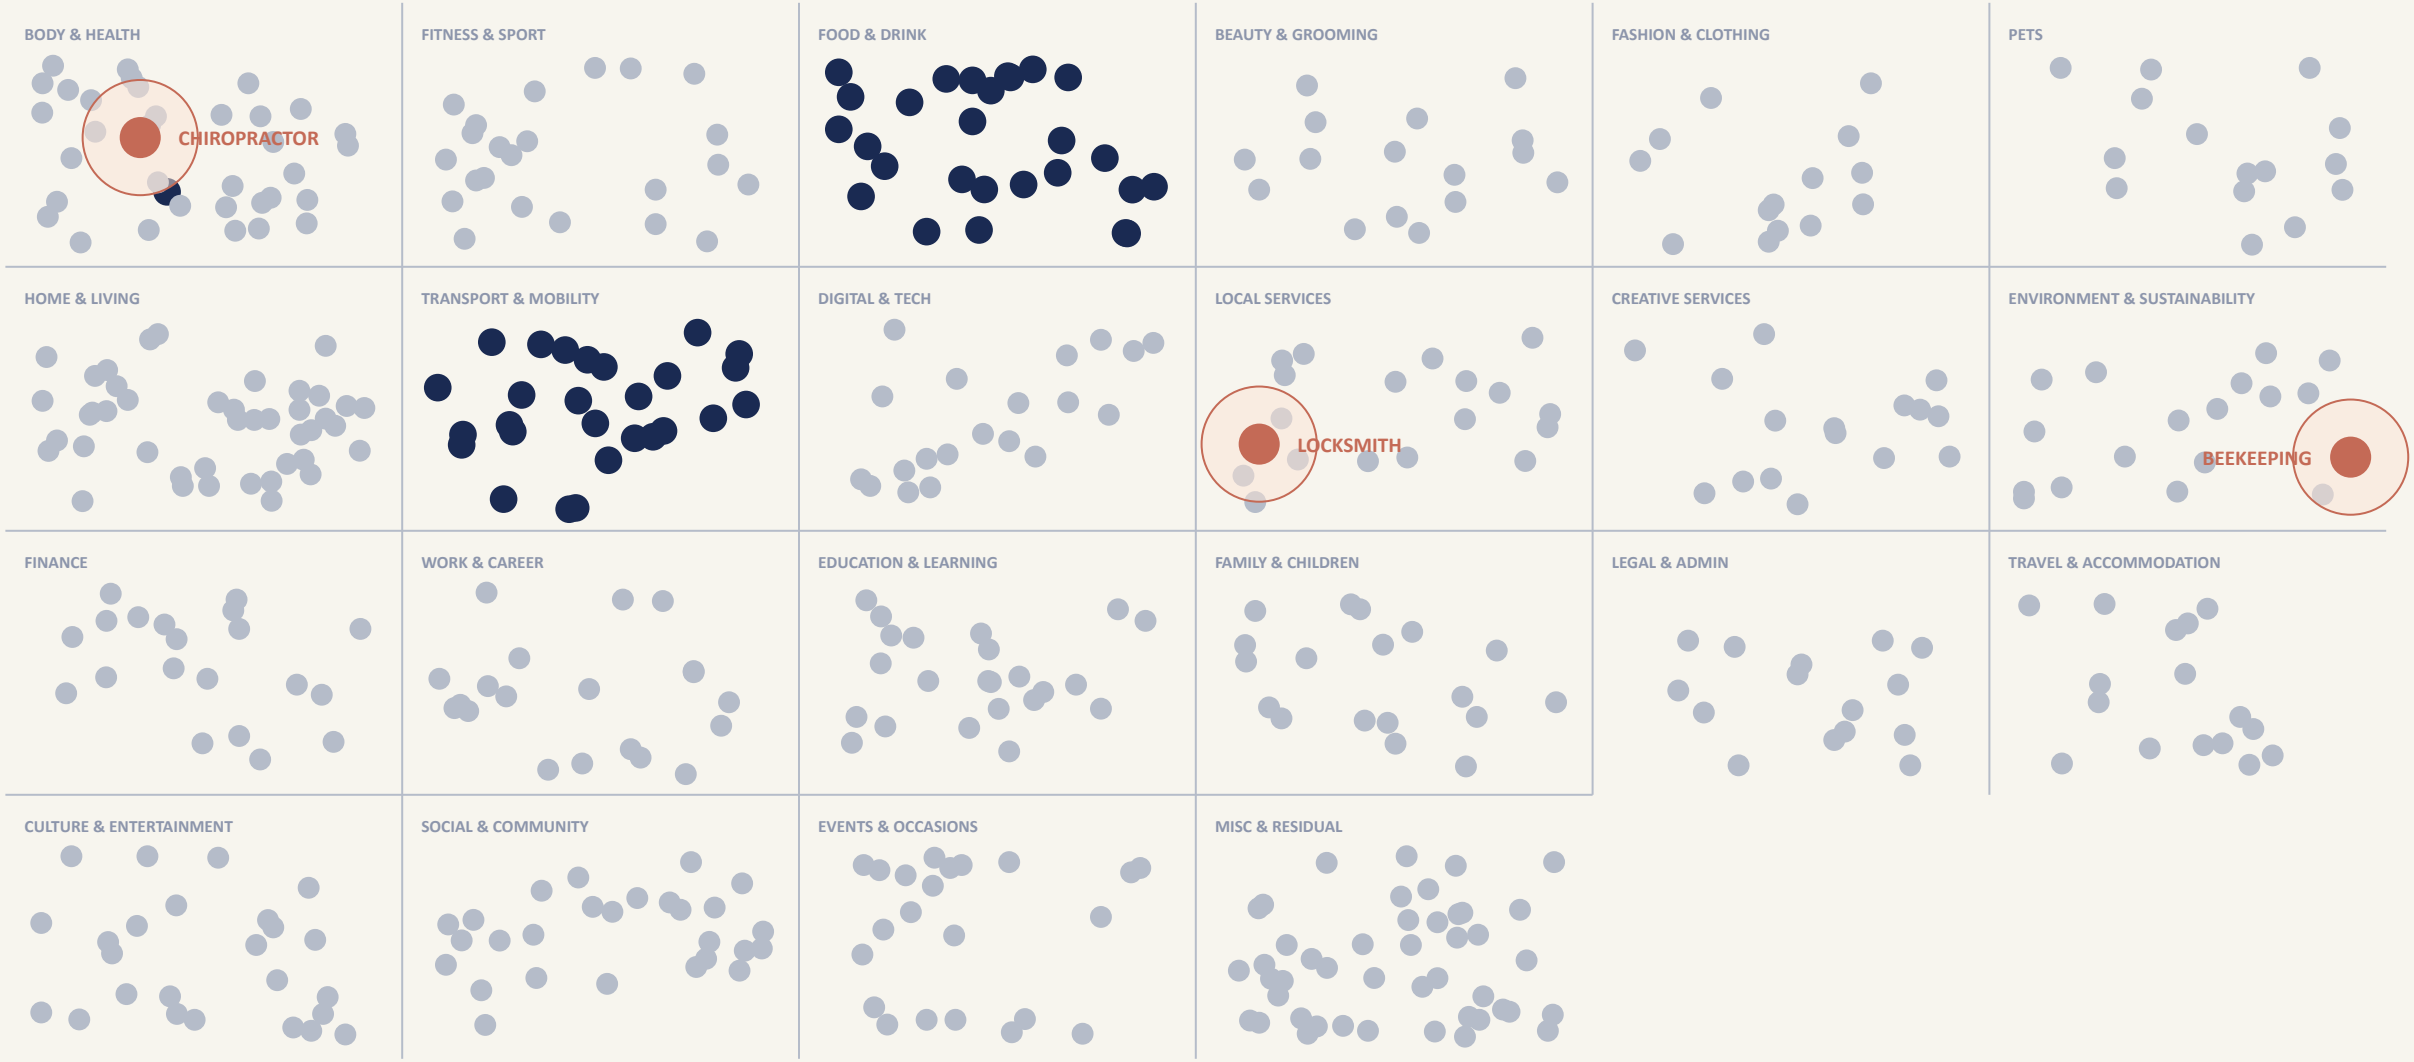

These nodes have NO consolidated answer-aggregator. The graph is the first place they get one.

User asks for something not yet in the seed list.

"who else does pet ultrasound?" — new node forks from VET. The graph extends itself.

NEW NODES

+1

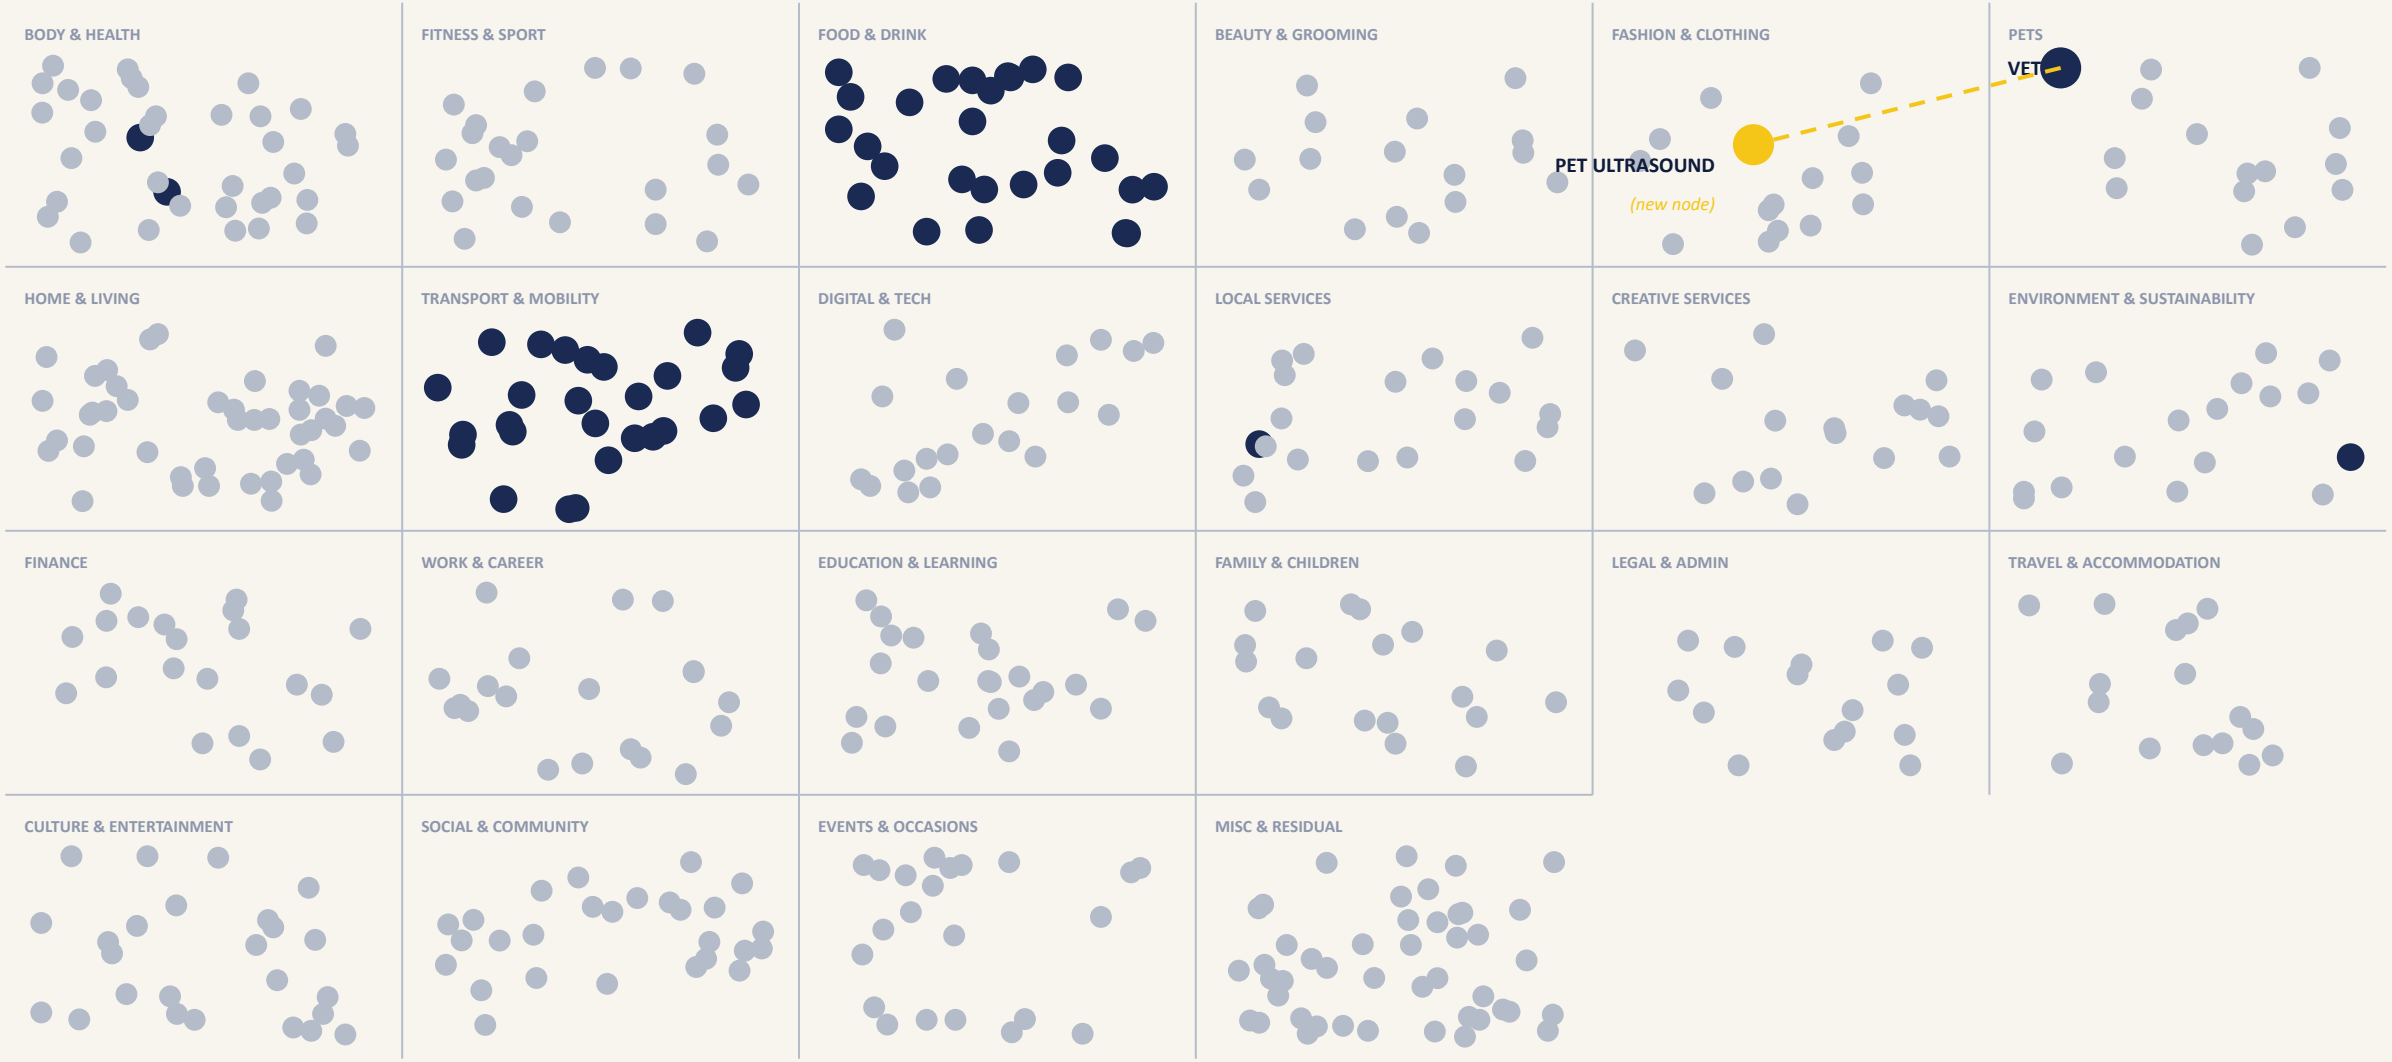

The taxonomy isn't fixed. The graph spawns nodes when a real query has no good match.

Provider data accrues fastest where incumbents have least.

LOCKSMITH had 0 edges on day 0. Now: 8 providers, 4 constraints, 12 cities.

LOCKSMITH EDGES

24

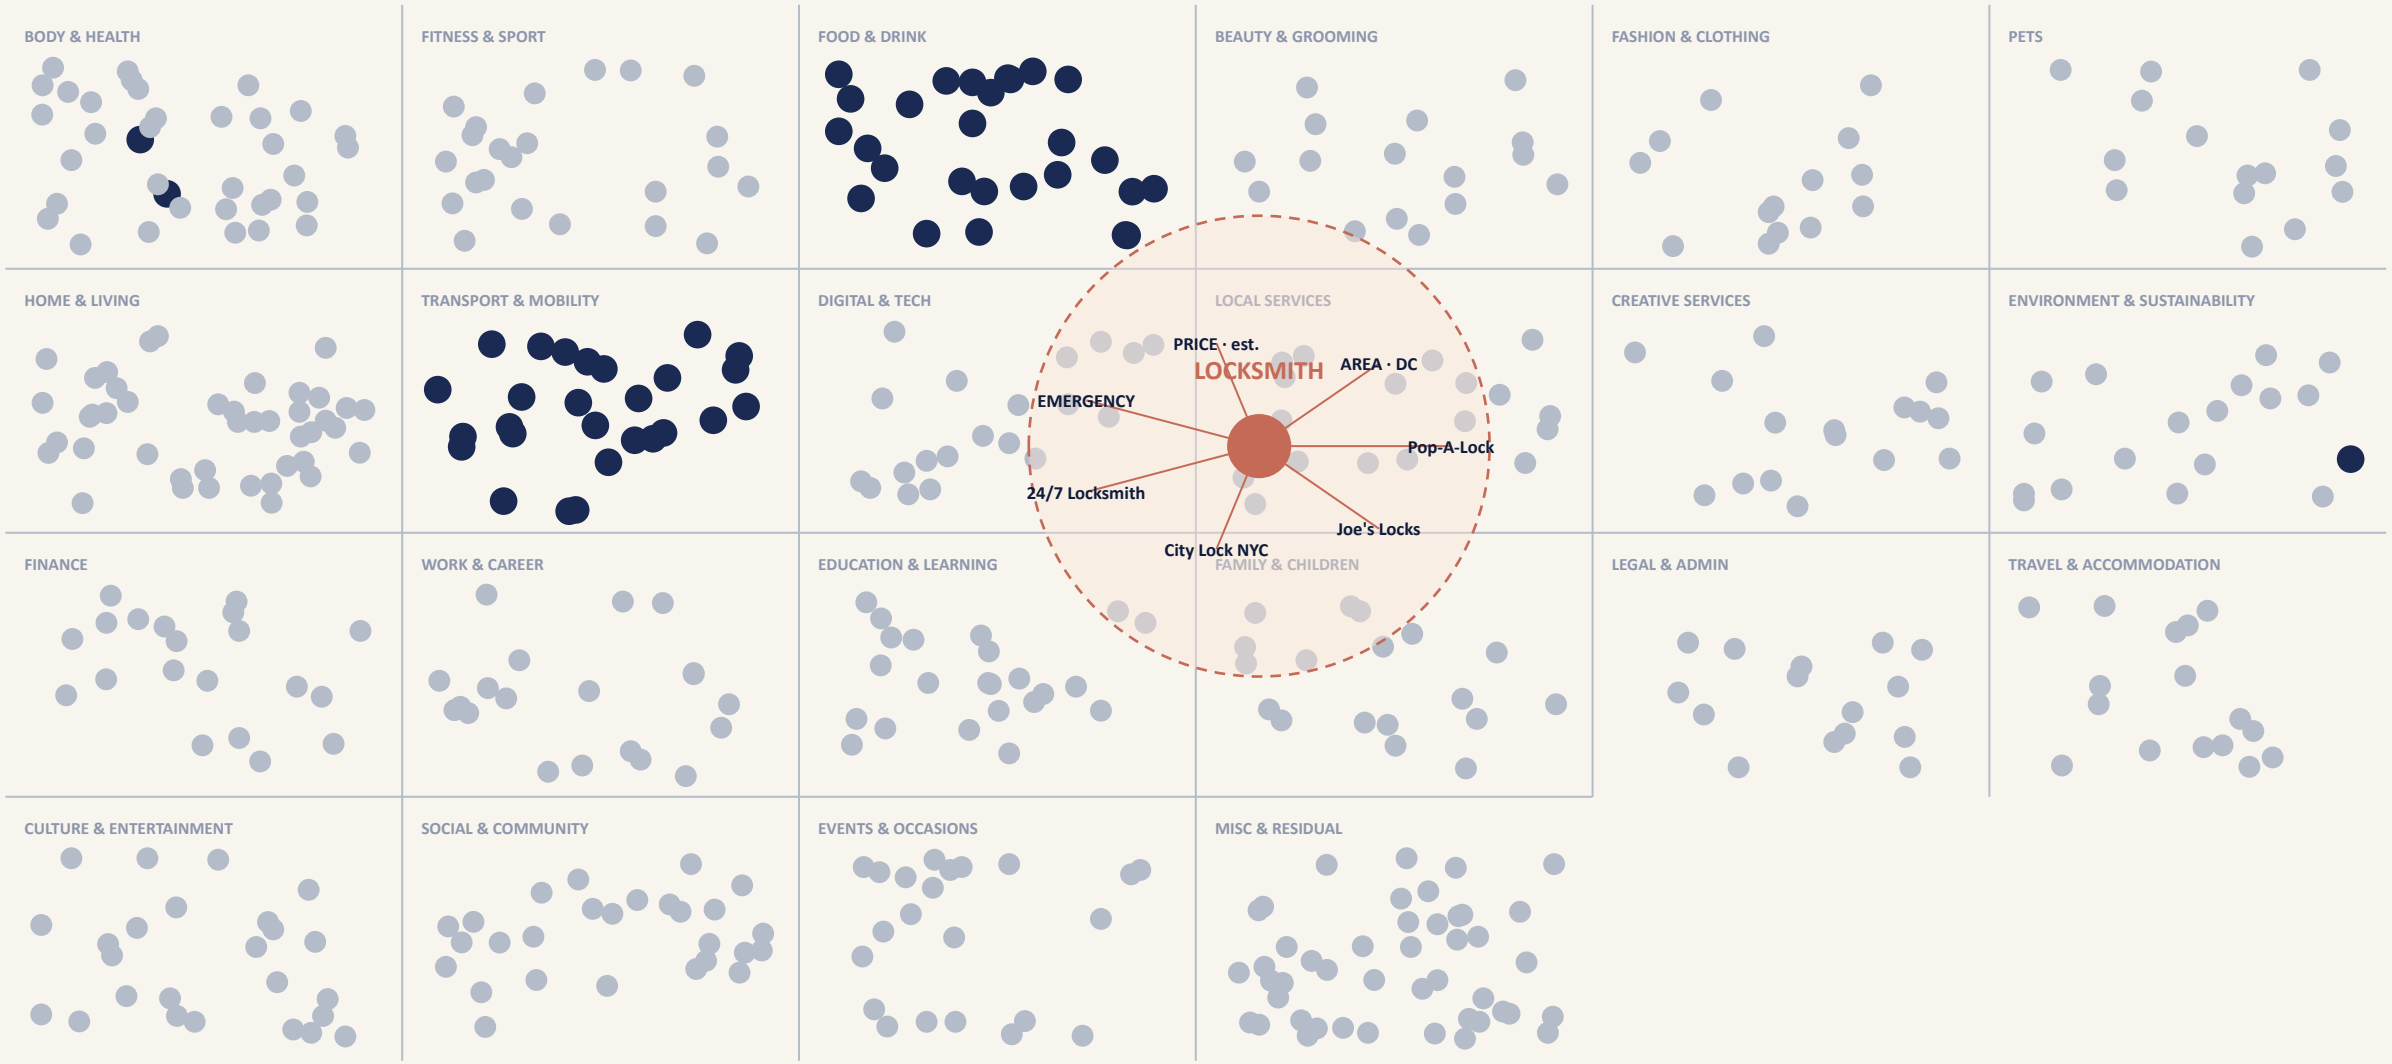

This is the slide that matters. Long-tail data is where the moat is built.

The flywheel turns asymmetrically.

Long-tail intents gain edges FASTER than head categories — because no one else is collecting this data.

GROWTH RATE

ASYMMETRIC

VERIFIED HEAD · 256 NODES

Slow growth.

Incumbents (DoorDash, Doctolib, Airbnb) already have most provider/constraint data.

Marginal value of each new edge: low.

EDGE GROWTH RATE

FRAGMENTED TAIL · 129 NODES

Compounding growth.

No incumbent has consolidated provider data here.

Every routing event is the FIRST observation. The graph is the only place the data exists.

EDGE GROWTH RATE

The asymmetry IS the defensibility.

Same node. Multiple language facets.

DOCTOR with Spanish-speaking, Mandarin-speaking, ASL-fluent provider edges.

LANGUAGE EDGES

+12

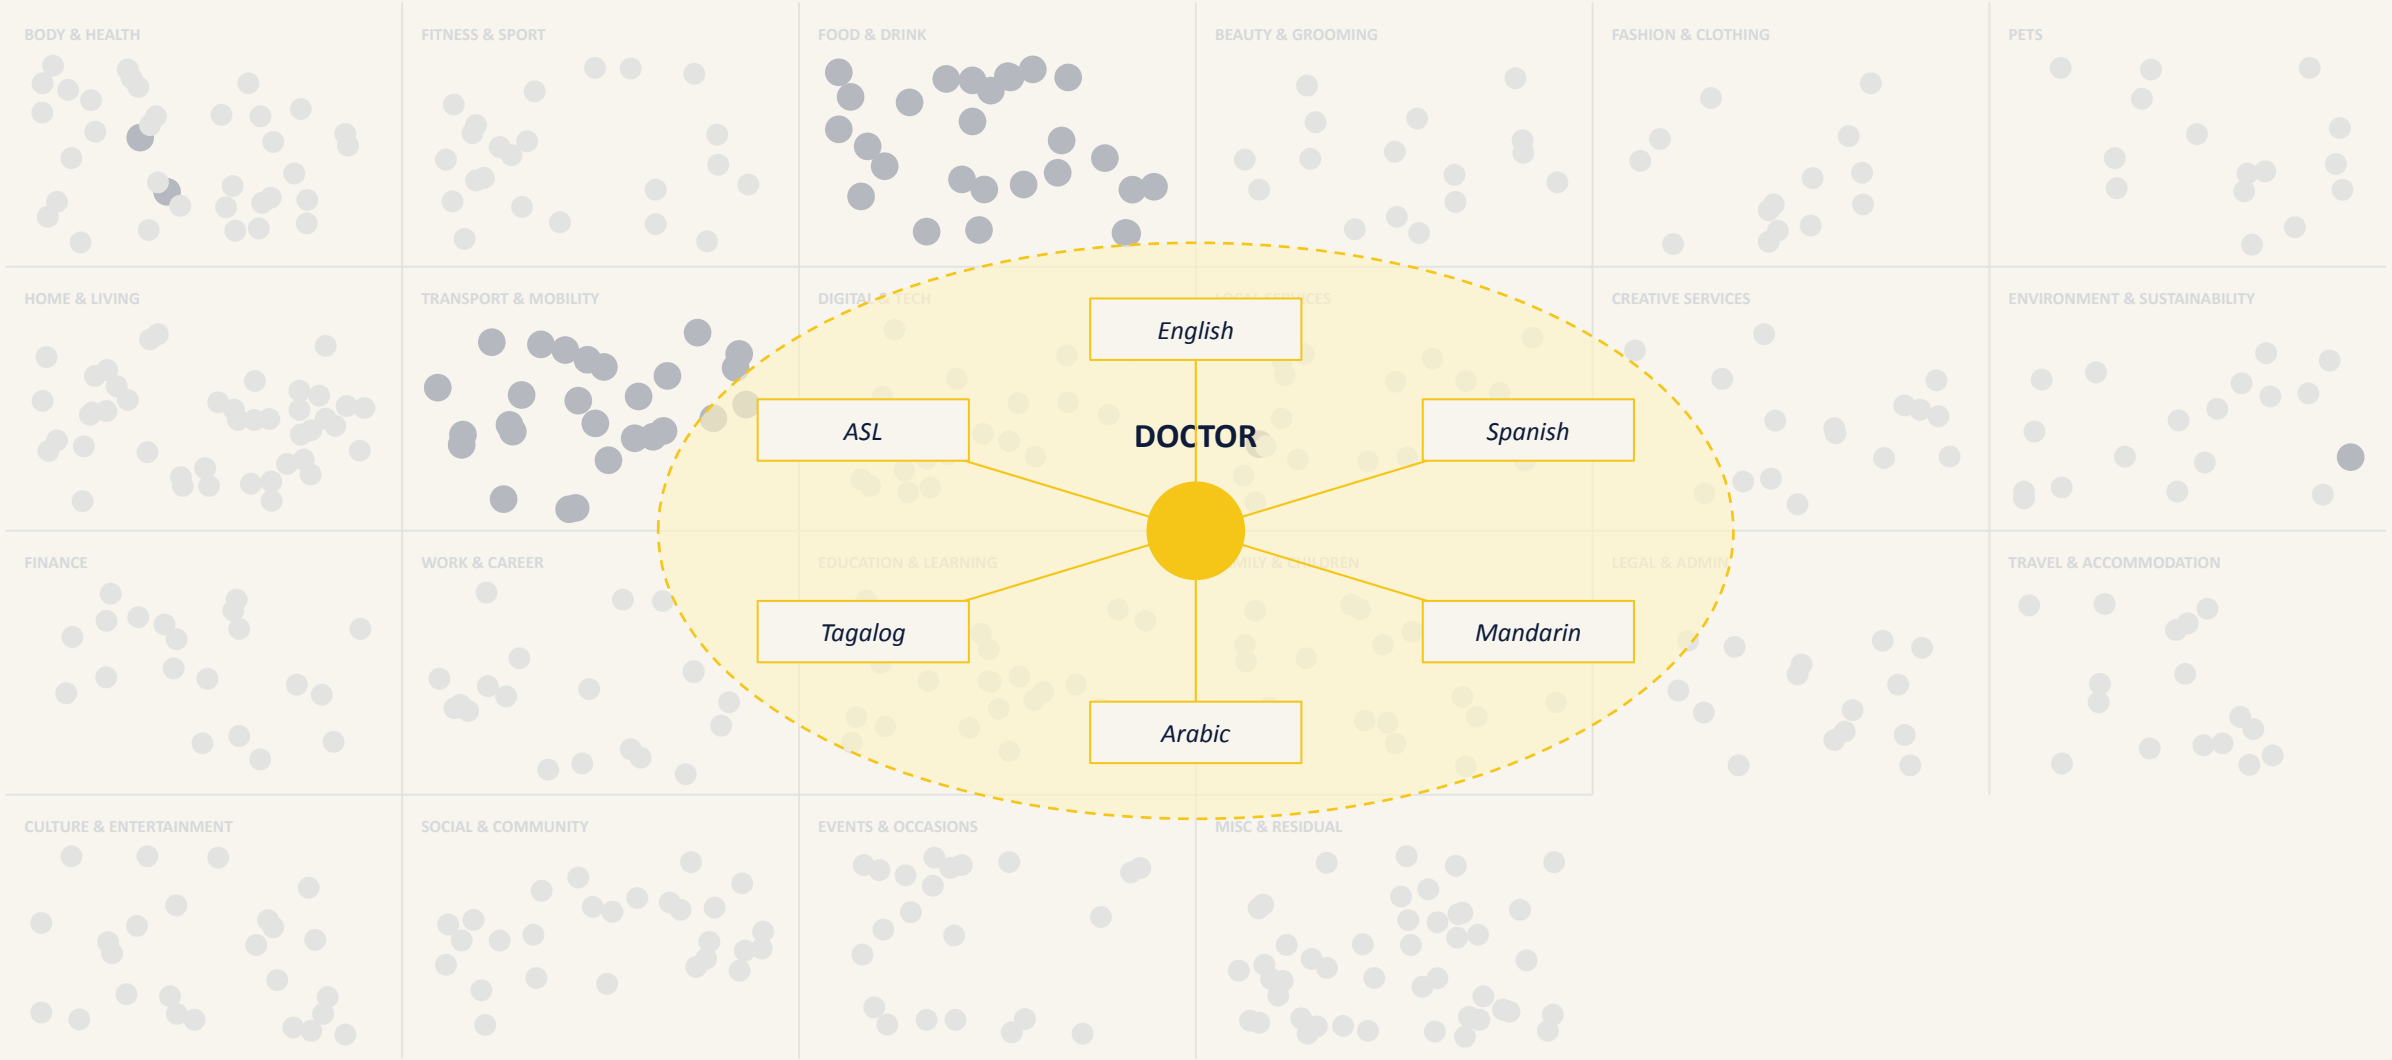

Constraint stacks compose. DOCTOR + Spanish + Aetna + open-Saturday → a single routable query.

What incumbents see vs. what we see.

Search engines index the head. We see the head AND the tail. Same primitive, more surface.

COVERAGE

100%

INCUMBENT VIEW · HEAD ONLY

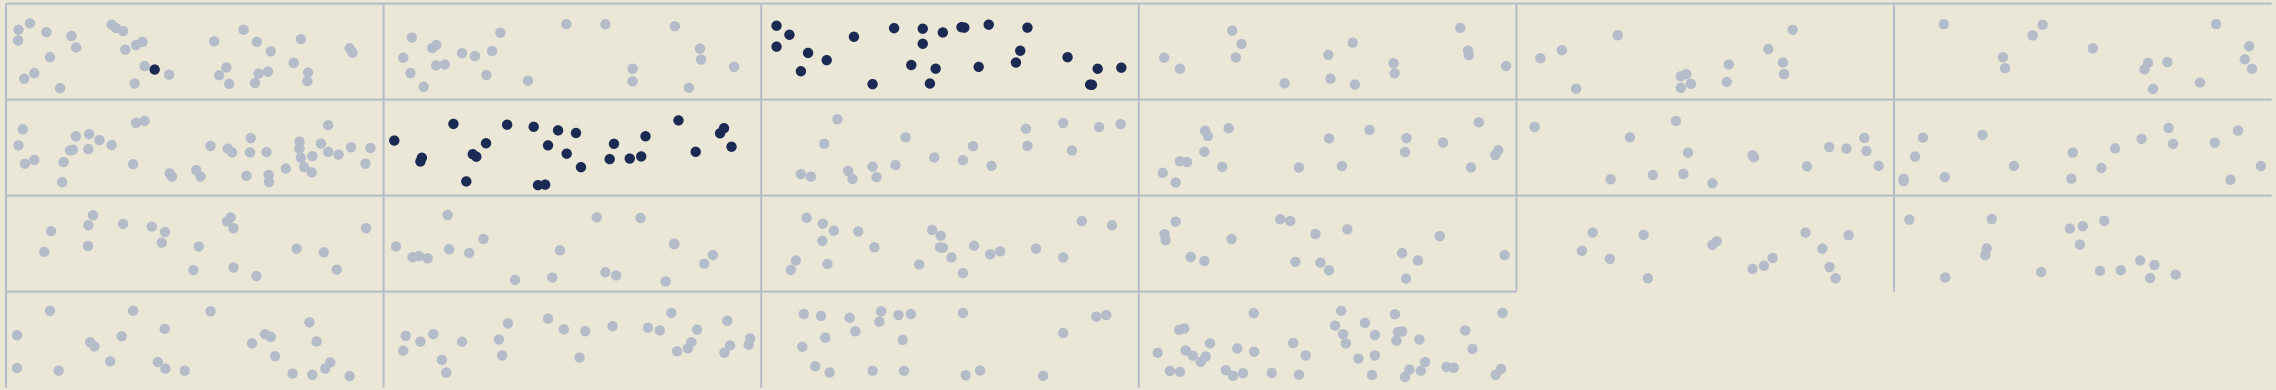

WHOELSE.AI VIEW · HEAD + TAIL

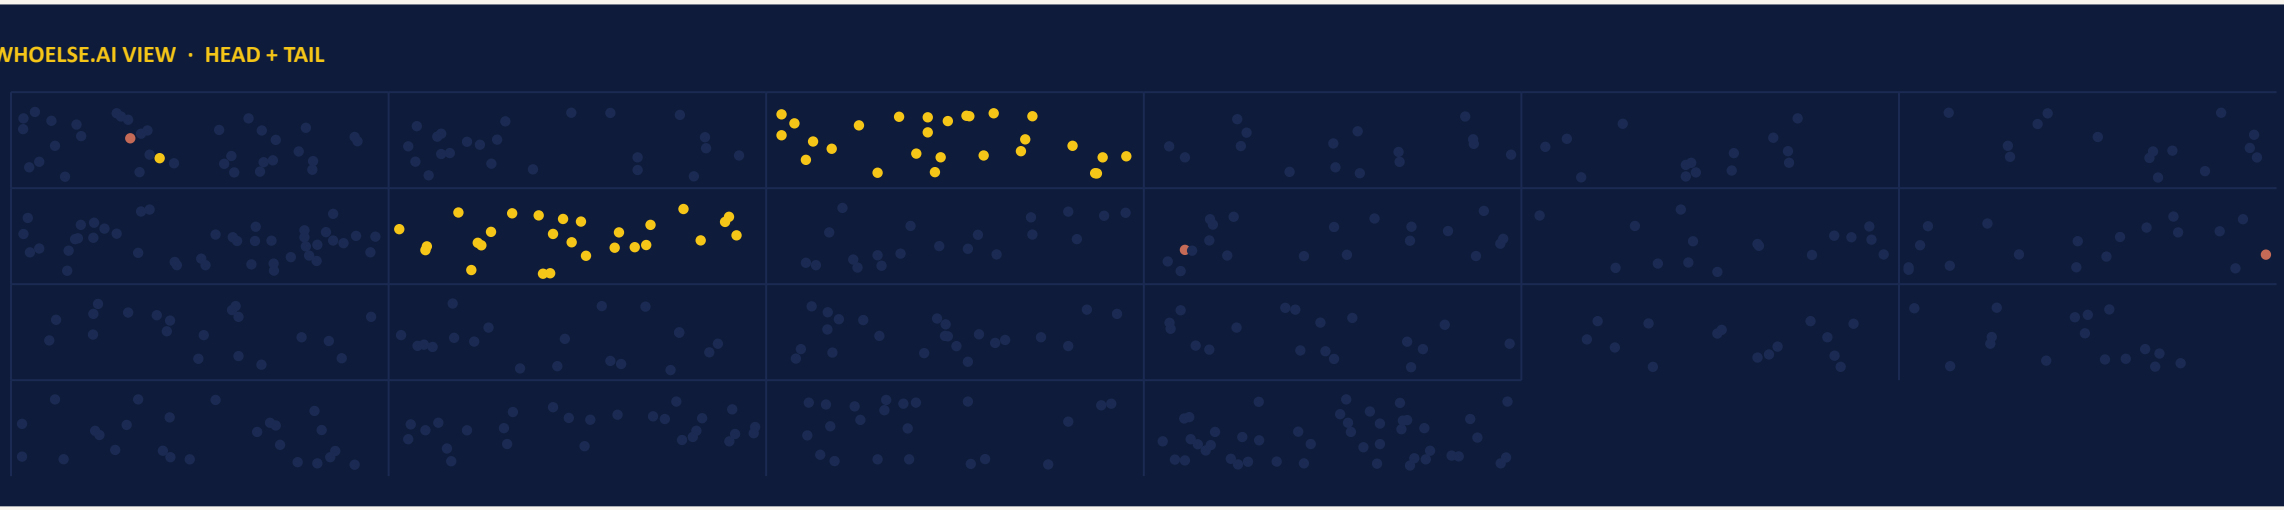

Incumbents see the lit head. We see the entire surface — including everything they've never indexed.

100 routing events / day.

Sparse activation. Cluster cores visible. Most nodes still dim.

QUERIES / DAY

100

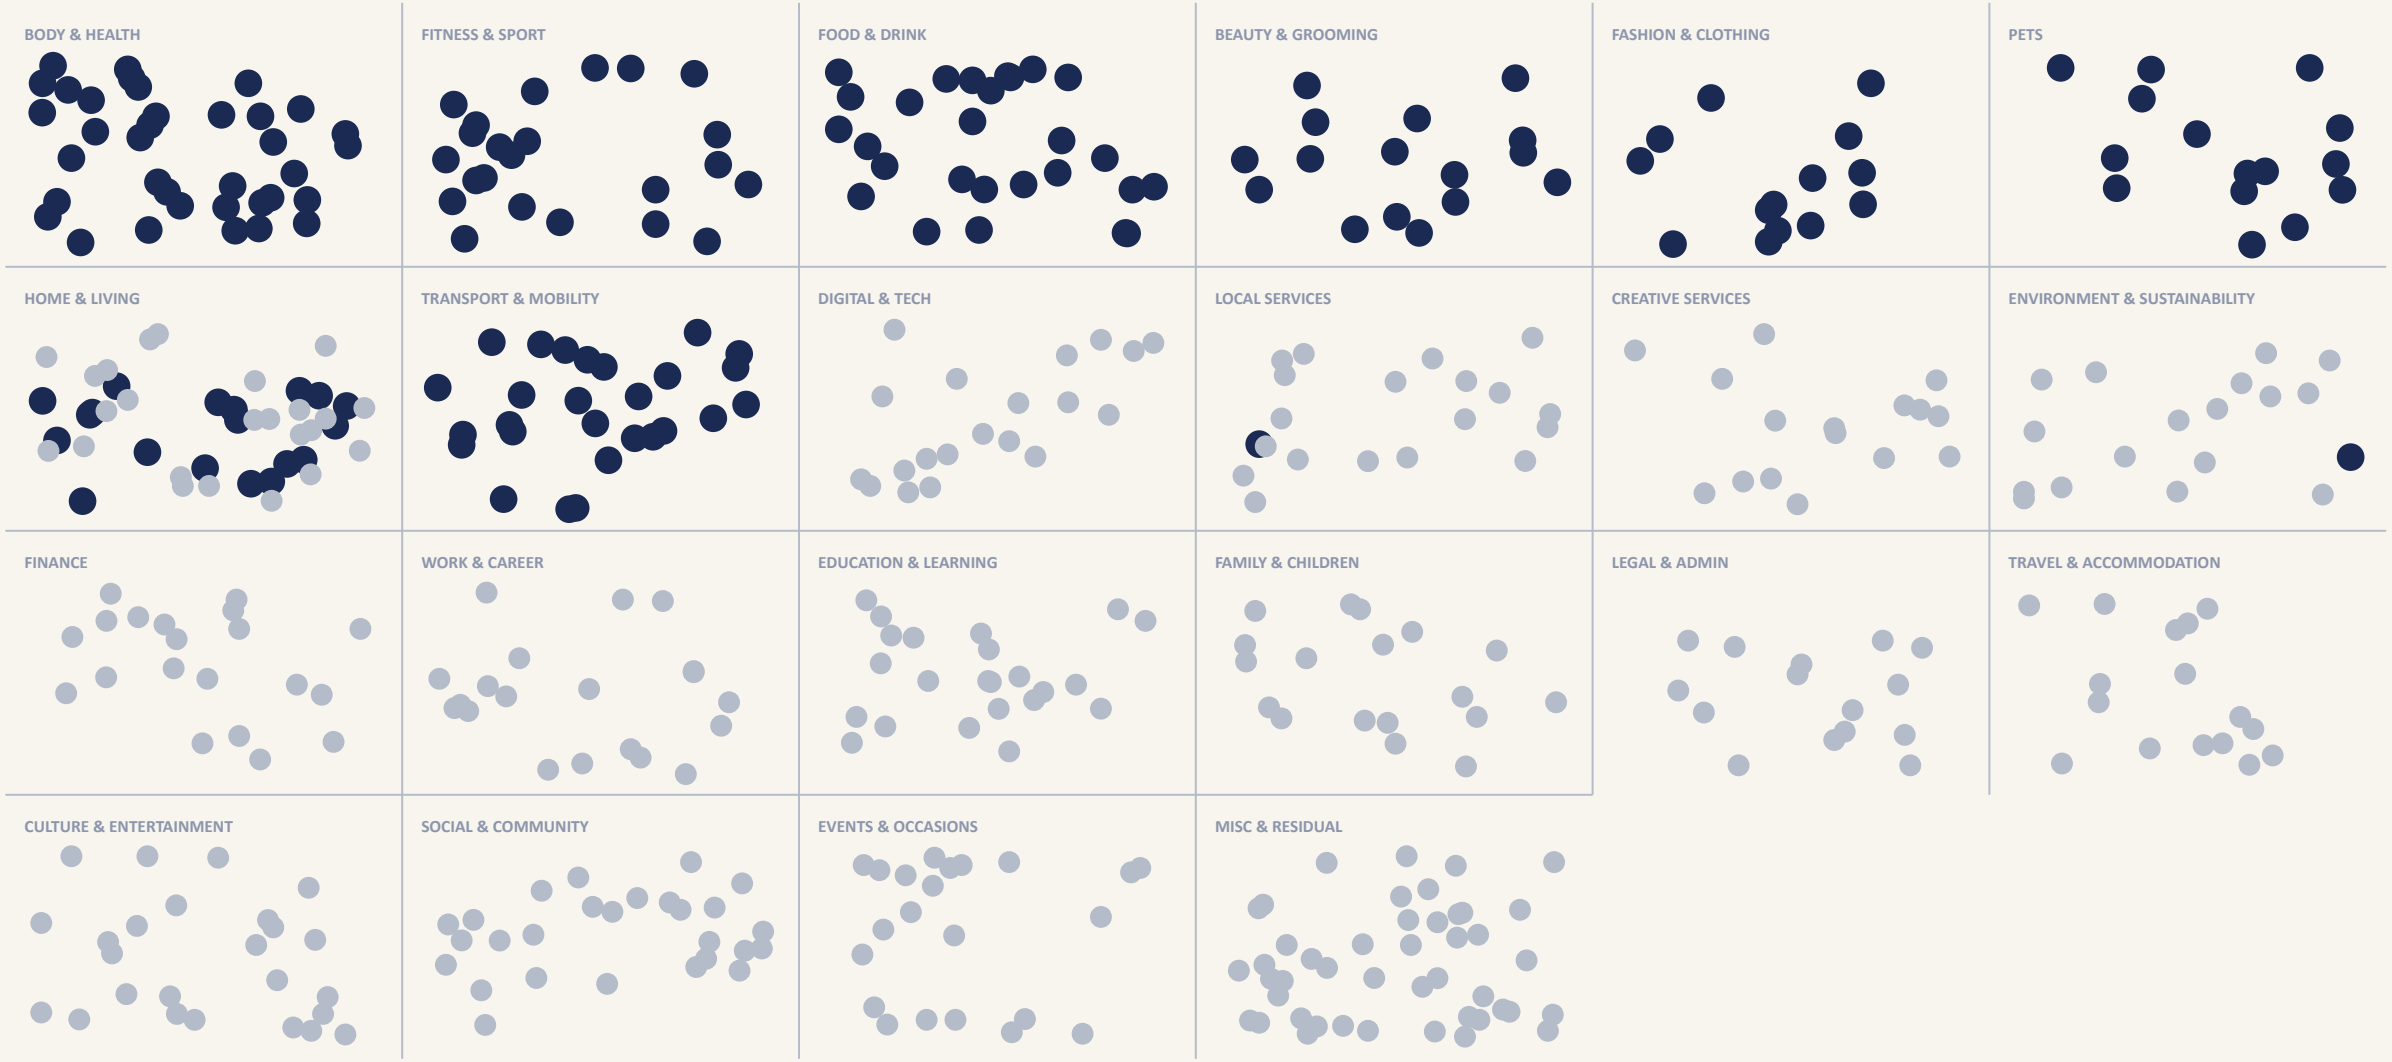

Early-stage density. The shape is visible but most nodes have <5 edges.

10,000 routing events / day.

Cluster structure firms up. Verified head heavily lit. Tail still gaining.

QUERIES / DAY

10K

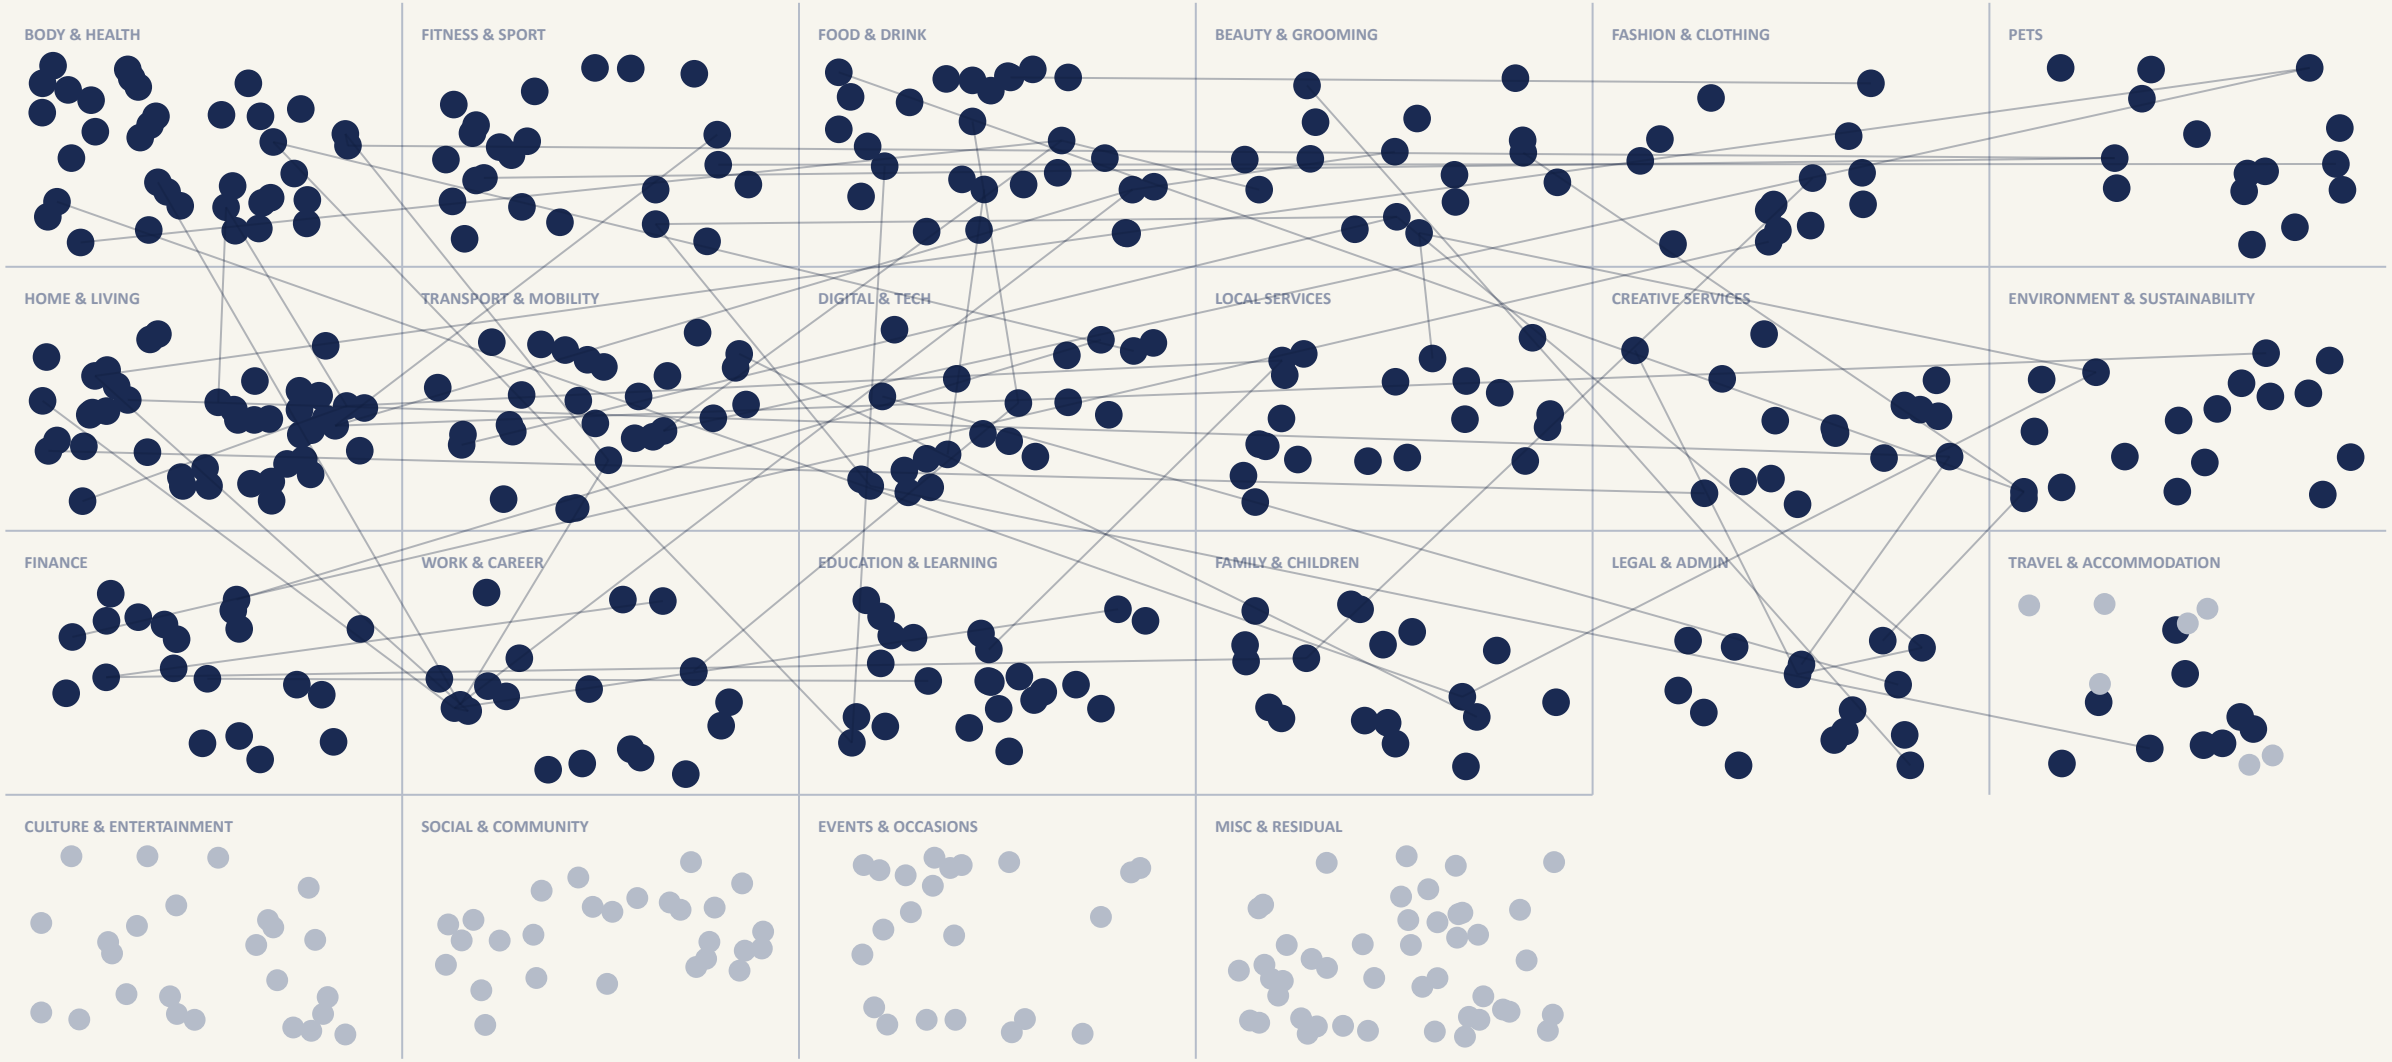

Edge density doubles roughly every order of magnitude in query volume.

1 million routing events / day.

Almost every node is active. Cluster interiors thick with edges. Tail catching up to head.

QUERIES / DAY

1M

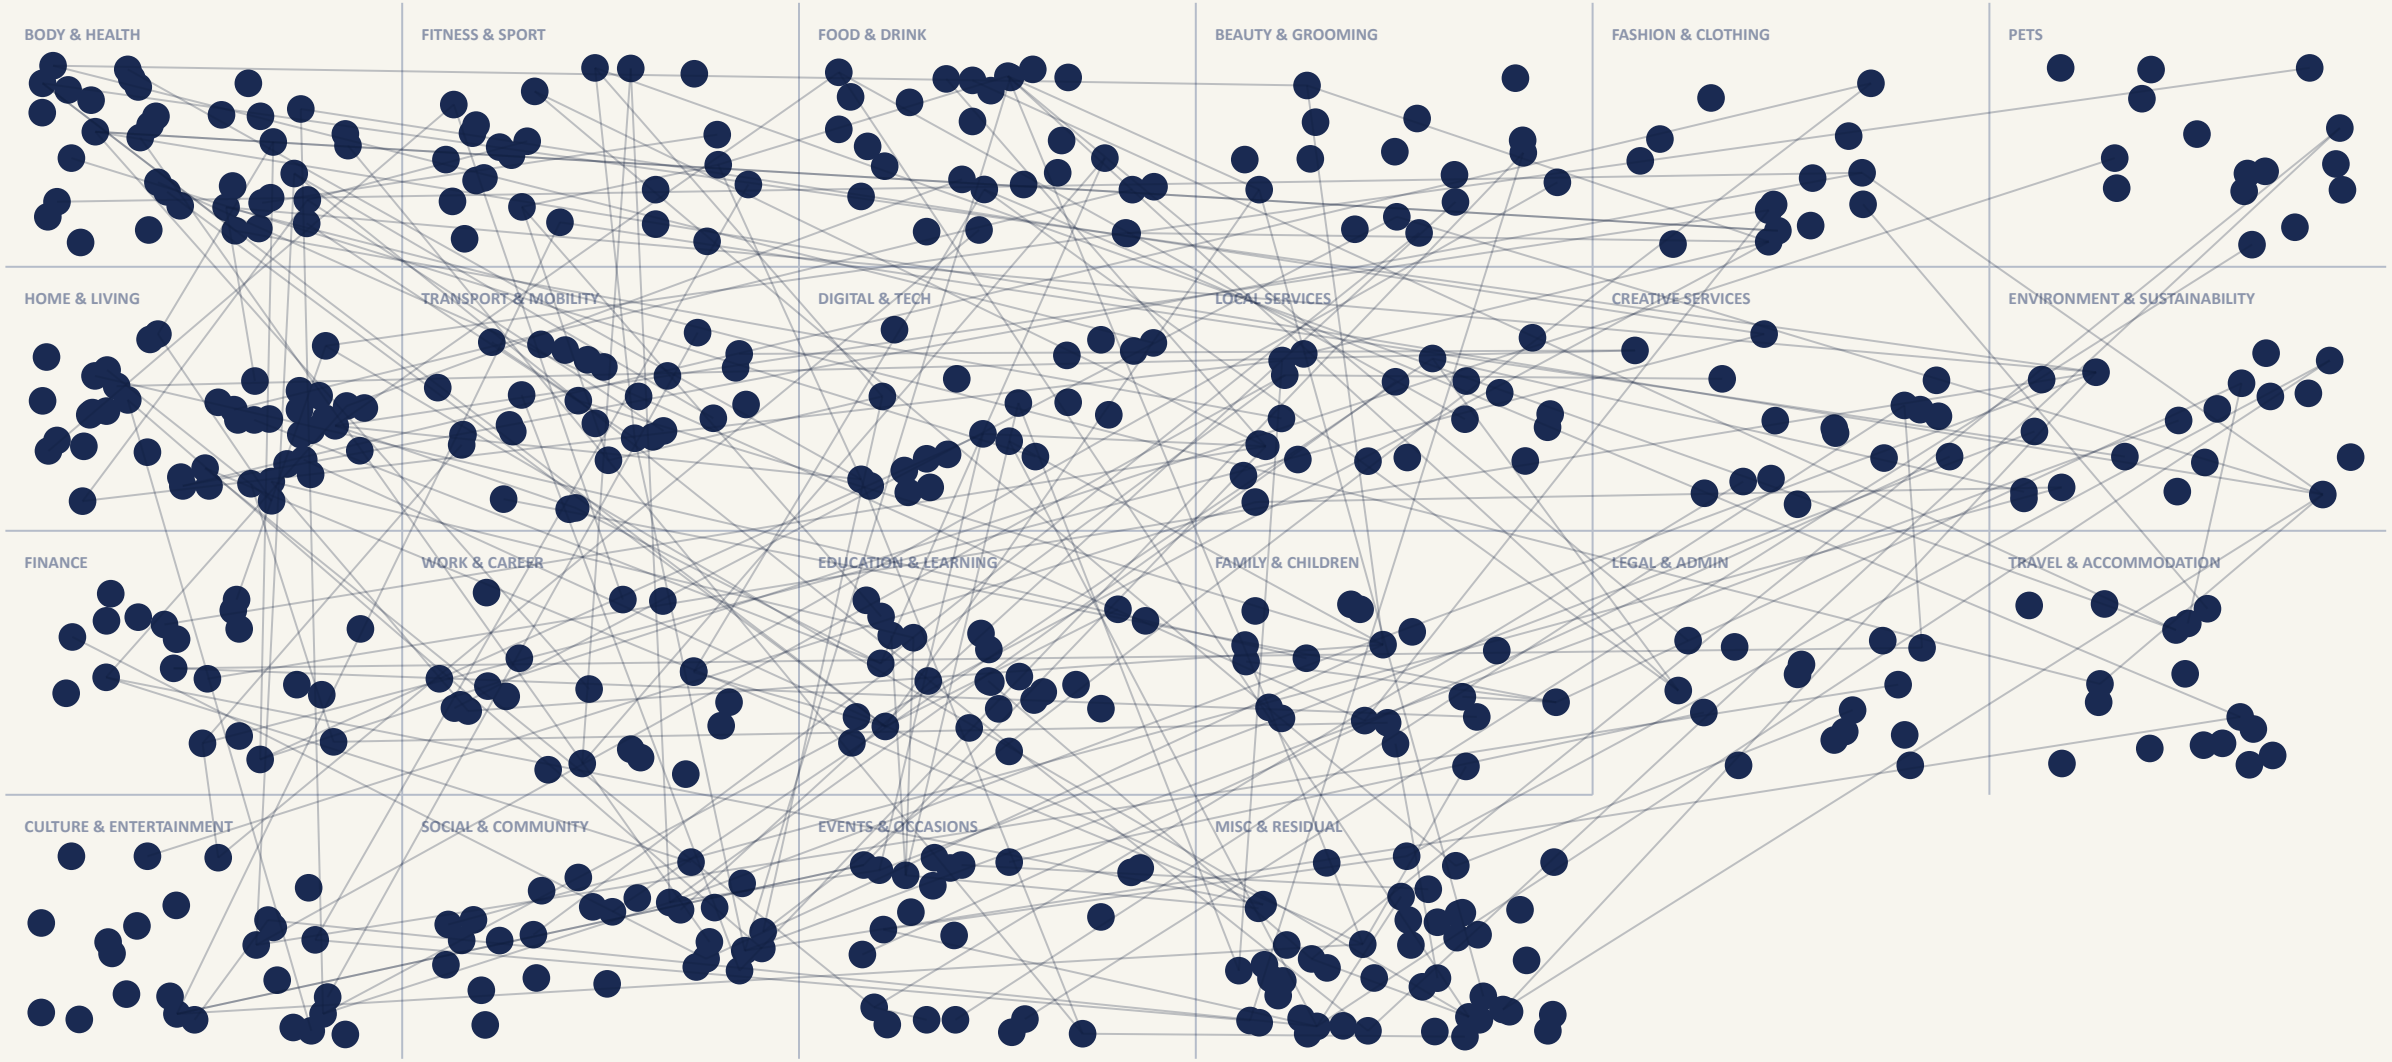

The graph stops looking like a list of nodes. It starts looking like a network.

100M routing events / day.

The graph is opaque with edges. Each cluster is a bright core.

QUERIES / DAY

100M

At scale, the lattice IS the product. The graph is the product.

One node at scale.

DOCTOR with 100k+ provider edges, 6 constraint families, 40+ language facets.

DOCTOR EDGES

100K+

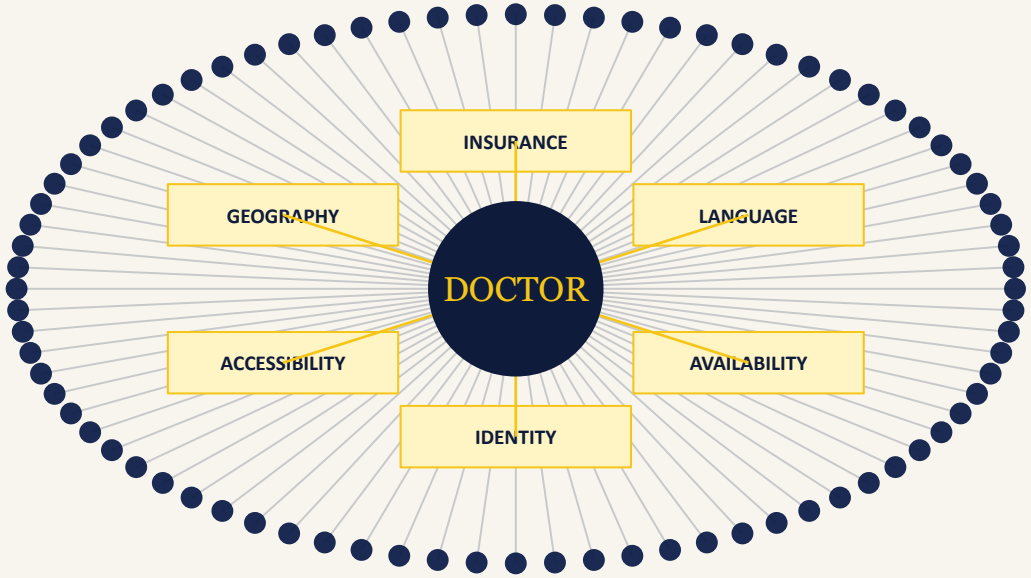

Every node looks like this at scale. 506 of these in one graph.

What the operator sees.

506 cluster cores. ~10 trillion edges. One graph.

GRAPH

MATURE

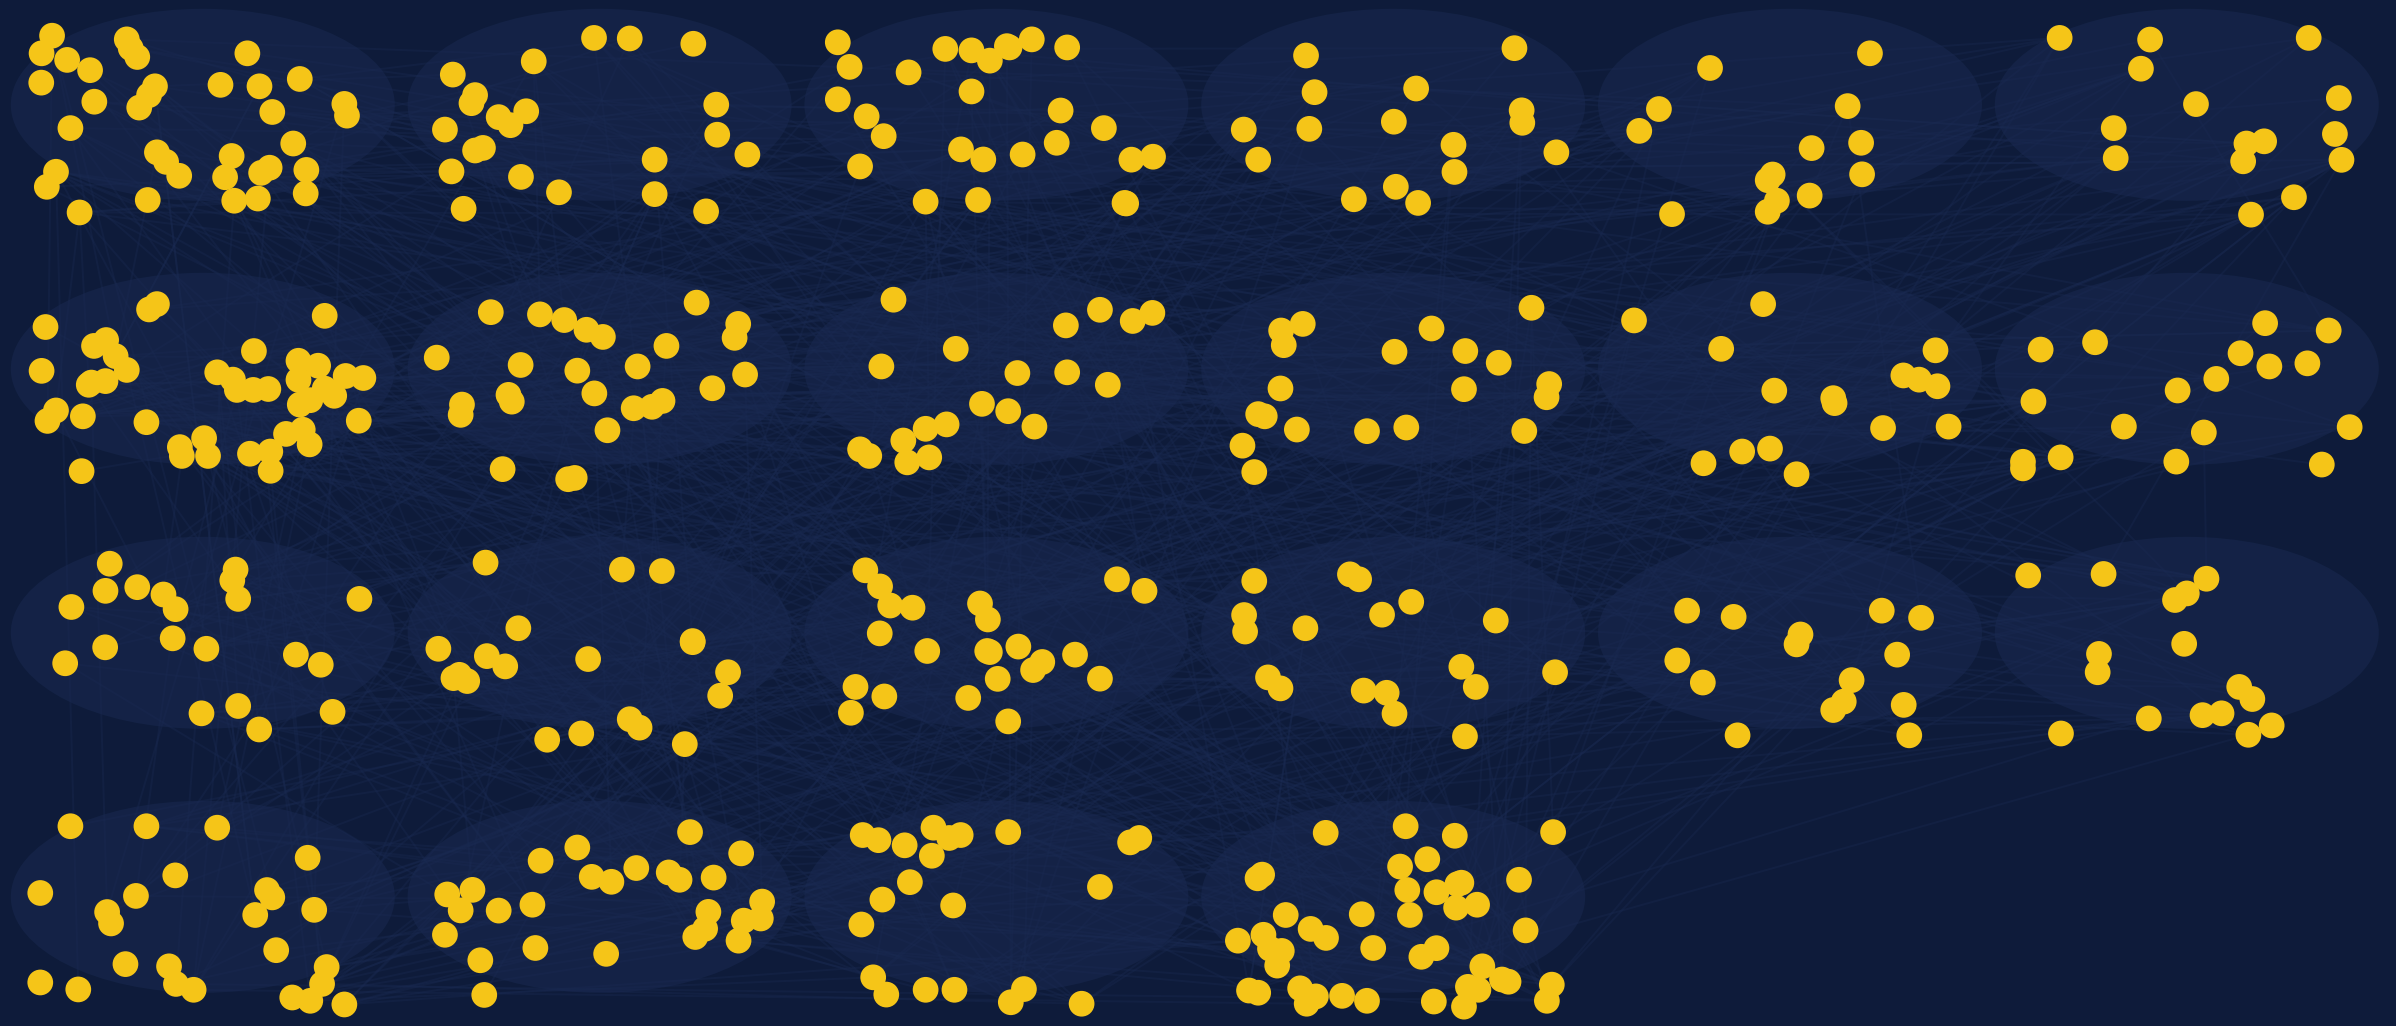

What every previous slide was building toward.

What you ship.

The graph is a queryable API. In, an intent. Out, a ranked list with constraints applied.

INPUT · NATURAL LANGUAGE

“who else delivers vegan food past midnight in DC?”

GRAPH · SHEAF OF EDGES

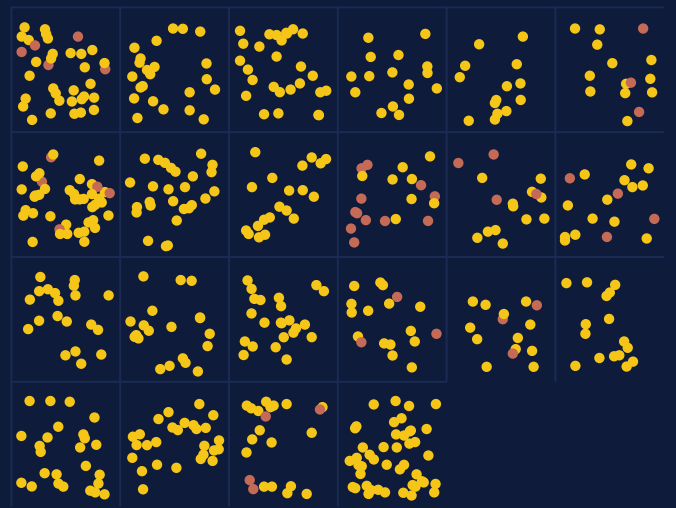

OUTPUT · RANKED PROVIDERS

1. Beyond Sushi — 0.8 mi
2. Plant Bar — 1.2 mi
3. Loving Hut — 2.1 mi
4. HipCityVeg — open till 1am
5. Fancy Radish — pickup only

The graph is the contract. Operators query against it. The taxonomy is the schema.

Why it can't be replicated.

Cold-start is structurally impossible. The graph is the data, the data is the graph.

01

Graph density compounds.

Each routing event adds one or more edges. Edges multiply over time. New entrants start at zero, against operators with billions.

02

Long-tail data is unique.

Verified categories (DoorDash, Doctolib) are knowable from public sources. Fragmented data (locksmith, calligrapher, beekeeper) exists ONLY in the graph.

03

DIN SPEC 2343 is open.

Anyone can use the taxonomy. Only one entity has the operational graph. The standard is a moat by enabling, not gating.

THE THESIS

We don't build apps. We grow a graph.

Every routing event makes the next one cheaper to answer. The graph is the asset. The graph is the moat. The graph is the company.

WHOELSE.AI

Watch it grow.
Then own the graph.

30 frames. One graph. From the first query to a 506-node lattice with ~10 trillion edges. The same flywheel runs at every scale.

TOBIAS MARTENS

tm@whoelse.ai · +1 202 820 2151

WASHINGTON DC · SEED 2026