

# Riak Search

John Muellerleile  
Basho Technologies

Erlang Factory 2010  
San Francisco

# Agenda

- Introduction
- A Search Tale
- Riak Search

# “I used to make stuff”

(but now I am a database hacker)

- You are not your primary key
- Having to roll your own indexes = not OK\*
- Time spent was not on *developing your app*

\* choosing to is different matter altogether

# A moment of weakness

“My way of joking is to tell the truth;  
it’s the funniest joke in the world”

George Bernard Shaw

# How do I query it?

## Fault-tolerance



# A Search Tale

DIY Search in Three Acts

# Act I: “I Love Lucene”

- “Life is good!”
- Dedicated server + backup
- Not growing quickly, relatively static data set
- Fast & predictable

# Act 2: “Cluster Luck”

- “I love a good challenge!”
- A few shards+backups, some clever scripts
- Performance is good enough
- Nobody notices when the indexing master fails, it’s just a couple of documents
- Amazon EC2 appreciates your business



# Act 3: “SNAFU”

- Lots of shards
- Operational nightmare
- Diminishing returns: indexing, querying
- High “hit by bus” factor
- Scripts no longer viewed as “clever”
- Amazon EC2 holiday cards are now hand-written

# Where does it hurt?

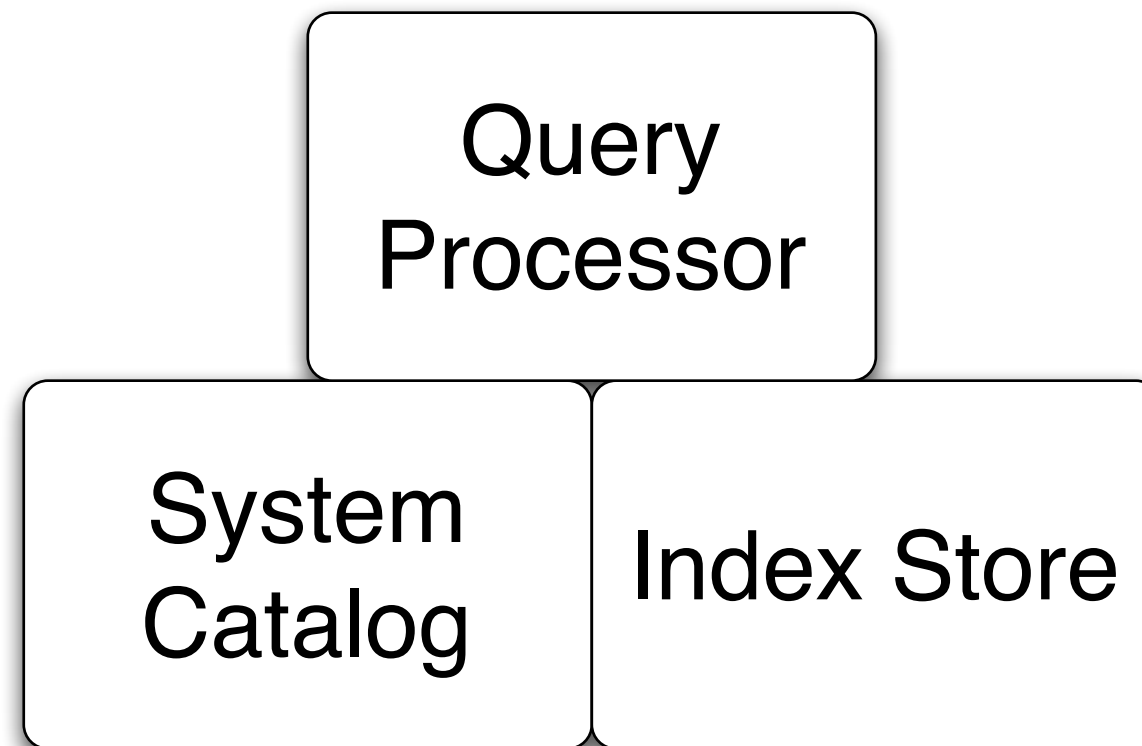
- Complexity: more shards, more problems
- Failure: cases range from “oops” to “*Epic*” (and are not mutually exclusive)
- Diminishing returns: indexing, querying
- Increasing cost: Ops, Dev, Opportunity

# Riak Search Goals

- Decentralized: no SPOF
- Distributed: only worry about local data
- Homogeneity: all nodes do the same thing
- Value: Adding nodes adds performance & capacity
- Flexibility: Query platform that's efficient and extensible

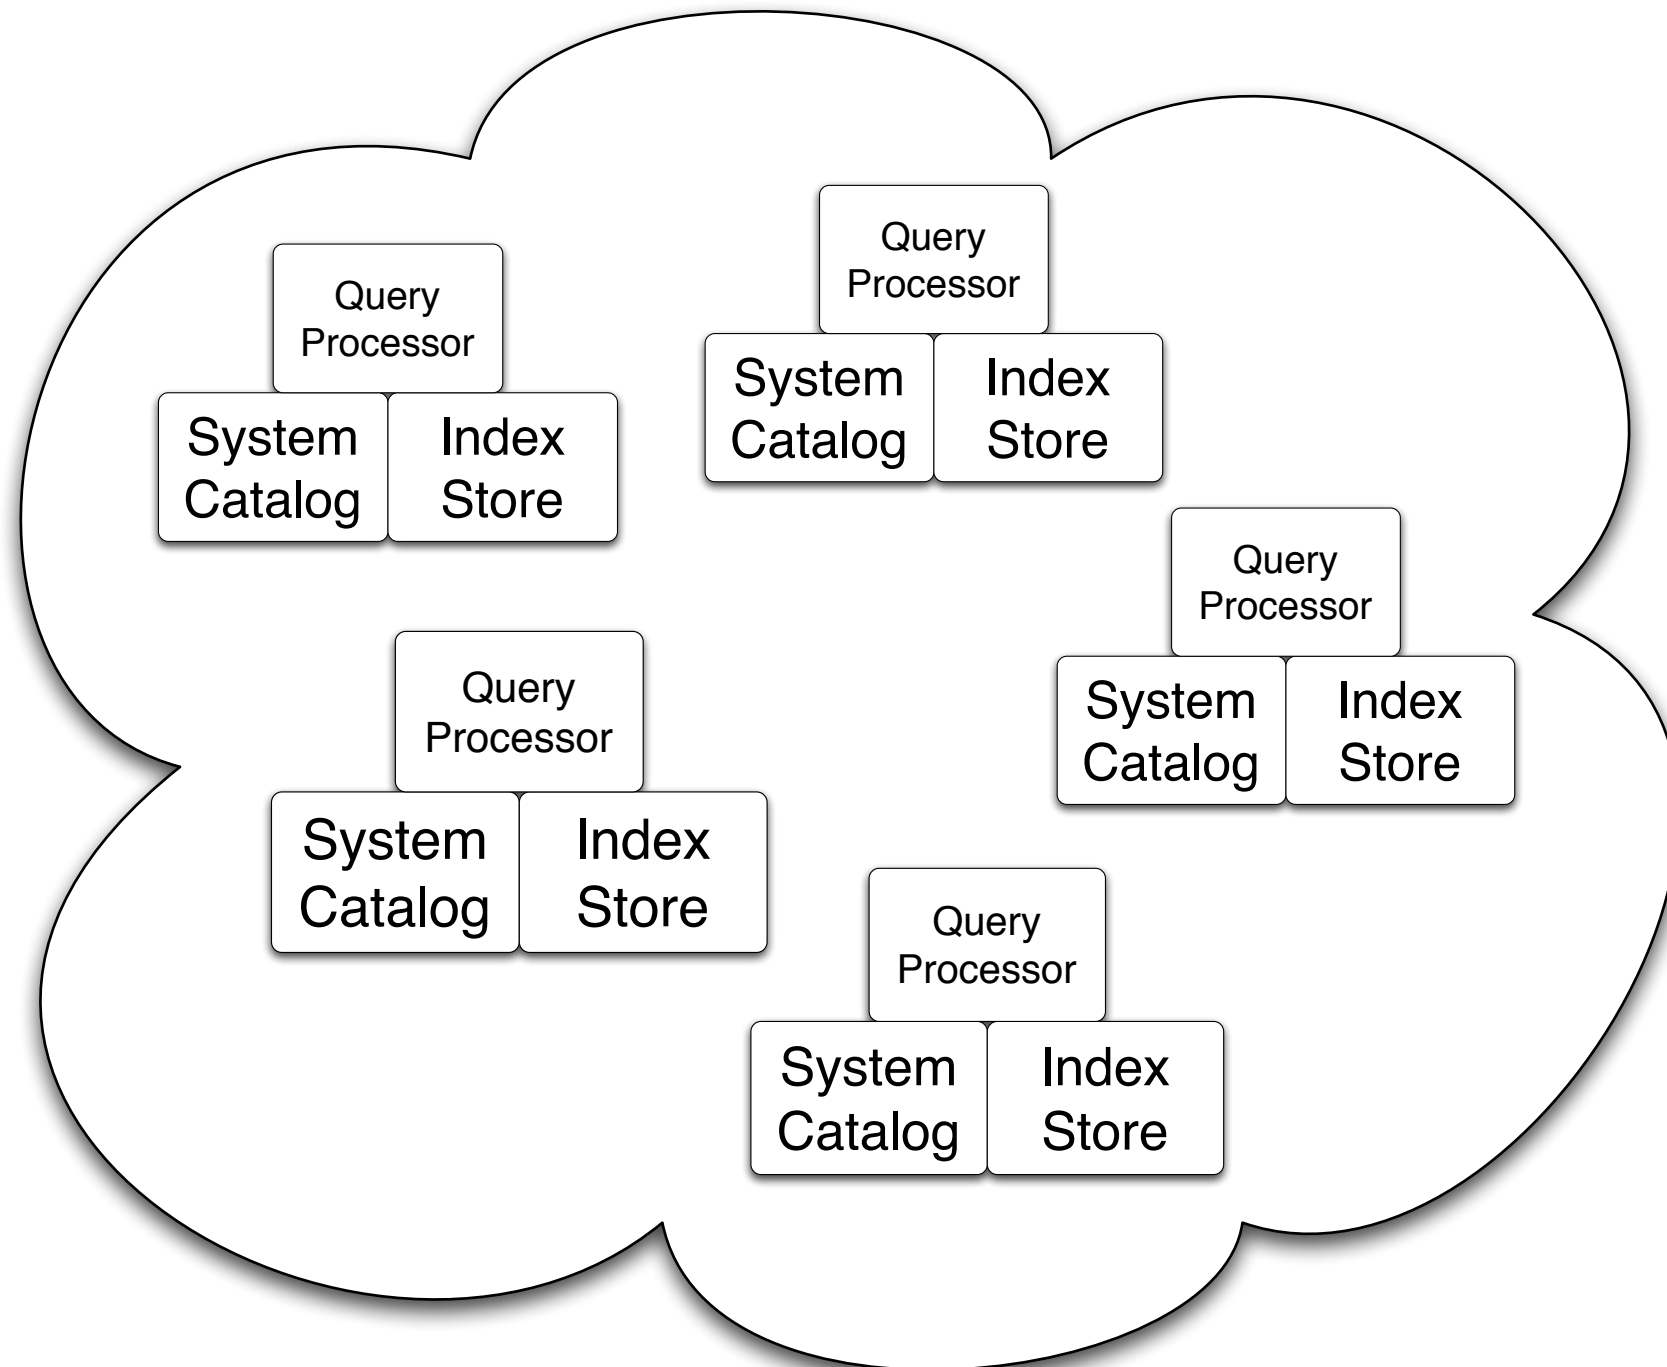
# Architecture

## Riak Search Node



Keep it simple & foster emergence

# No Unique Snowflakes



# Index

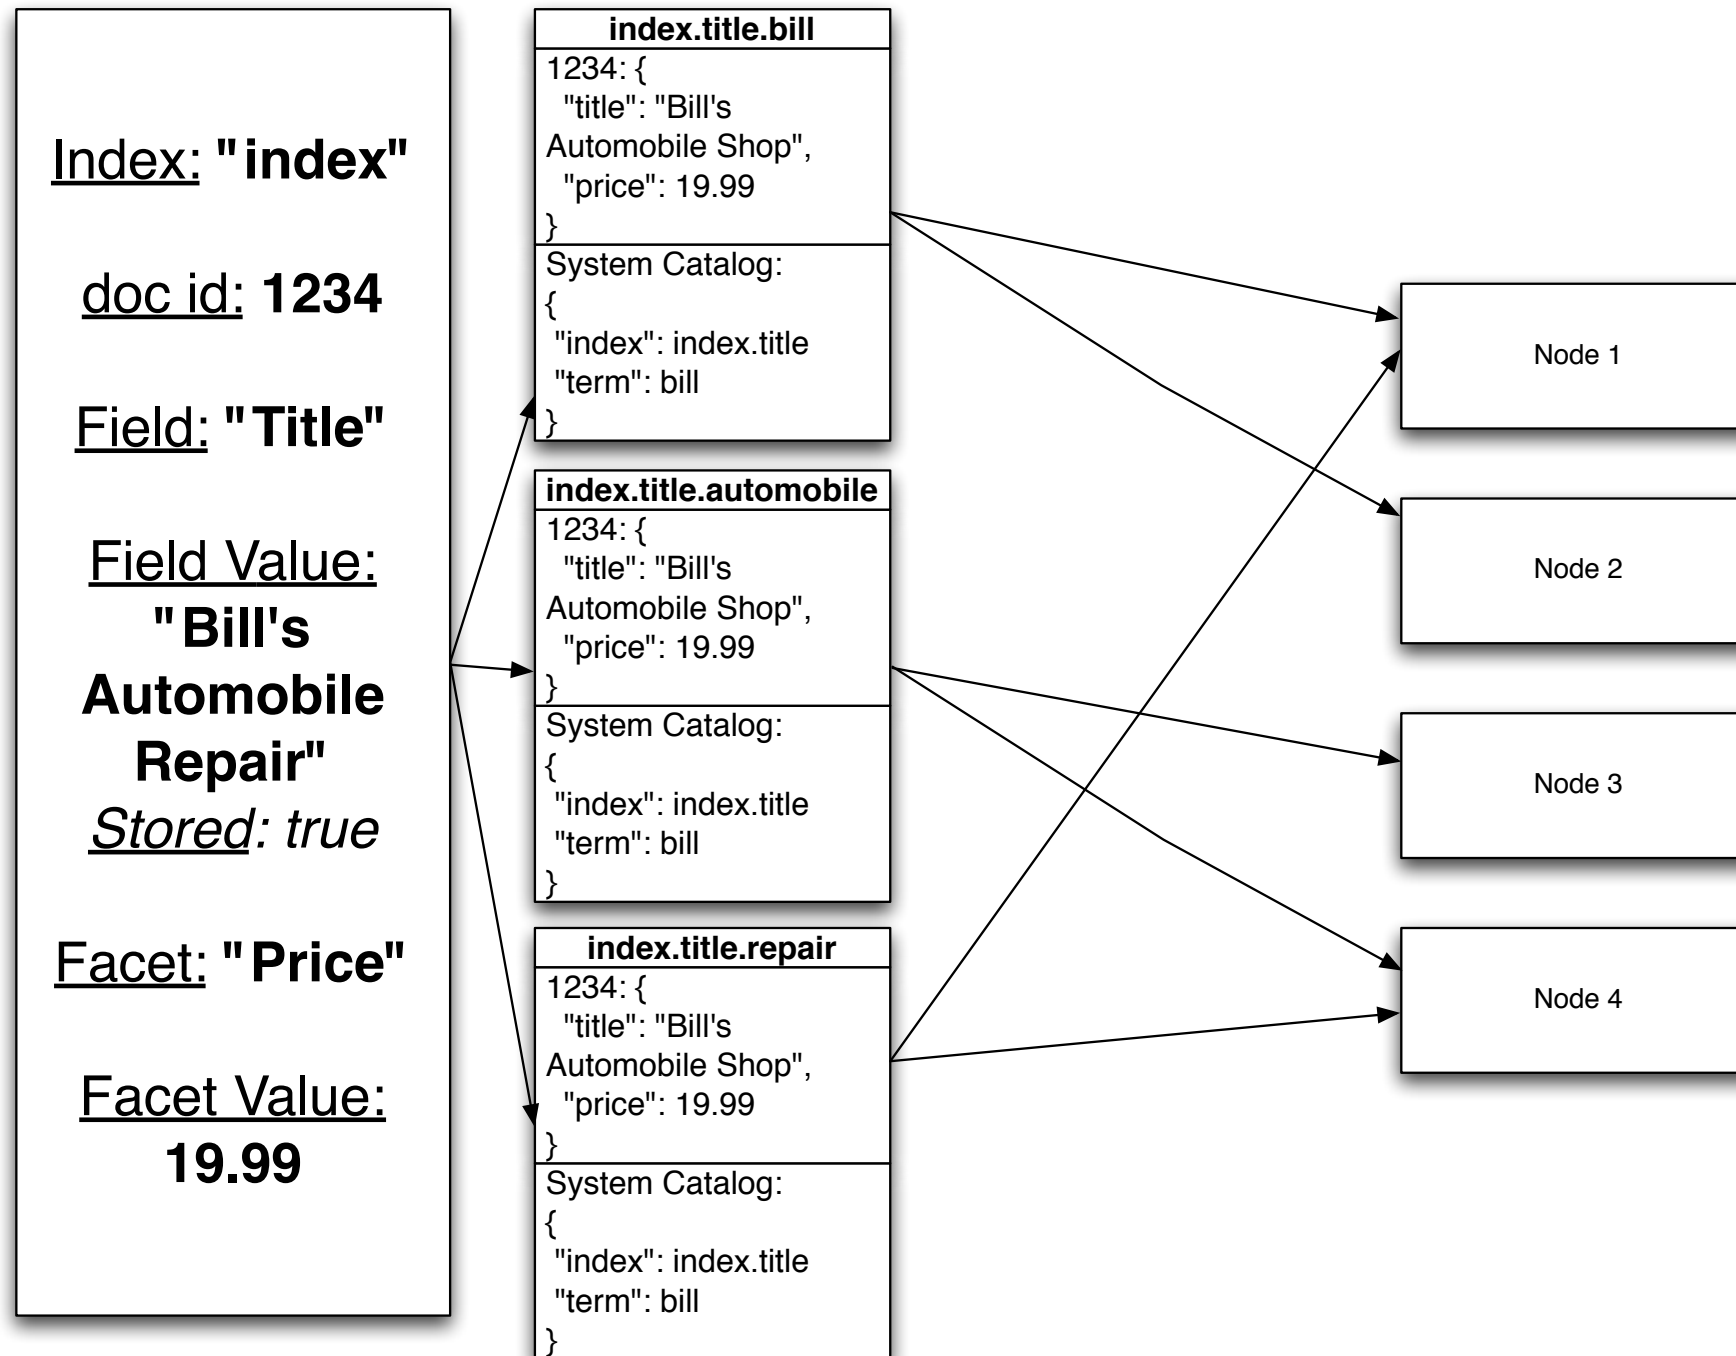
- “Column-*esque*”, use as many as you want
- Every column has a primary key
- Every key points to at least a “doc id”
- Can also store “facets”

# System Catalog

- Searchable index of node-local columns
- Can contain arbitrary metadata
- Built to be small & fast; always in-memory
- e.g., query:  
`+index:"products"`  
`+term:auto*`  
`+category:"books"`  
**result:**  
`products.automobile,`  
`products.auto, products.automata,`  
`products.autoclave, [...]`

# Indexing Crash Course

Tokenize ➡ Term ➡ Column ➡ Store





# Query Processor

- Parser
- Query Intermediate Language
- Planner
- Executor

# Parser

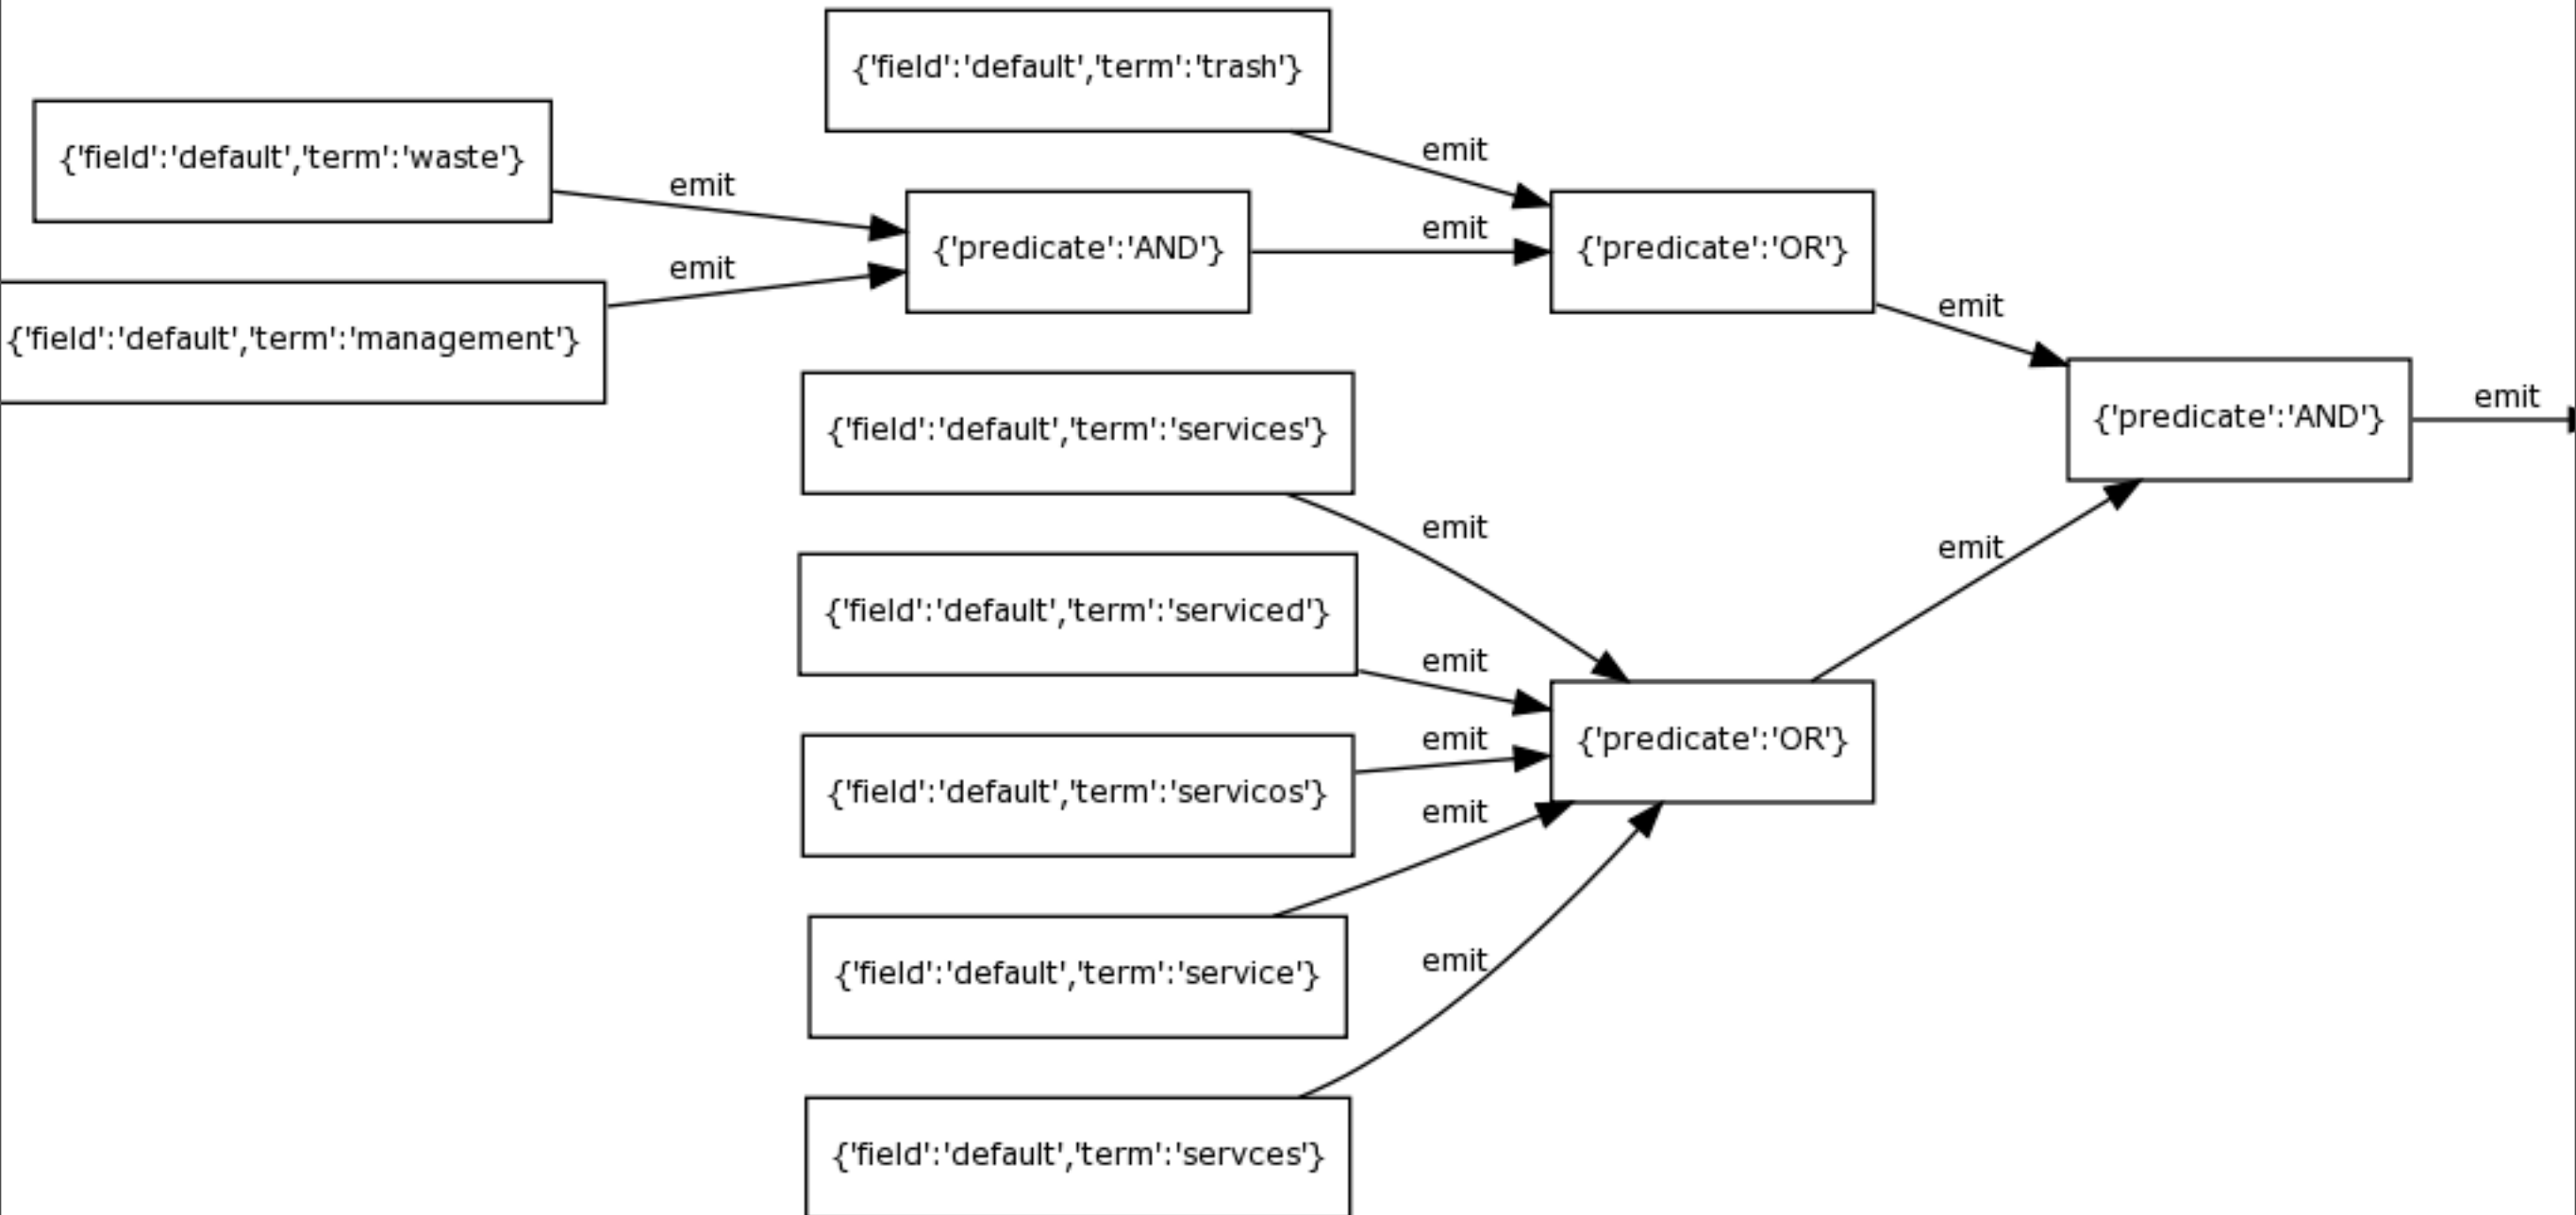
- Parsers produce networks of “*primitives*” the planner and executor can deal with
- Implement multiple query languages
- Use different query languages with the same data in many cases
- Lucene syntax is implemented

# QIL

## Query Intermediate Language

- QIL is used to express data flows
- A network of simple components:  
*producers, filters and accumulators*

# You like DAGs?



# QIL Vertices

- Producers: stream node-local data  
(`get_all, ...`)
- Filters: *AND, OR, NOT, ...*
- Accumulators: collect and/or stream results

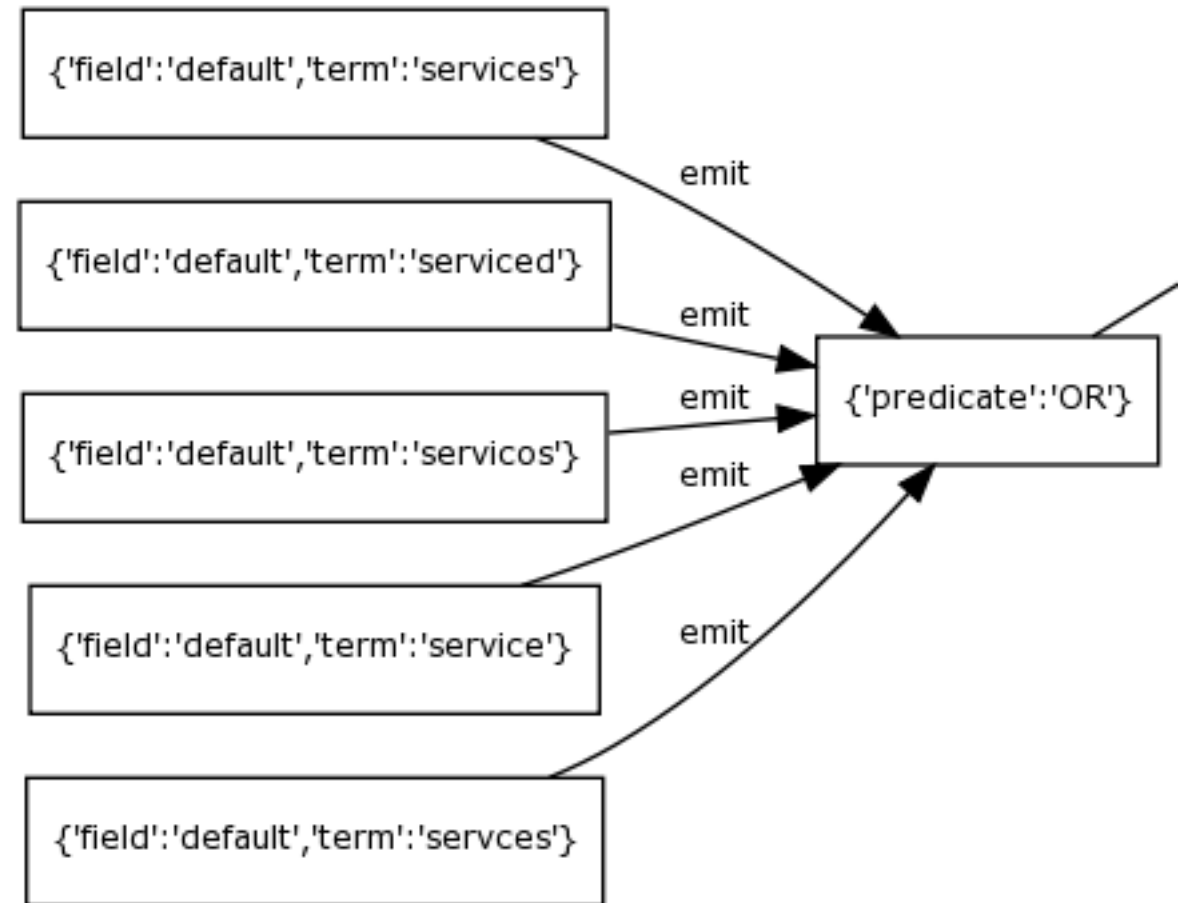

# QIL Edges

- A stream of tuples
- One format

# Expansions

## A Fuzzy Decomposition

services~0.8



# Expansions

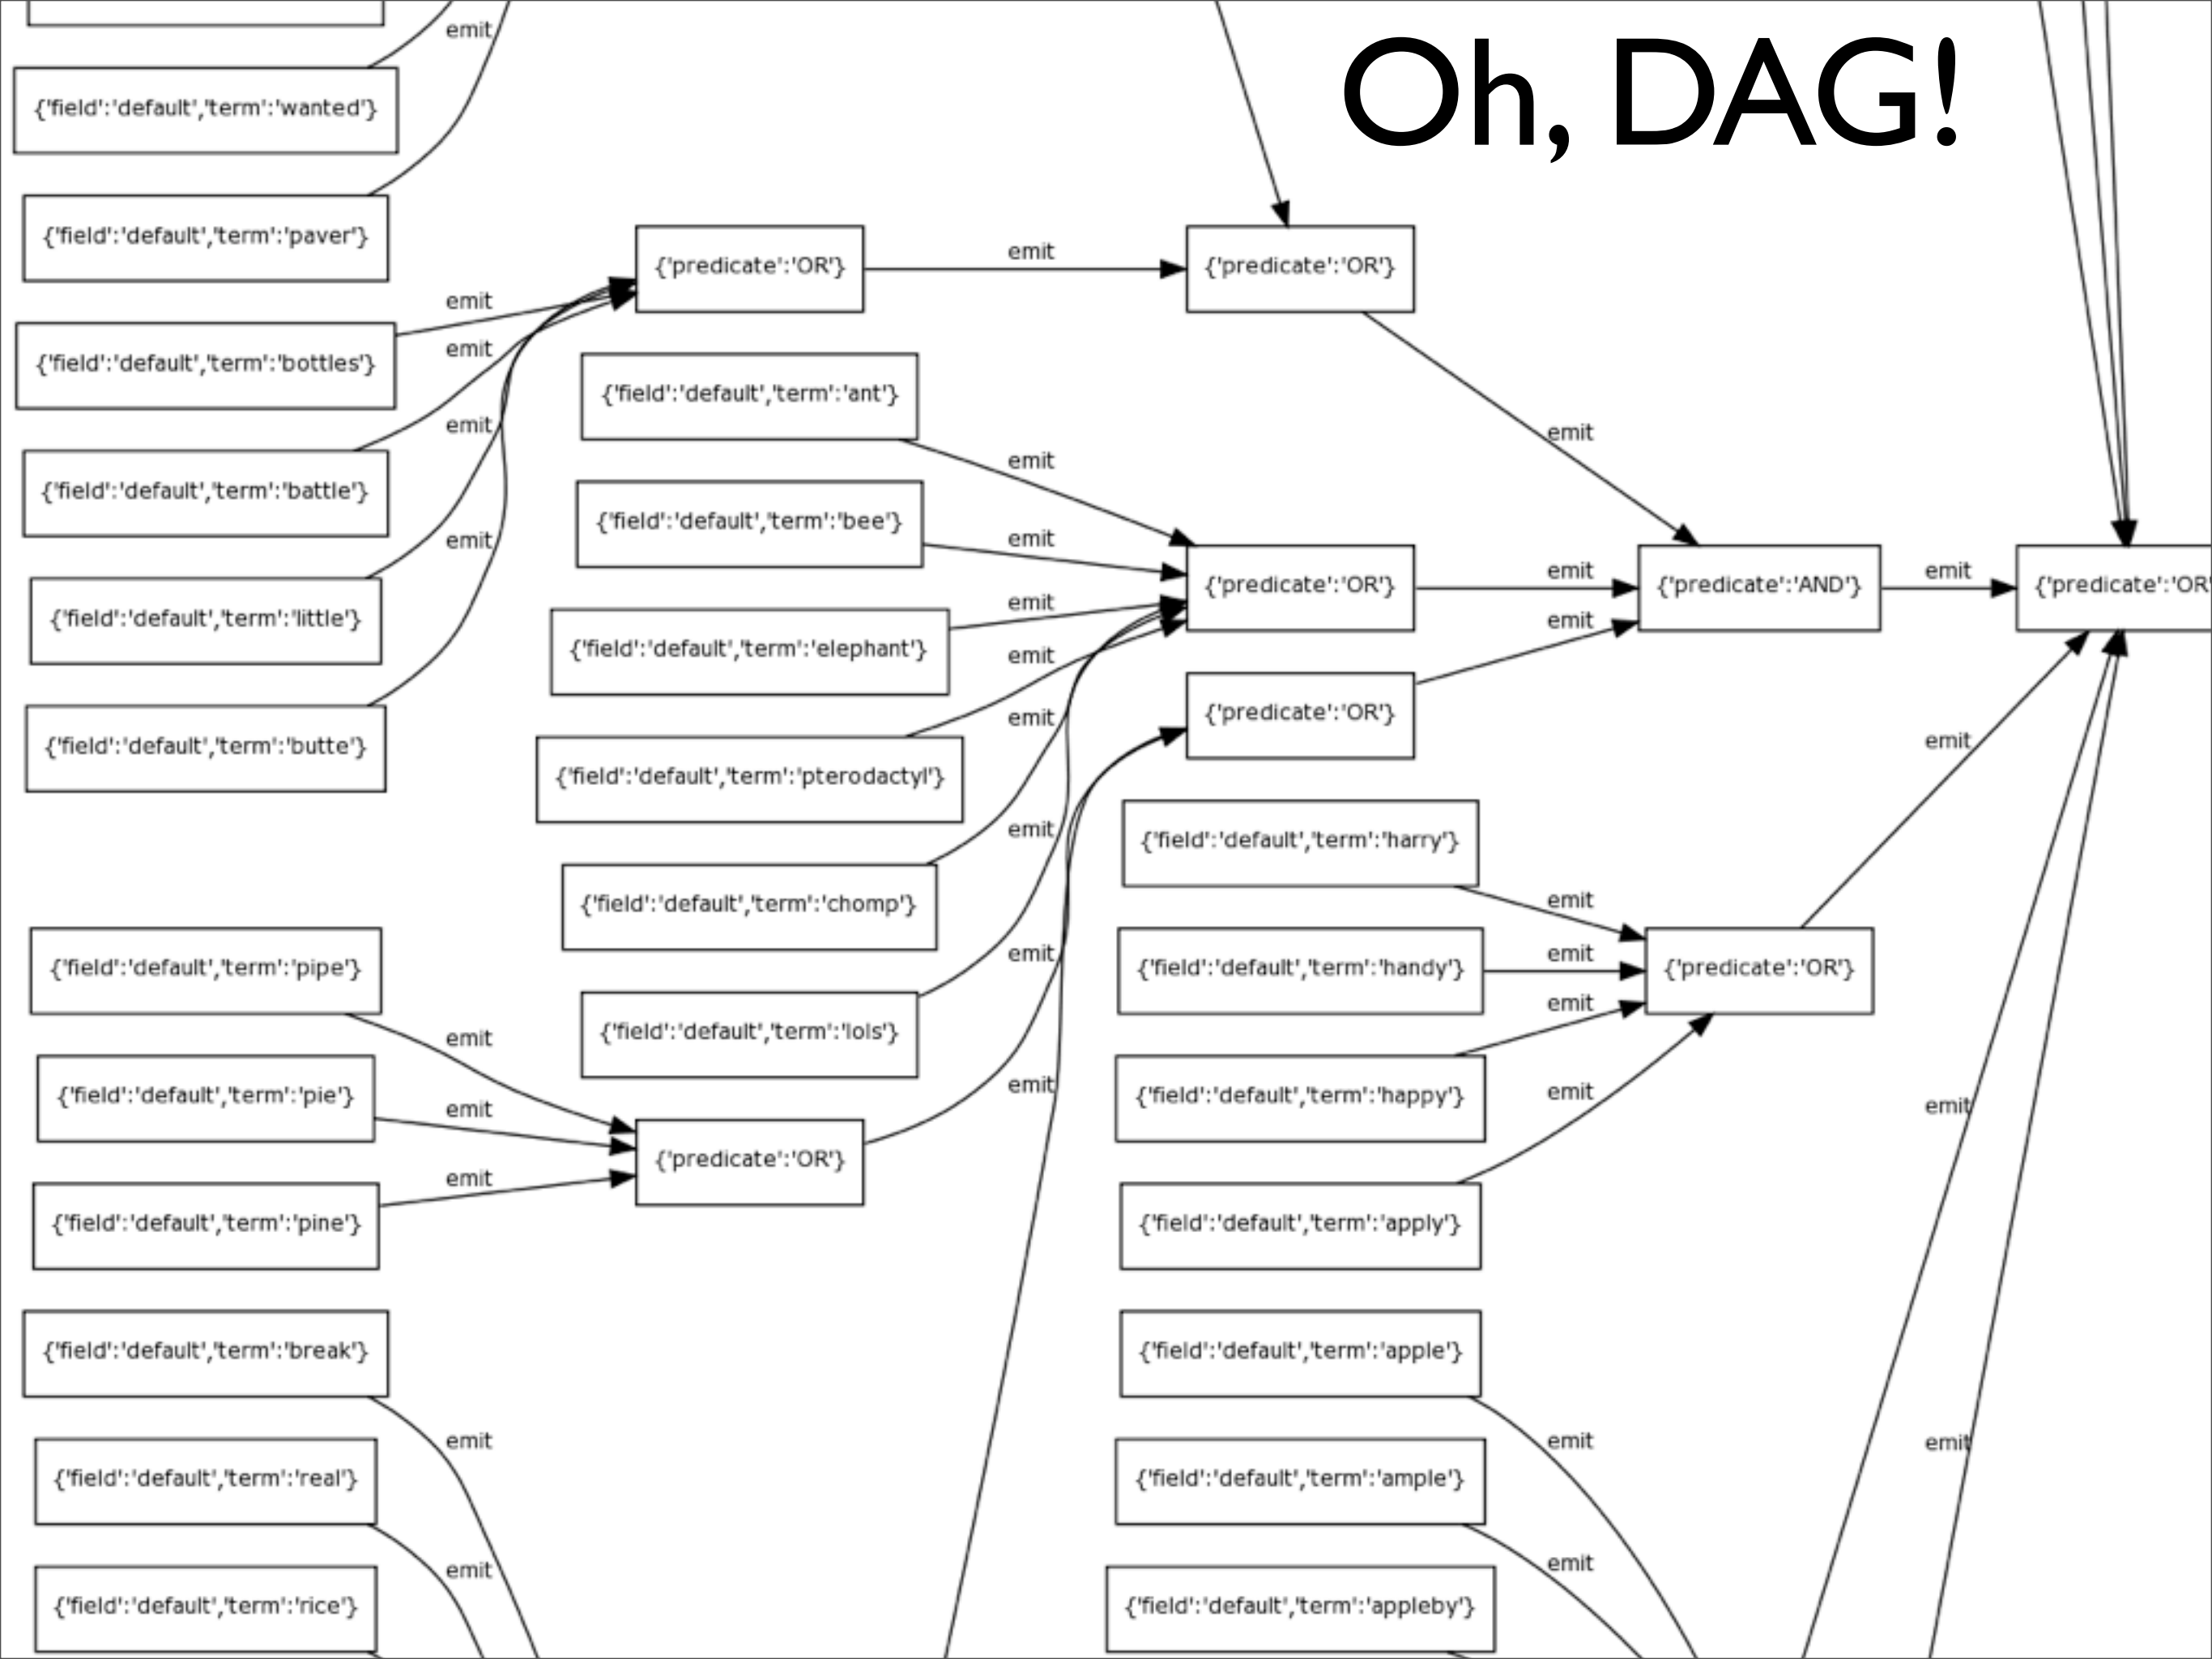
- *Range* ➡ “OR” of all existing indexes in the range
- *Wildcard, regex* ➡ “OR” of all existing columns that match
- Expansion is an easy way to express a *relationship*, not just string variations (related terms, synonyms, etc.)
- Expansions can be applied ad-hoc



# Keep It Simple & Foster Emergence

- All components produce and/or expect the same messages
- Easy to reason about, extend
- *Networks of primitives built up by expansion produce complex behaviors*

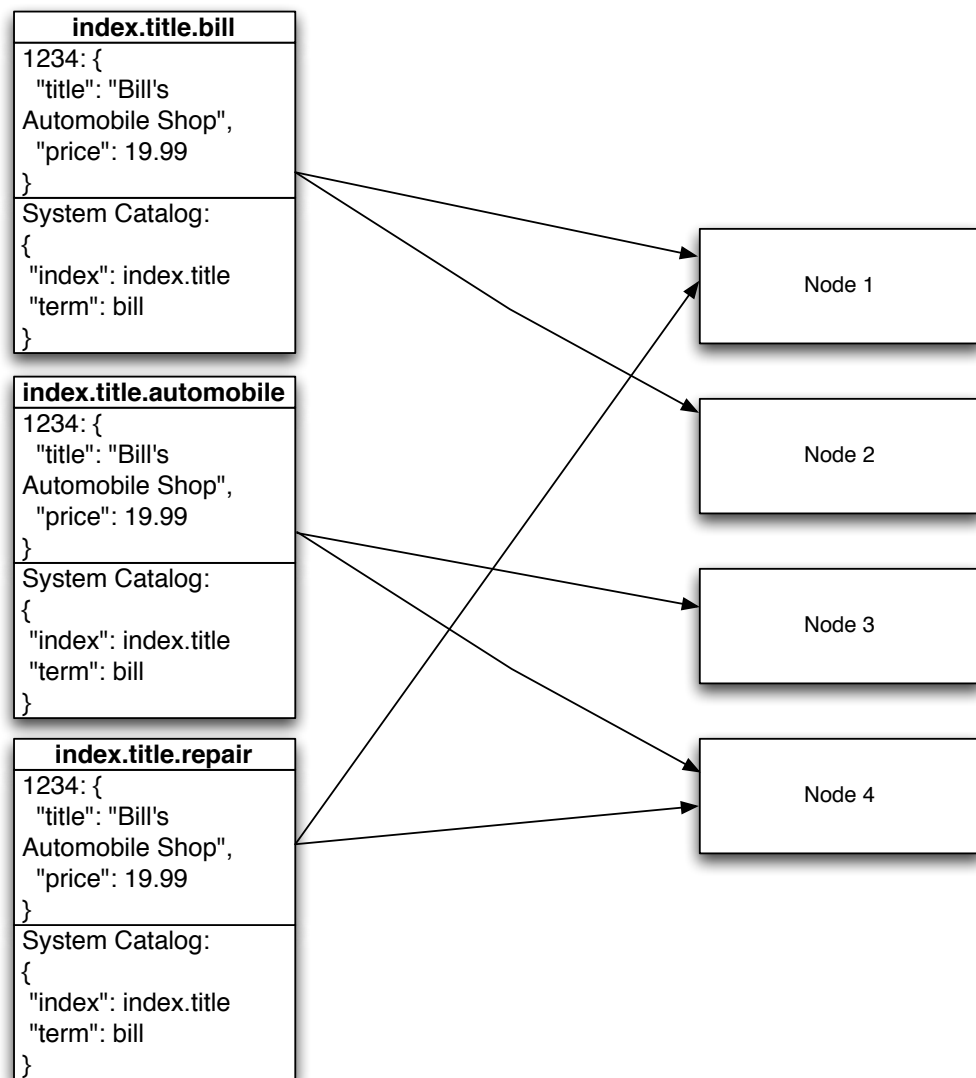
# Oh, DAG!



title:bill

+title:auto

+price:[18.00 TO 20.00]



**1. Locate columns**

**2. Plan query**

2a. optimize for co-located node-local data

2b. optimize filter process placement

**3. Execute plan**

3a. start accumulators

3b. start filters

3c. start producers

# QIL-ler Features

- Designed to be extended, customized
- Built to pipeline
- Lots of opportunity for optimization,  
*e.g., logically factor: favor node-coincident indexes for expansions, even partially, to reduce unnecessary tuple streams*

# Why Erlang?

- Queries execute as process networks
- Each process an automaton, shared nothing
- Stability in a storm of short lived processes
- Sane ways to reason about and manage them
- Erlang: a perfect fit!

# Project Status

- Prototype being supported for a small group of users, beta program is closed
- Production version has been in the works for awhile, though no availability date (yet)
- It will be open source (Apache)
- *I can't wait to get it into your hands!*

# Thanks!

- direct: johnm at basho dot com
- twitter: @jrecursive
- inquiry: riak@basho.com
- Basho Technologies: http://basho.com/
- search team: @rklophaus @kevsmith  
@hobbyist @dizzyco @jon\_meredith