

(The art of not shooting yourself in the foot using Elasticsearch)

Using Elasticsearch as the Primary Data Store

ElasticNL Amsterdam

2019-10-08

@ElasticNL

Volkan Yazici

<https://vlkan.com>

@yazicivo

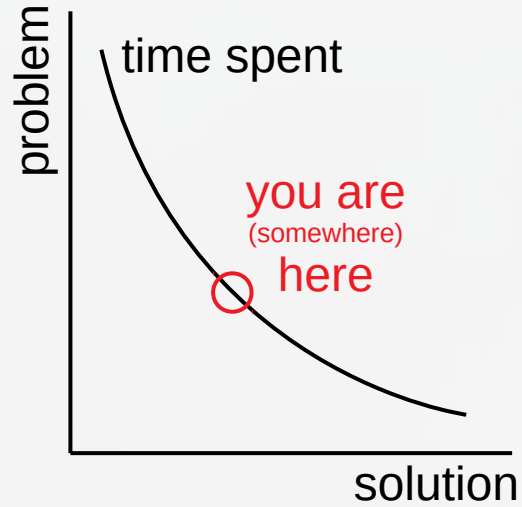
Poll time!

- Recently purchased an item online?
- Elasticsearch users?
- Elasticsearch users with 10+ node clusters?
- Updating Elasticsearch indices real-time?
- Has ever lost data in Elasticsearch?

Who am I?

- Volkan Yazici (volkan.com) – @yazicivo – github/vy)
- Java plumber in the domain of search (bol.com, since 2014)
- interested in networking & concurrency
 - OpenJDK Project Loom (aka. fibers/coroutines for JVM)
 - Reactive Streams (Reactor, RxJava)
- F/OSS contributor
 - log4j2-logstash-layout
 - HRRS (HTTP Request Replay Suite)
 - reactor-pubsub
- BS in math, MS and PhD in CS

Disclaimer





bol.com

9+ million active¹ clients²

17+ million articles²

200k+ sellers²

1500+ employees²

62+ million visits/month²

¹ Customers who ordered an item in the last 365 days.

² As of October 2018.

E-commerce search

- Search
 - Matching
 - Ranking
 - Faceting
- Guidance
 - Suggestions
 - Auto-corrections
 - Recommendations

The screenshot shows the Bol.com search results for 'wijn glazen'. The page features a navigation bar with the Bol.com logo, a search bar containing 'wijn glazen', and various utility links like 'Inloggen' and 'Winkelwagen'. Below the navigation, there are promotional banners and a 'Kies een categorie' section with five options: 'Wijn glazen', 'Rode wijnglazen', 'Witte wijnglazen', 'Glazen', and 'Champagne gla'. The main content area is titled ''wijn glazen' in Alle Artikelen' and shows 1,268 results. Two product listings are visible:

- Leonardo Puccini Witte Wijn glas - 6 Stuks**: Vaatwasserbestendig | 6 stuks | 560 ml | Glas. Price: 24,-. Rating: 4.5 stars (39 reviews). Description: 'Het ontwerp van deze Puccini wijnglazen voldoet niet alleen aan de huidige tijdgeest, maar de brede kelk zorgt er ook voor dat de... Meer'. Status: 'Op voorraad' (In stock).
- Leonardo Puccini Rode wijnglas - 0,75 l - 6 stuks**: Vaatwasserbestendig | 6 stuks | 750 ml | Glas. Price: 30,96. Rating: 4.5 stars (26 reviews). Description: 'Het ontwerp van deze Puccini wijnglazen voldoet niet alleen aan de huidige tijdgeest, maar de brede kelk zorgt er ook voor dat de... Meer'. Status: 'Op voorraad' (In stock).

Who is using search?

- Customers
- Sellers
 - via web
 - via API
- Bots
 - search engines (Google, Bing, etc.)
 - competitors
- Internal services

Search input

- **Product attributes** (title, EAN, ISBN, color, etc.)
- **Seller offers** (price, availability, deliverability)
- **Derived content** (for ranking)
 - Sale popularity
 - Price quality
 - Customer feedback (reviews, etc.)
- **Configuration** (faceting, value translations, etc.)

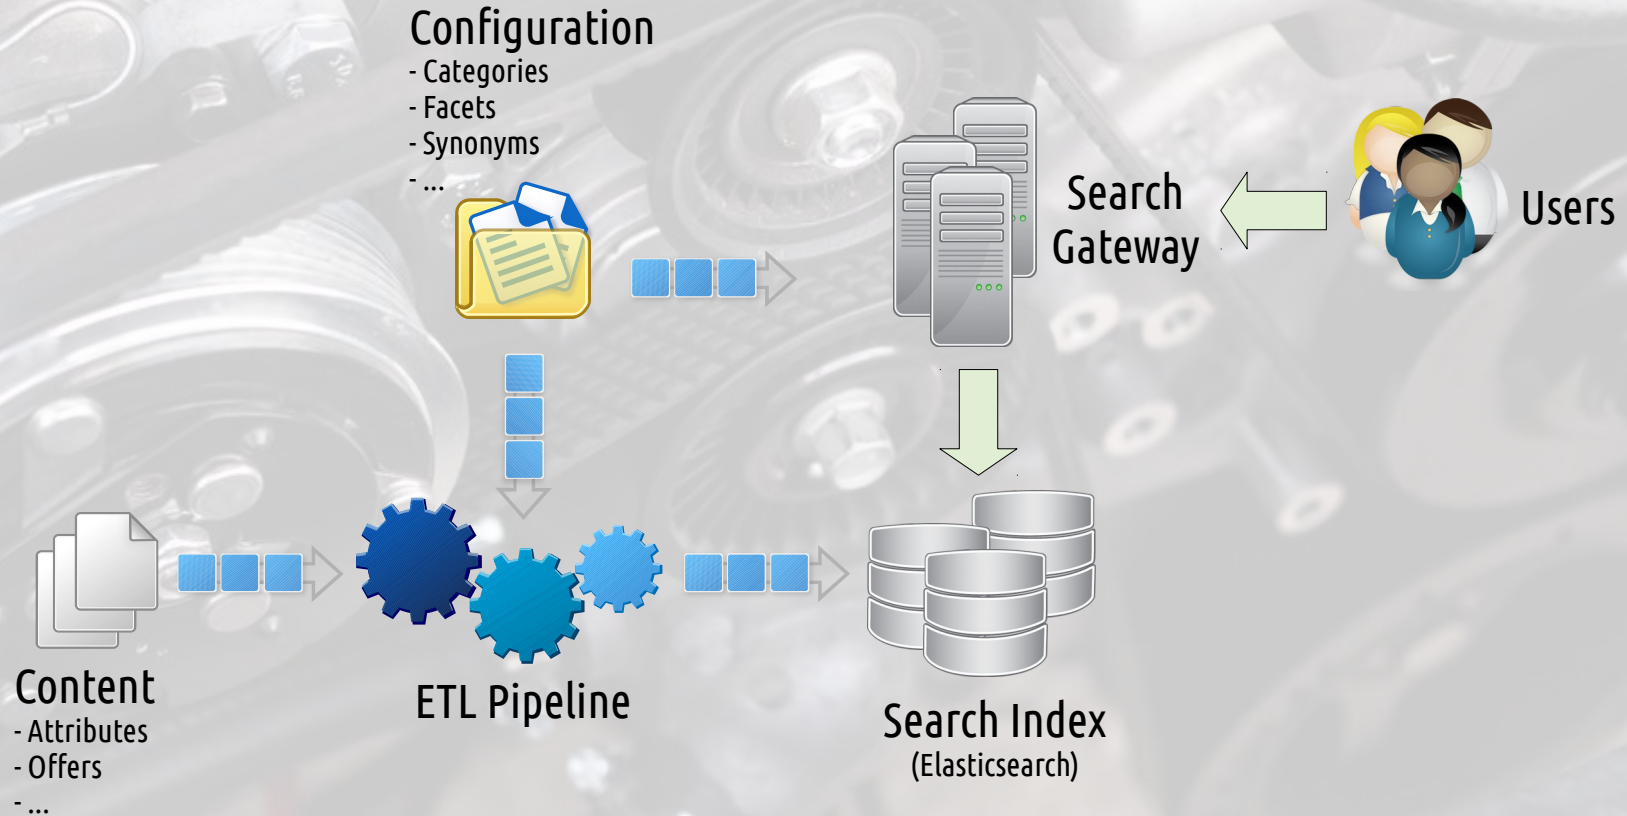


volume
&
frequency

Search output

- **Hits** (products and offers)
- **Facets**
- **Auto-corrections**
- **Redirects** (huge SEO impact)

Architecture overview



Data arrival latency

Source	Past	Present	Future
Attributes	1/24h	streaming	streaming
Offers	streaming	streaming	streaming
Facets	1/24h	1/24h	streaming
Indexing	1/24h	1/5h	streaming

Performance

- Search
- ETL
- Caching
(see [Varnishing Search Performance](#))





ETL
(Extract, Transform, Load)

Low-volume (~8/day)
traffic triggering
multiple (millions!)
product updates

High-volume
(millions/day) traffic
each triggering single
product updates

Configuration

- Categories
- Facets
- Synonyms
- ...



```
if (attrs.gpc.family_id == 1234 &&  
    attrs.gpc.chunk_id == 5678) {  
    doc.category = "books"  
}
```

```
if (attrs.disk_capacity_bytes != null) {  
    doc.disk_capacity_gigabytes =  
        attrs.disk_capacity_bytes / 1e9  
}
```

category == "smart phones"

```
if (attrs.disk_capacity_bytes != null) {  
    doc.disk_capacity_terabytes =  
        attrs.disk_capacity_bytes / 1e12  
}
```

category == "computers"



Content

- Attributes
- Offers
- Discounts
- ...

```
{  
  gpc: {  
    family_id: ...,  
    chunk_id: ...  
  },  
  disk_capacity_bytes: ...,  
  ...  
}
```



ETL Pipeline

exposure rules (signaling
search gateway to when/where
to expose these facets)

Why ETL at all?

Strategy	Advantages	Disadvantages
Without ETL	Changes take immediate effect	Latency and throughput hurts Aggregations become impractical
With ETL	Optimal query-time performance	Need to bake affected products

Content stream

- Sources
 - Content
 - Offer
 - Ranking
 - ...


- High-volume traffic

- ETL'ing is expensive
(due to tens of thousands of configurations)

```
if (attrs.gpc.family_id == 1234 &&
    attrs.gpc.chunk_id == 5678) {
  doc.category = "books"
}
```


Configuration stream

- Business screens
 - Configuration snapshots
 - Query on any field
- Low-volume traffic
- Retrospective changes



```
if (attrs.gpc.family_id == 1234 &&  
    attrs.gpc.chunk_id == 5678) {  
    doc.category = "books"  
}
```

H1-HG
Design

1. ~~~~~
2. ~~~~~
3. ~~~~~

Equal | Down

H1-Headline

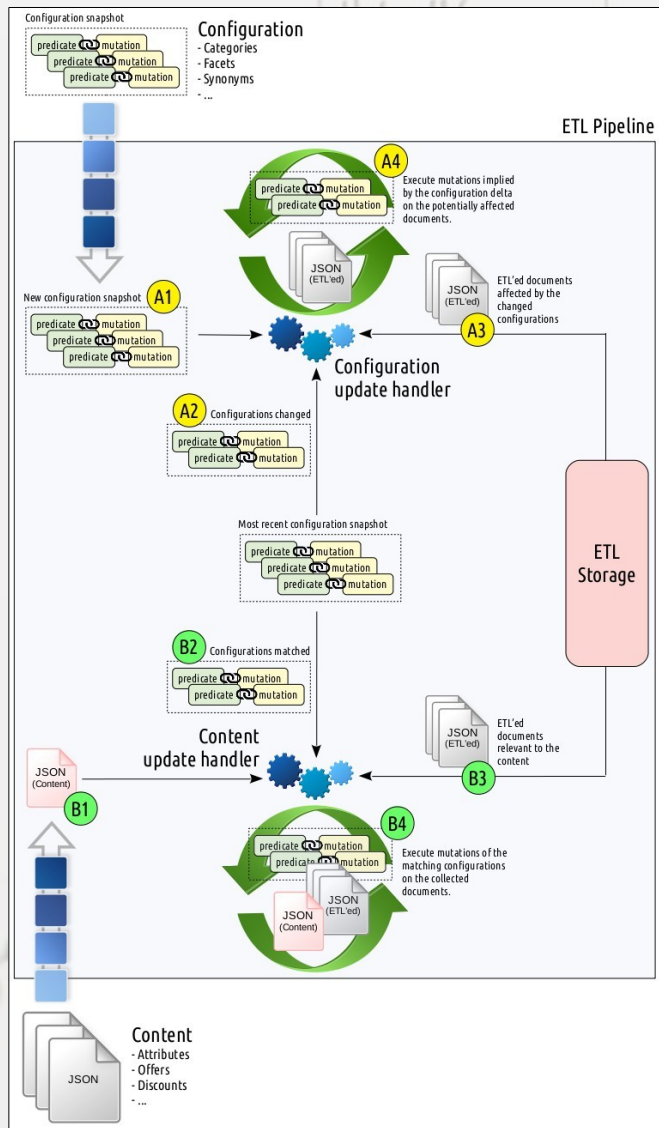
ETL Operational Overview



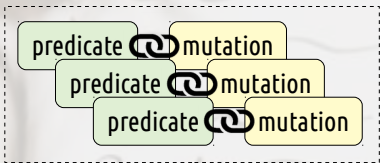
Video
Module

Menu






Configuration snapshot



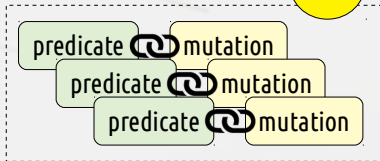
Configuration

- Categories
- Facets
- Synonyms
- ...



New configuration snapshot

A1

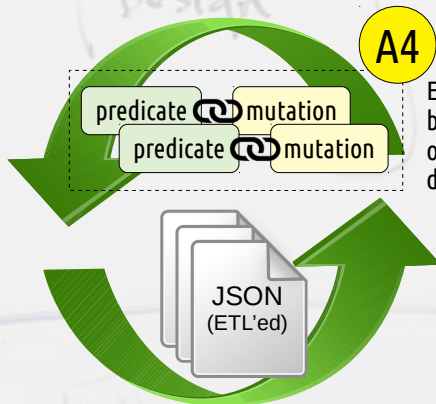



```
if (attrs.gpc.family_id == 1234 &&  
    attrs.gpc.chunk_id == 5678) {  
    doc.category = "books"  
}
```

A2 Configurations changed



Most recent configuration snapshot



Execute mutations implied by the configuration delta on the potentially affected documents.




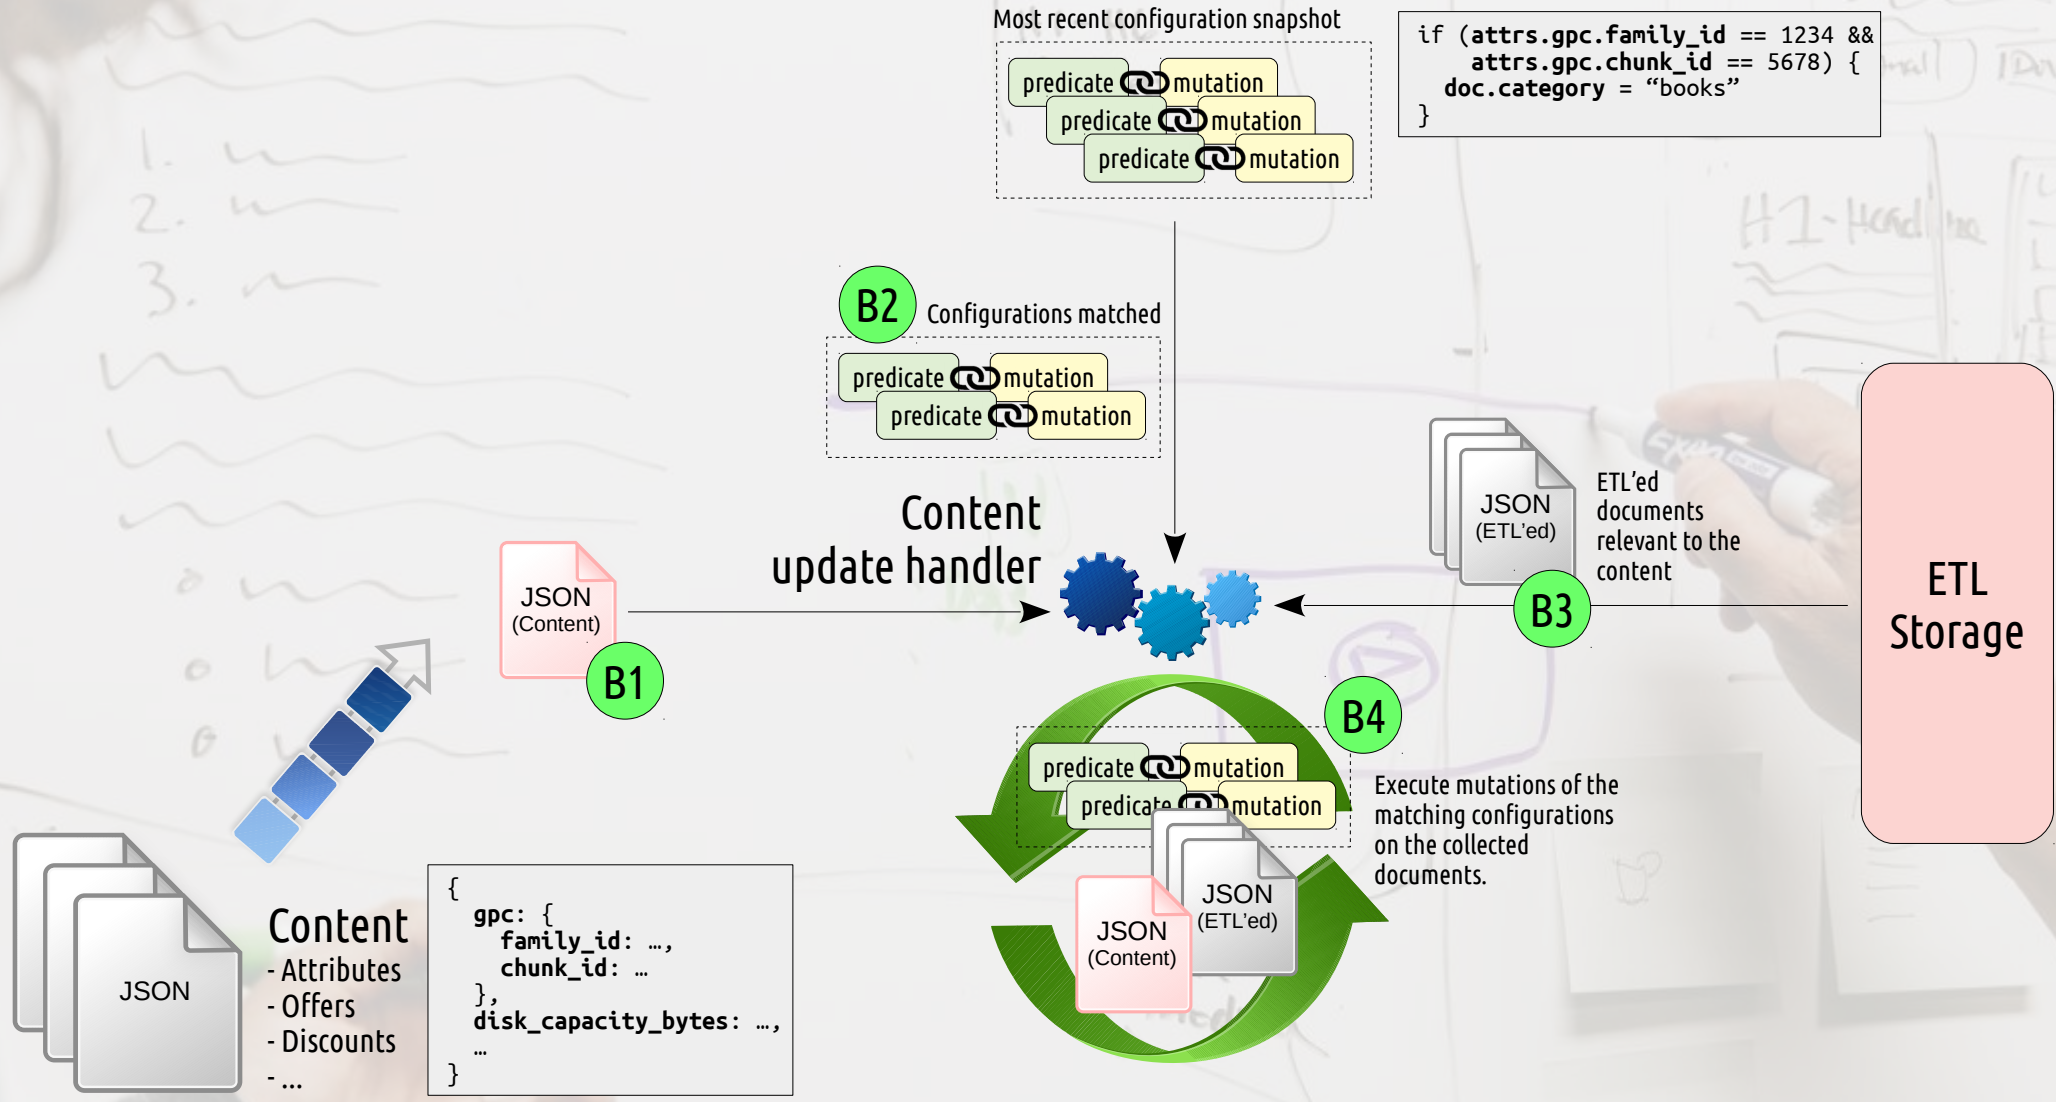
ETL'ed documents affected by the changed configurations

A3



Configuration update handler





Old ETL

- One giant PL/SQL troop marching 1/24h
- “Baseline” taking ~12h
- Failures hurt a lot
- Difficult to
 - innovate
 - debug
 - observe
- At the edge of software limits
 - e.g. max column count
 - multiple threads in PL/SQL
 - optimizer hints getting broken as
 - upgrades take place
 - data size change

Battle of ETL Storage Solutions

Storage Solution	Distributed?	Sharded?	Required Indices	Integrity Measure
PostgreSQL	No	No	One ¹	Transactions
PostgreSQL (partitioned)	No	Yes ²	One ¹	Transactions
MongoDB	Yes	Yes ³	Some ⁴	Transactions/CAS ⁵
Elasticsearch	Yes	Yes	None	CAS ⁶

- 1) PostgreSQL jsonb index covers all fields.
- 2) PostgreSQL partitioning is not sharding in distributed sense, but still serves a similar purpose.
- 3) MongoDB sharding requires manual configuration.
- 4) MongoDB requires an explicit index for each whitelisted field allowed in ETL configuration predicates.
- 5) MongoDB updateMany() or findAndModify() can be leveraged for the desired integrity.
- 6) Elasticsearch _version field can be leveraged to implement a CAS (compare-and-swap) loop.

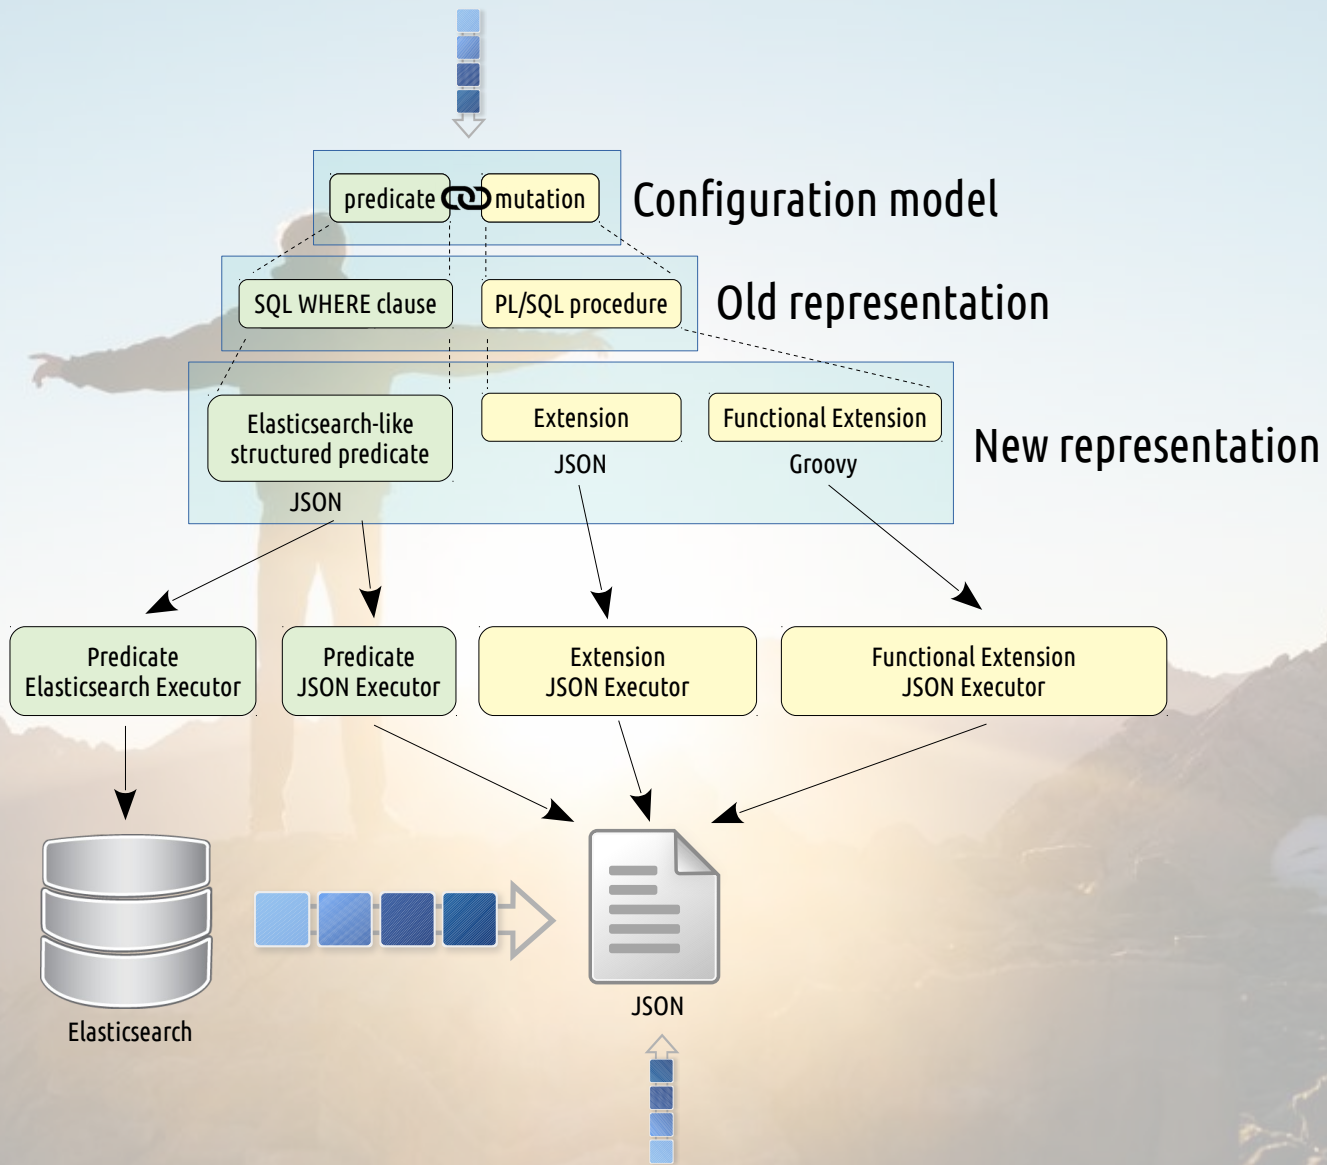
A silhouette of a person standing on a rocky mountain peak with their arms outstretched, looking towards a bright sun over a mountain range. The scene is bathed in a warm, golden light, suggesting a sunrise or sunset. The text "New ETL" is overlaid on the image.

New ETL

Storage solution winner: **Elasticsearch**



- Versatile query support
- Implicit indexing
- Scales good for reads, ok'ish for writes
- Easy to maintain
- Extensive experience



TL;DR

Google-like search != e-commerce search
(though both employ full-text search)

ETL = the art of cooking content (for search)

ETL rules necessitate search as well
(due to excessive faceting)

Elasticsearch is a good candidate for storage in ETL

Thank you!

(Questions?)

Volkan Yazici

<https://vlkan.com>

@yazicivo @ElasticNL