

Varnishing Search Performance

Changing the (search) engine of a racecar going 300 km/h



bol.com^{eo}
Volkan Yazıcı
<http://vlkan.com/>

Varnishing Search Performance

Changing the (search) engine of a racecar going 300 km/h

SERIOUSLY!



bol.com^{eo}

Volkan Yazıcı

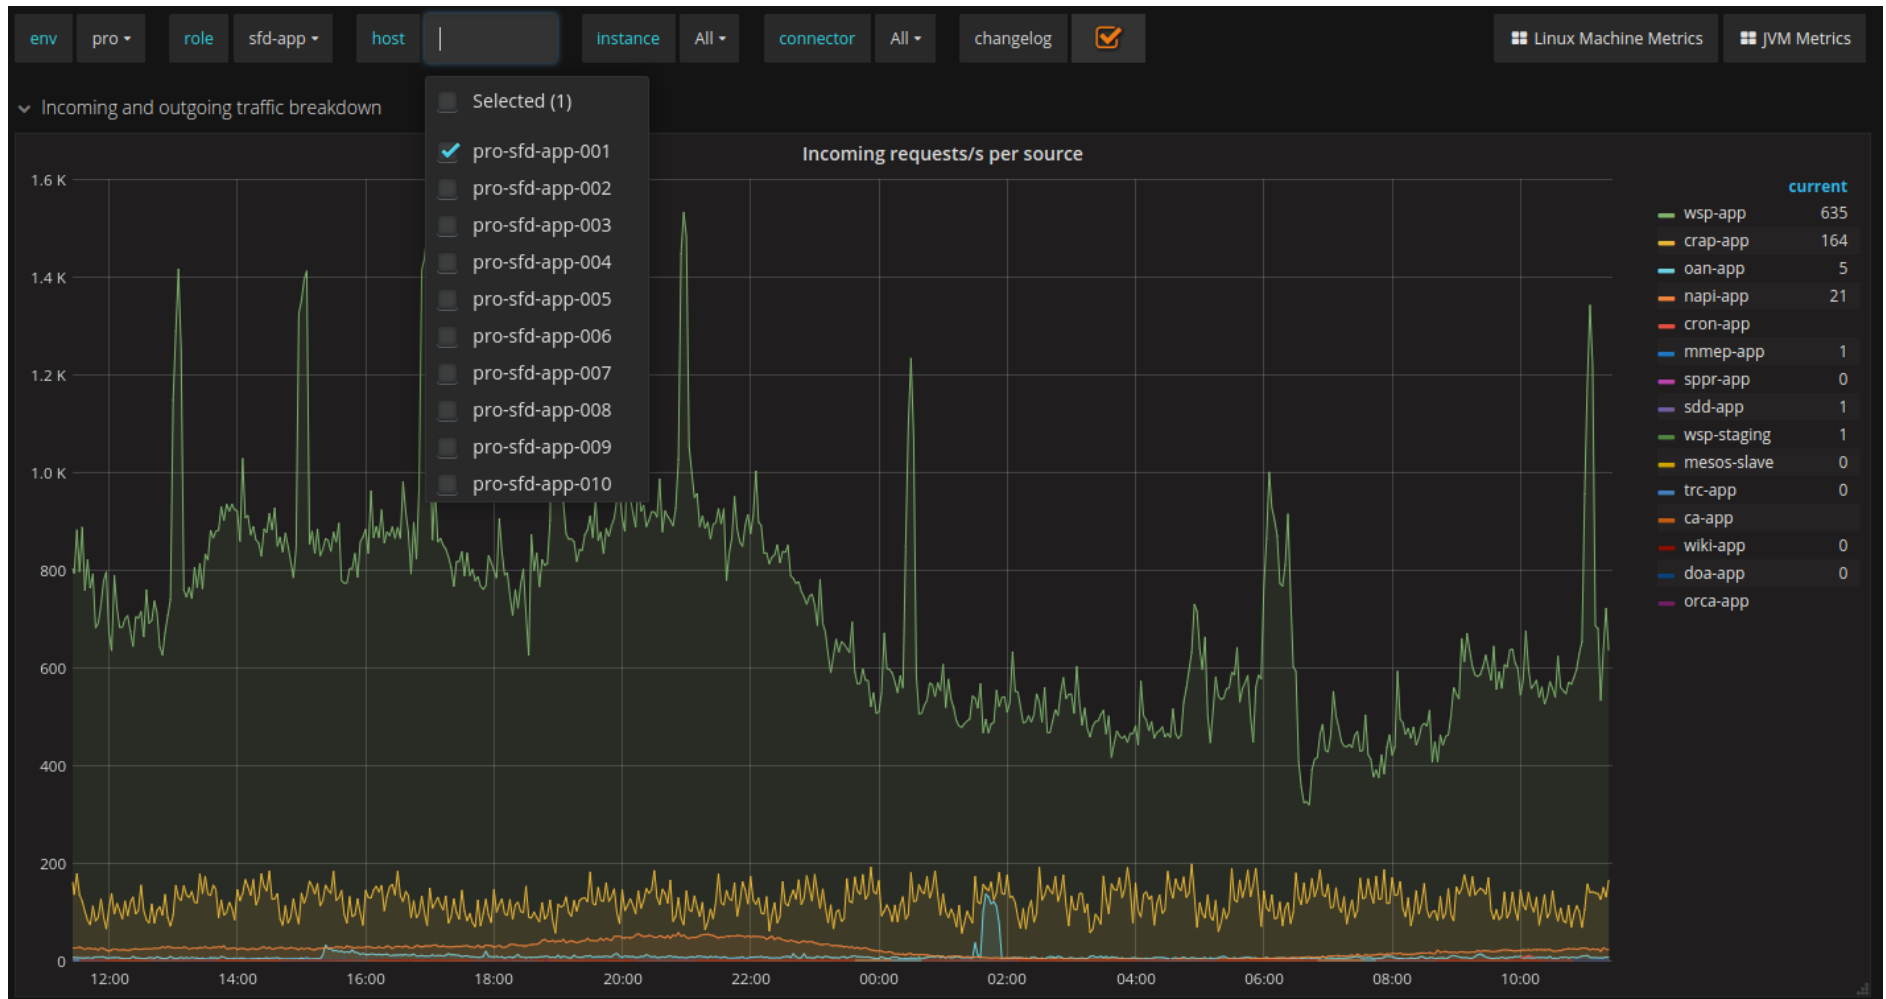
<http://vlkan.com/>

Secret Agenda


I will try to convince you Elasticsearch **performance**
and **caching** are indeed difficult subjects. So that I can
justify our complex caching **setup**.

Who uses **search** anyway?

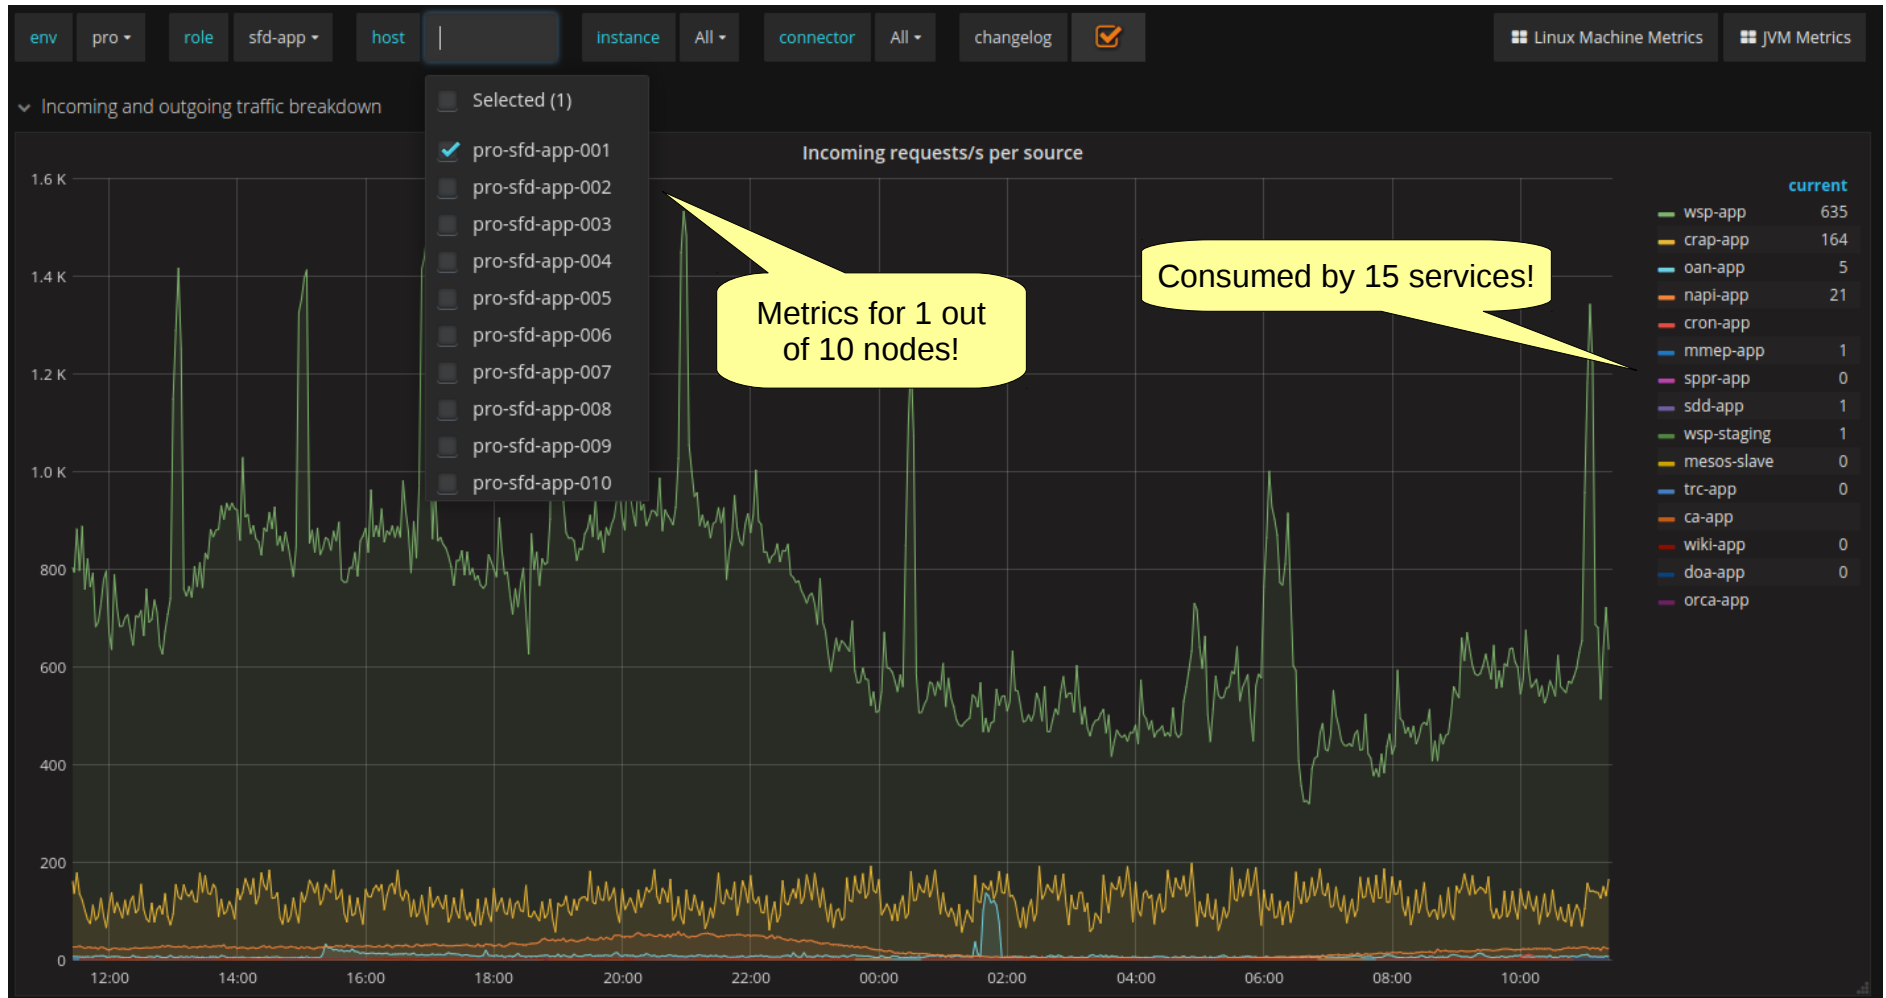
Who uses search anyway?



Who uses search anyway?



Who uses search anyway?



The Untold Story of **Performance** in Elasticsearch



The Untold Story of **Performance** in Elasticsearch

FAST

You have a couple of
constant_score
filters on analyzed fields.

SLOW



The Untold Story of **Performance** in Elasticsearch

FAST

You have a couple of
constant_score
filters on analyzed fields.

Add some **boost**
and **should** spice.

SLOW



The Untold Story of **Performance** in Elasticsearch

FAST

You have a couple of **constant_score filters** on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't foreseen.

SLOW



The Untold Story of **Performance** in Elasticsearch

FAST

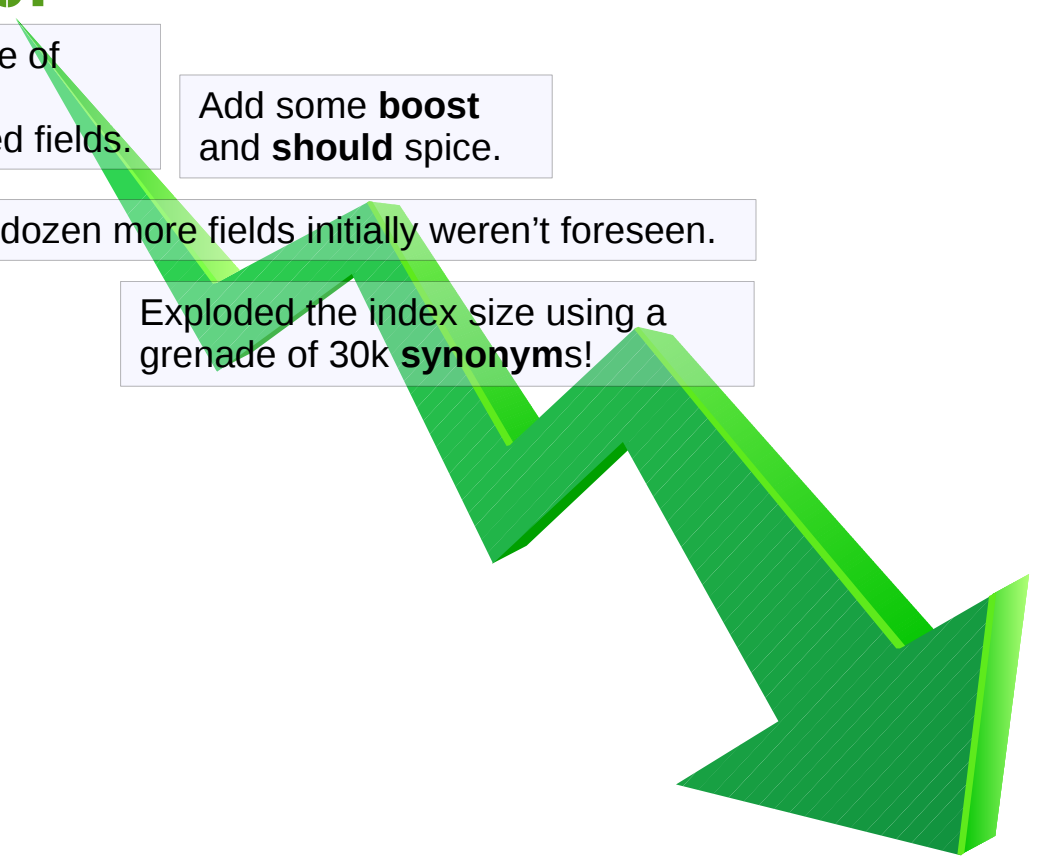
You have a couple of **constant_score filters** on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't foreseen.

Exploded the index size using a grenade of 30k **synonyms!**

SLOW



The Untold Story of Performance in Elasticsearch

FAST

You have a couple of **constant_score filters** on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't foreseen.

Exploded the index size using a grenade of 30k **synonyms!**

Needed to group hits by color, size, etc. Now you have a **collapse**.

SLOW



Capcom

Monster Hunter World - PS4

PlayStation 4 | PEGI-leeftijd: 16

★★★★★ (14)

Welkom in een nieuwe wereld. Word jager en dood woeste monsters in een levend, ademend ecosysteem waarin je de omgeving om je heen... [Meer](#)

In verschillende varianten verkrijgbaar:



Xbox One
€ 77,00



Xbox One
download
€ 69,99

The Untold Story of Performance in Elasticsearch

FAST

You have a couple of **constant_score** filters on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't foreseen.

Exploded the index size using a grenade of 30k **synonyms**!

Needed to group hits by color, size, etc. Now you have a **collapse**.

Decided to add 20 **aggregations**.

SLOW



Capcom

Monster Hunter World - PS4

PlayStation 4 | PEGI-leeftijd: 16

★★★★★ (14)

Welkom in een nieuwe wereld. Word jager en dood woeste monsters in een levend, ademend ecosysteem waarin je de omgeving om je heen... [Meer](#)

In verschillende varianten verkrijgbaar:



Xbox One
€ 77,00



Xbox One
download
€ 69,99

Games platform ^

Er zijn verschillende merken die consoles... >

- PlayStation 4 (19)
- Xbox One (10)
- Nintendo ZDS (8)

Meer

Merk ^

Er zijn verschillende merken die cons... >

Game editie v

Wat mag het kosten? ^

Kost gemiddeld € 260,-

Opslagcapaciteit

Consoles hebben verschillende... >

- 1 TB (20)
- Geen opslag
- 4 GB (10)

Meer

Nieuw of tweedehands v

Levertijd v

Serie

De consoles kunnen worden uitgebra... >

- Xbox One S (2)
- The Legend of C...

Opties

Consoles bieden verschillende opt... >

- WIFI (48)
- Inclusief game
- Bluetooth (30)

Meer

Select ledenvoordeel v

Beste getest v

Wat vinden anderen? v

The Untold Story of Performance in Elasticsearch

FAST

You have a couple of **constant_score** filters on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't foreseen.

Exploded the index size using a grenade of 30k **synonyms**!

Needed to group hits by color, size, etc. Now you have a **collapse**.

Decided to add 20 **aggregations**.

Hey! There aren't 20 facets here!

SLOW

The screenshot shows a facet menu for 'Games platform'. It lists various console brands like PlayStation 4, Xbox One, and Nintendo 2DS. Below this, there are filters for 'Opslagcapaciteit' (Storage capacity) with a range from 10 to 480 GB, and 'Serie' (Series) with options like 'Xbox One S' and 'The Legend of Zelda'. There are also sections for 'Aanraders' (Recommendations) and 'Opties' (Options) like 'WIFI' and 'Bluetooth'. A yellow speech bubble points to the menu with the text 'Hey! There aren't 20 facets here!'.

The image shows a product listing for 'Monster Hunter World - PS4' by Capcom. It includes a PS4 game cover, a star rating of 5 stars (14 reviews), and a description: 'Welkom in een nieuwe wereld. Word jager en dood woeste monsters in een levend, ademend ecosysteem waarin je de omgeving om je heen... Meer'. Below the description, it says 'In verschillende varianten verkrijgbaar:' and shows two variants: 'Xbox One' for €77,00 and 'Xbox One download' for €69,99. A red dashed circle highlights these two variants.

The Untold Story of Performance in Elasticsearch

FAST

You have a couple of **constant_score** filters on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't fores

Exploded the index size using grenade of 30k **synonyms**!

Needed to group hits by color, size, etc. Now you have a **collapse**.

Decided to add 20 **aggregations**.

Doubled the amount of aggregations.

SLOW

Games platform
Er zijn verschillende merken die consoles...
 PlayStation 4 (19)
 Xbox One (10)
 Nintendo ZDS (8)
Meer

Merk
Er zijn verschillende merken die cons...
 Sony (19)
 Nintendo (17)
 Microsoft (11)
Meer

Game editie
Wat mag het kosten?
Kost gemiddeld € 260,-
10 tot 480

Oplagcapaciteit
Consoles hebben verschillende...
 1 TB (20)
 Geen opslag
 4 GB (10)
Meer

Serie
De consoles kunnen worden uitgebra...
 Xbox One S (2)
 The Legend C

Aanraders
Uitleg & kiezen
 Bestbeoordeeld (9)
 Nieuw verschenen (1)
 Te reserveren (1)

Opties
Consoles bieden verschillende opt...
 WIFI (48)
 Inclusief game
 Bluetooth (30)
Meer

Select ledenvoordeel
Beste getest
Wat vinden anderen?



Capcom
Monster Hunter World - PS4
PlayStation 4 | PEGI-leeftijd: 16
★★★★★ (14)
Welkom in een nieuwe wereld. Word jager en dood woeste monsters in een levend, ademend ecosysteem waarin je de omgeving om je heen... [Meer](#)

In verschillende varianten verkrijgbaar:



Xbox One
€ 77,00



Xbox One
download
€ 69,99

The Untold Story of Performance in Elasticsearch

FAST

You have a couple of **constant_score** filters on analyzed fields.

Add some **boost** and **should** spice.

Needed a dozen more fields initially weren't fores

Exploded the index size using grenade of 30k **synonyms**!

Needed to group hits by color, size, etc. Now you have a **collapse**.

Decided to add 20 **aggregations**.

Doubled the amount of aggregations.

Just realized you're serving 5% of calls on production behind an A/B test.

SLOW

Games platform
Er zijn verschillende merken die consoles...
 PlayStation 4 (19)
 Xbox One (10)
 Nintendo 2DS (8)
Meer

Merk
Er zijn verschillende merken die cons...
 Sony (19)
 Nintendo (17)
 Microsoft (11)
Meer

Game editie
Wat mag het kosten?
Kost gemiddeld € 260,-
10 tot 480
Oplagcapaciteit
Consoles hebben verschillende...
 1 TB (20)
 Geen opslag
 4 GB (10)
Meer

Serie
De consoles kunnen worden uitgebra...
 Xbox One S (2)
 The Legend C

Aanraders
Uitleg & kiezen
 Bestbeoordeeld (9)
 Nieuw verschenen (1)
 Te reserveren (1)

Opties
Consoles bieden verschillende opt...
 WIFI (48)
 Inclusief game
 Bluetooth (30)
Meer

Select ledenvoordeel
Beste getest
Wat vinden anderen?

Capcom
Monster Hunter World - PS4
PlayStation 4 | PEGI-leeftijd: 16
★★★★★ (14)
Welkom in een nieuwe wereld. Word jager en dood woeste monsters in een levend, ademend ecosysteem waarin je de omgeving om je heen... [Meer](#)
In verschillende varianten verkrijgbaar:
 Xbox One € 77,00
 Xbox One download € 69,99

EVER MAKE A MISTAKE IN LIFE?



**LET'S MAKE THEM BIRDS. YEAH,
THEY'RE BIRDS NOW.**

EVER MAKE A MISTAKE IN LIFE?

SLOW QUERY



LET'S MAKE THEM **CACHED** YEAR,

THEY'RE **CACHED** NOW.

The Untold Story of **Caching** with Elasticsearch

The Untold Story of **Caching** with Elasticsearch

SIMPLE




COMPLEX

The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.



COMPLEX

The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.

Reproducing production problems
becomes more cumbersome.

COMPLEX



The Untold Story of **Caching** with Elasticsearch

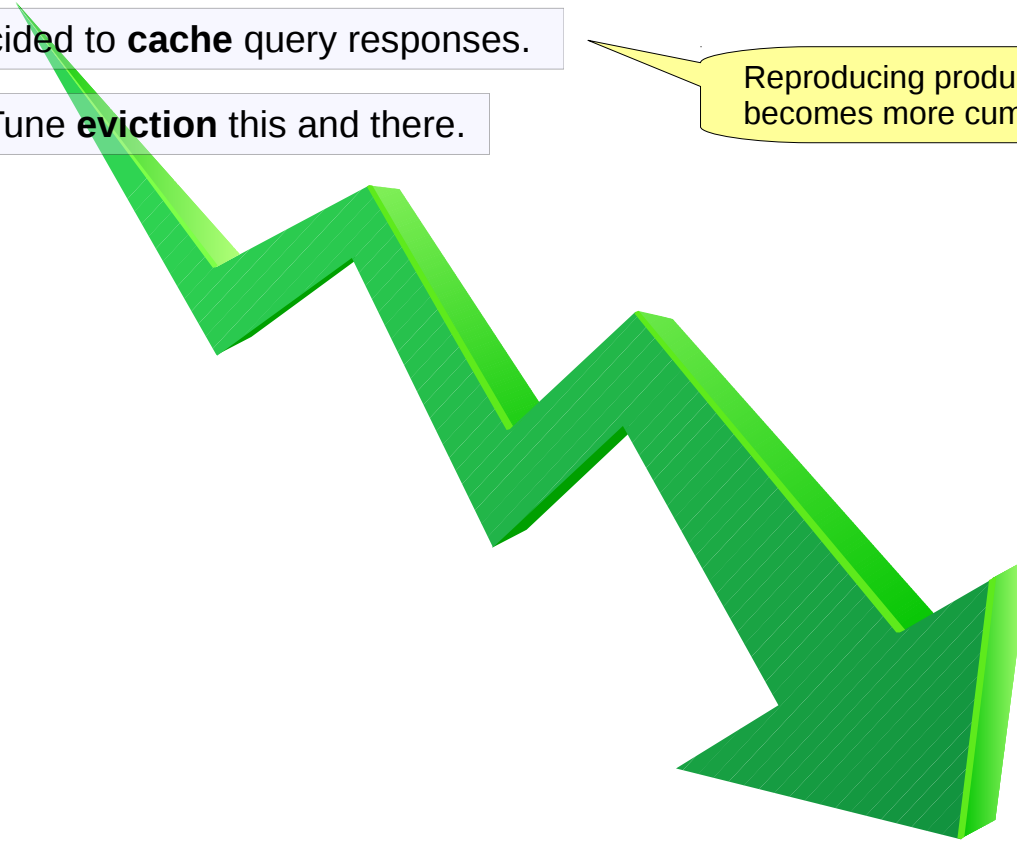
SIMPLE

You decided to **cache** query responses.

Tune **eviction** this and there.

Reproducing production problems
becomes more cumbersome.

COMPLEX



The Untold Story of **Caching** with Elasticsearch

SIMPLE

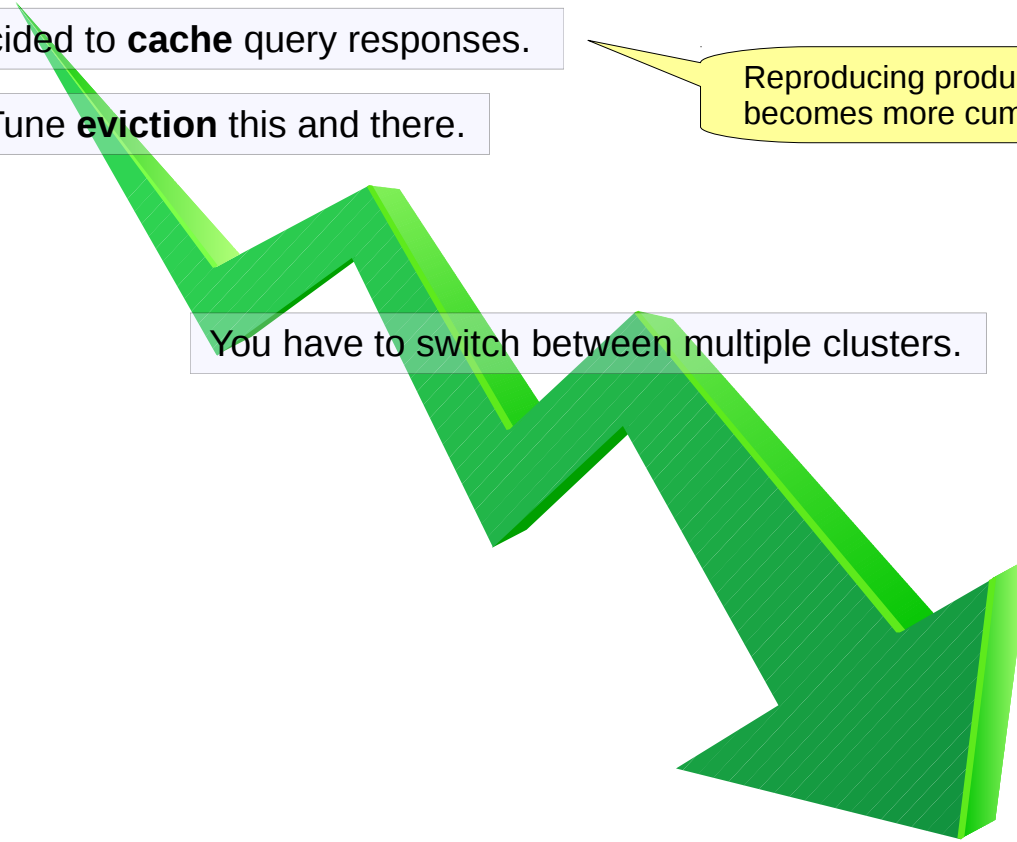
You decided to **cache** query responses.

Tune **eviction** this and there.

Reproducing production problems
becomes more cumbersome.

You have to switch between multiple clusters.

COMPLEX



The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.

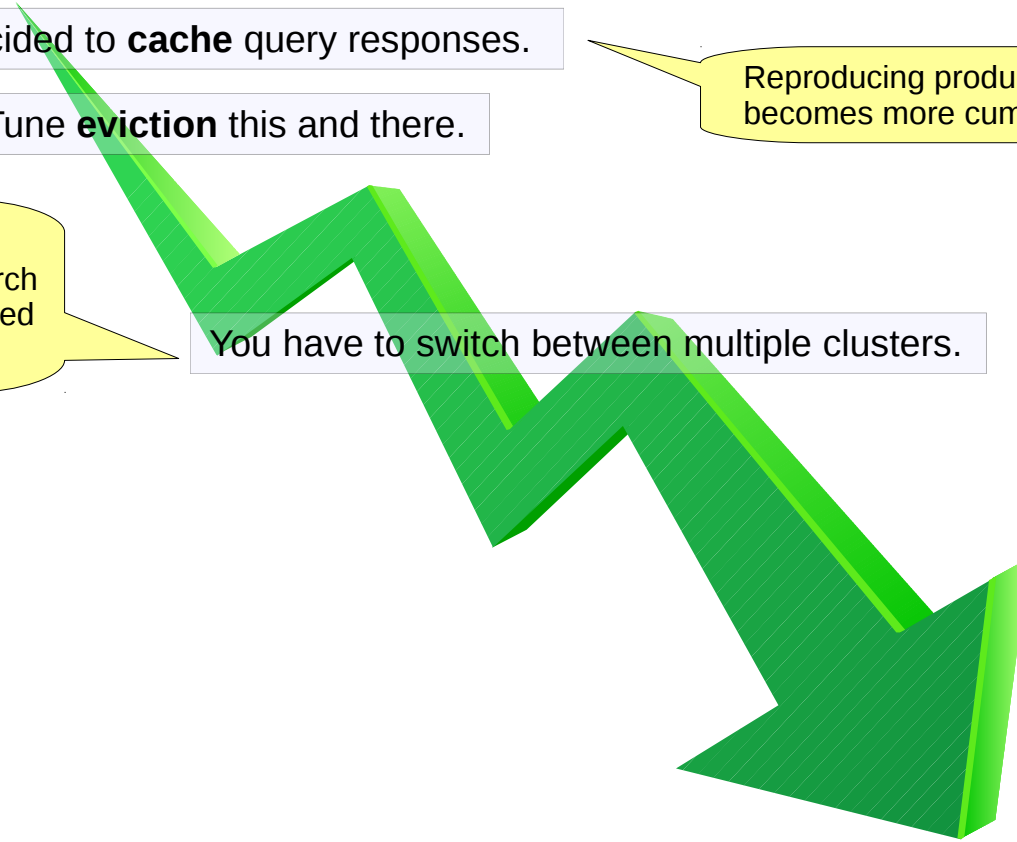
Tune **eviction** this and there.

Reproducing production problems becomes more cumbersome.

Management interfaces:
- which clusters are used for search
- which clusters are getting indexed
- in-progress index operations

You have to switch between multiple clusters.

COMPLEX



The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.

Tune **eviction** this and there.

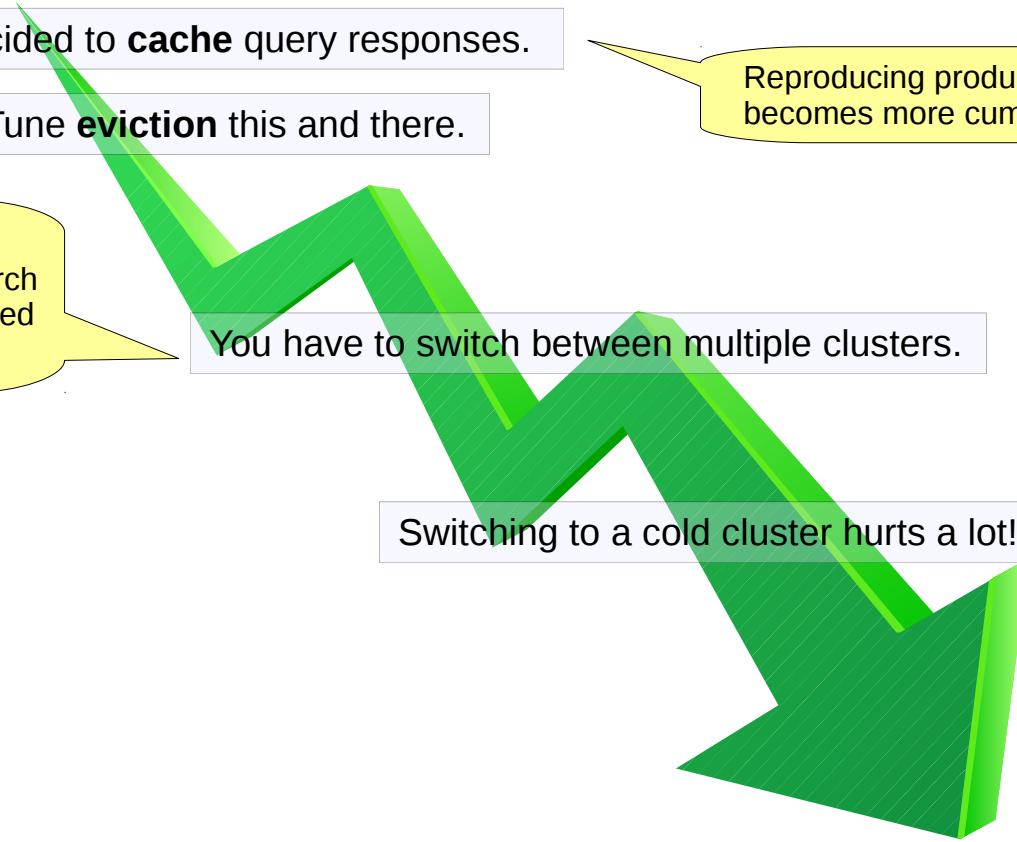
Reproducing production problems becomes more cumbersome.

Management interfaces:
- which clusters are used for search
- which clusters are getting indexed
- in-progress index operations

You have to switch between multiple clusters.

Switching to a cold cluster hurts a lot!

COMPLEX



The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.

Tune **eviction** this and there.

Reproducing production problems becomes more cumbersome.

Management interfaces:

- which clusters are used for search
- which clusters are getting indexed
- in-progress index operations

You have to switch between multiple clusters.

Switching to a cold cluster hurts a lot!

How to warm-up?

- replay traffic
- gradually move traffic (session stickiness)

COMPLEX

The Untold Story of **Caching** with Elasticsearch

SIMPLE

You decided to **cache** query responses.

Tune **eviction** this and there.

Reproducing production problems becomes more cumbersome.

Management interfaces:

- which clusters are used for search
- which clusters are getting indexed
- in-progress index operations

You have to switch between multiple clusters.

Switching to a cold cluster hurts a lot!

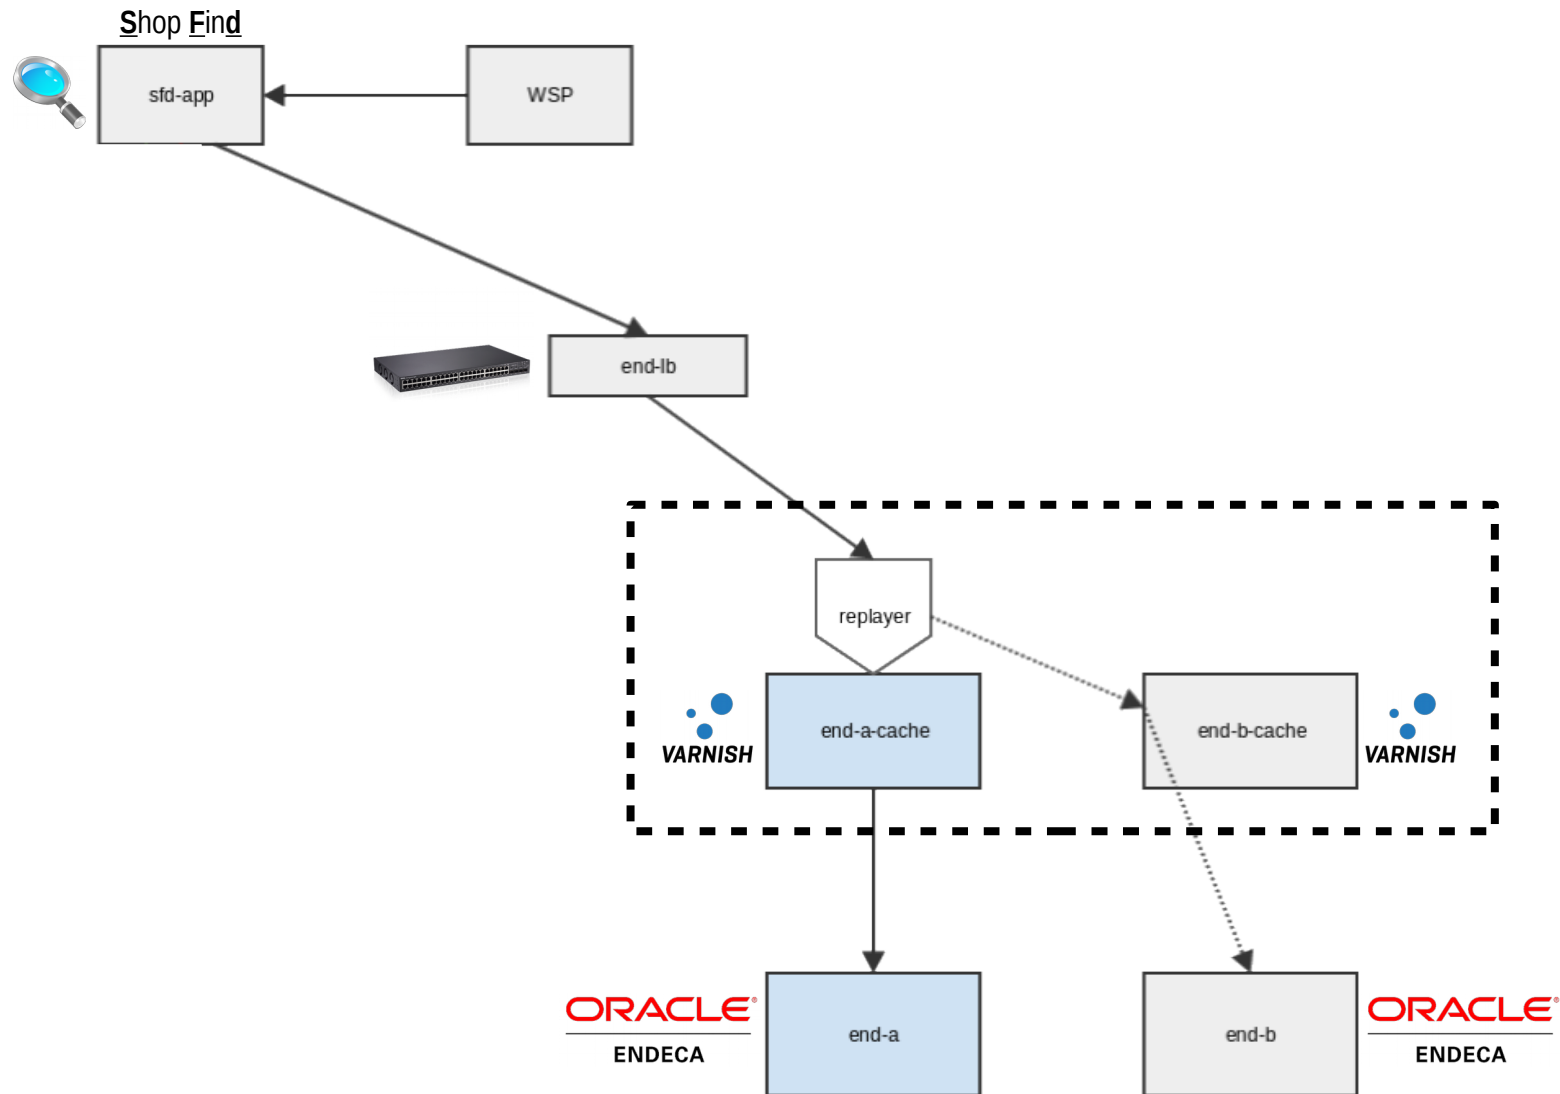
How to warm-up?

- replay traffic
- gradually move traffic (session stickiness)

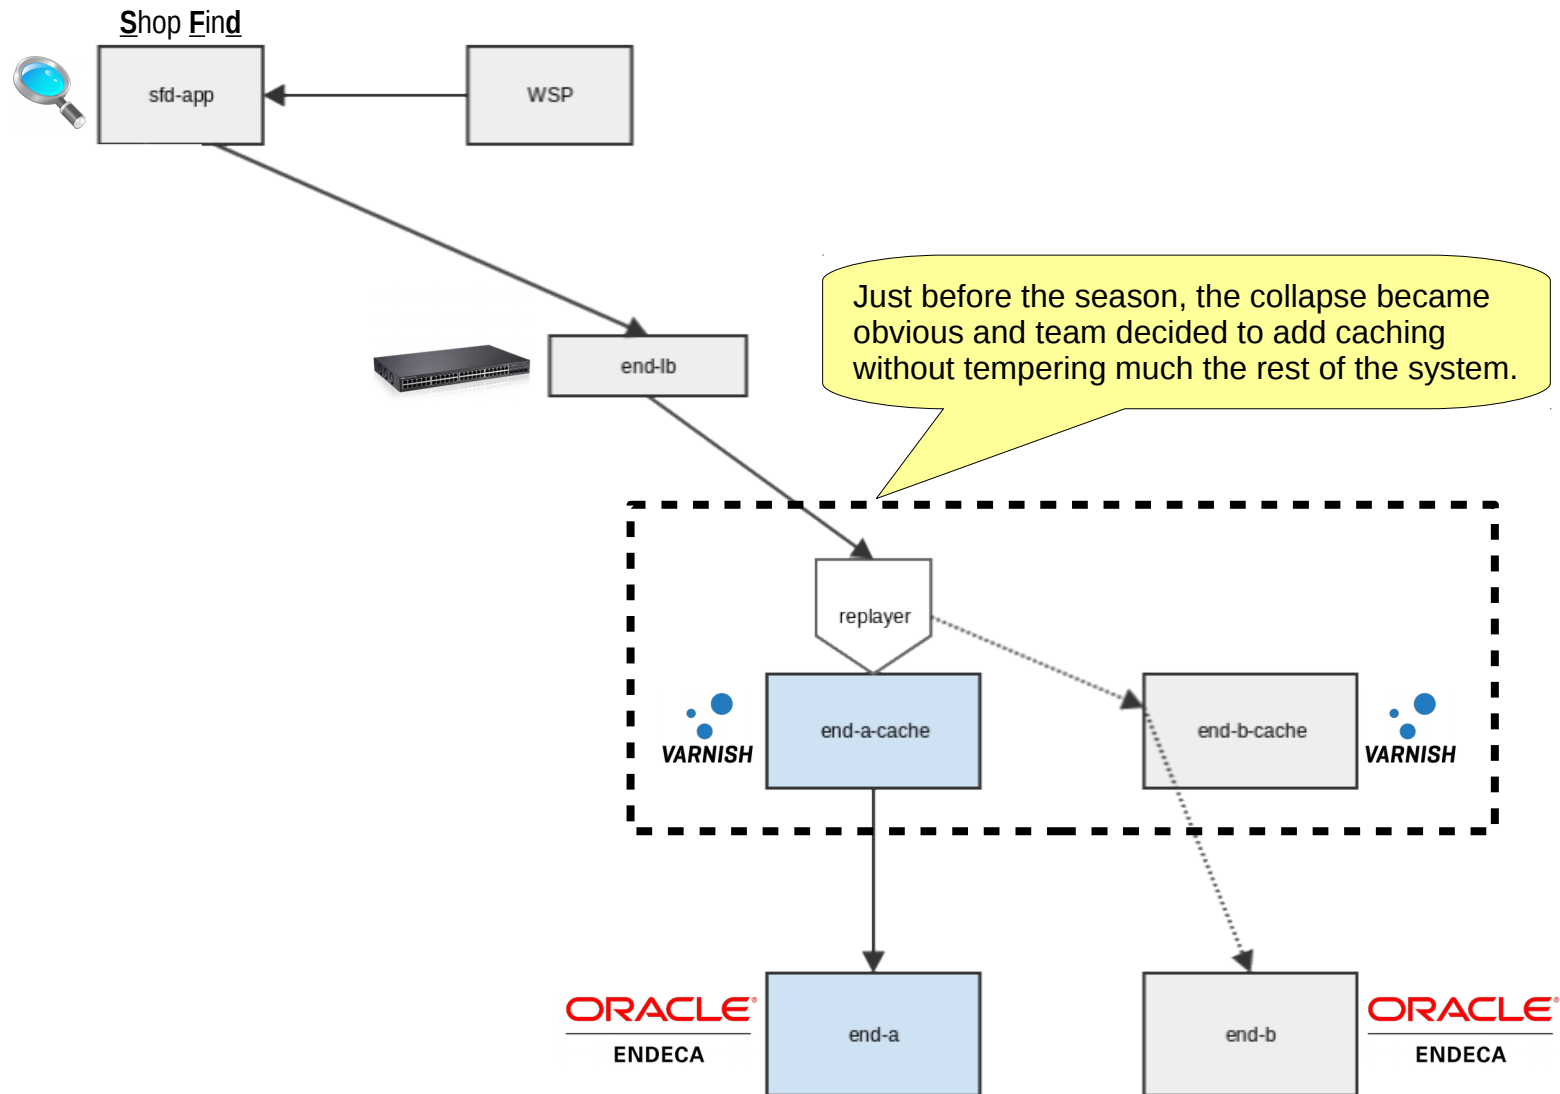
You need to migrate from a legacy system.

COMPLEX

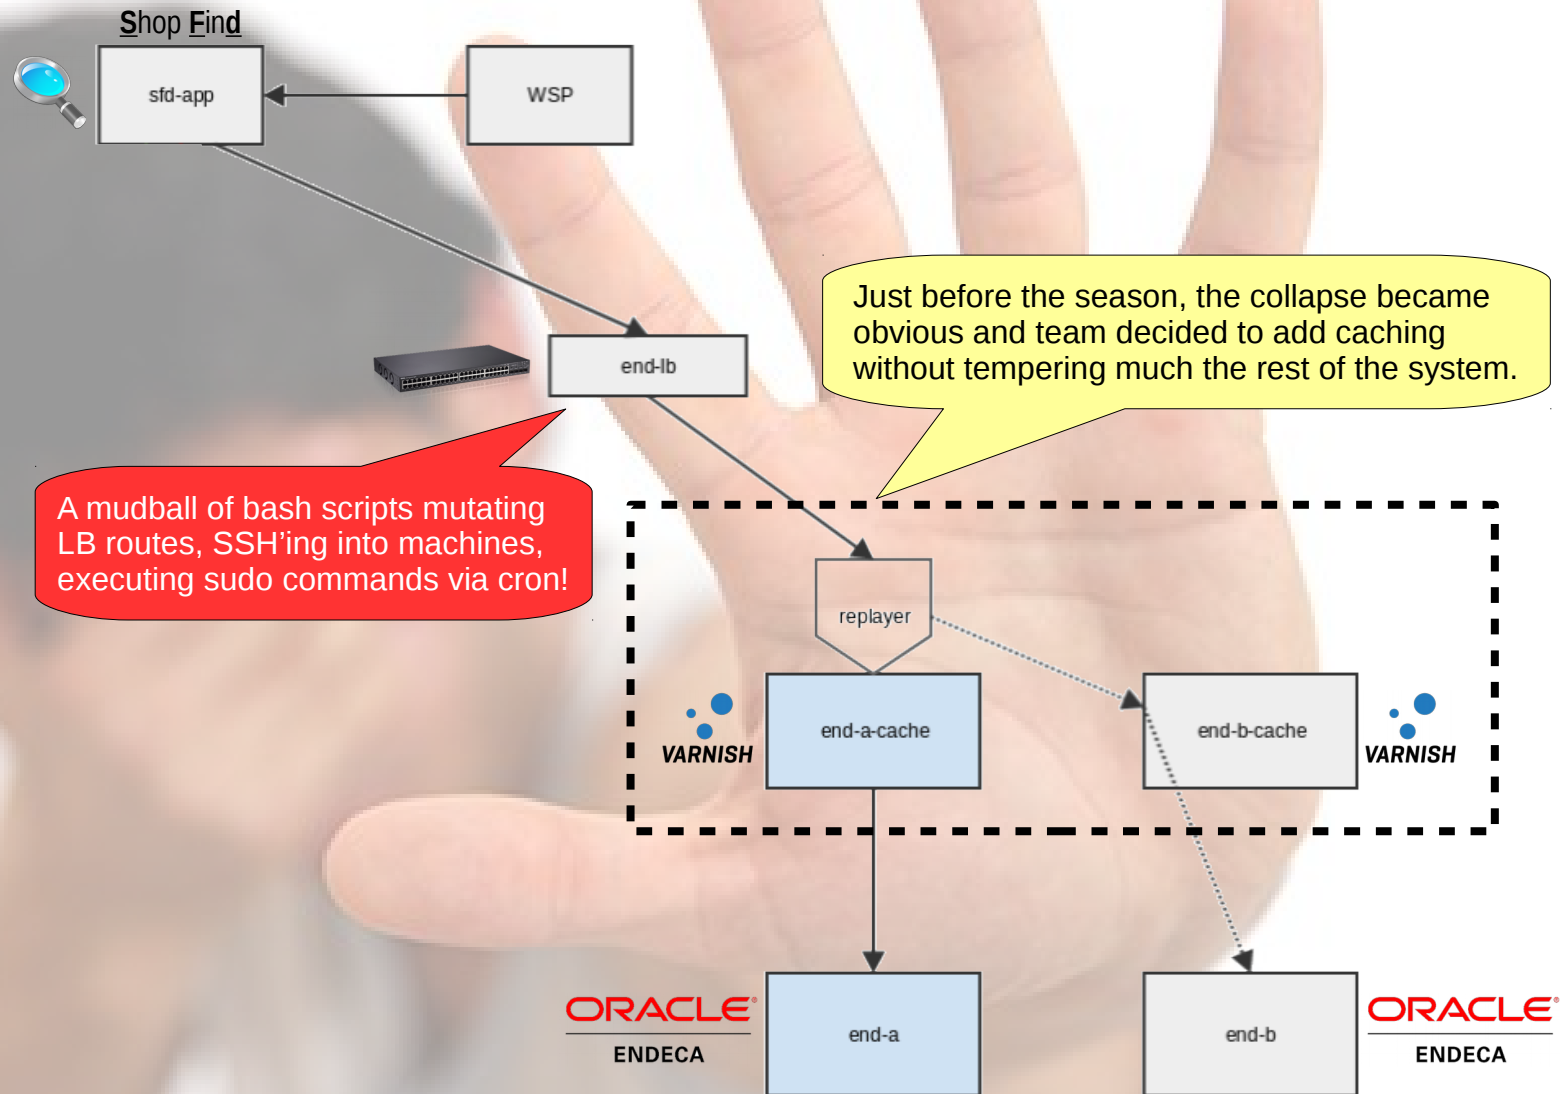
Hail the Legacy!



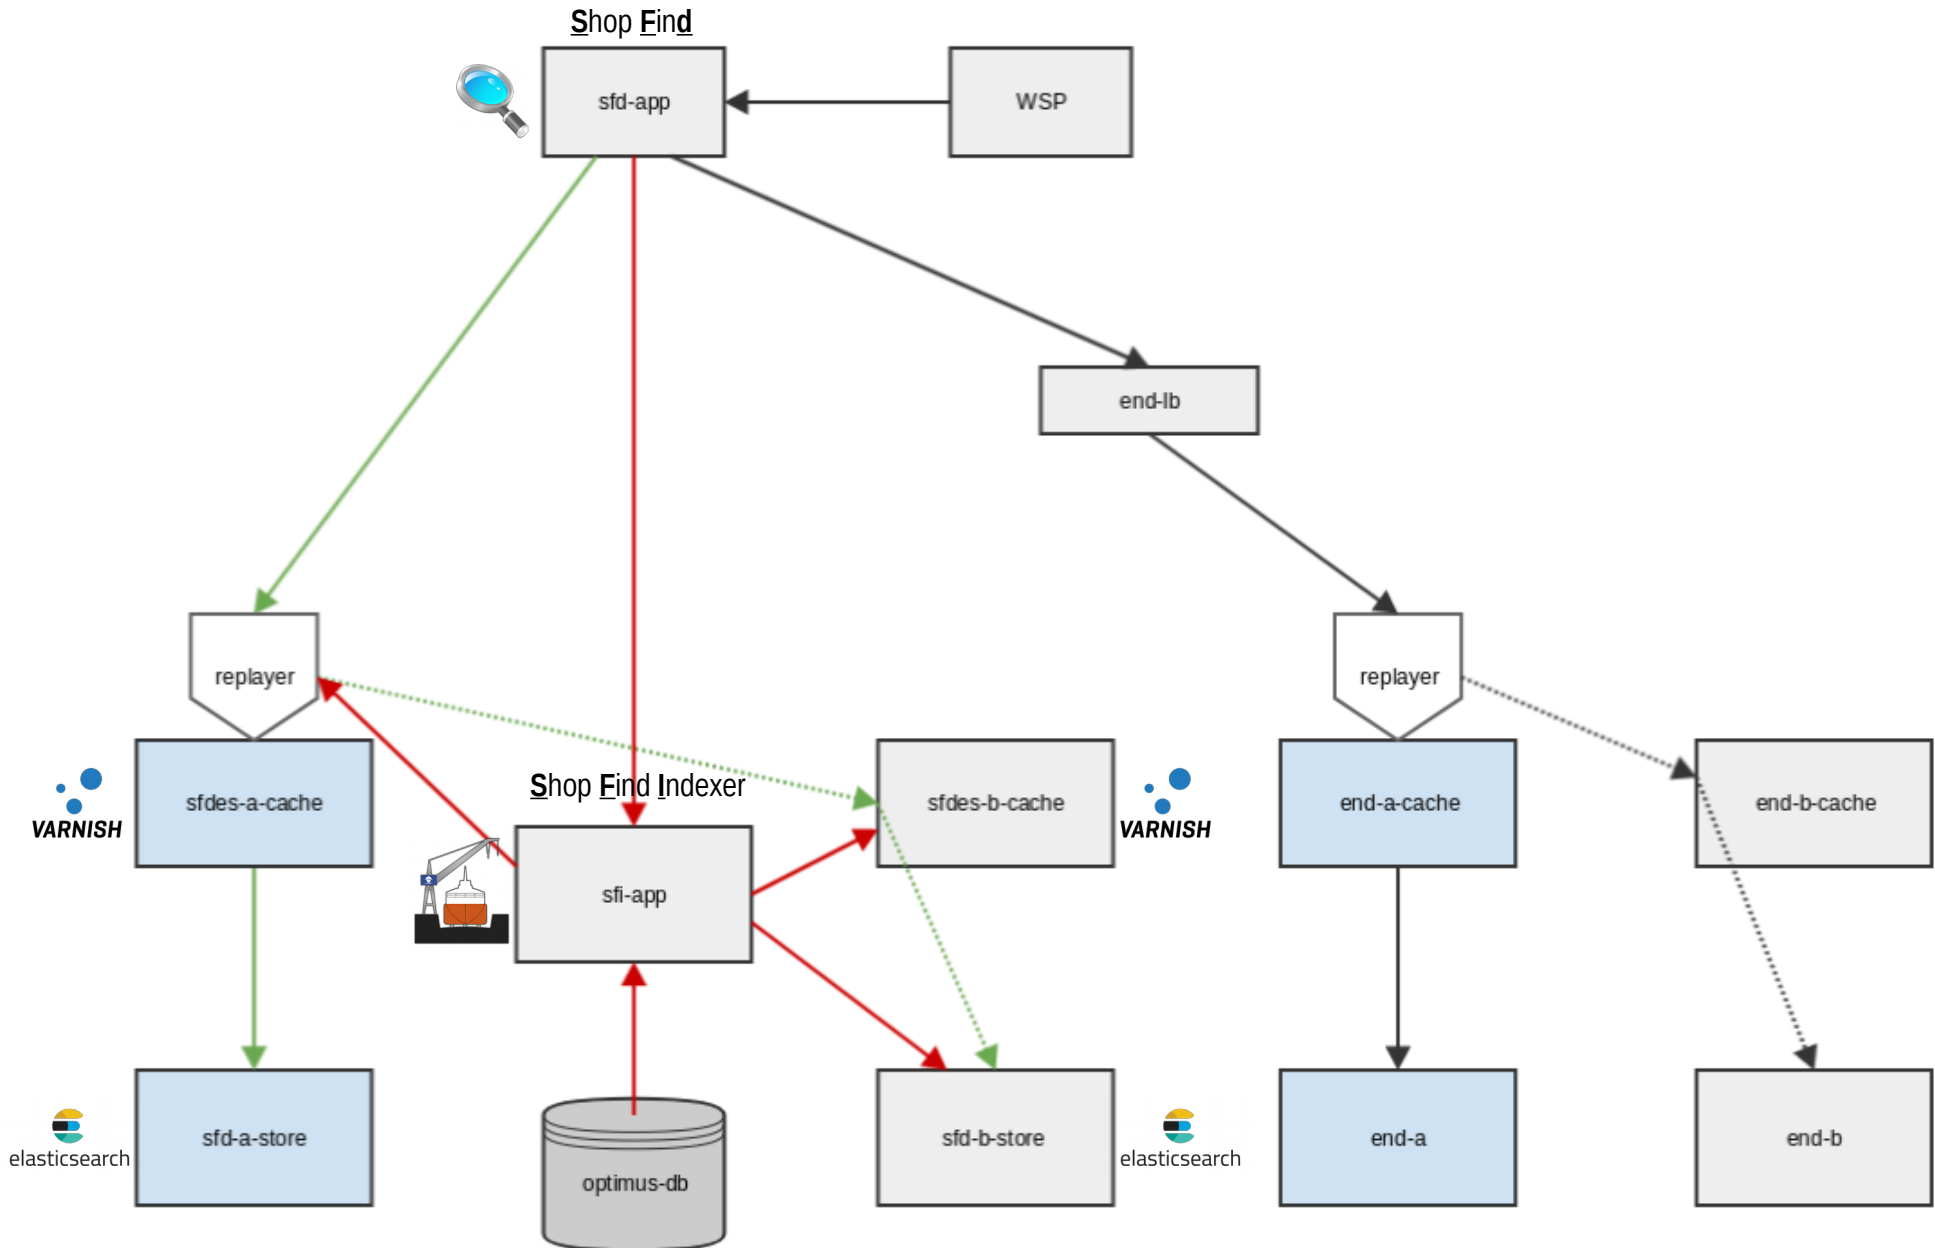
Hail the **Legacy!**



Hail the Legacy!











Long live the **king!**




Long live the **king!**

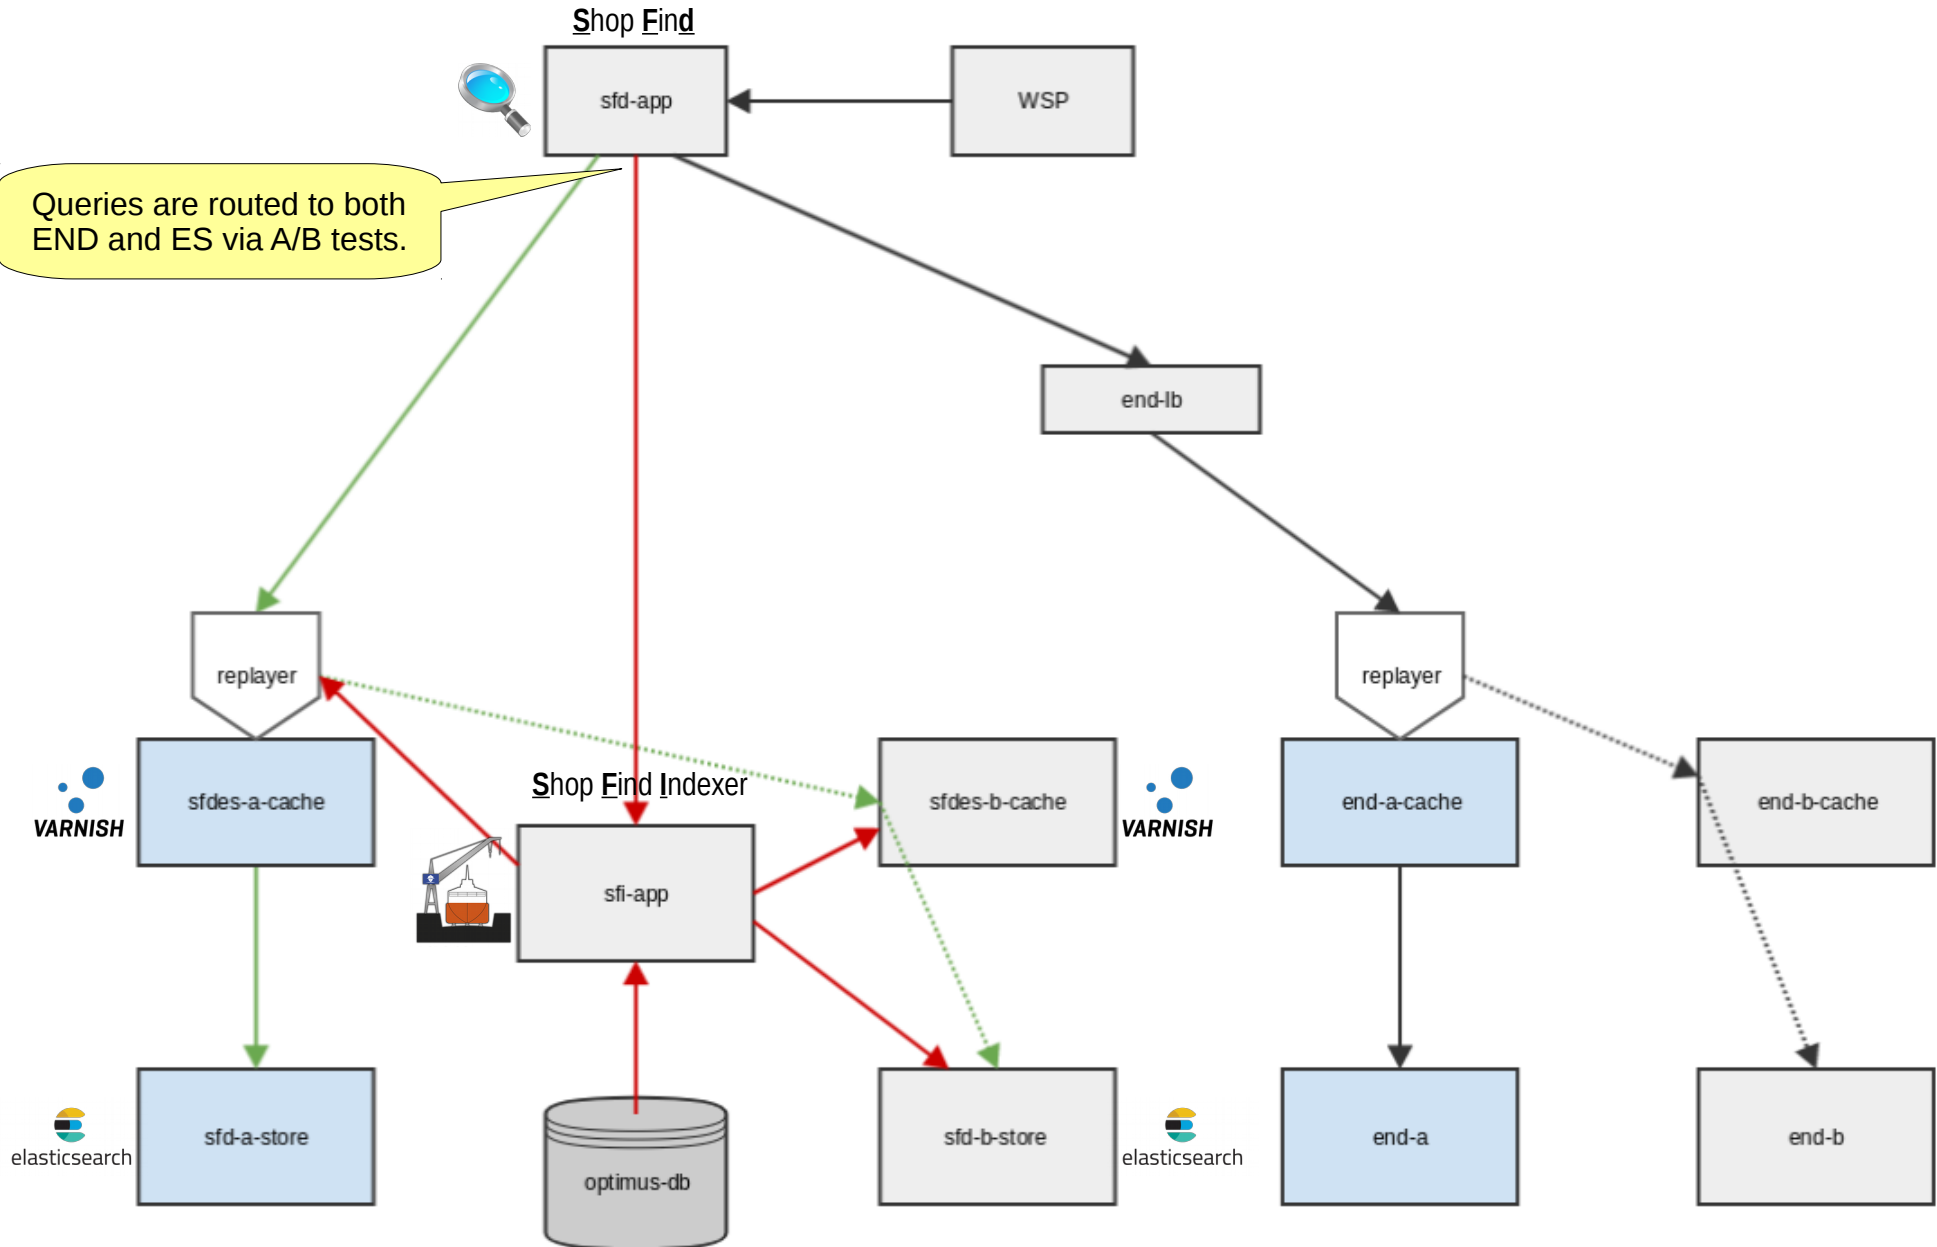
Shop Find



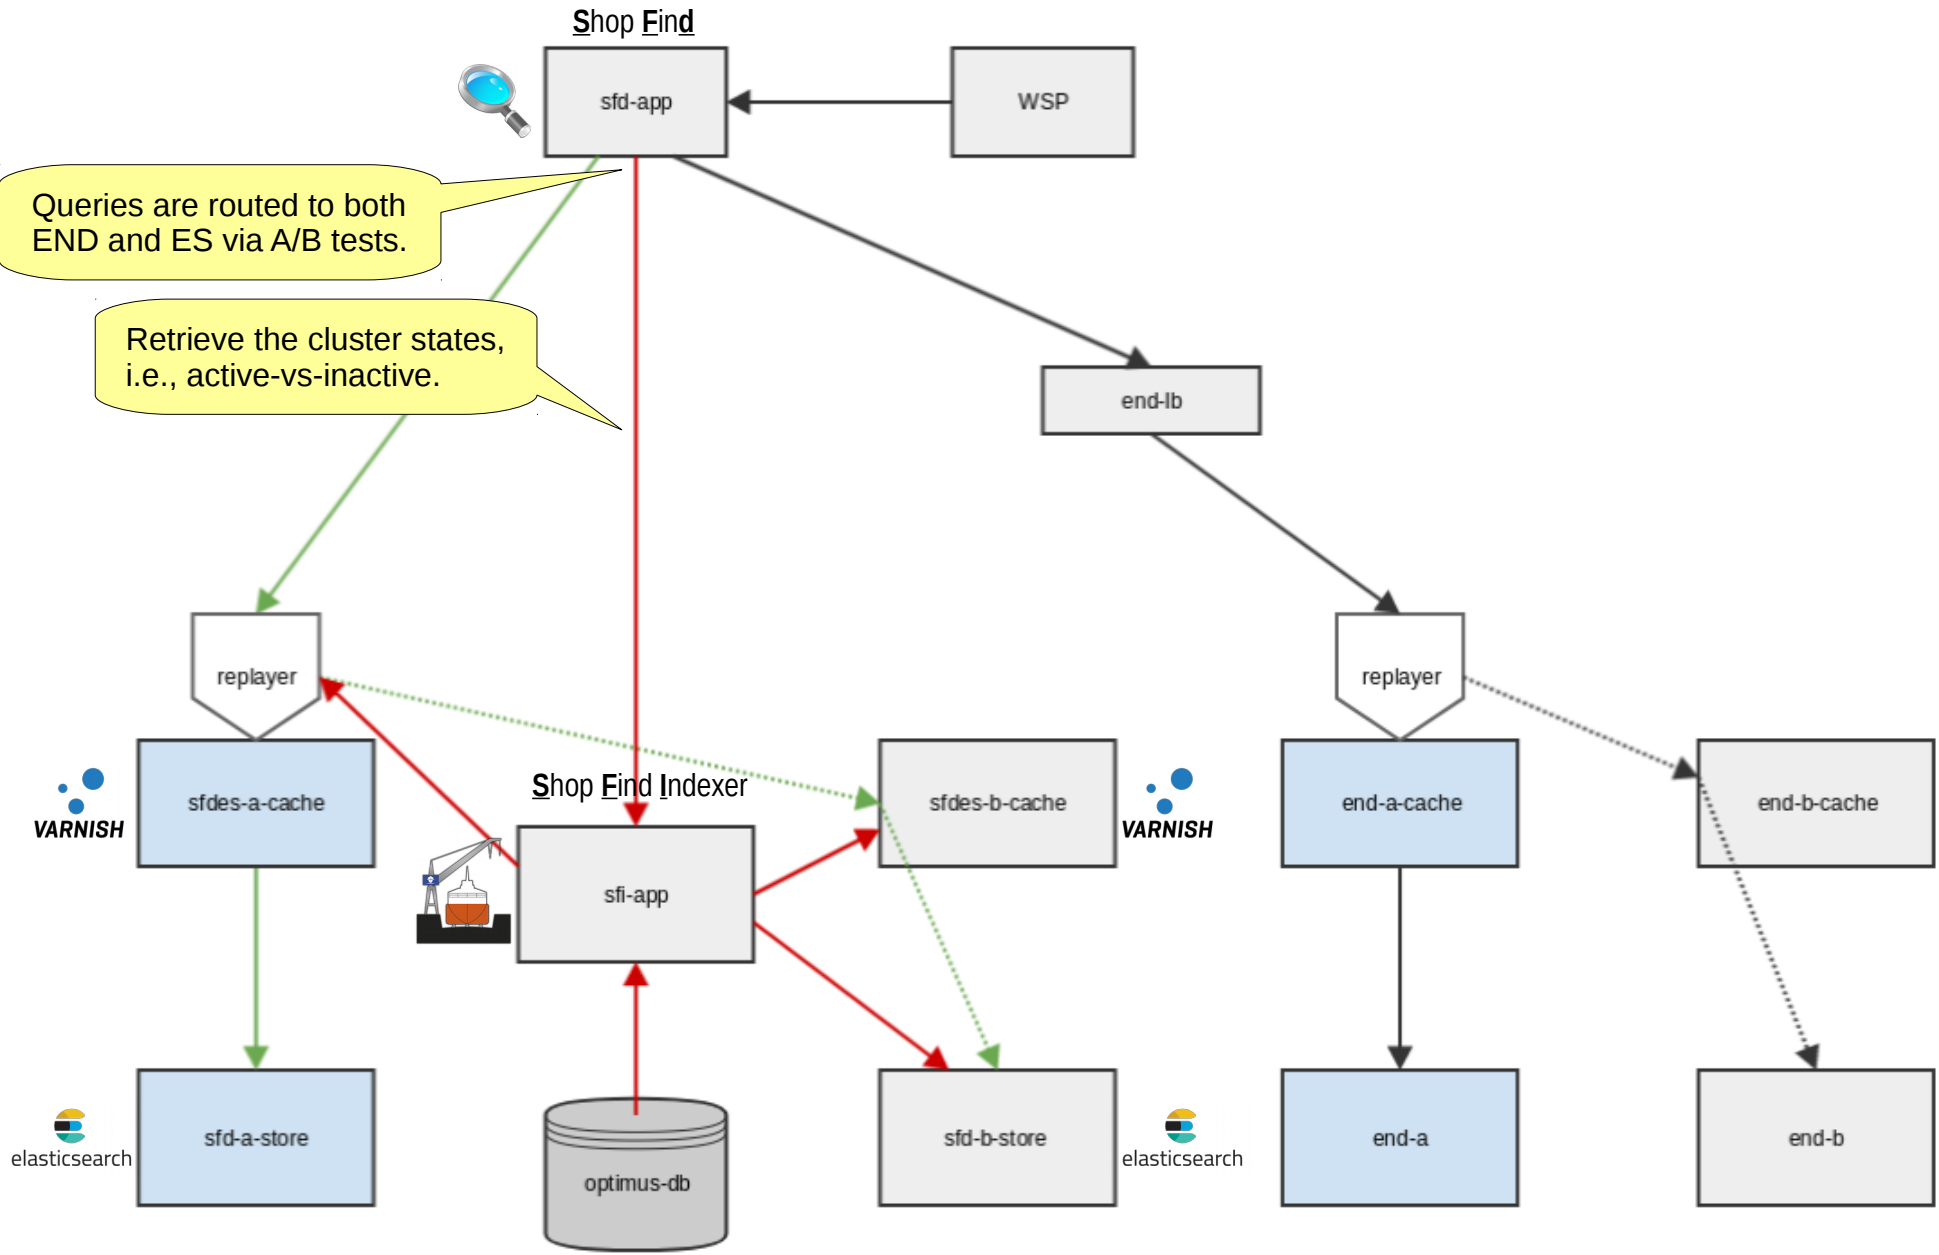
Shop Find Indexer



Queries are routed to both END and ES via A/B tests.



Long live the **king!**

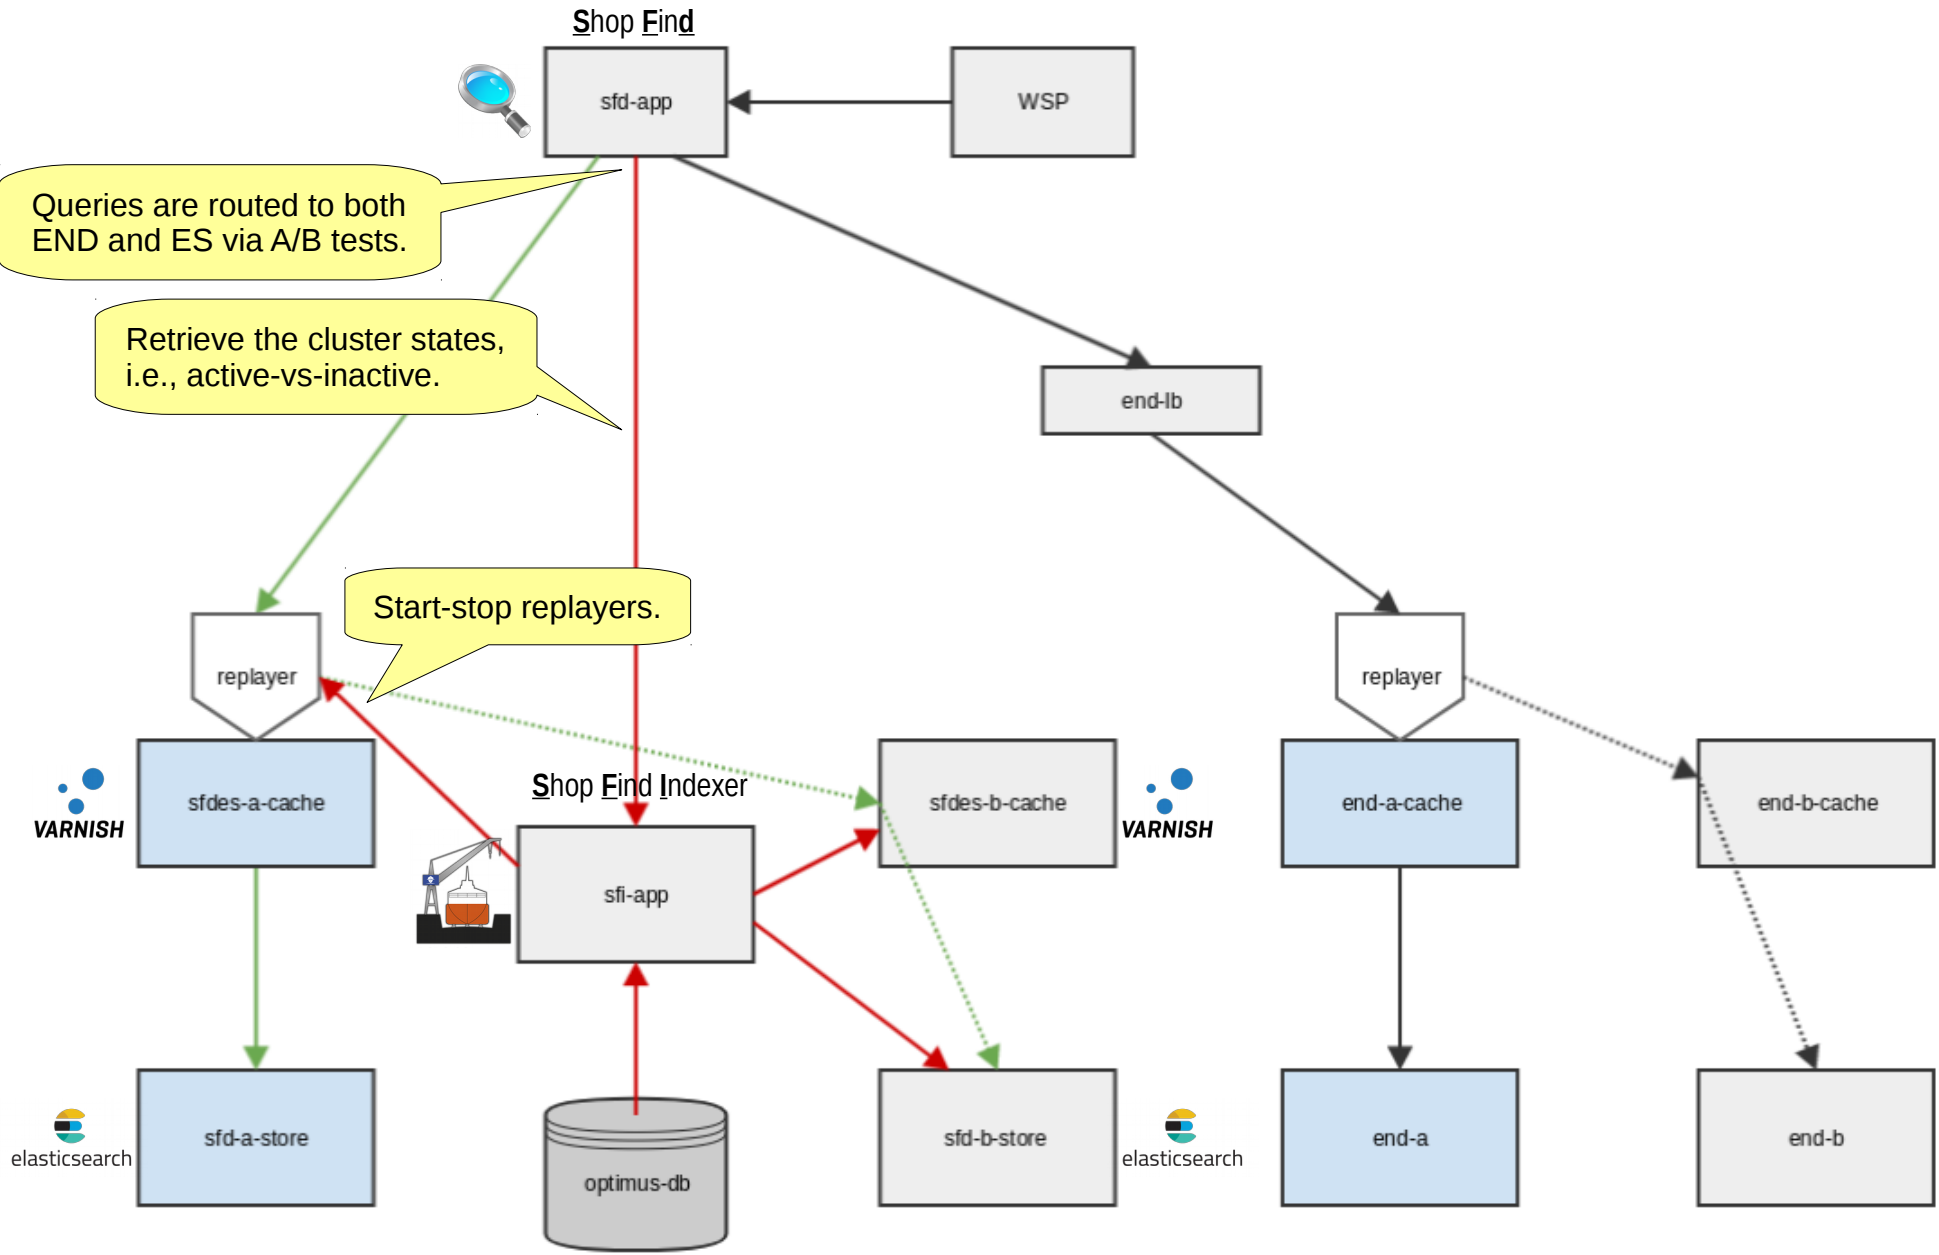


Queries are routed to both END and ES via A/B tests.

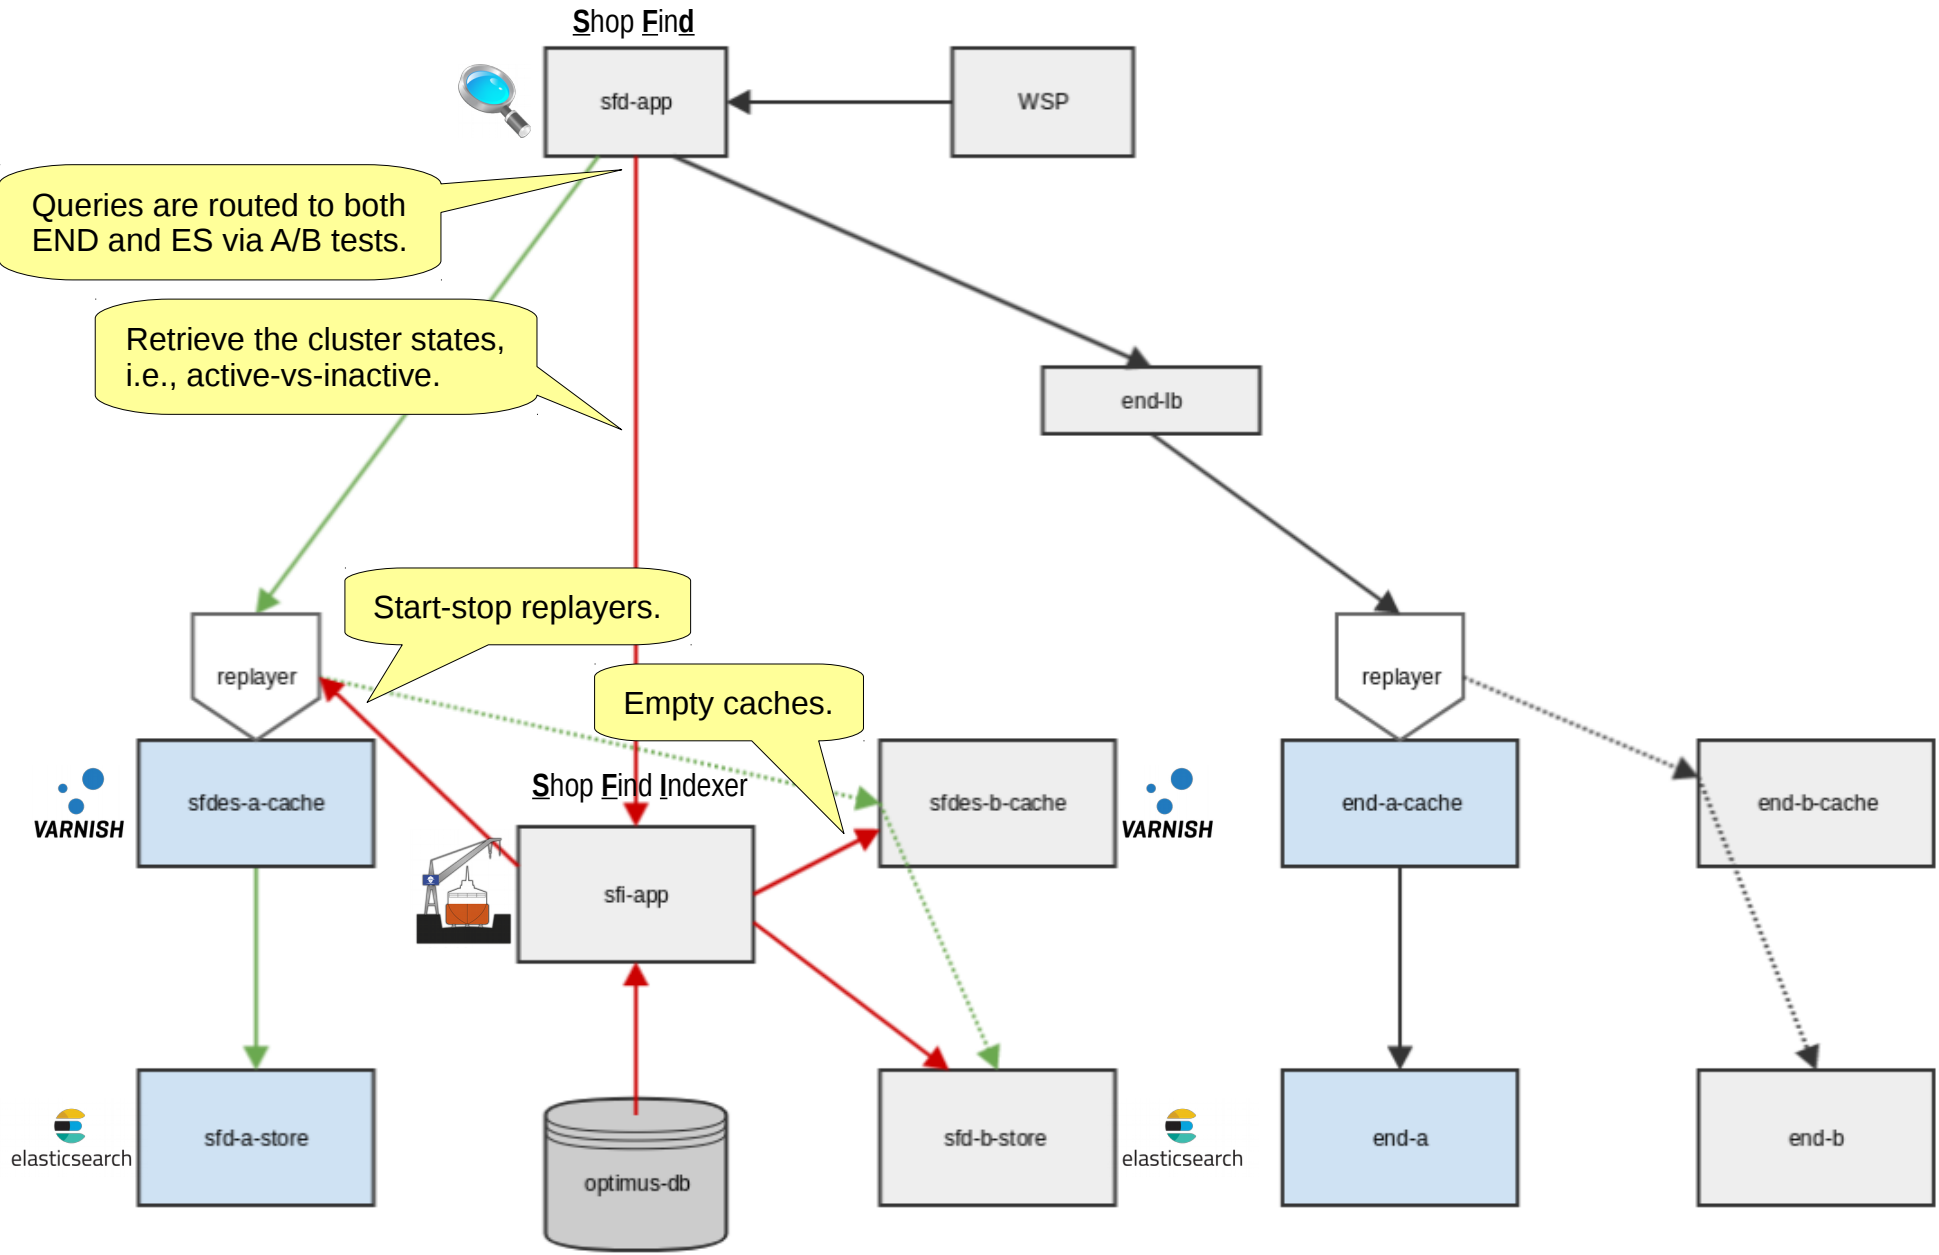
Retrieve the cluster states, i.e., active-vs-inactive.



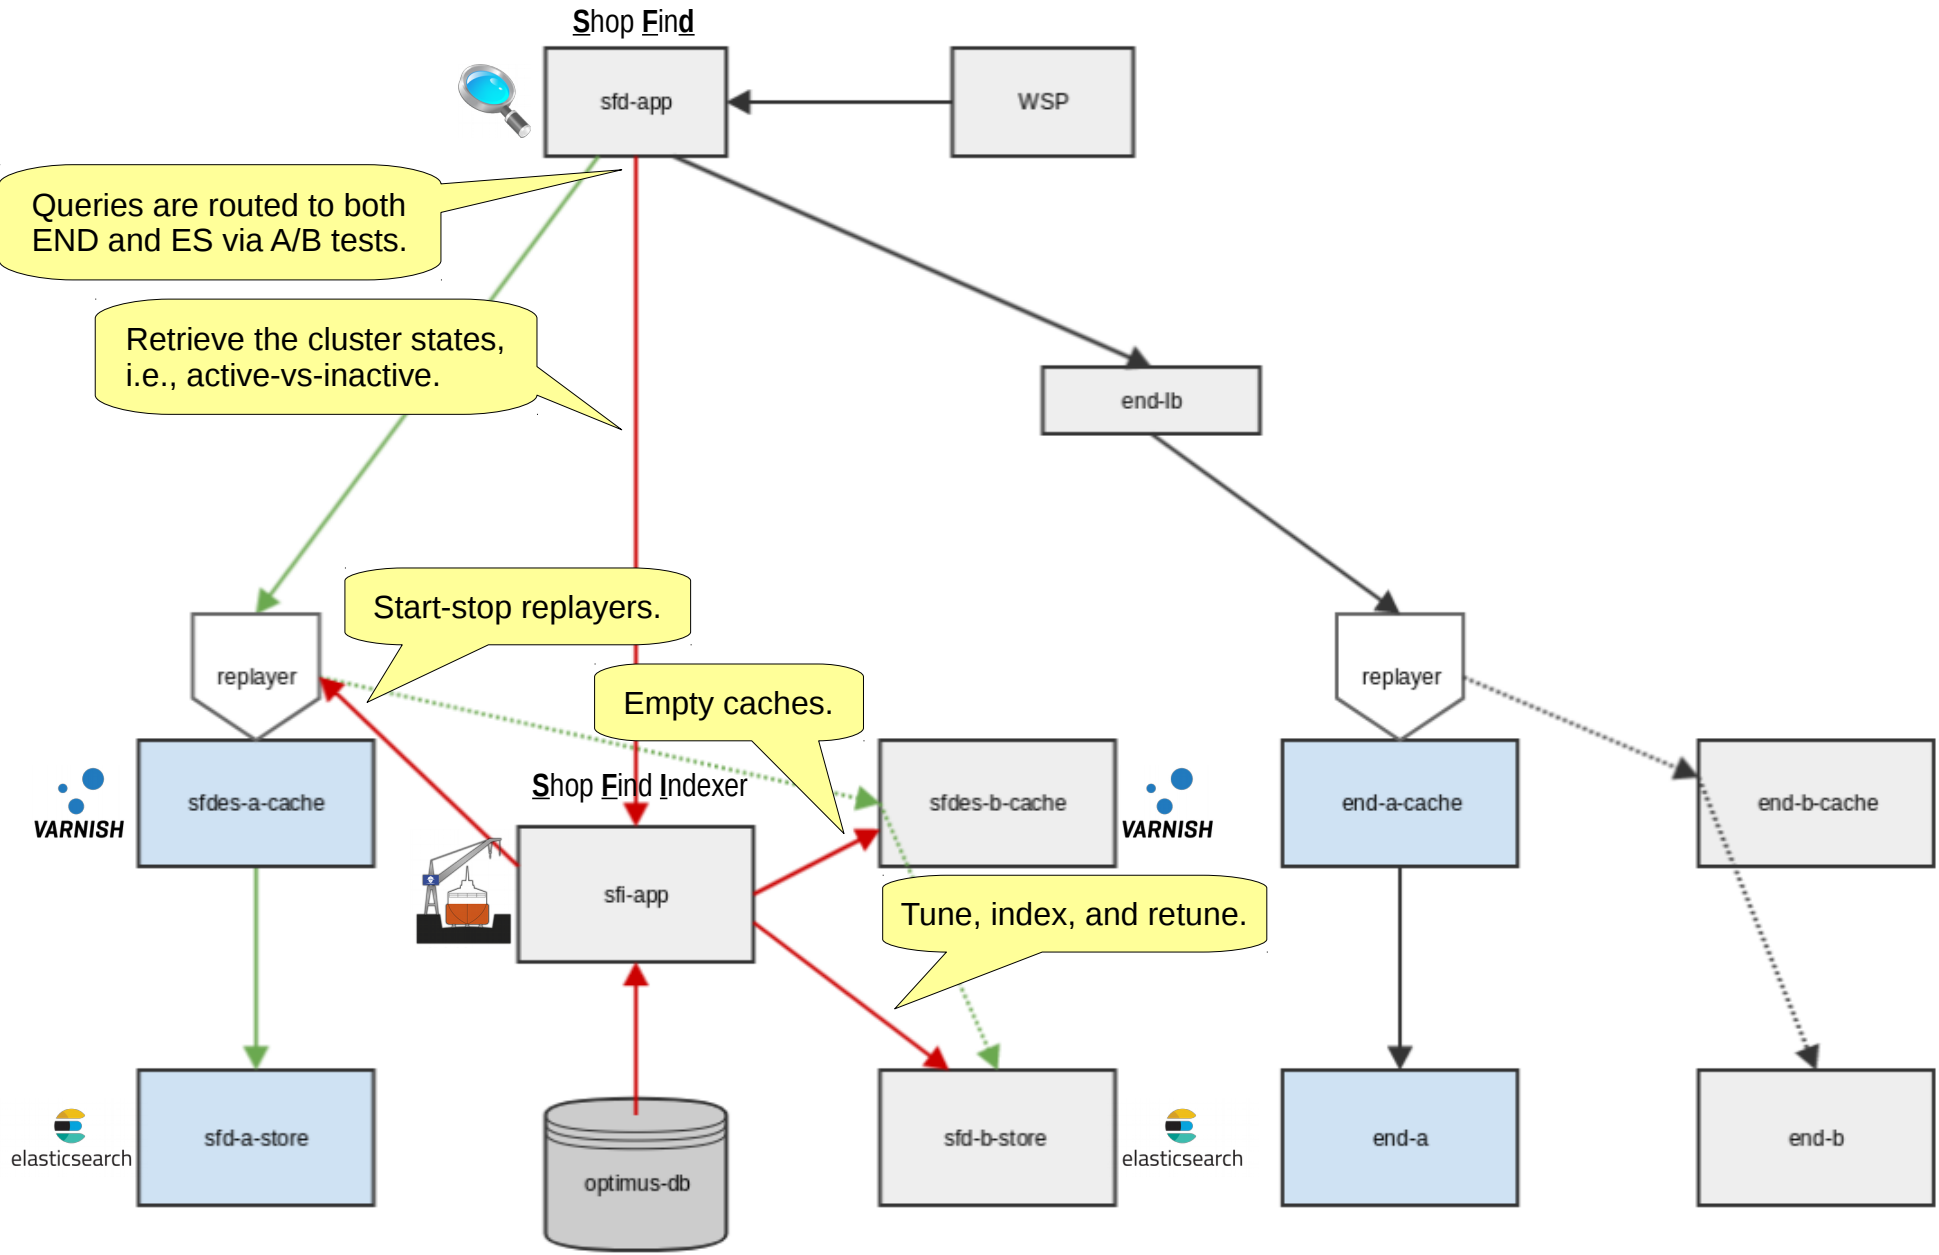
Long live the **king!**



Long live the **king!**

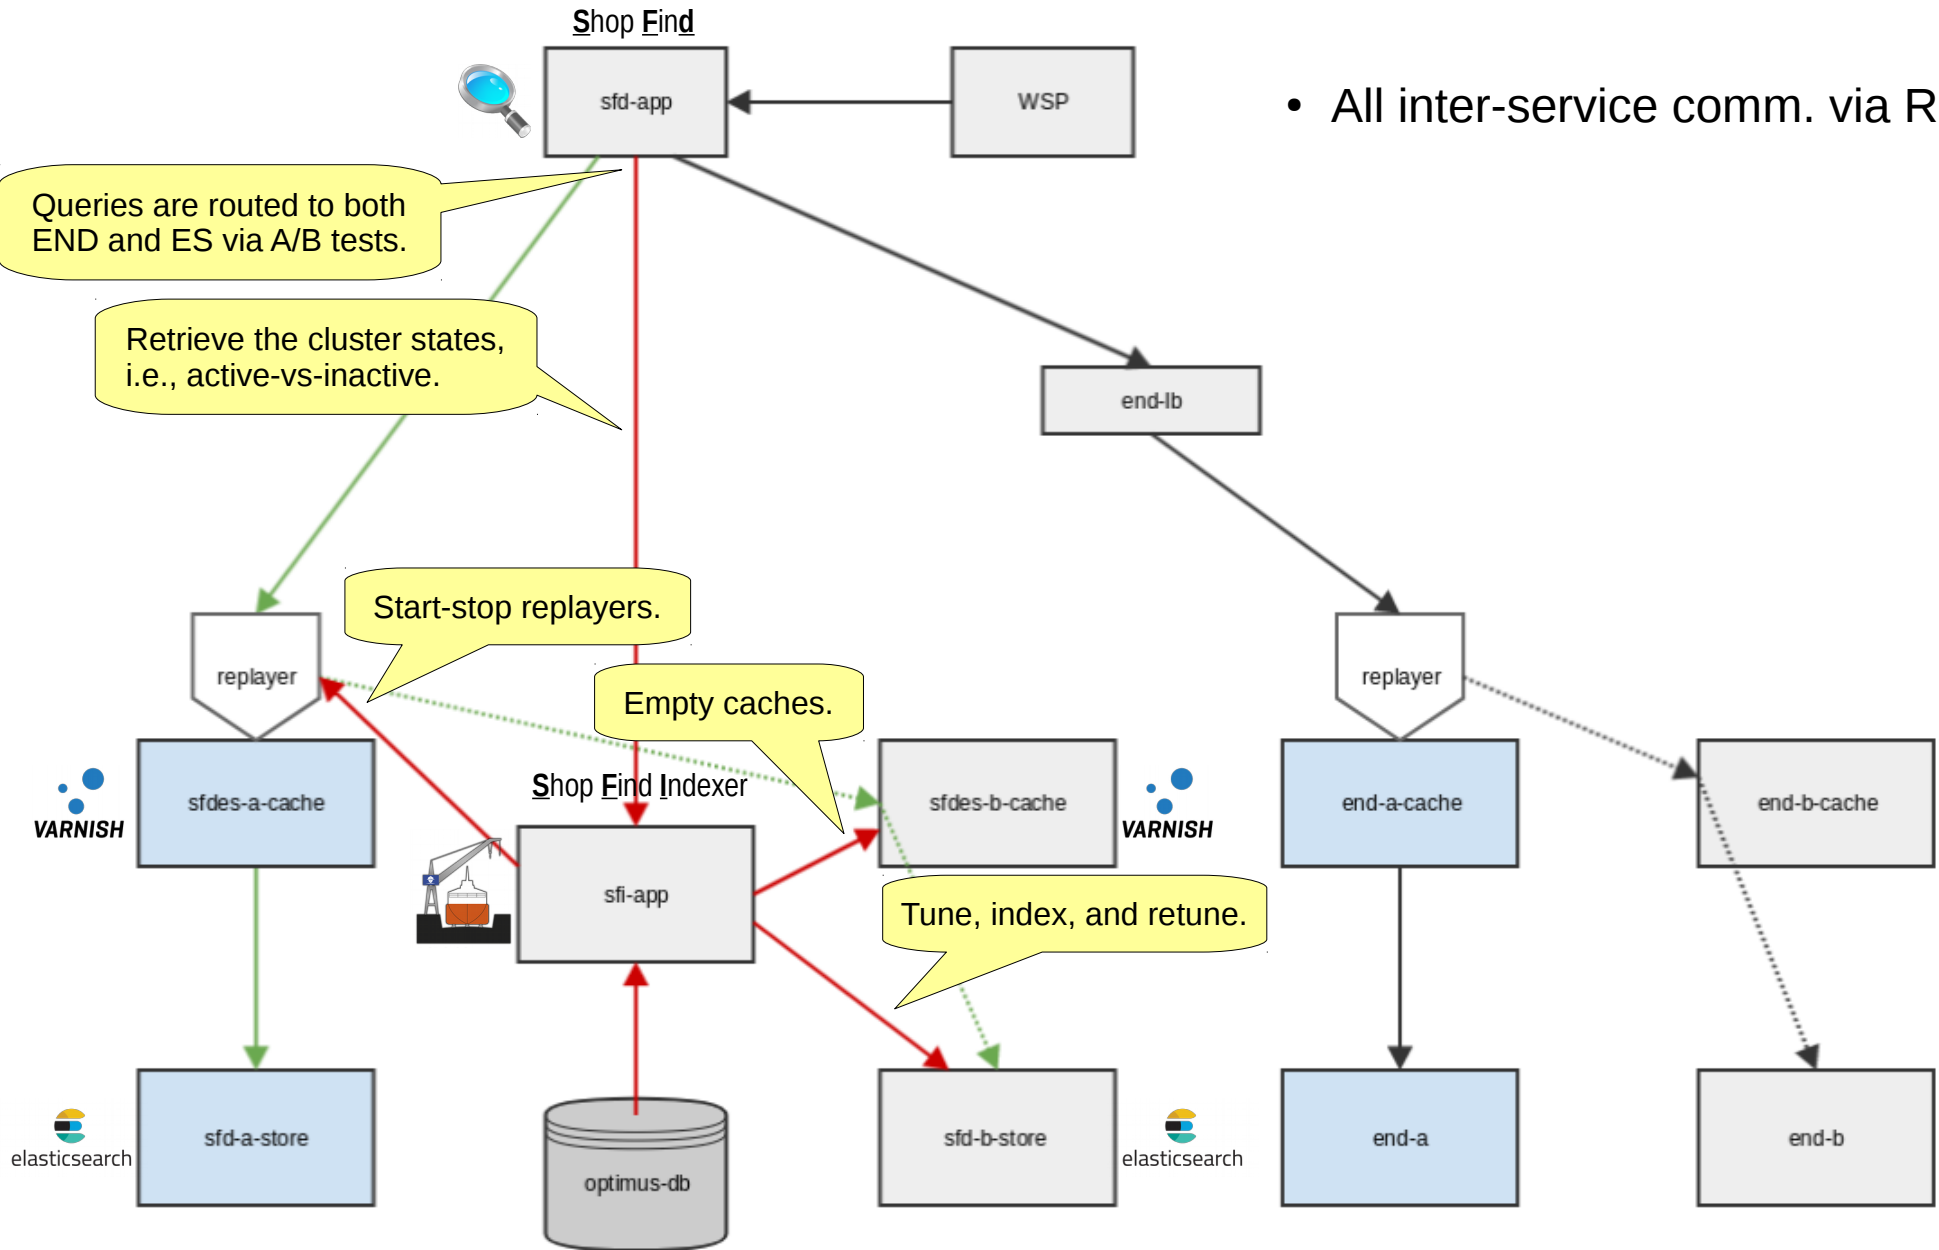


Long live the **king!**

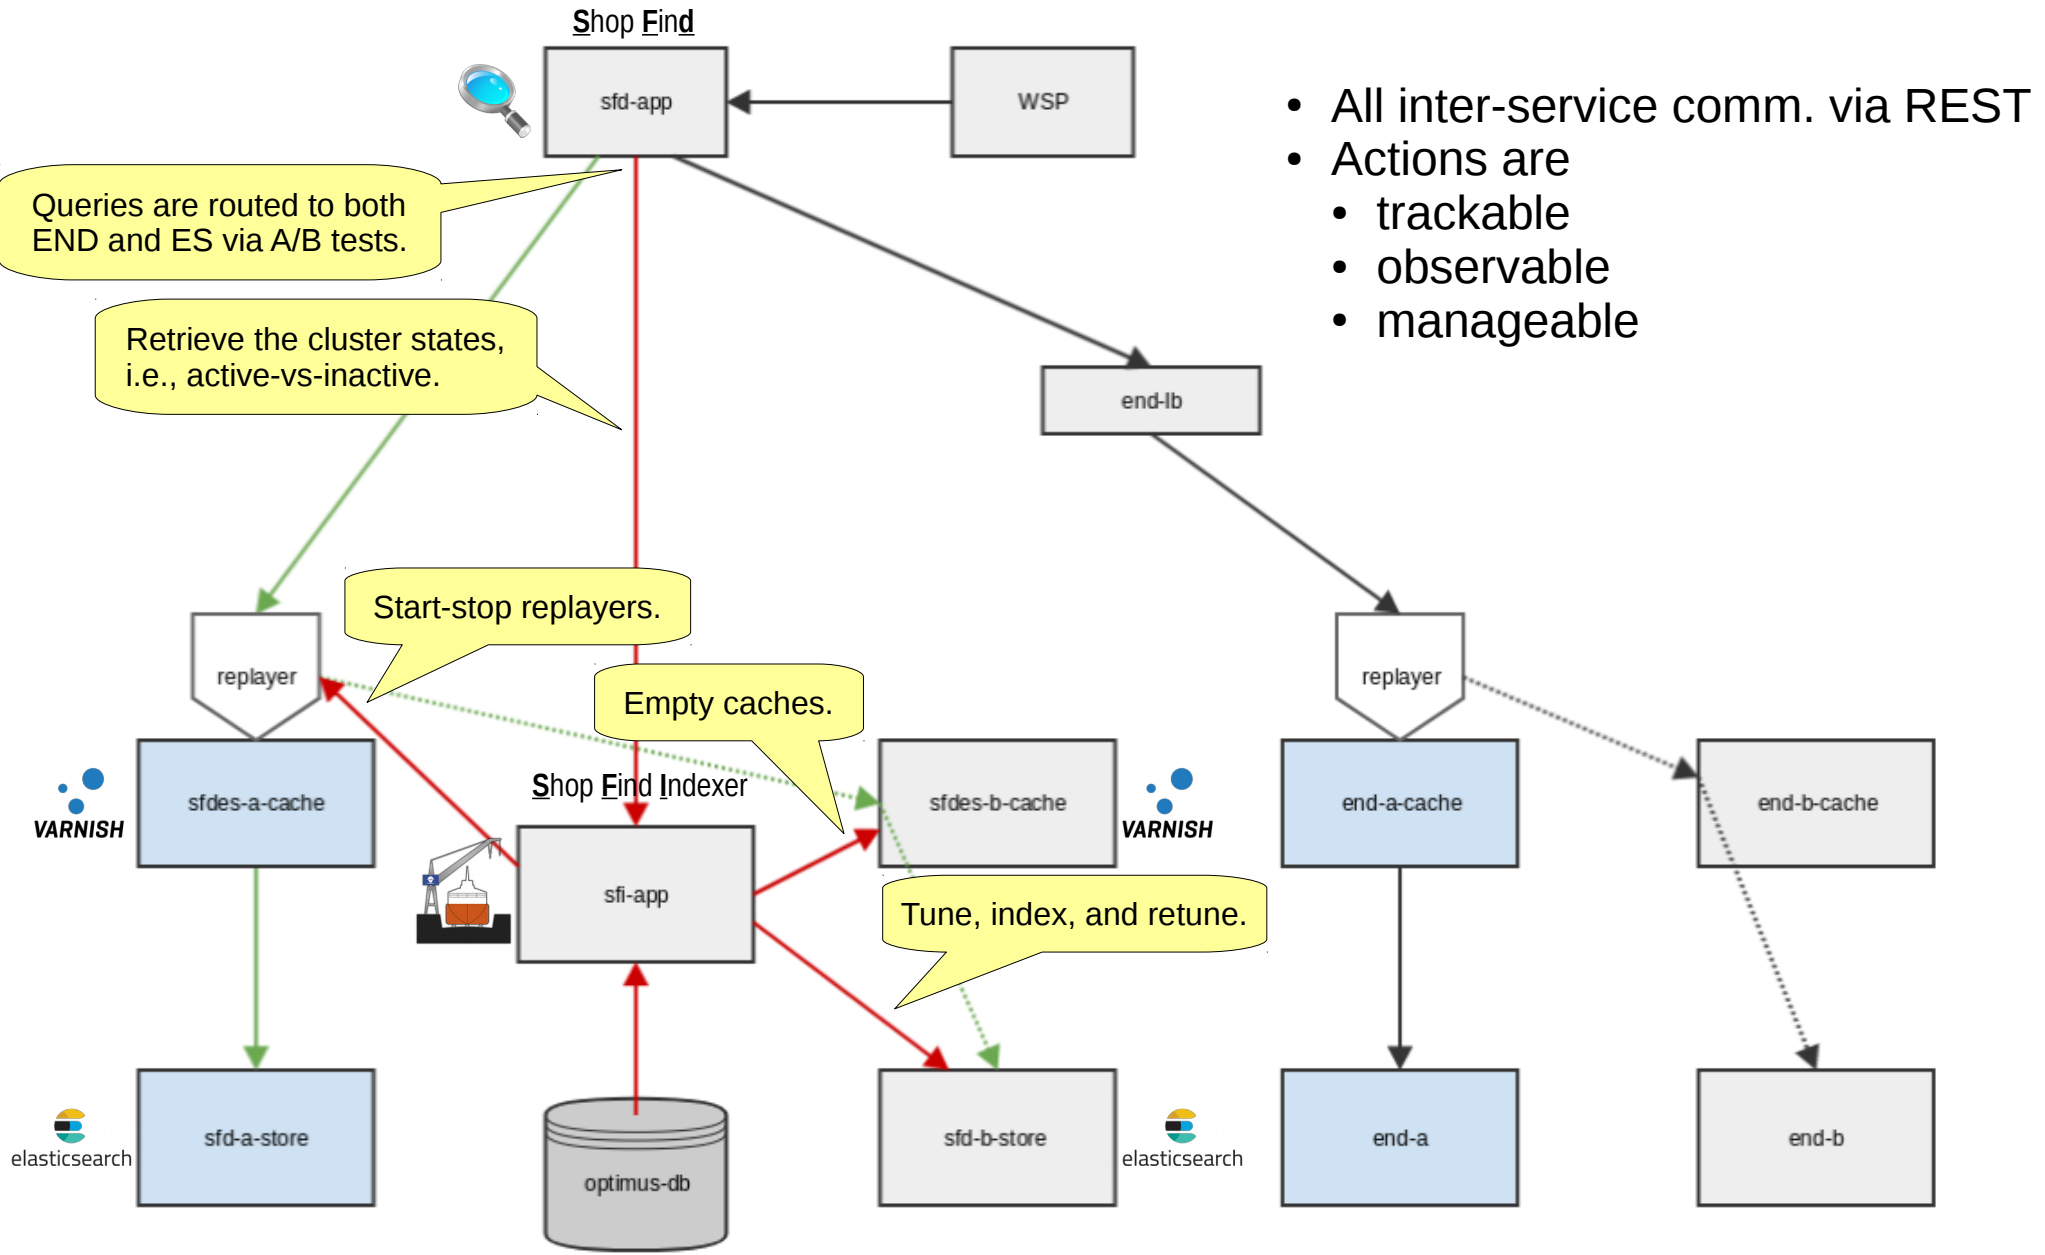


Long live the **king!**

- All inter-service comm. via REST



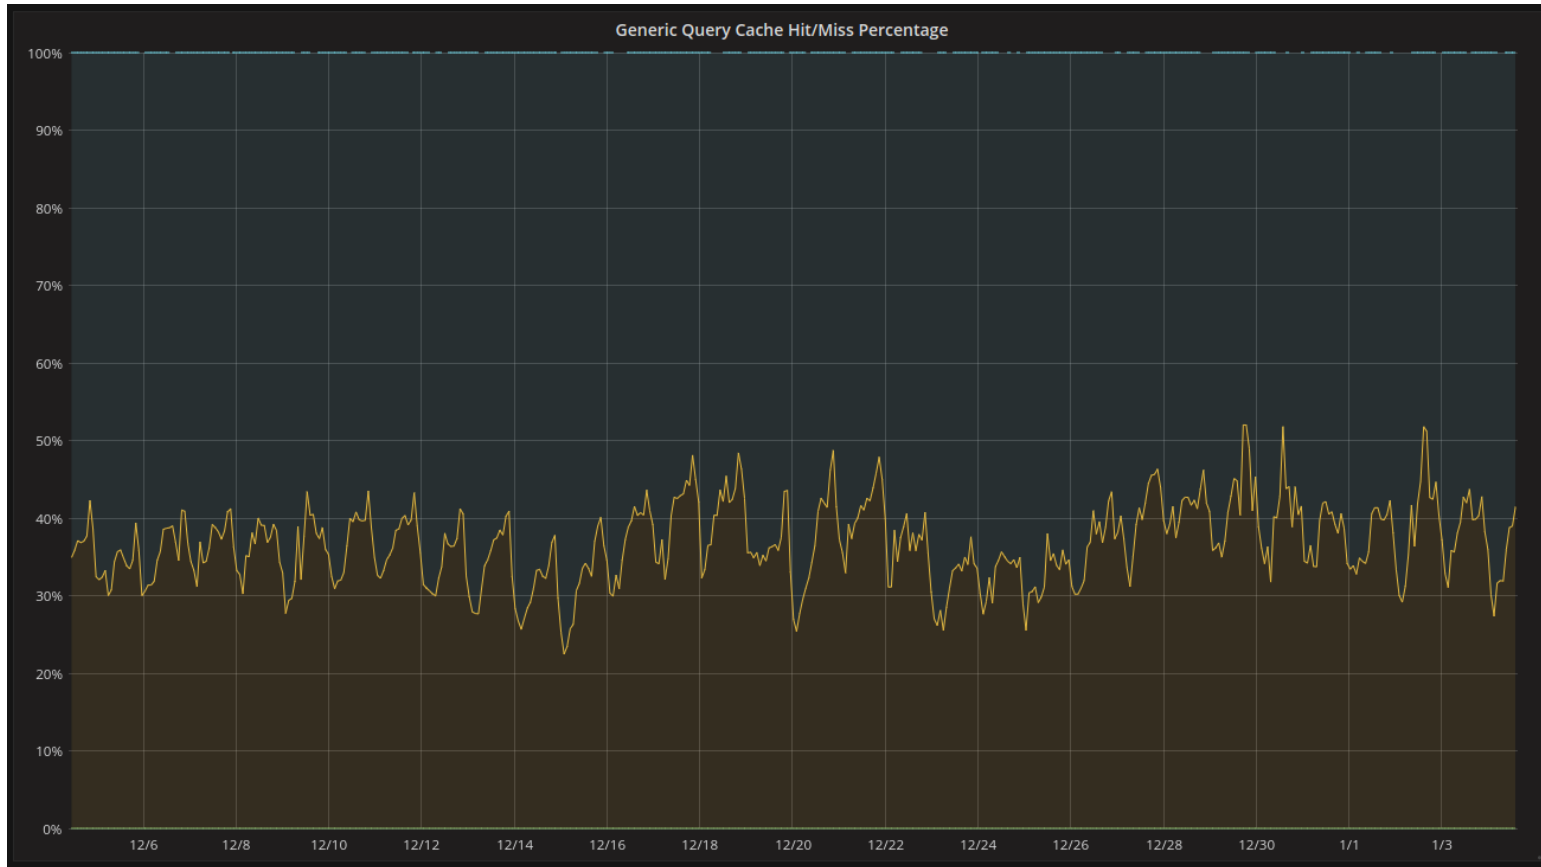
Long live the **king!**



- All inter-service comm. via REST
- Actions are
 - trackable
 - observable
 - manageable

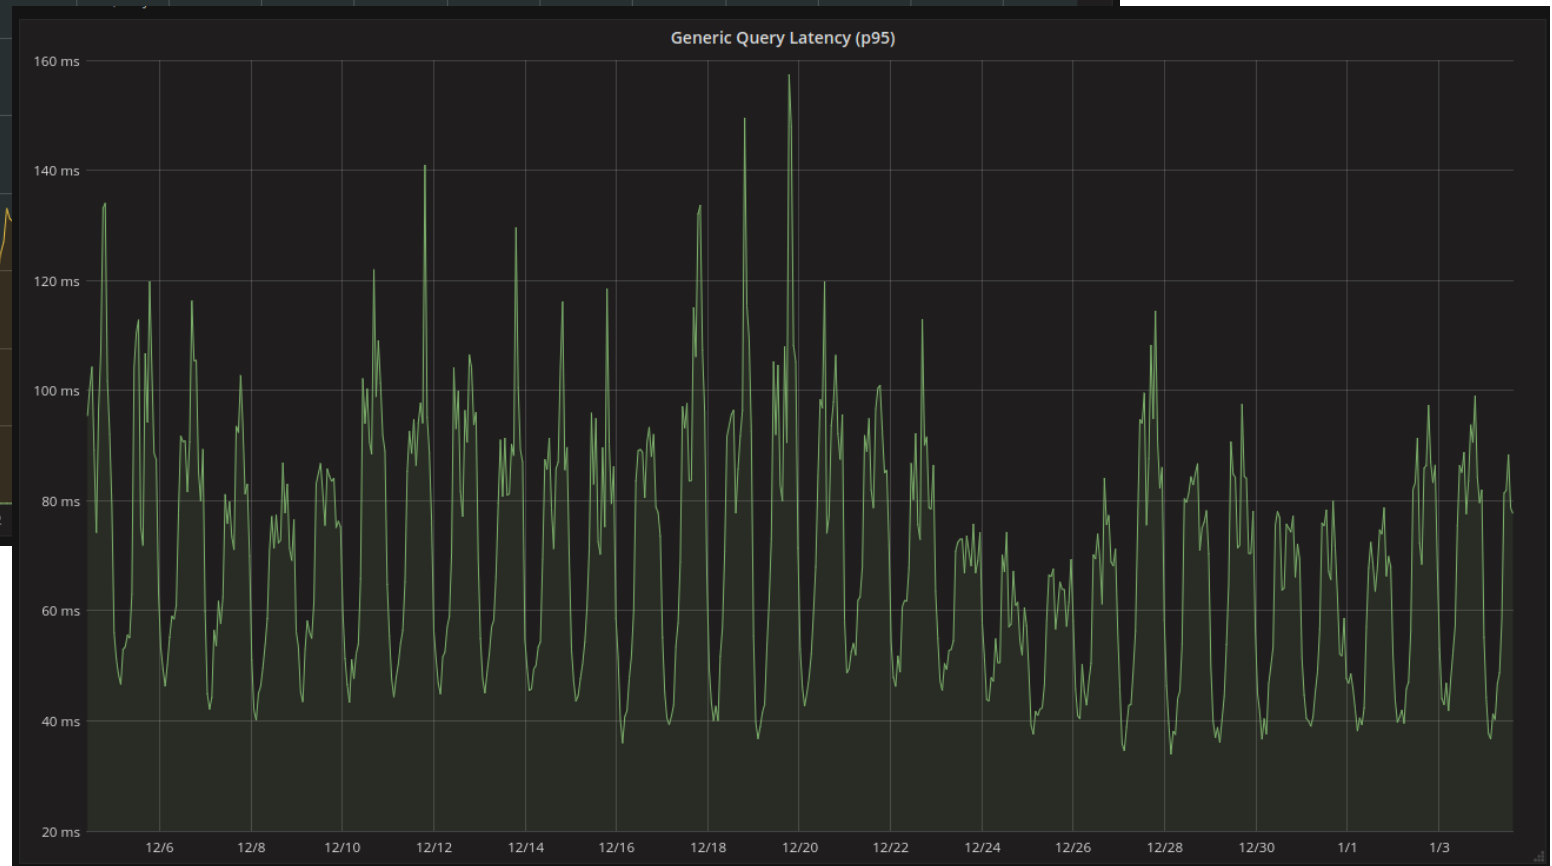
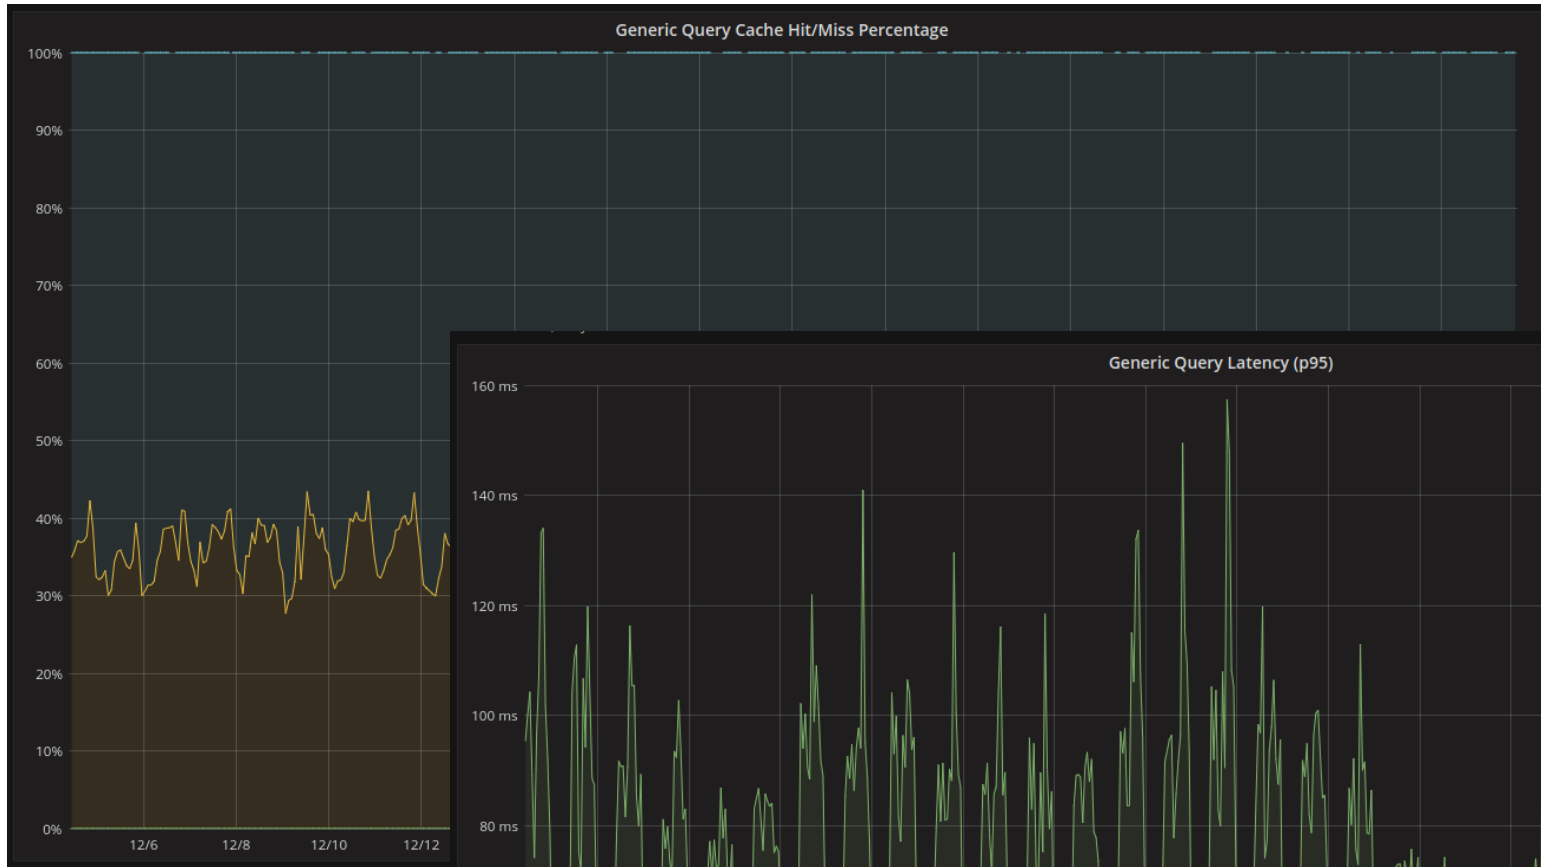
Did it work? (1/2)

(for generic queries)




Did it work? (1/2)

(for generic queries)





Did it work? (2/2)

(for single product queries)



Did it work? (2/2)

(for single product queries)



Questions?

(Do you still want to have caching?)