

Addressing IPv6

A CDN perspective

@jta joão taveira Araújo

RIPE 74



clickbait



Addressing IPv6

A CDN perspective

@jta joão taveira Araújo

RIPE 74



Addressing IPv6

A CDN perspective

@jta joão taveira arújo

your mileage may vary



RIPE 74




GitHub



 *shopify*

vimeo



Pinterest

stripe

imgur

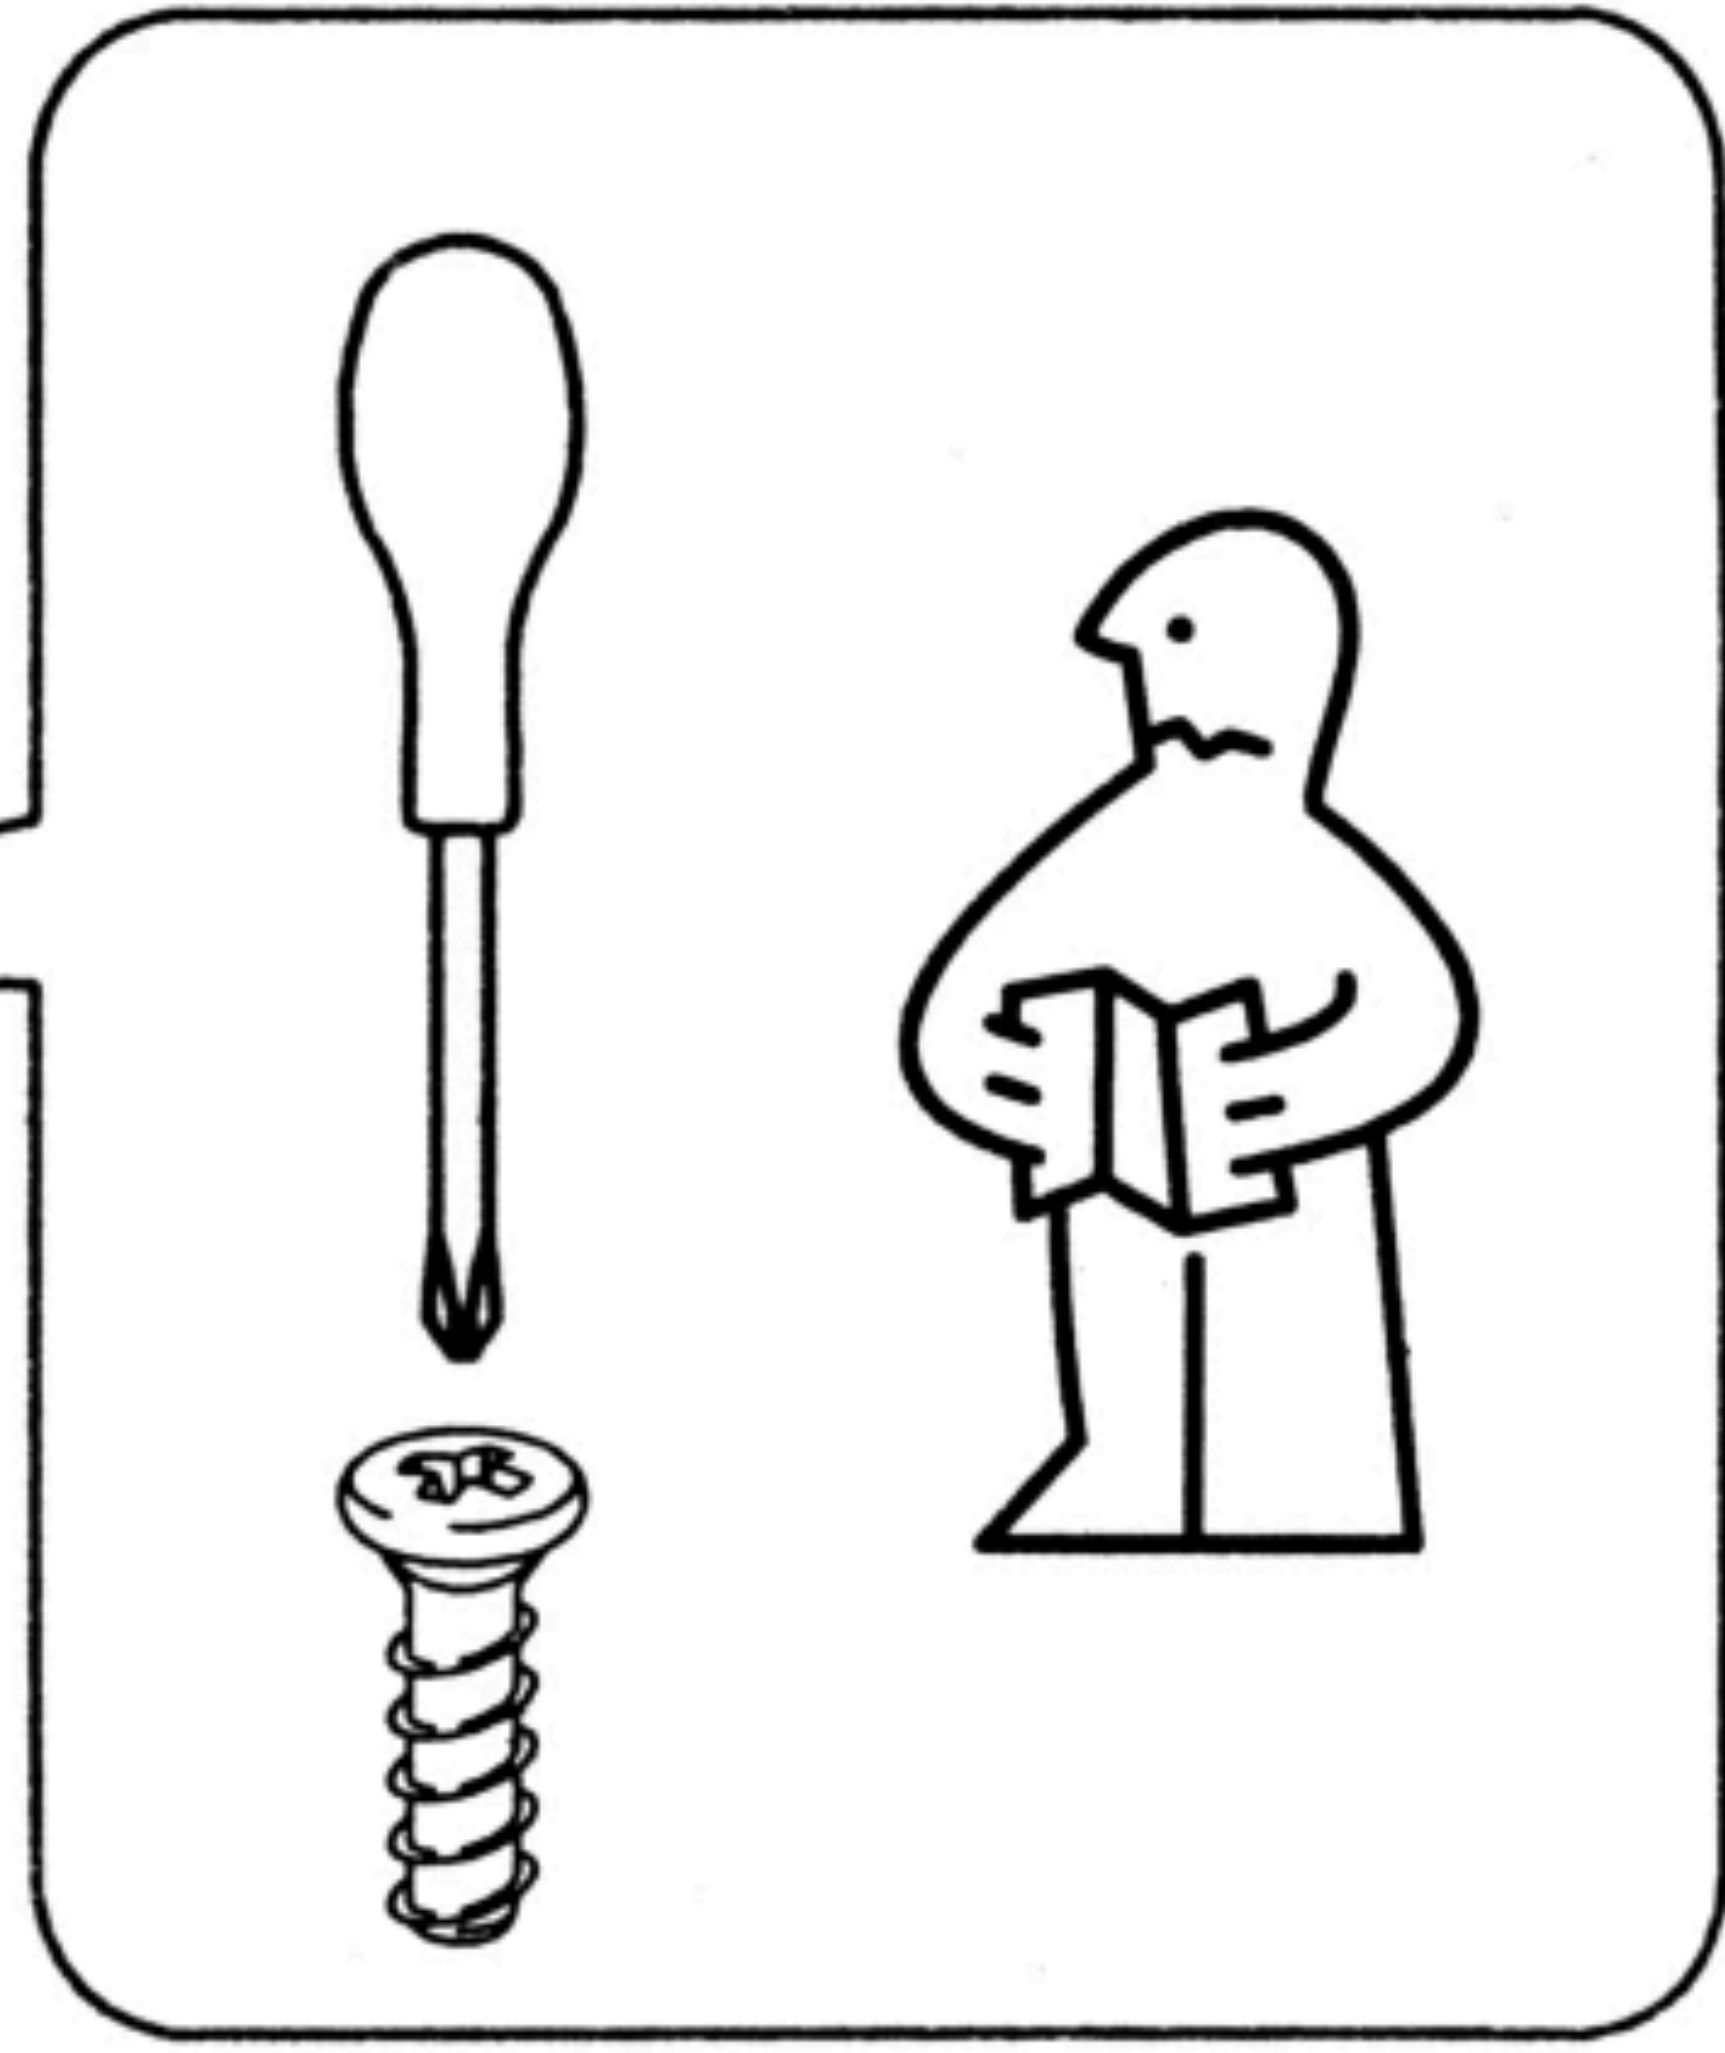
@fastly

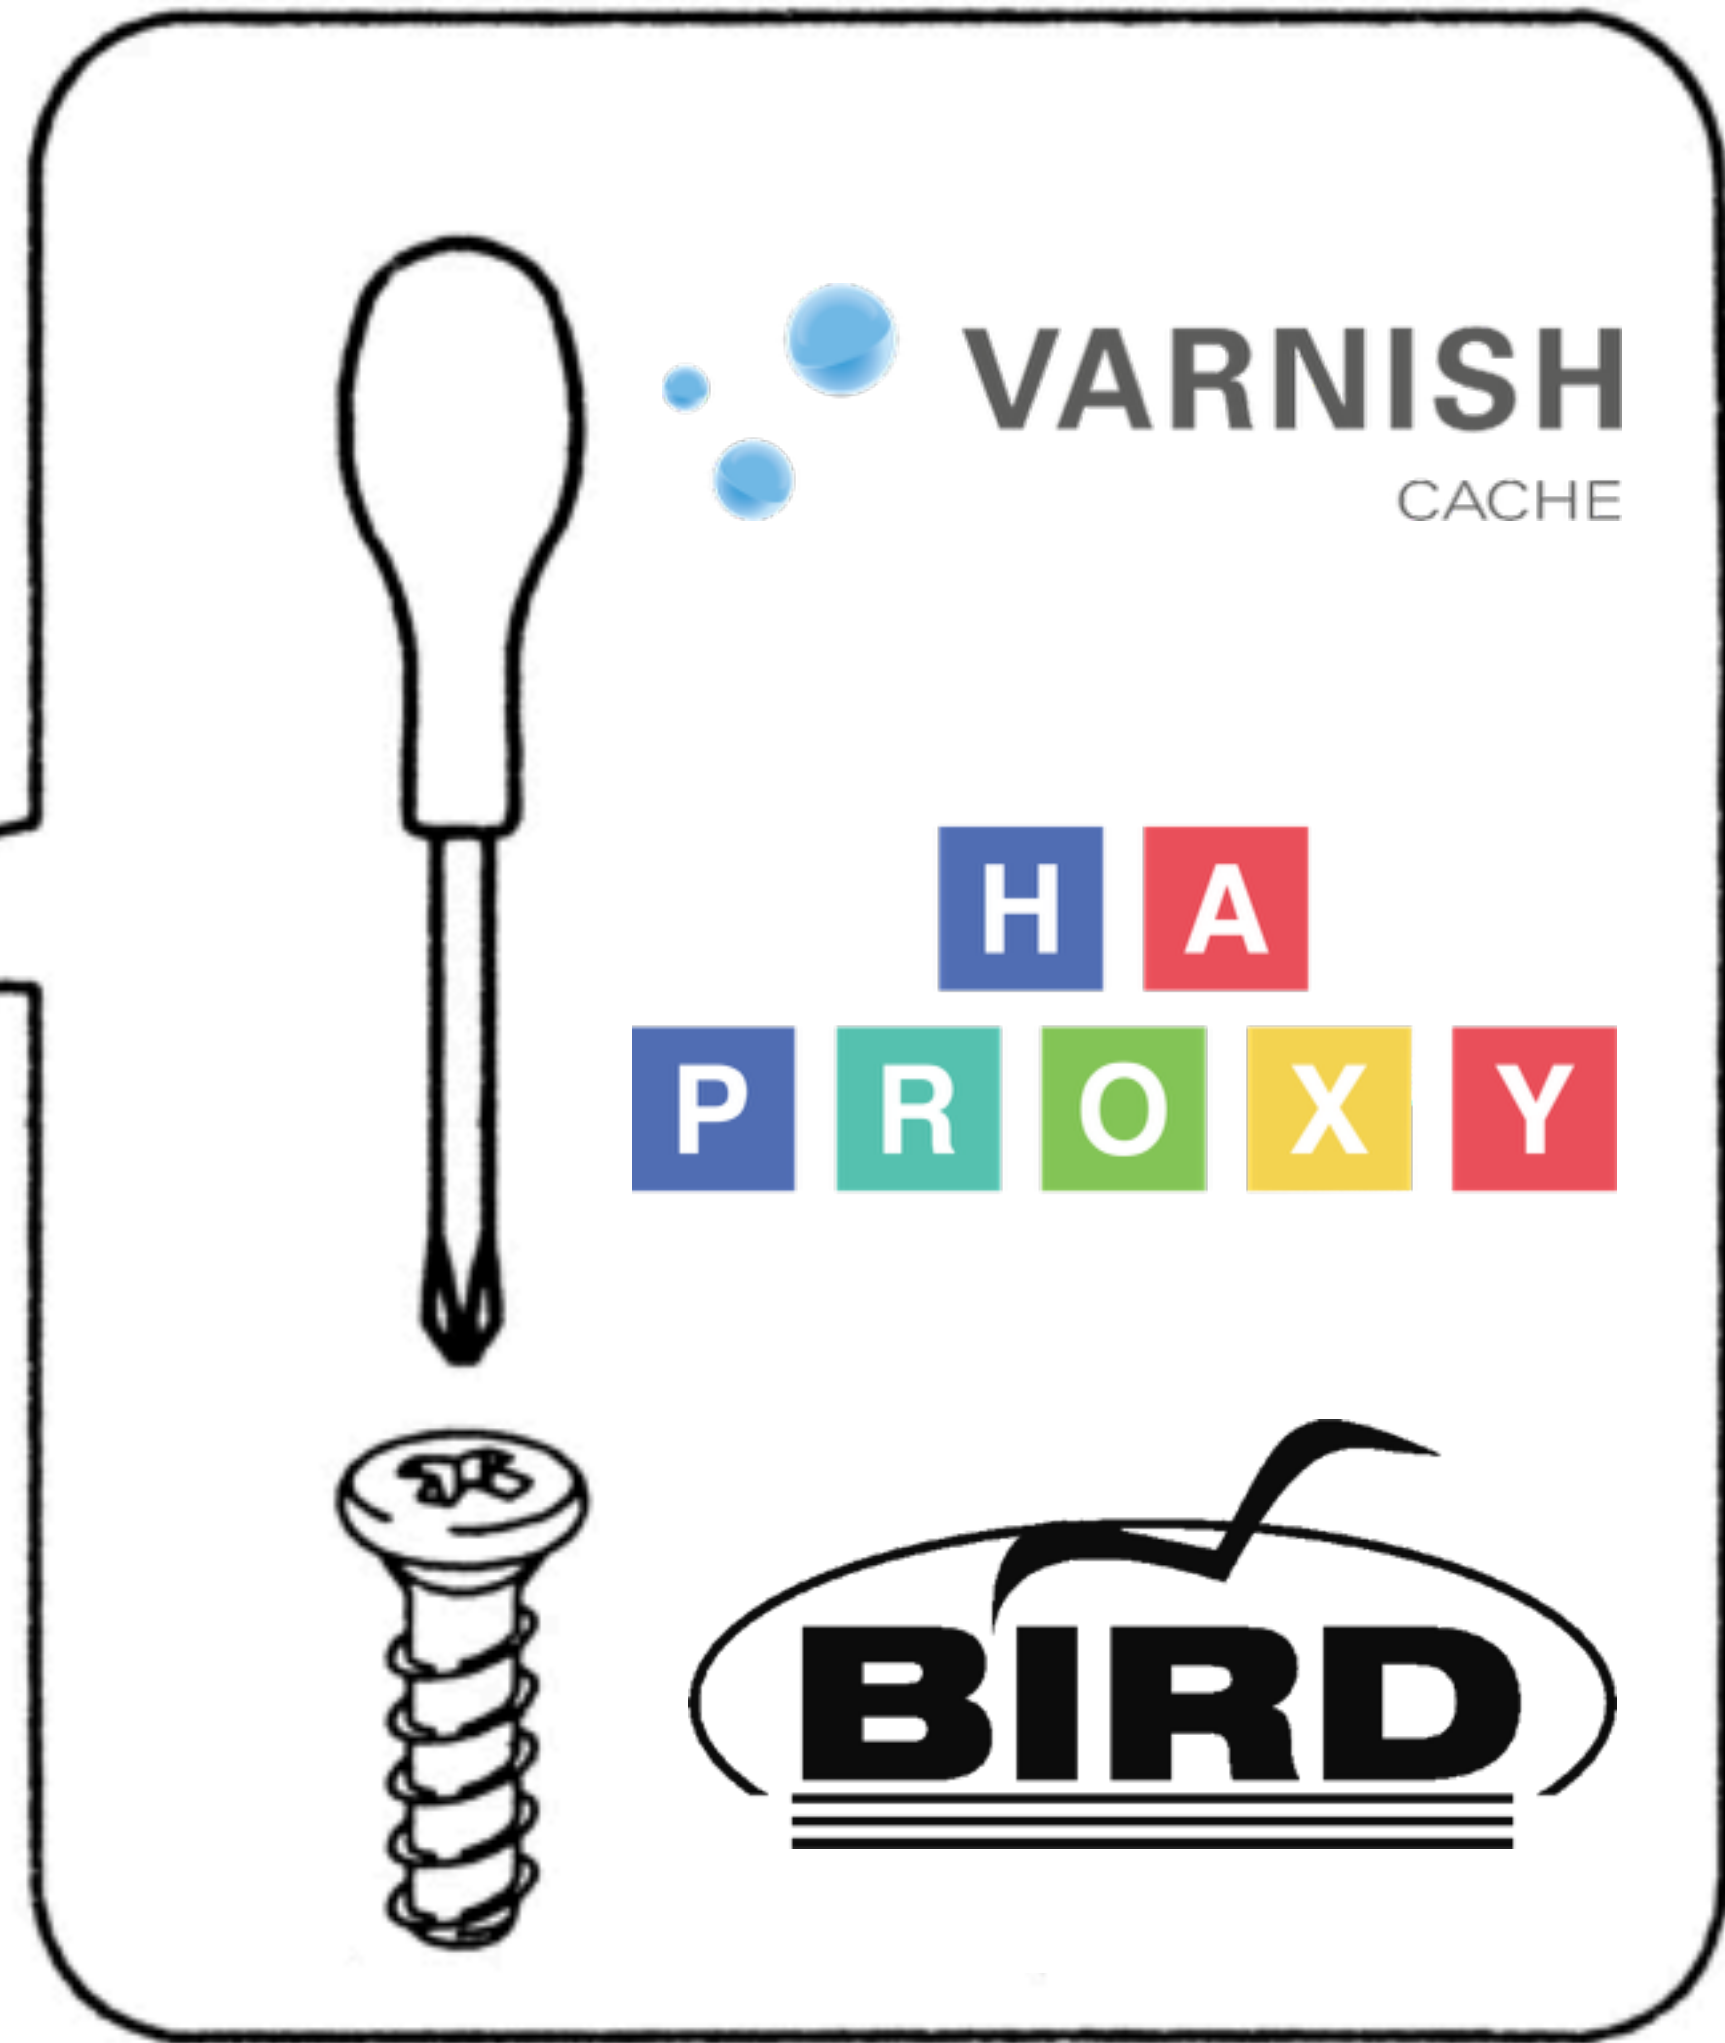
 reddit

theguardian

 airbnb







SFO


LHR






SFO
LHR

Unicast model




2012




clients

Points of Presence




SJC




LAX




Unicast model




2012



Prefix announcements




SJC





LAX

Unicast model




2012




DNS



IP address




SJC




LAX


Unicast model




2012




IP allocation





2012




IP allocation



2012

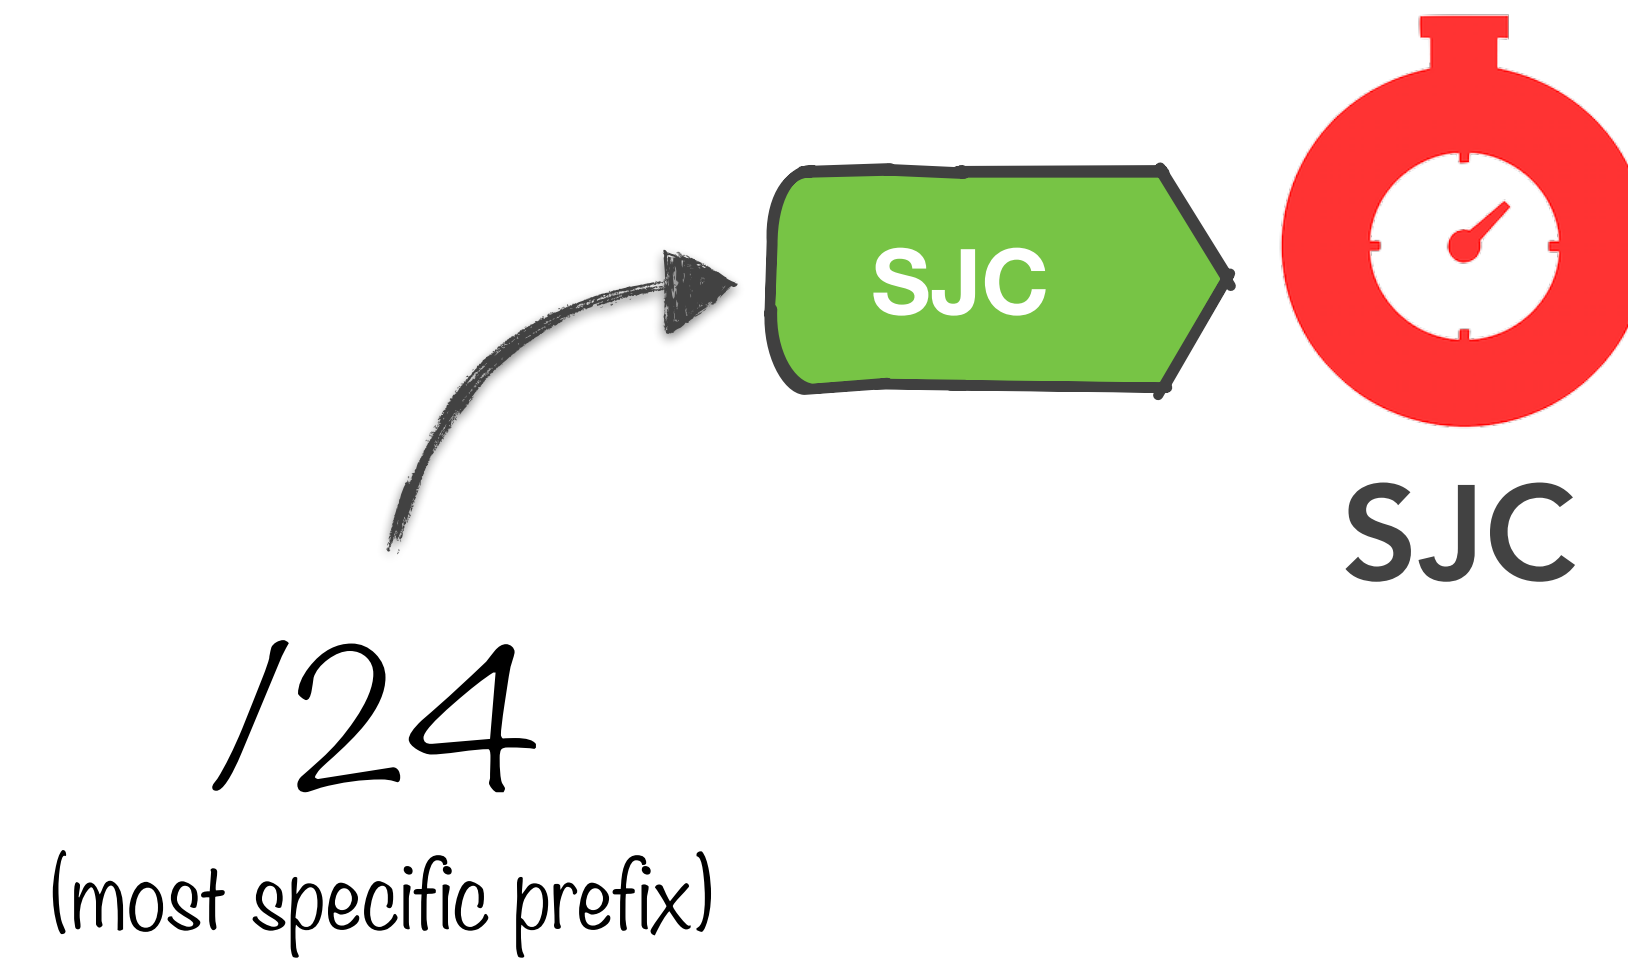



IP allocation




2012



Hosts



IP allocation



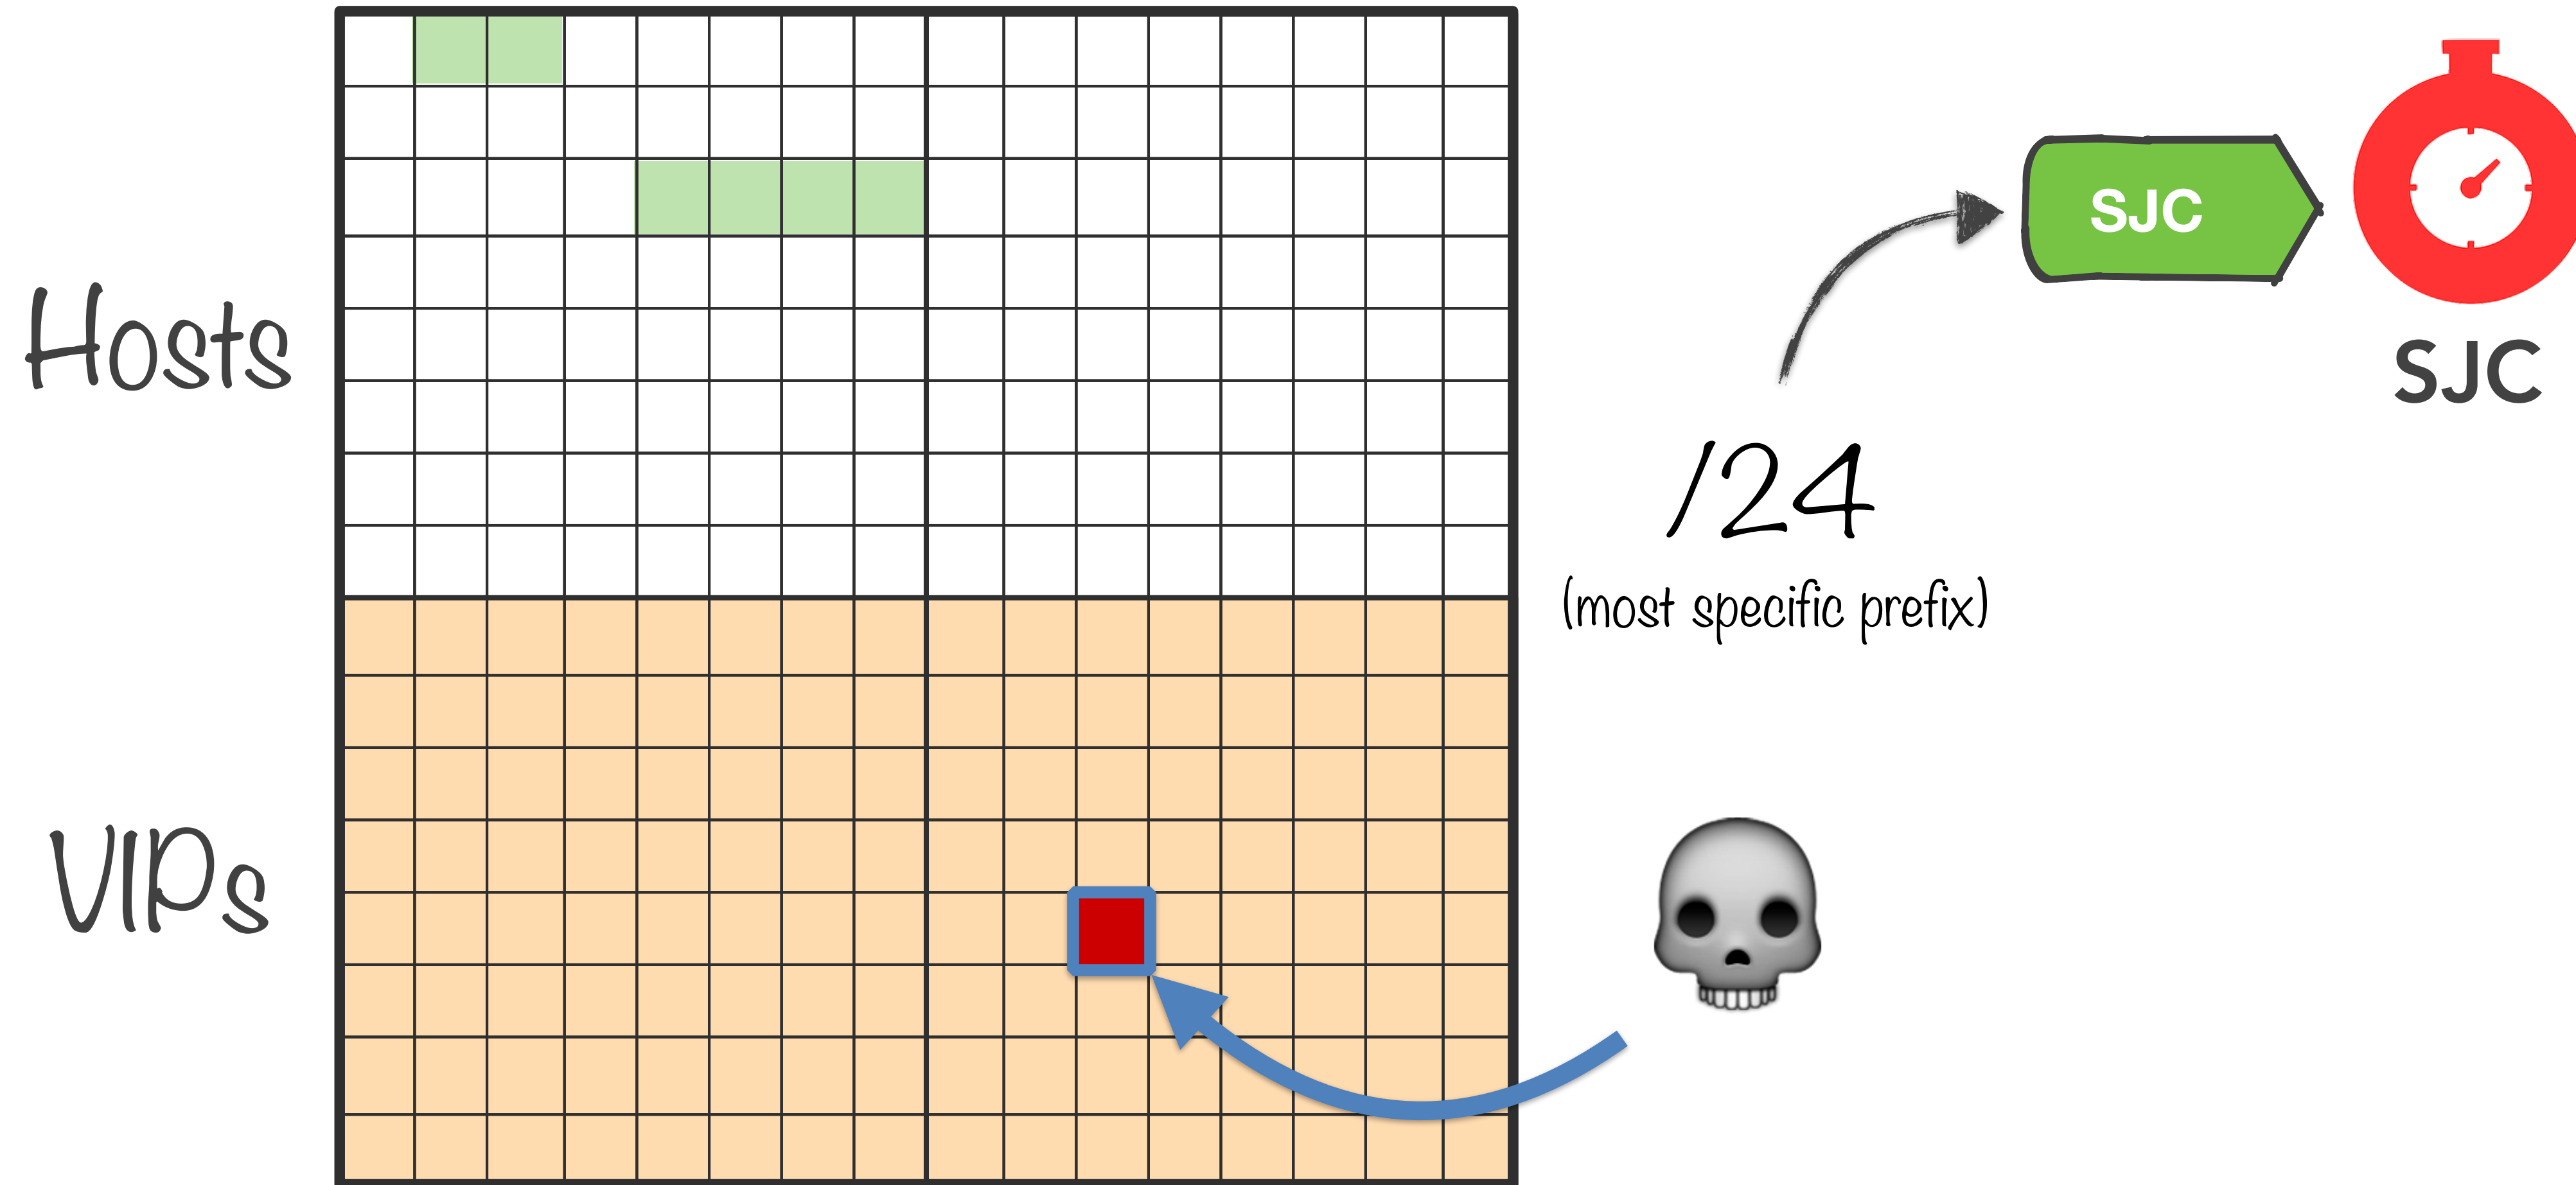
2012




Fate-sharing






2012



Everything is fine




2013

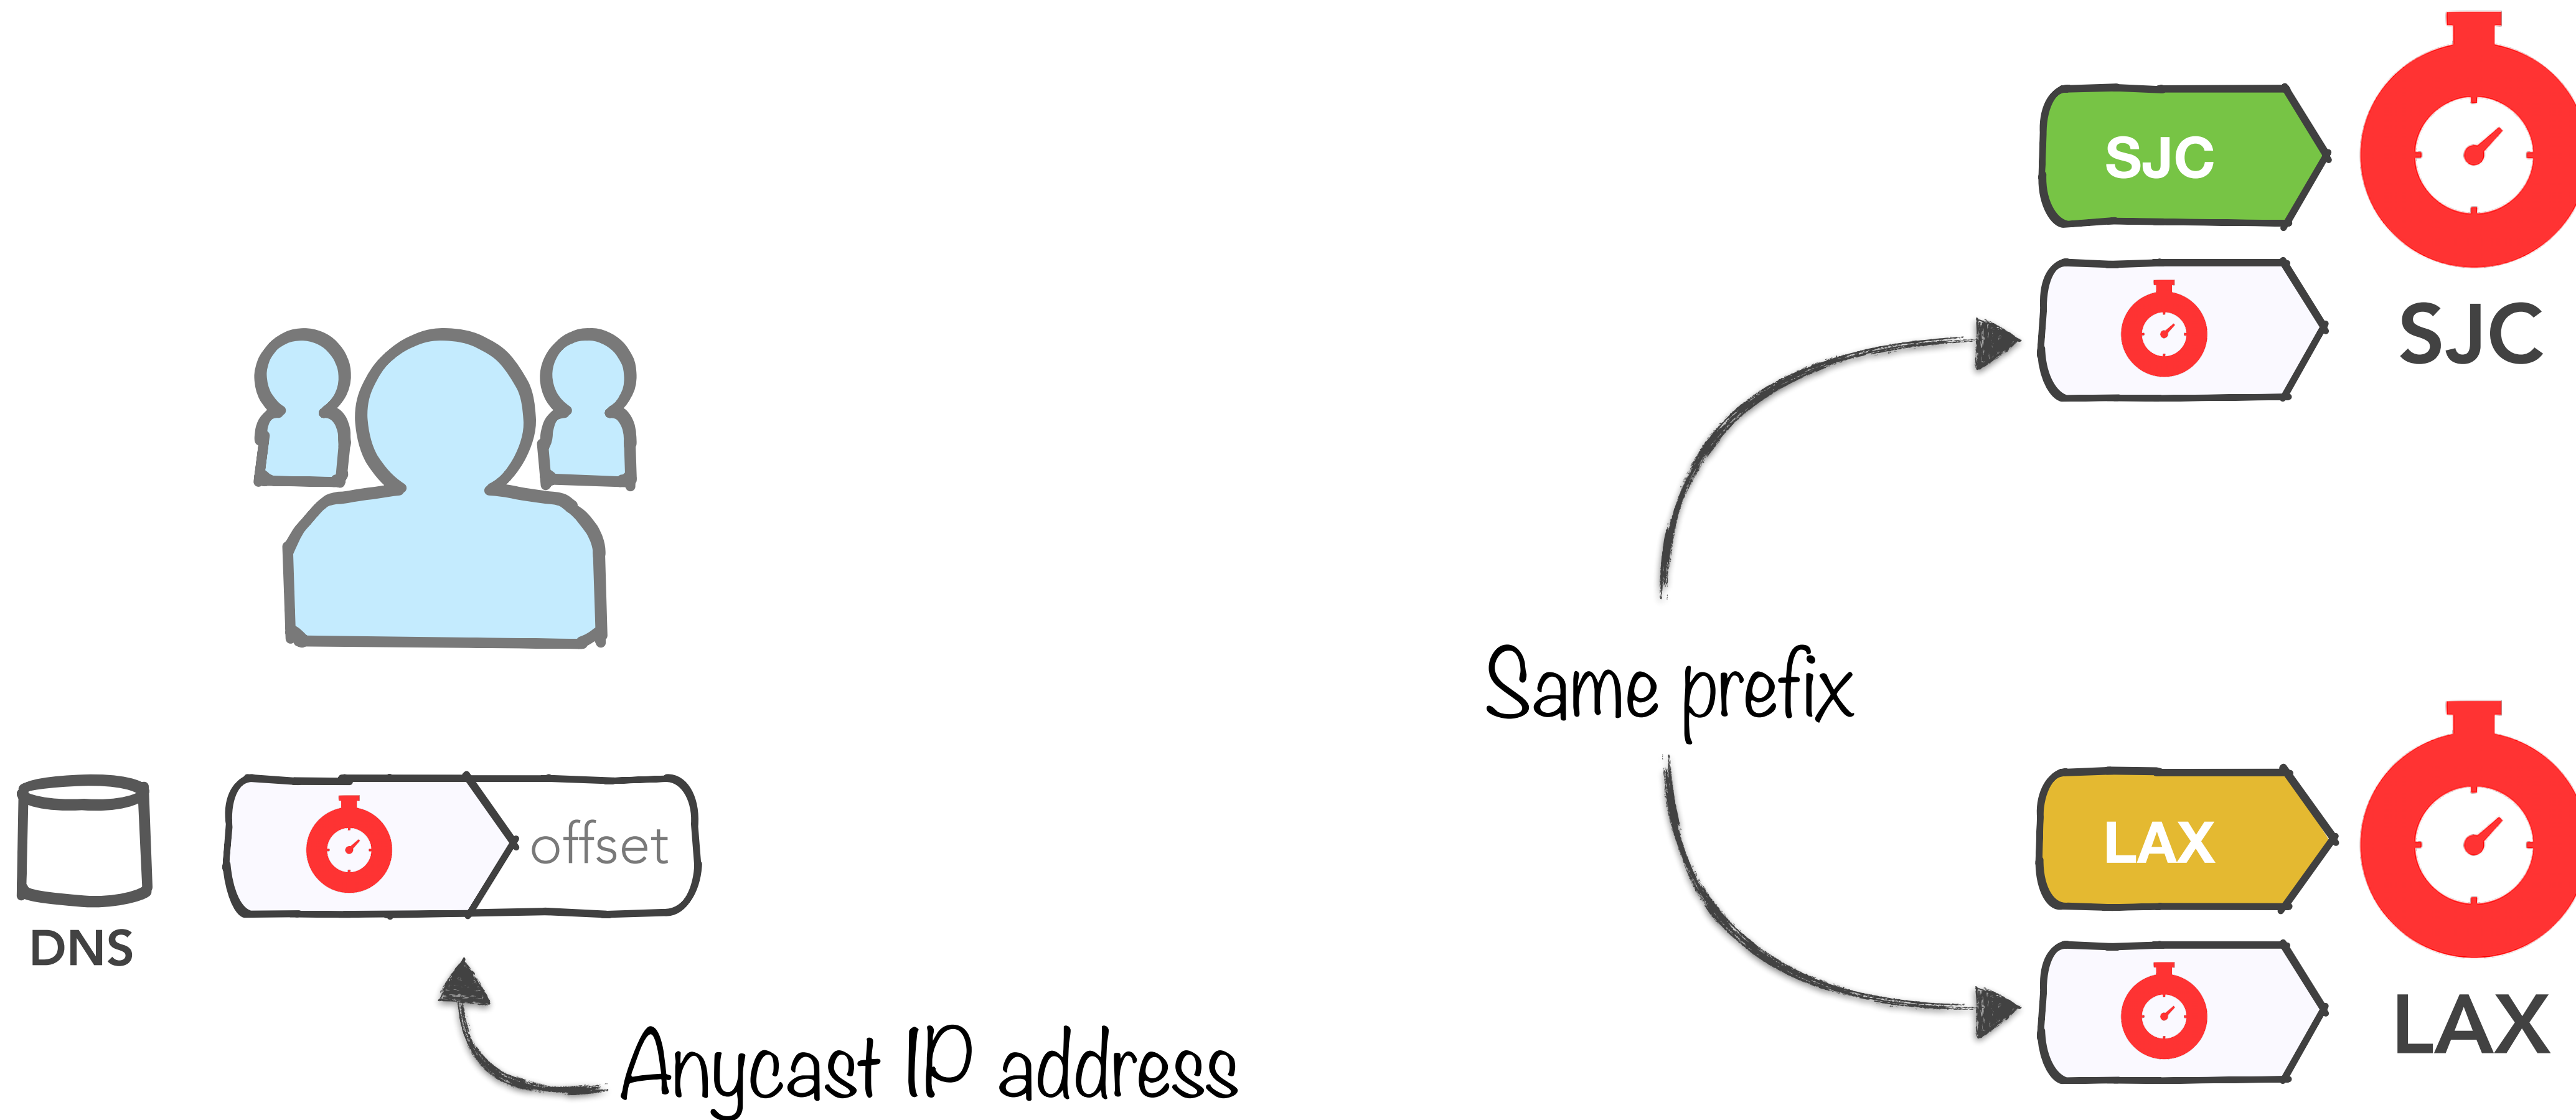
-  fate-sharing between address types
-  unicast has poor fallback properties
-  /24 per POP will run out at some point

 anycast to support apex domains and in-house DNS

Anycast model






2014




Anycast model






2014



SJC




DNS




LAX


Anycast model




2014




Anycast model



2014




Everything is fine



2014



- 🔥 fate-sharing between address types
- 🔥 unicast has poor fallback properties
- 💣 /24 per POP will run out at some point
- 🔥 anycast is hard to get right
- 🔥 inbound path control is terrible
- 🔥 overhead of running concurrent models



IP addressing



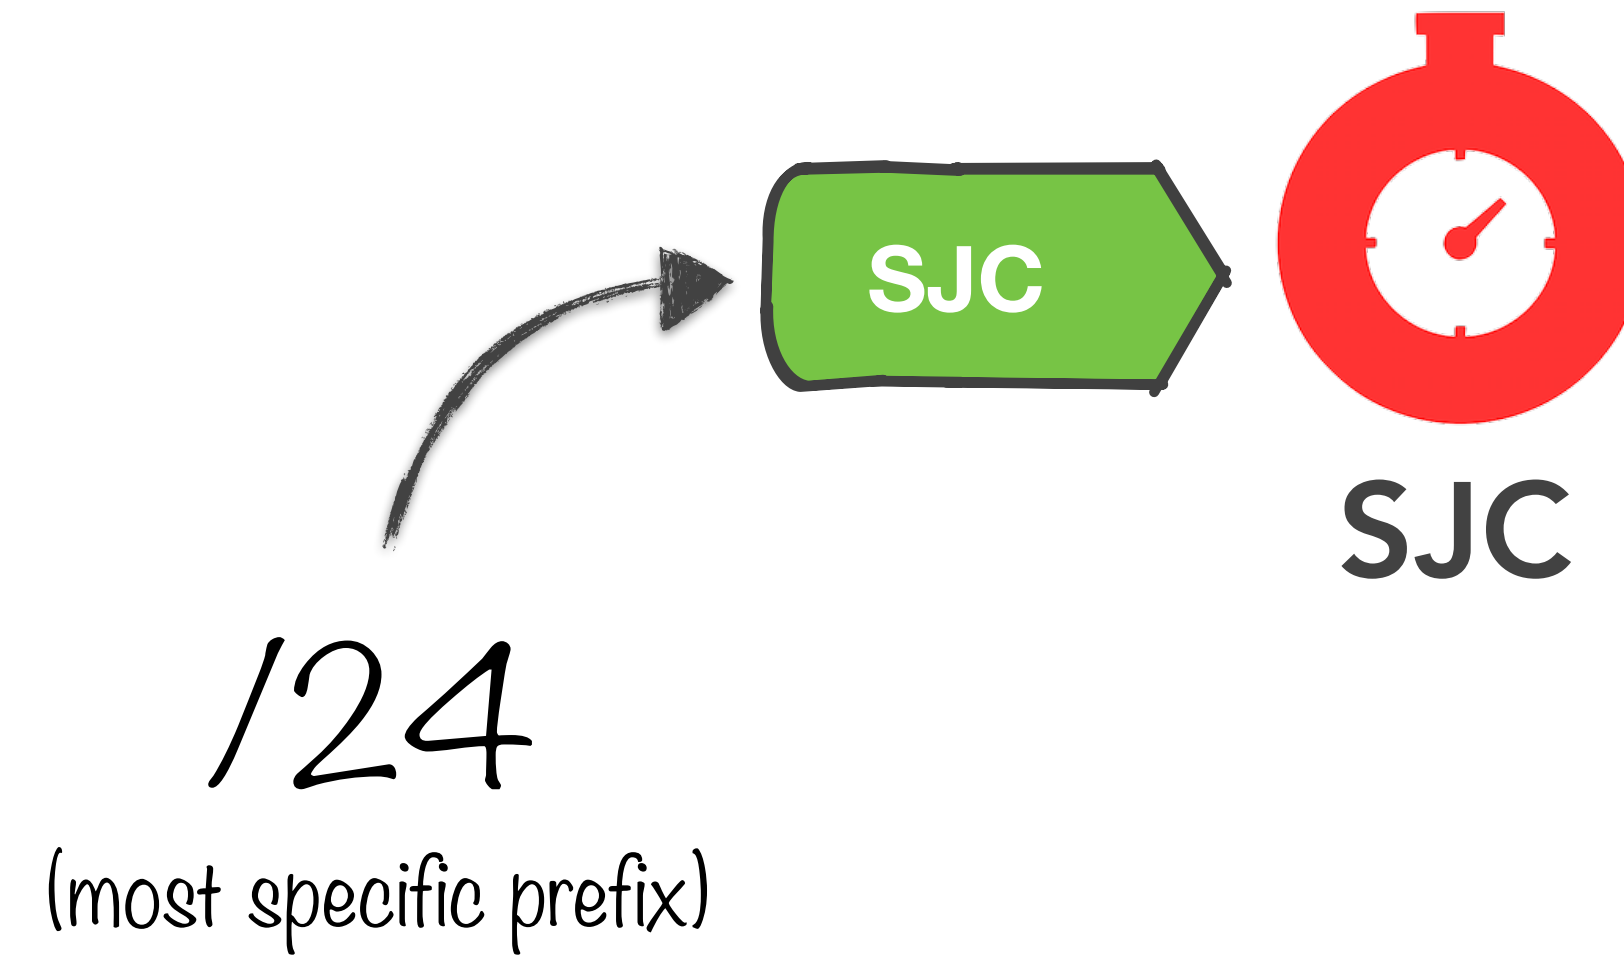

2013



IP addressing





2014




IP addressing










2015



Everything is terrible



2015

-  fate-sharing between address types
-  unicast has poor fallback properties
-  ~~but growth in IPv4 allocation matches same point~~
24 growth in IPv4 allocation matches same point
-  anycast is hard to get right
-  inbound path control is terrible
-  overhead of running concurrent models
-  you should probably do IPv6



The good news

no first-mover advantage

at least two competitors already offered IPv6

limited demand for IPv6

more valued features: caching, purging, logging, stats, VCL


already lost the very few customers who cared about IPv6

wasn't affecting our retention rate or growth

no need to rush, so clean slate




Everything is terrible




2015

- 🔥 fate-sharing between address types
- 🔥 unicast has poor fallback properties
- 🔥 outgrown IPv4 allocation scheme
- 🔥 anycast is hard to get right
- 🔥 inbound path control is terrible
- 🔥 overhead of running concurrent models

Decouple address types




Decouple address types



VIPs


one-to-many mapping, service abstraction



Infrastructure

one-to-one mapping to a physical endpoint

Everything is terrible



2015



~~fate sharing between address types~~



unicast has poor fallback properties



outgrown IPv4 allocation scheme



anycast is hard to get right




inbound path control is terrible



overhead of running concurrent models



Everything is terrible



2015



~~fate sharing between address types~~



unicast has poor fallback properties



outgrown IPv4 allocation scheme



anycast is hard to get right



inbound path control is terrible

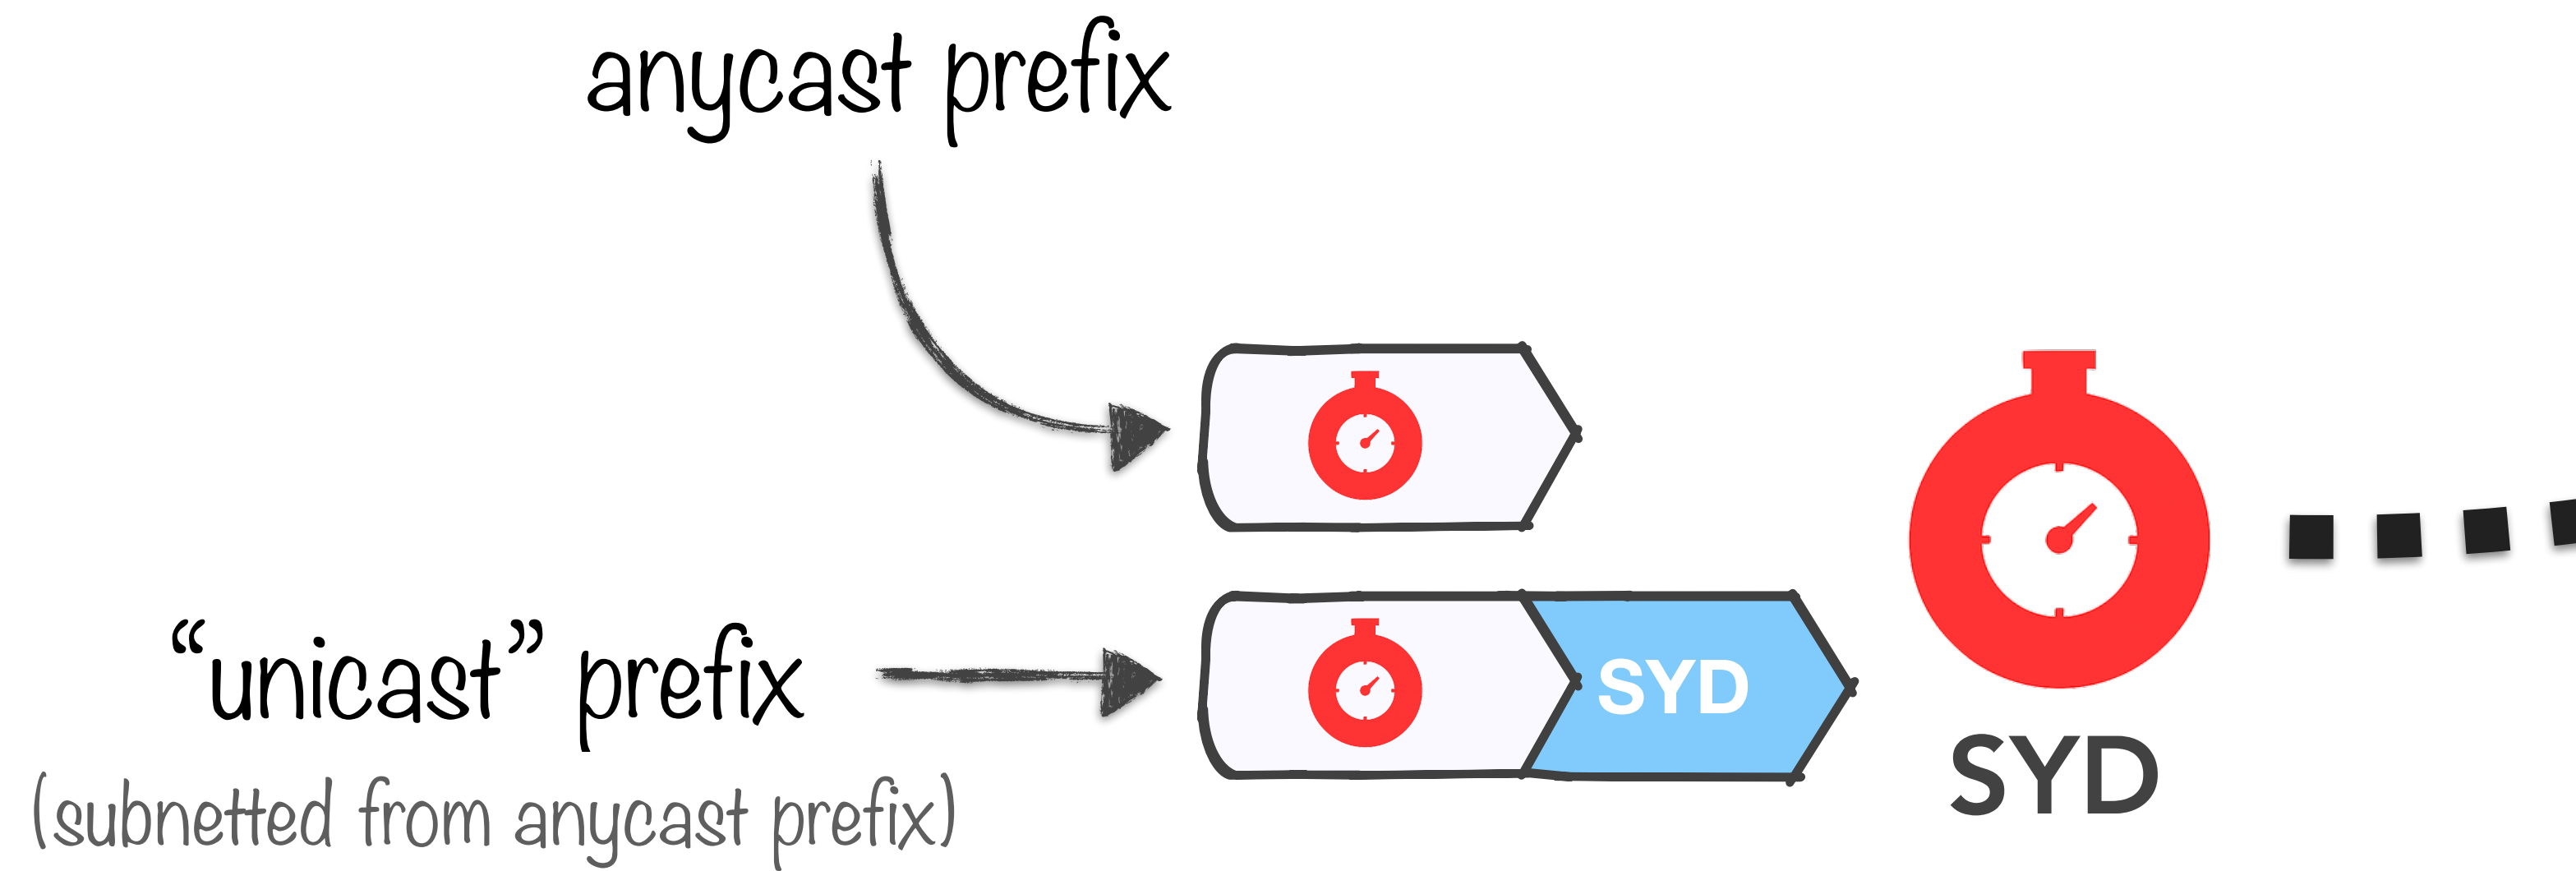


overhead of running concurrent models




Backing anycast

 ~2015




Backing anycast

📅 ~2015



Backing anycast

📅 ~2015




Backing anycast

 ~2015



Everything is terrible



2015



~~fate sharing between address types~~



unicast has poor fallback properties



outgrown IPv4 allocation scheme



anycast is hard to get right




inbound path control is terrible



overhead of running concurrent models



Everything is terrible



2015



~~fate sharing between address types~~



~~unicast has poor fallback properties~~



outgrown IPv4 allocation scheme



~~anycast is hard to get right~~




inbound path control is terrible



overhead of running concurrent models



Everything is terrible



2015



~~fate sharing between address types~~



~~unicast has poor fallback properties~~



outgrown IPv4 allocation scheme



~~anycast is hard to get right~~




inbound path control is terrible




~~overhead of running concurrent models~~





VIPs


one-to-many mapping, service abstraction




Infrastructure

one-to-one mapping to a physical endpoint


VIP allocation




2016




132




Per provider planes




2016



VIP allocation



2016



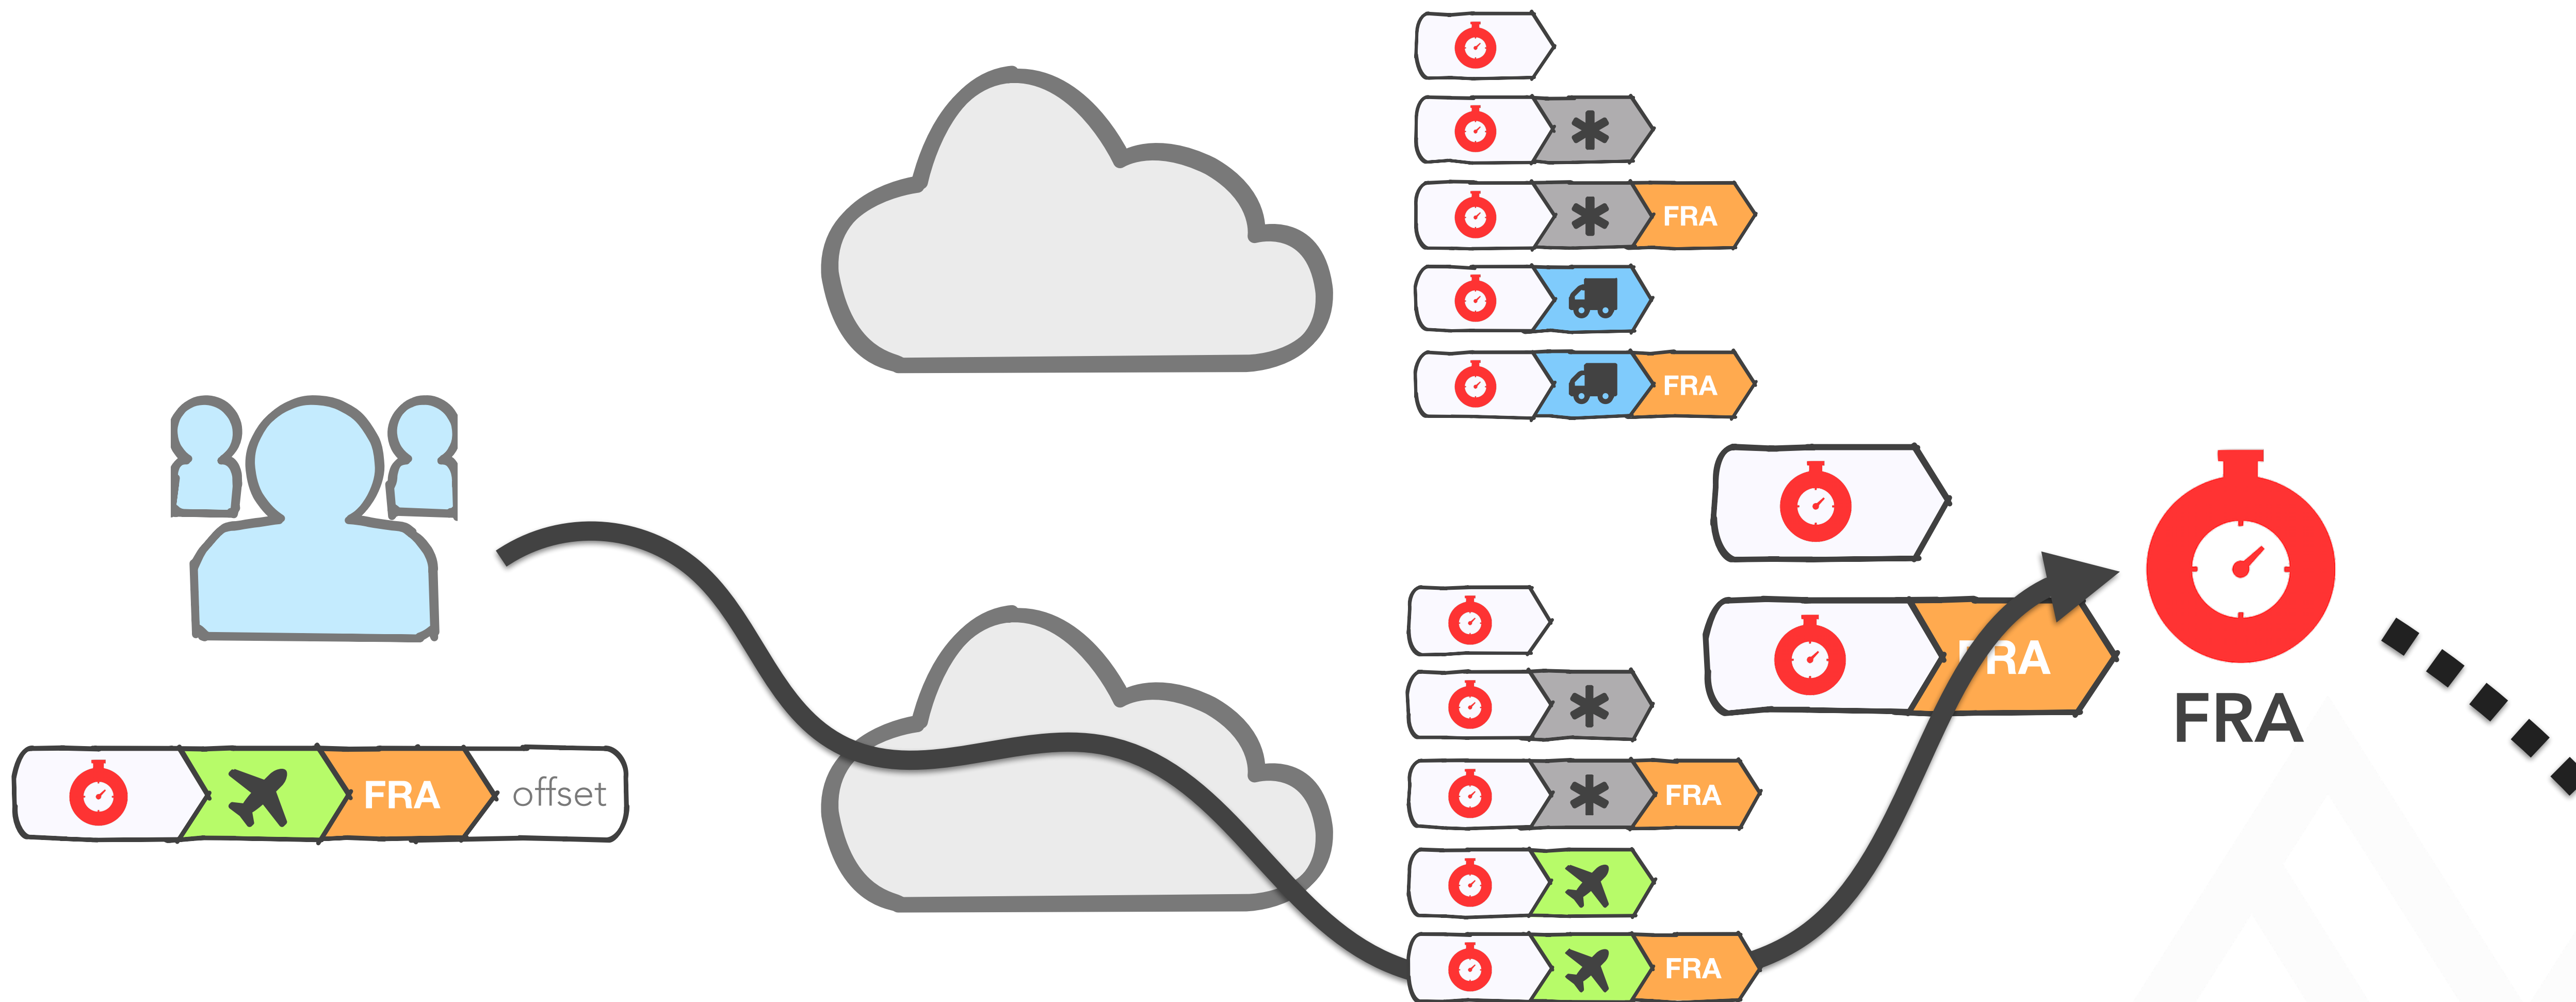
148




VIP allocation



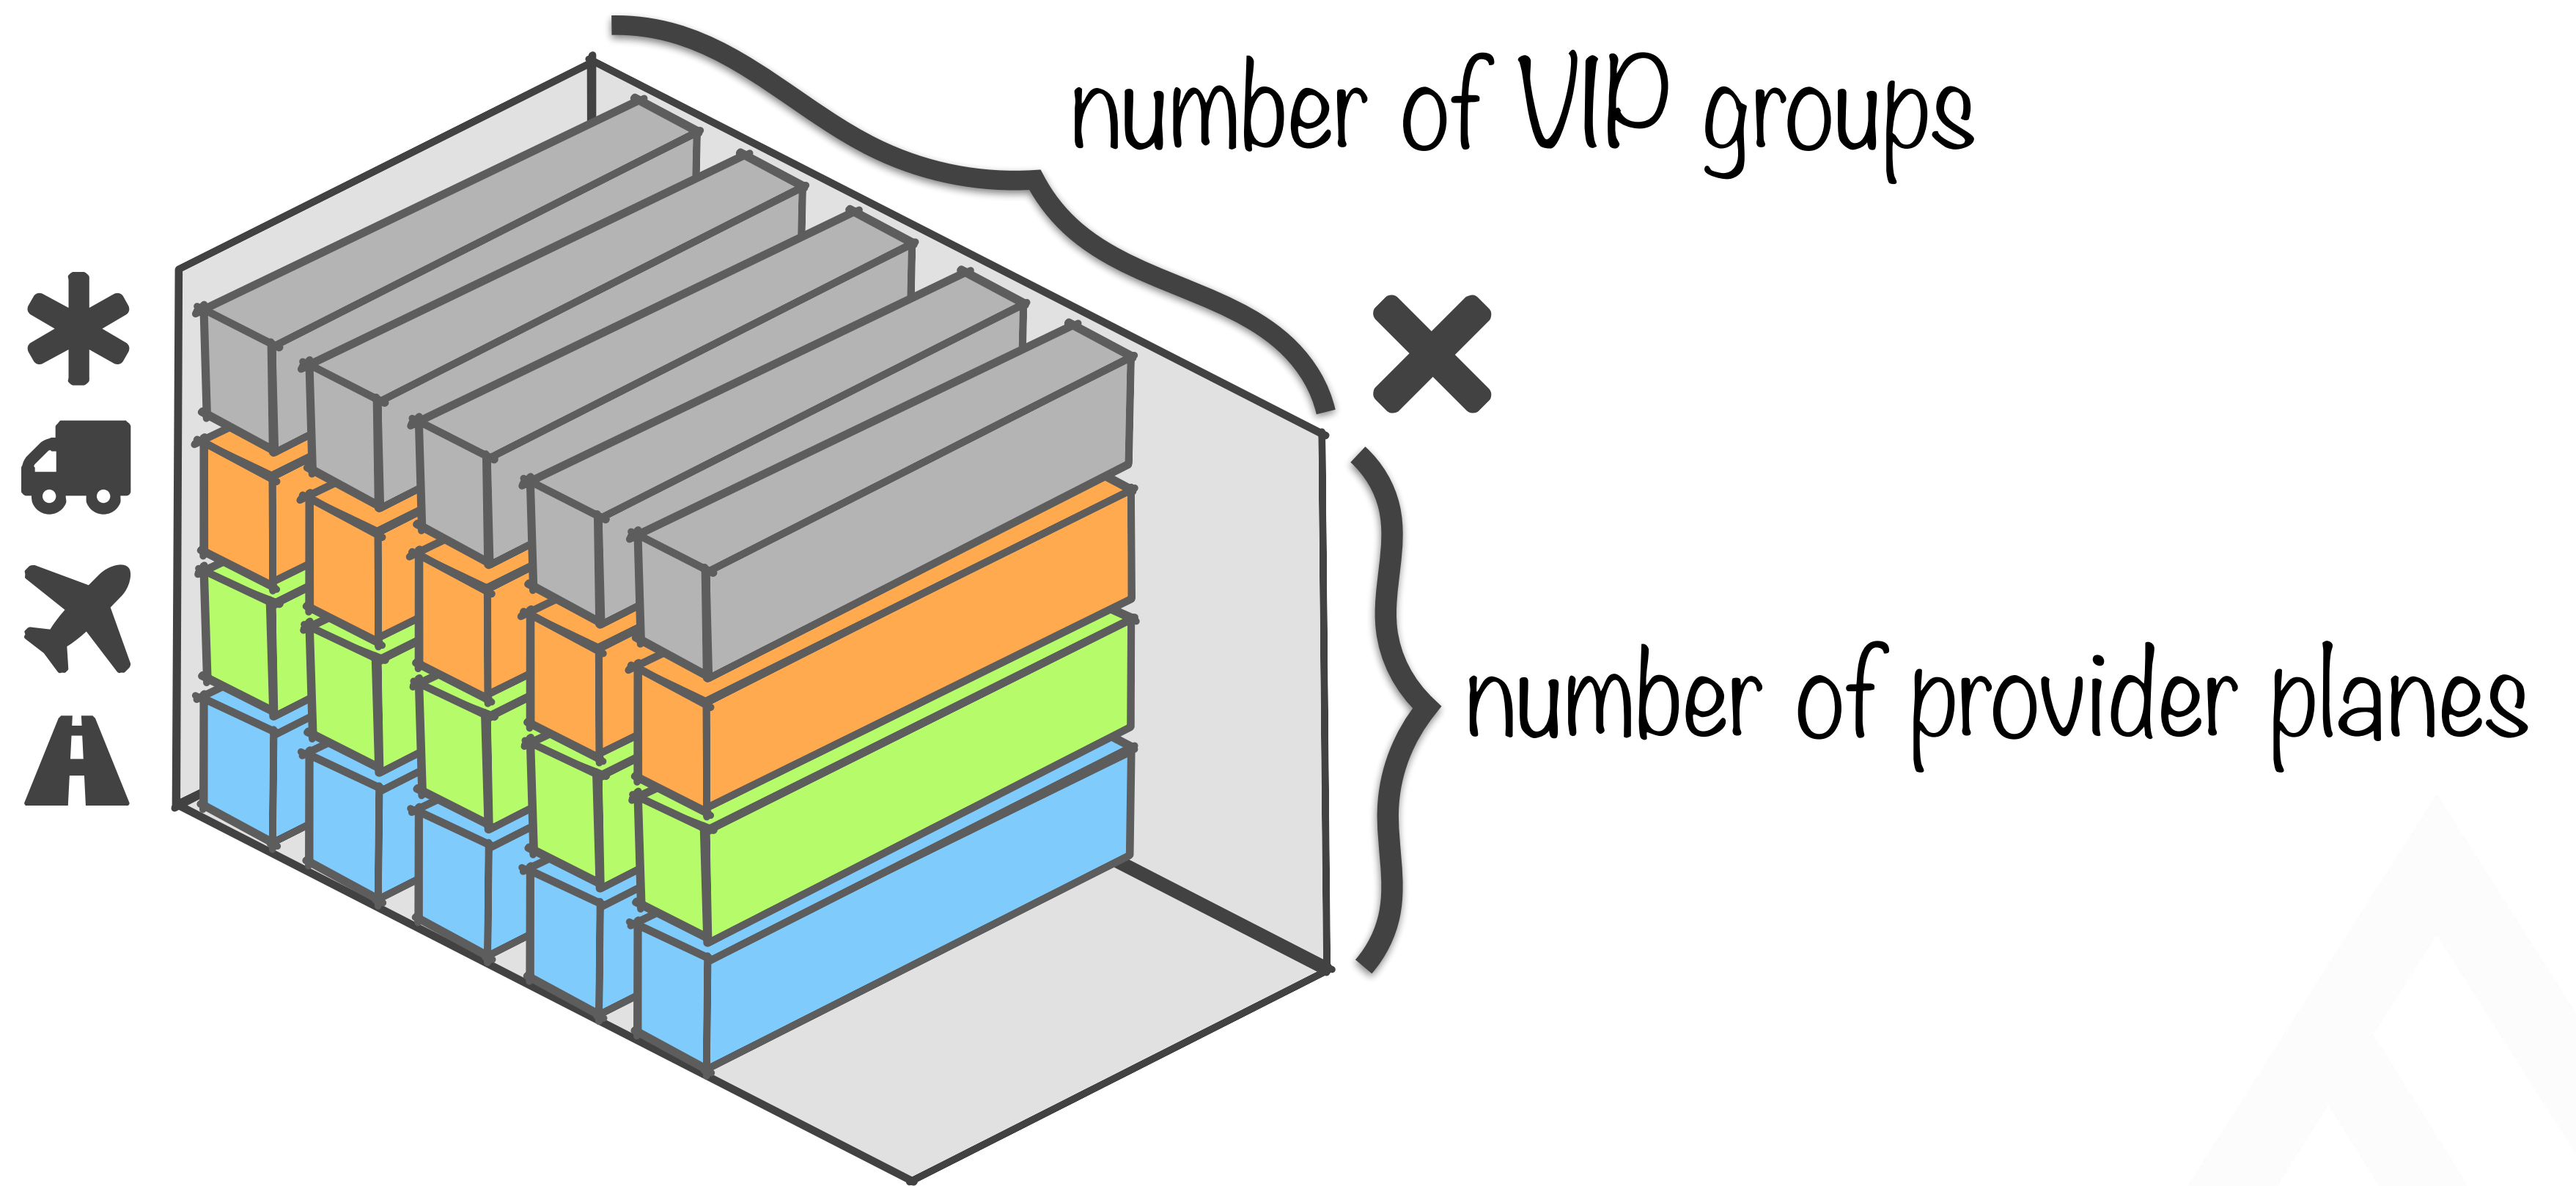
2016




Impact on global table



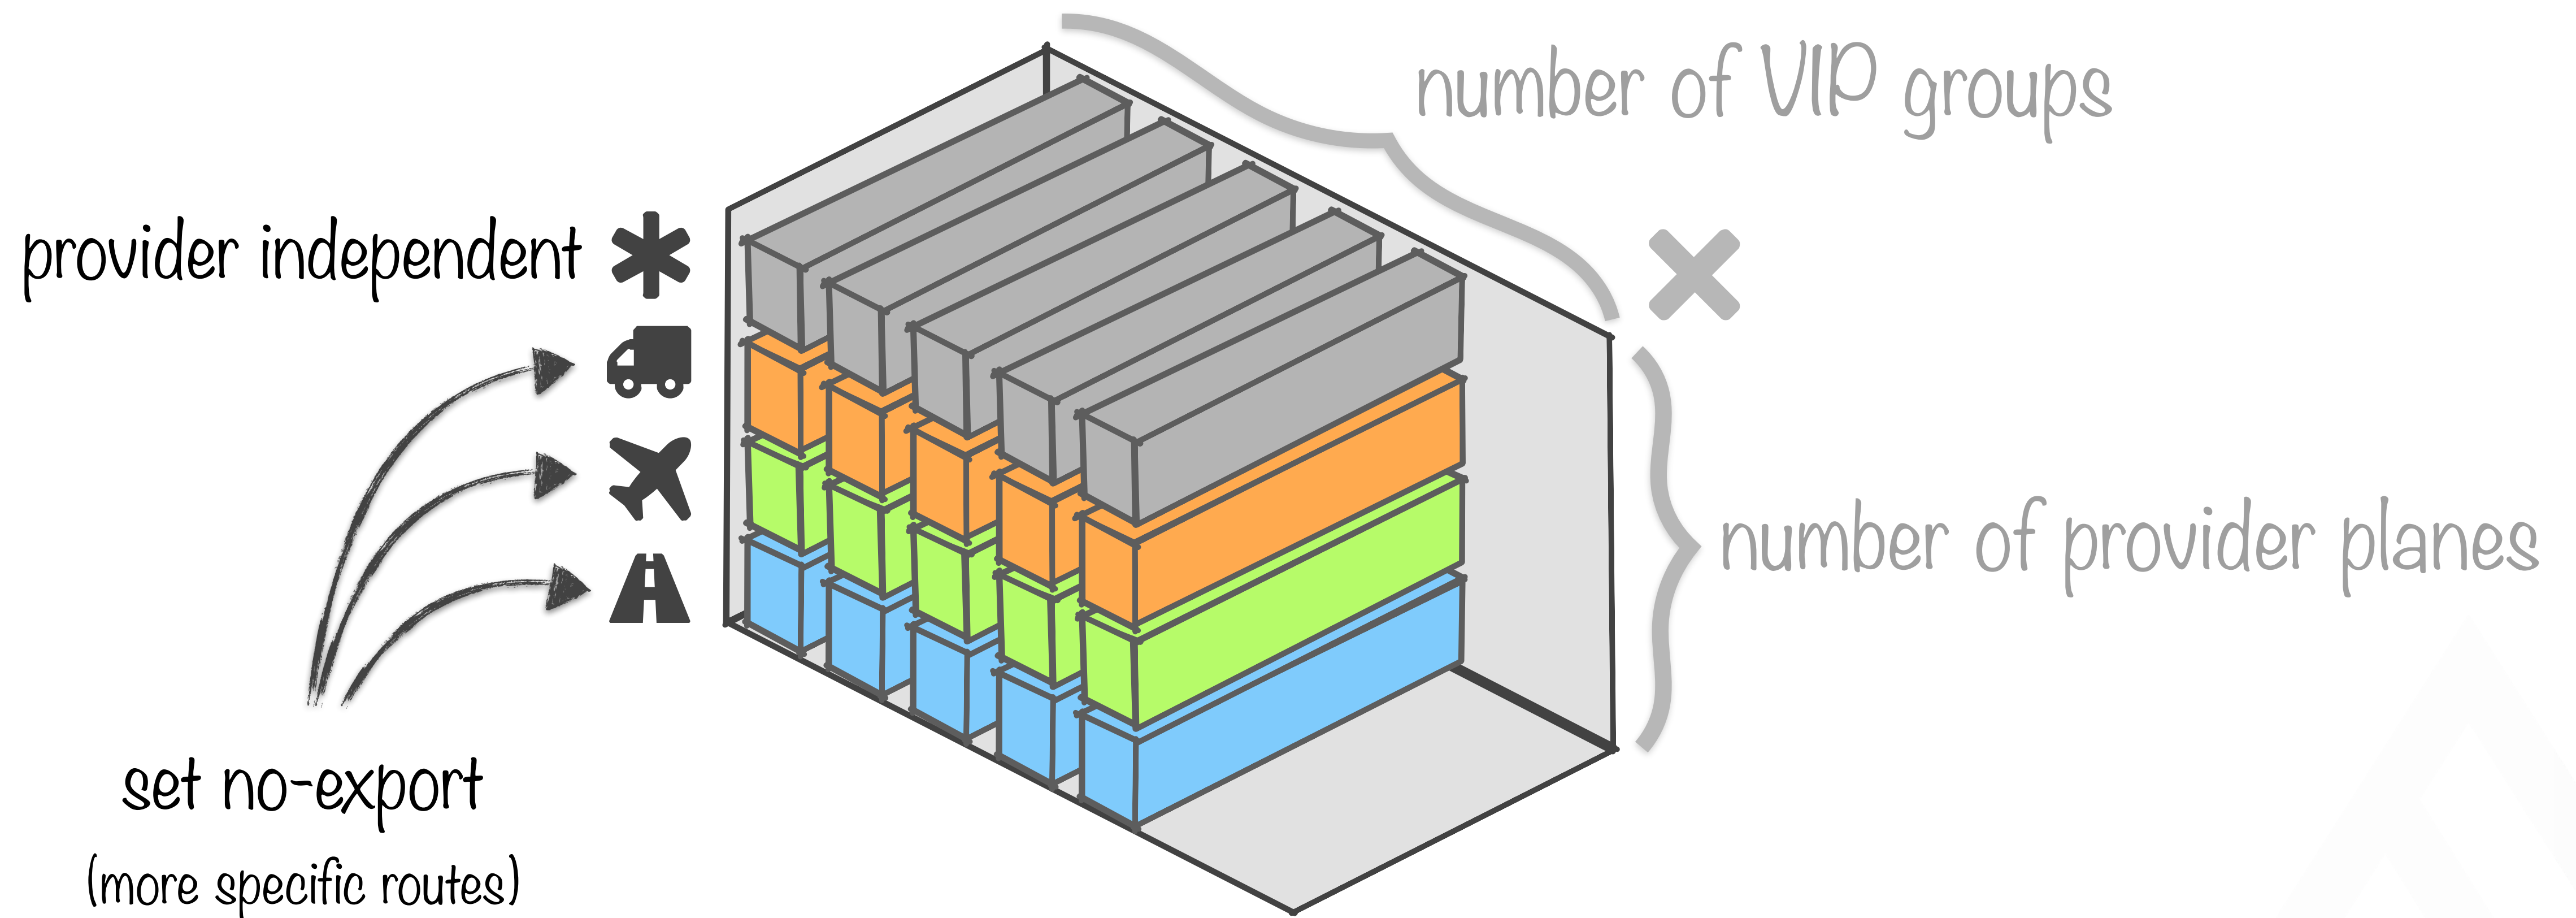
2016




Impact on global table



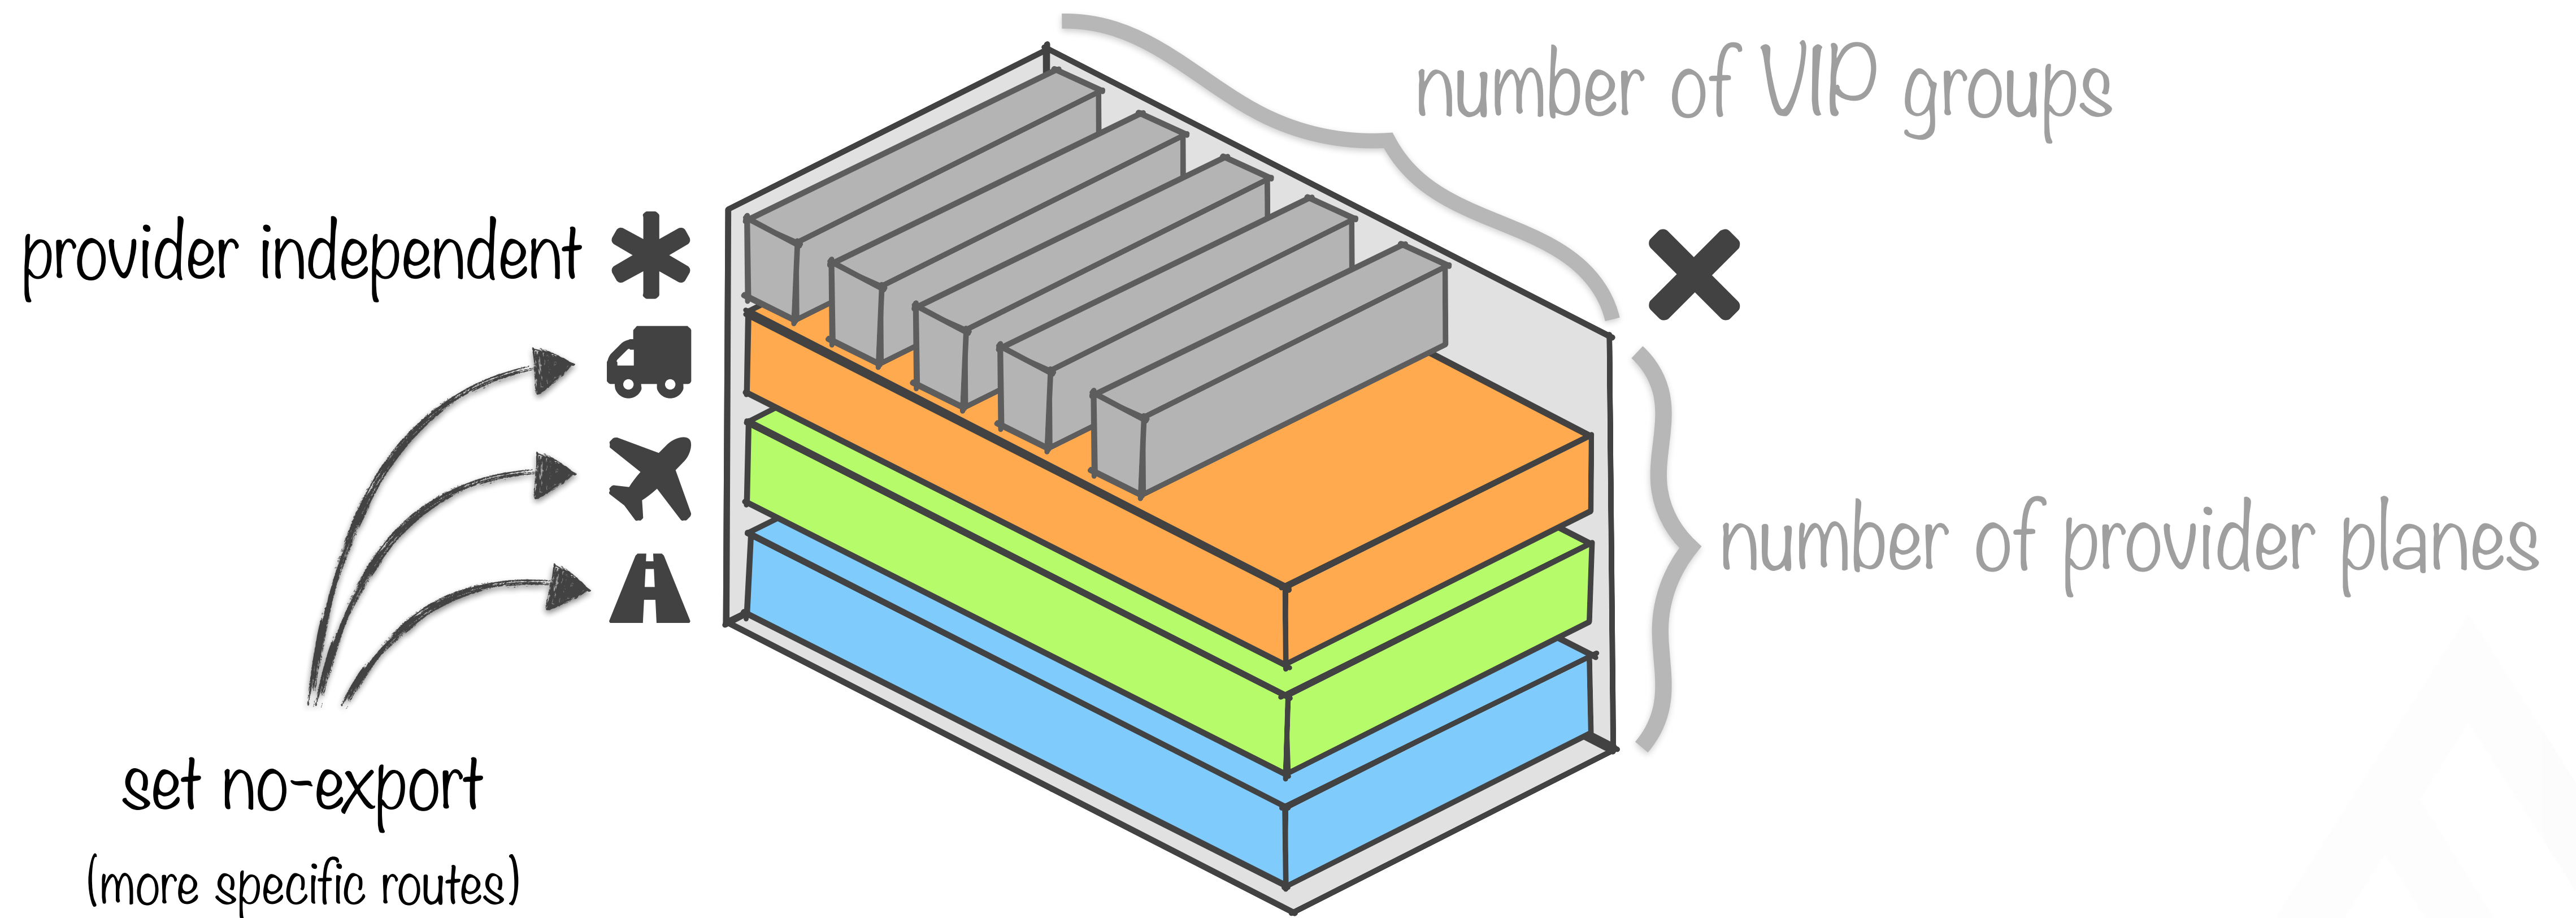
2016




Impact on global table



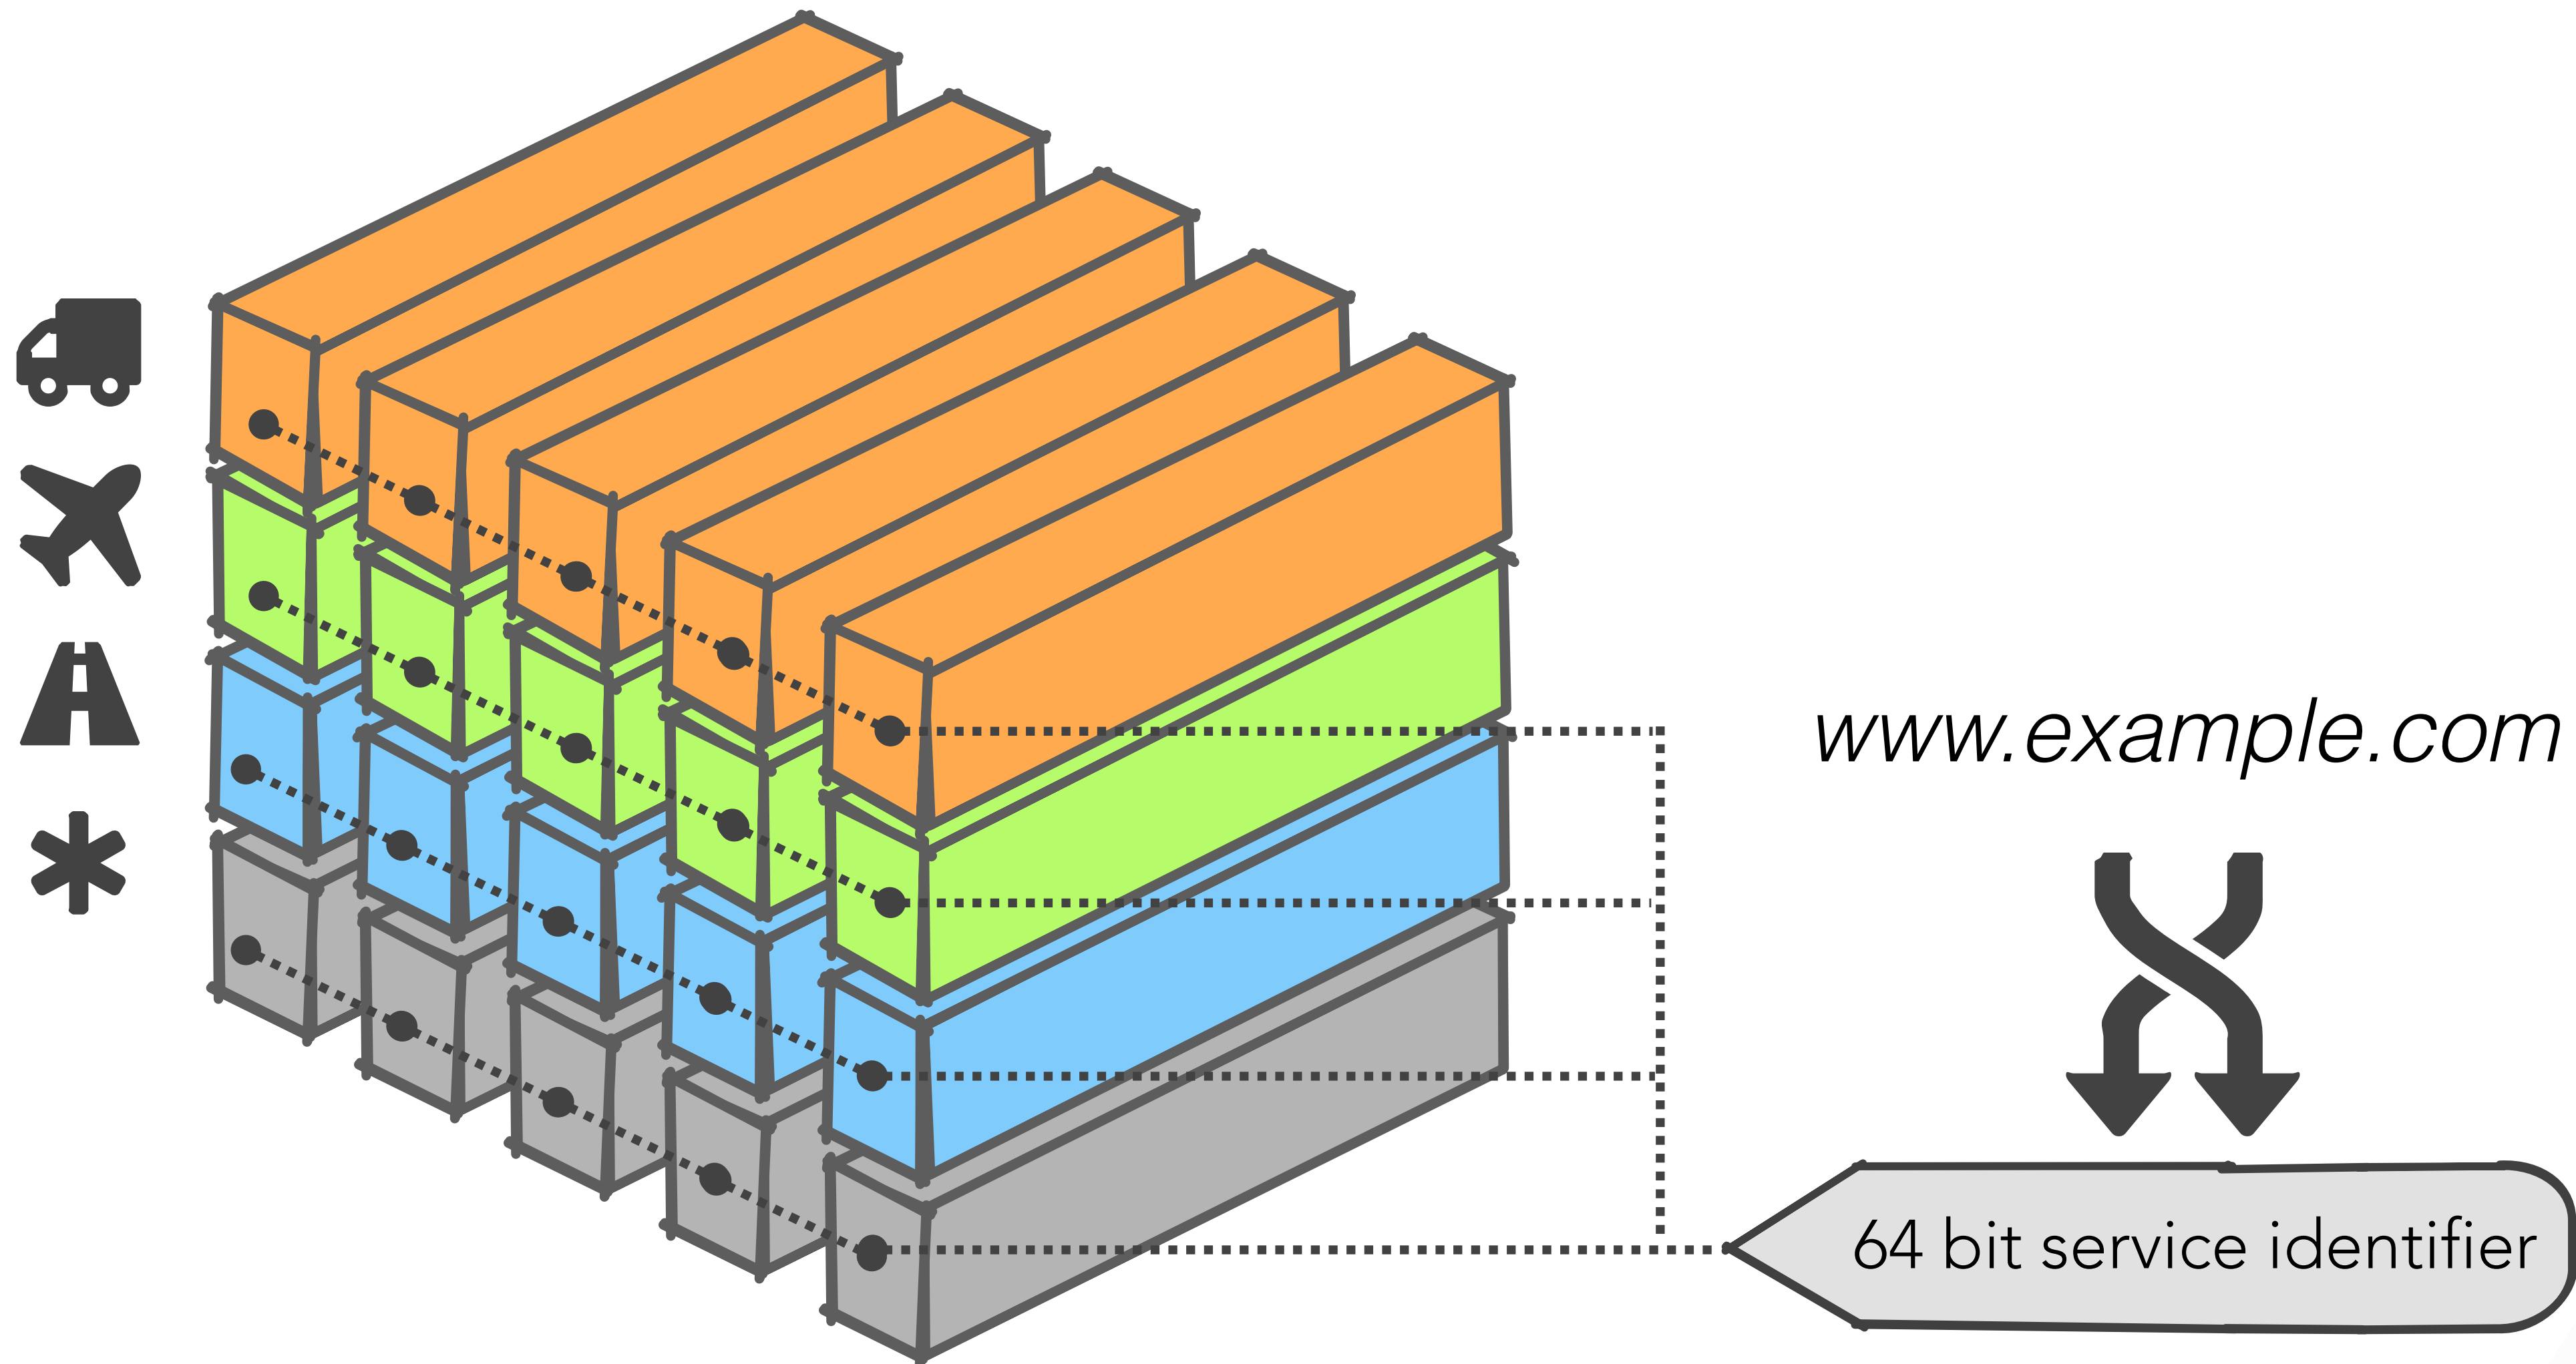
2016




Locator / Identifier




2016





VIPs


one-to-many mapping, service abstraction




Infrastructure

one-to-one mapping to a physical endpoint


Infrastructure allocation




2016




Infrastructure allocation



2016



Impact on global table




2016

number of announcements $\sim \sum_{i \in \text{POPs}} |\text{providers}_i|$



Impact on global table



2016


number of announcements $\sim \sum_{i \in \text{POPs}} |\text{providers}_i|$



each infrastructure prefix in a POP is a /40



Impact on global table




2016

number of announcements $\sim \sum_{i \in \text{POPs}} |\text{providers}_i|$ up to 16 /44s

each infrastructure prefix in a POP is a /40

Everything is terrible



2015



~~fate sharing between address types~~



~~unicast has poor fallback properties~~



outgrown IPv4 allocation scheme



~~anycast is hard to get right~~




inbound path control is terrible



~~overhead of running concurrent models~~



Everything is terrible



2016



~~fate sharing between address types~~



~~unicast has poor fallback properties~~



outgrown IPv4 allocation scheme



~~anycast is hard to get right~~




~~inbound path control is terrible~~









overhead of running concurrent models




Everything is terrible




2016

-  ~~fate sharing between address types~~
-  ~~unicast has poor fallback properties~~
-  outgrown IPv4 allocation scheme
-  ~~anycast is hard to get right~~
-  ~~inbound path control is terrible~~
-  overhead of running concurrent models


Locator names




2016




vip.




ntt.vip.



sjc.global.vip.




fra.inf.




peering.fra.inf.


Locator names




2016




vip.




ntt.vip.



sjc.global.vip.




fra.inf.




peering.fra.inf.


Locator names




2016



vip.



ntt.vip.



sjc.global.vip.




fra.inf.



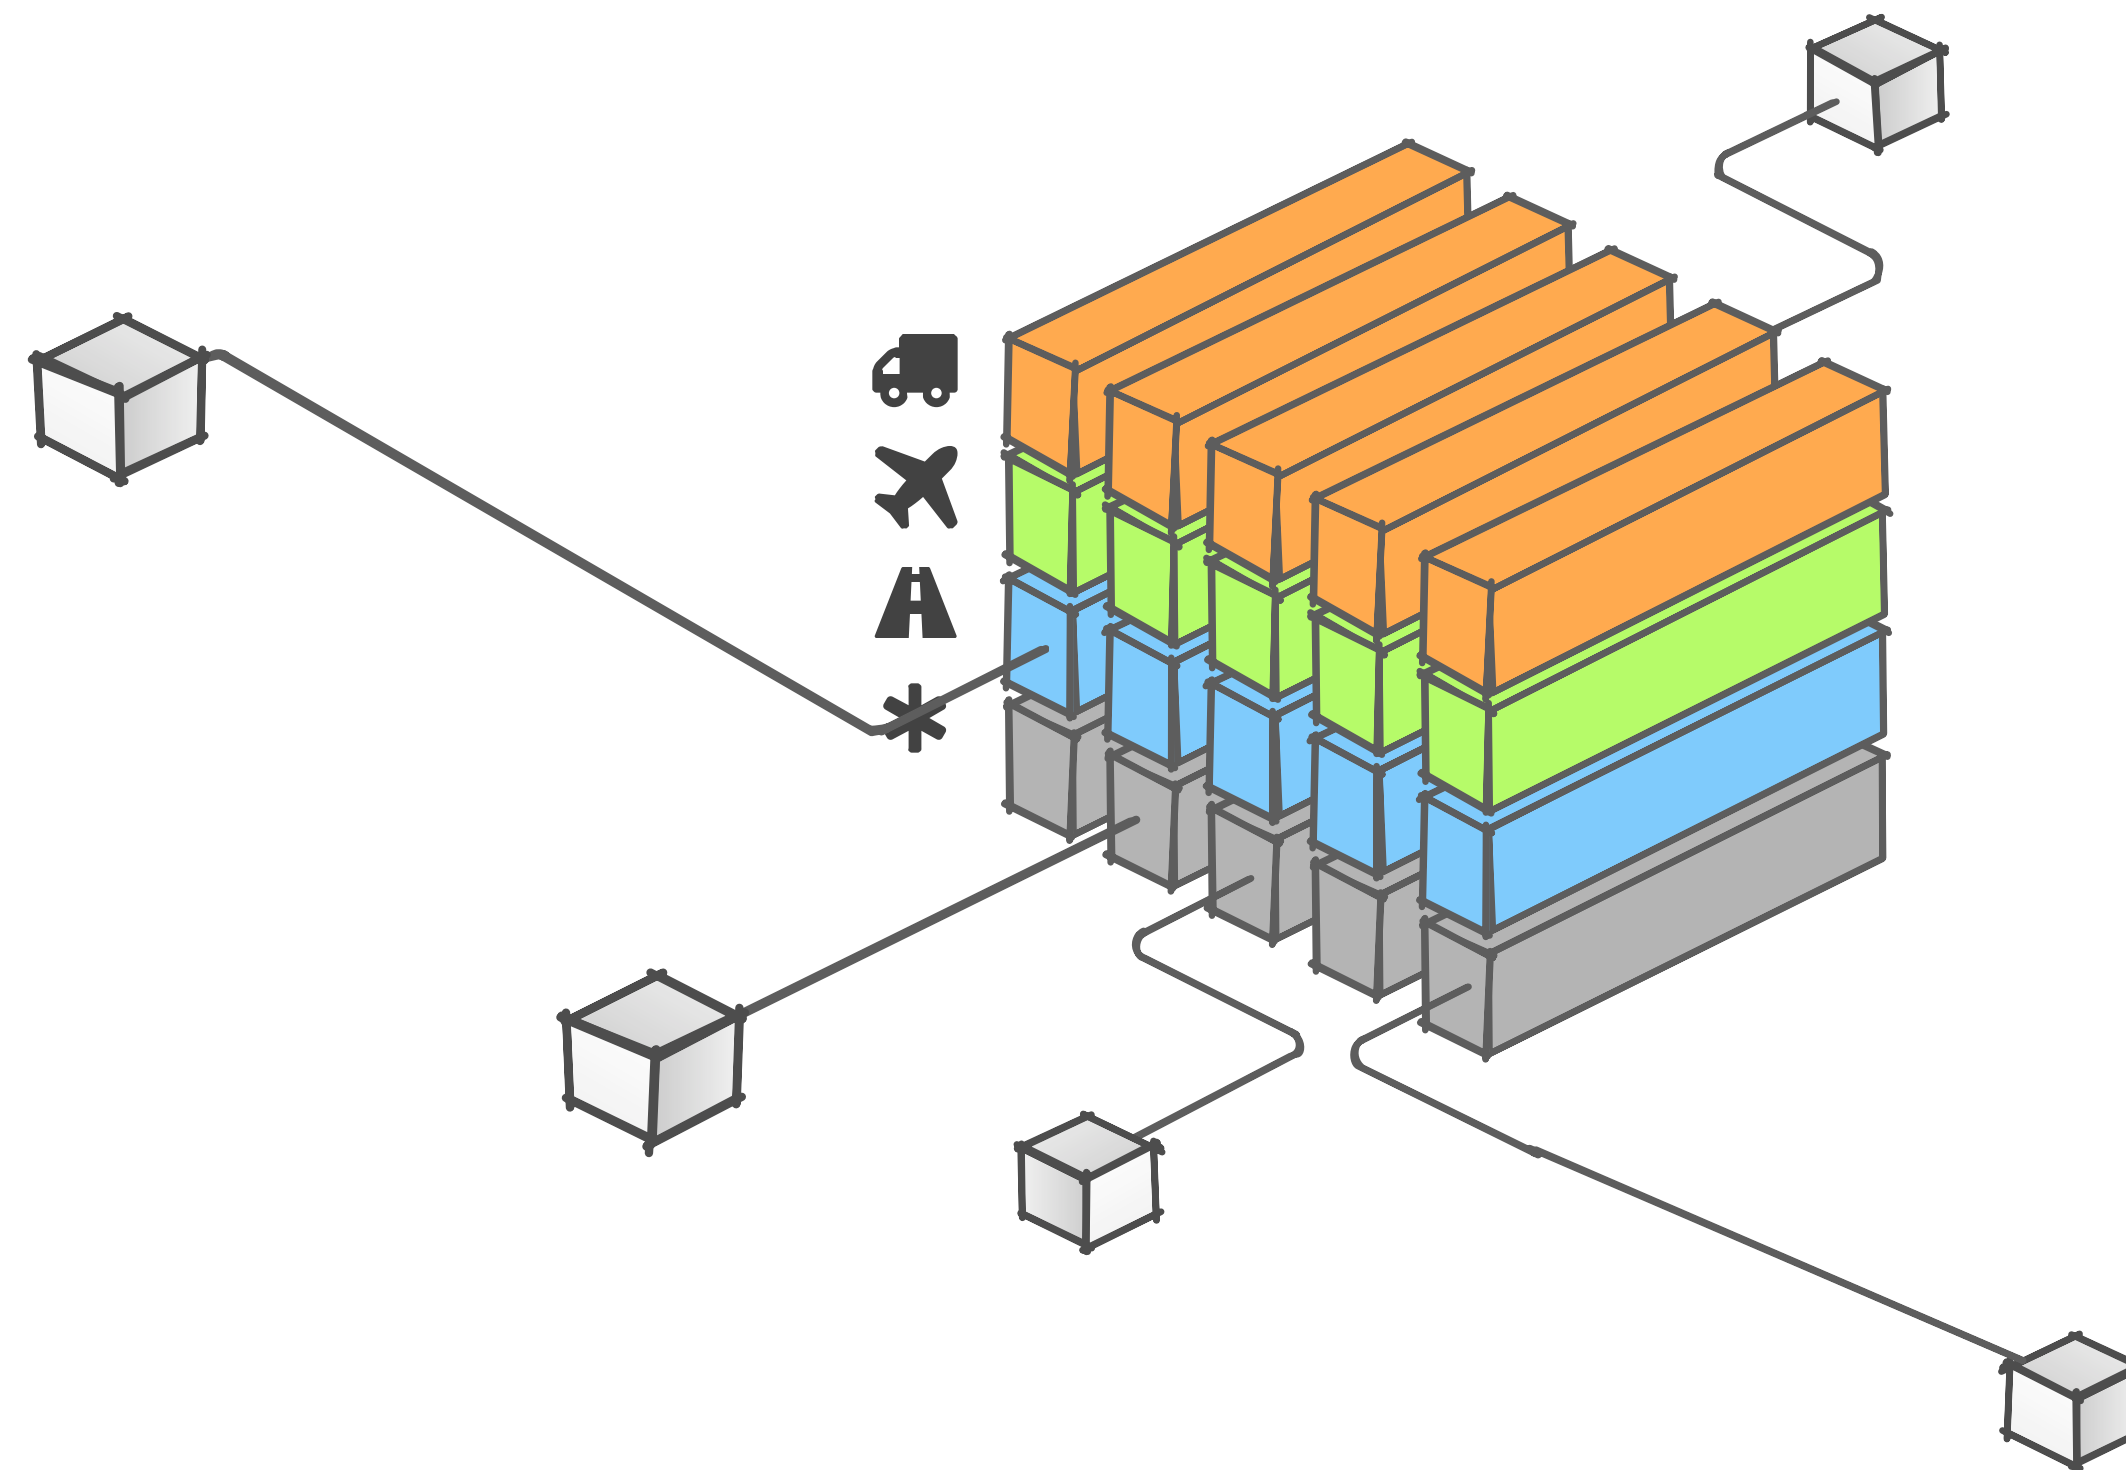
peering.fra.inf.

Anchoring IPv4



2016

```
→ bird-export.example  
if locator ~ "*cogent.vip" && provider != "cogent" then reject;  
if locator ~ ".*cogent.vip" && provider == "cogent" then set_no_export()
```



Summary

- ✓ decoupled address types
- ✓ graceful fallback
- ✓ VIP prefix mobility
- ✓ fine-grained inbound path control
- ✓ unified model based on locator names
- ✗ takes a long time

Intellectual heritage

ILNP mobility, multi-homing, inbound TE

MP-TCP resource pooling

re-ECN information asymmetry in connectivity markets



WIP

ILNP locators expose path diversity

MP-TCP pool path diversity at transport and above

re-ECN e2e metrics drive path selection



WIP

ILNP

MP-TCP

re-ECN




EU research from ~ 10 years ago







LHR



VARNISH
CACHE




H A
P R O X Y




A speech bubble with a black outline and rounded corners. It contains three rows of logos and text:

- Top row: The flag of Norway (red field with a blue Scandinavian cross) on the left, and the logo for "VARNISH CACHE" on the right. The logo features three blue circles of varying sizes to the left of the text "VARNISH" in a bold, sans-serif font, with "CACHE" in a smaller font below it.
- Middle row: The flag of France (vertical stripes of blue, white, and red) on the left, and the text "HA PROXY" on the right. "HA" is in white letters on blue and red squares, and "PROXY" is in white letters on blue, green, yellow, and red squares.
- Bottom row: The flag of the Czech Republic (white field with a red triangle at the bottom left and a blue triangle at the top left) on the left, and the logo for "BIRD" on the right. The logo features a stylized bird silhouette above the word "BIRD" in a bold, sans-serif font, with three horizontal lines below it.



VARNISH
CACHE



H A
P R O X Y



BIRD





Either the questions don't matter

ILNP mobility, multi-homing, inbound TE

MP-TCP resource pooling

re-ECN information asymmetry in connectivity markets



Either the ideas don't work

ILNP locators expose path diversity

MP-TCP paths exposed to transport/app

re-ECN e2e metrics drive path selection





Questions

[@jta](#) João Taveira Araújo

RIPE 74