All good things come in threes

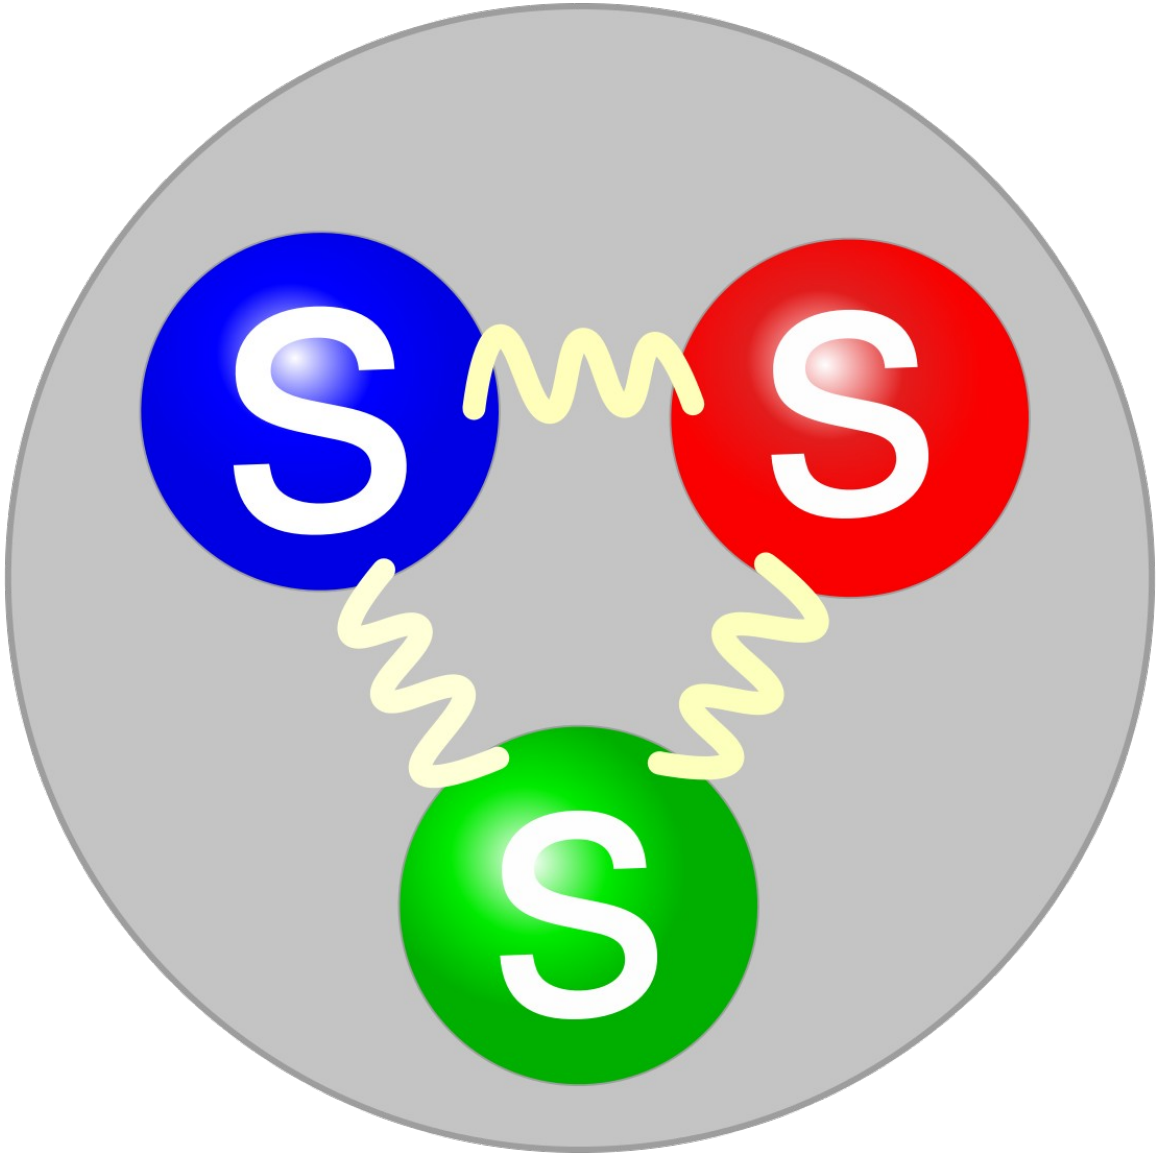

Jean-Marie de Boer

13 March 2024

image source: wikipedia.org

SURFdrive / Research Drive

- Owncloud based sync&share
- ~ 100.000 total users
- ~ 2 PB in data
- 24 racks
- Some 40 connected institutions
- 10 years in operation
- > 99.9% uptime

Some key values

- **Independence / transparancy**
  - No big tech
  - Open source
- **Privacy / security**
  - ISO standards
  - Certified datacenters
  - Auditting / pentesting
- **Performance**
  - Scale out
- **Availability**
  - Eliminate SPOF’s

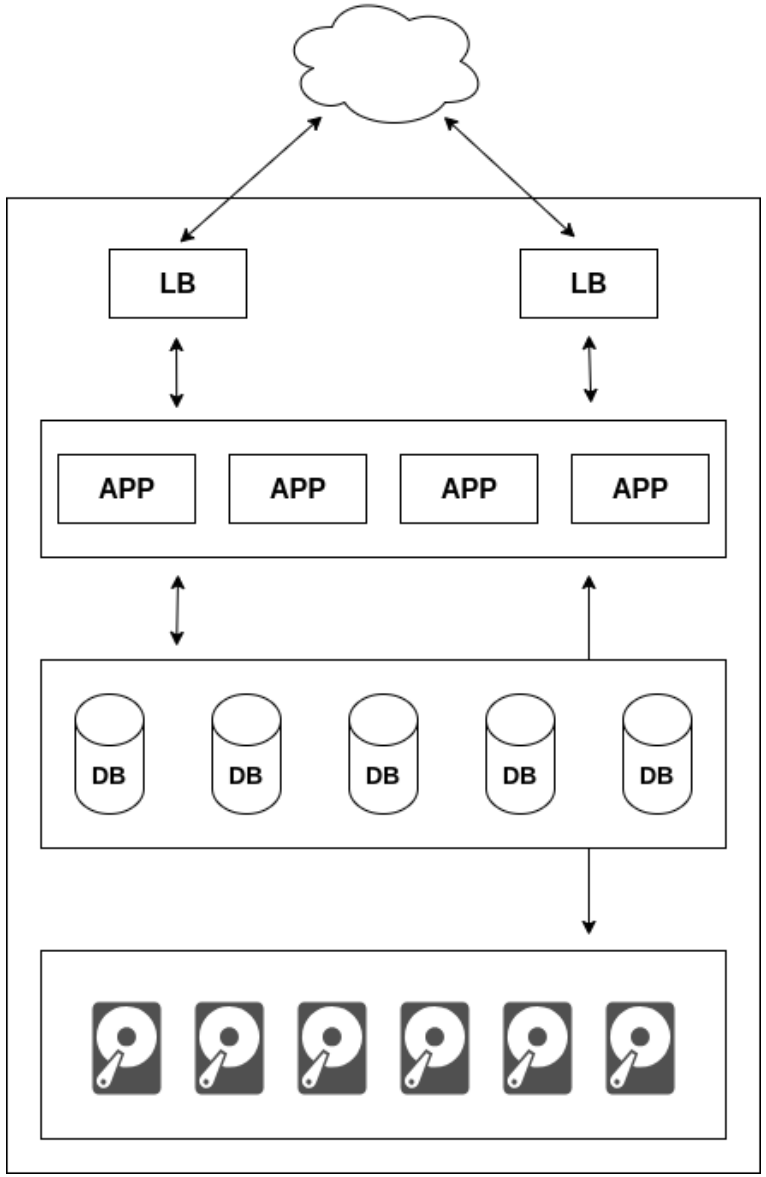

A BIG SPOF

image source: www.agsarchitects.net

A BIG SPOF

image source: i0.wp.com

A BIG SPOF

image source: www.agsarchitects.net

Why three datacenters?

- **We wanted active-active-active**
  - Failovers are bad ;)
  - Restoring from backup is hardly feasible
- **Quorum is easy**
  - Avoid split-brain
- **Less resources needed**
  - Counterintuative but true:

  When using 2 sites, if 1 is lost, the other needs 100% capacity → total is 200% capacity

  When using 3 sites, if one is lost, the other two need 50% each → total is 150% capacity

SURF

Aligning the organisation

- **Slow decision-making process**
  - Management needs convincing
- **Must be tendered**
- **Many requirements exist**
  - ISO standards
  - Technical (such as network)
- **We have a winner**
  - Equinix AM1 and AM8


AM8
8.8 KM
Science Park
12.8 KM
6.8 KM
AM1
SURF
Leaflet | © OpenStreetMap contributors

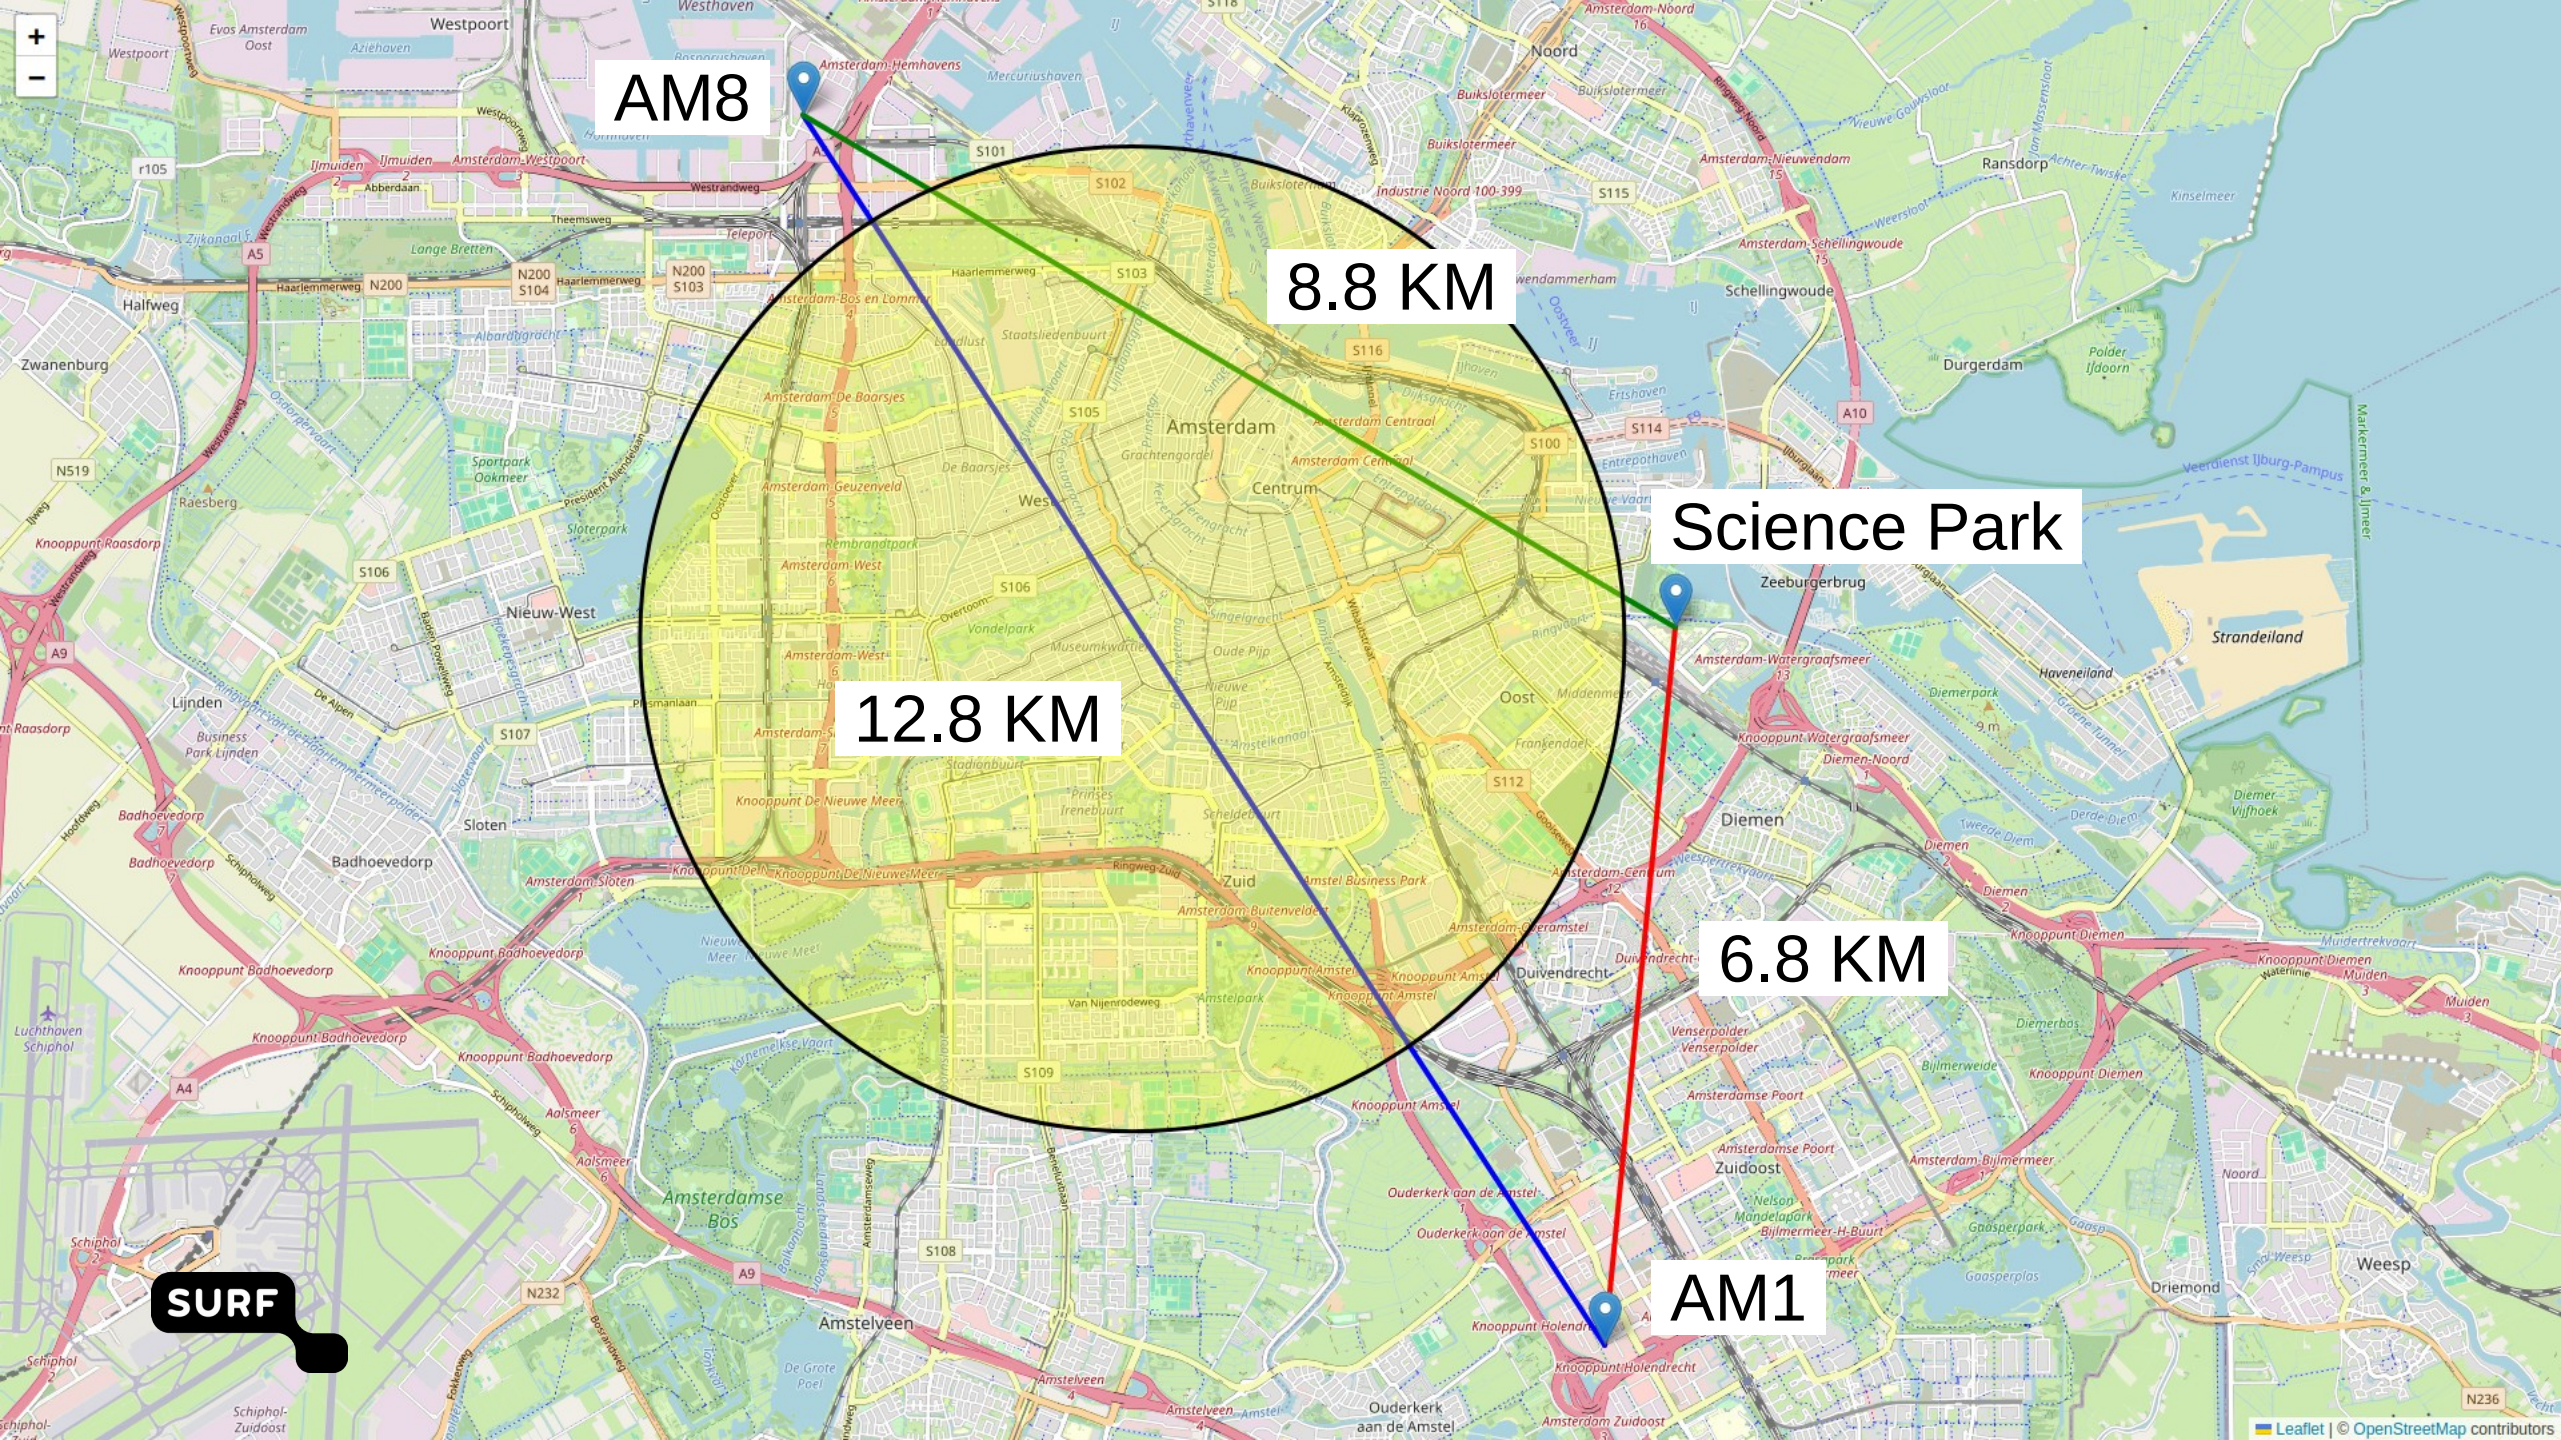
AM8
8.8 KM
Science Park
12.8 KM
6.8 KM
AM1
SURF
Leaflet | © OpenStreetMap contributors

Three datacenter stack

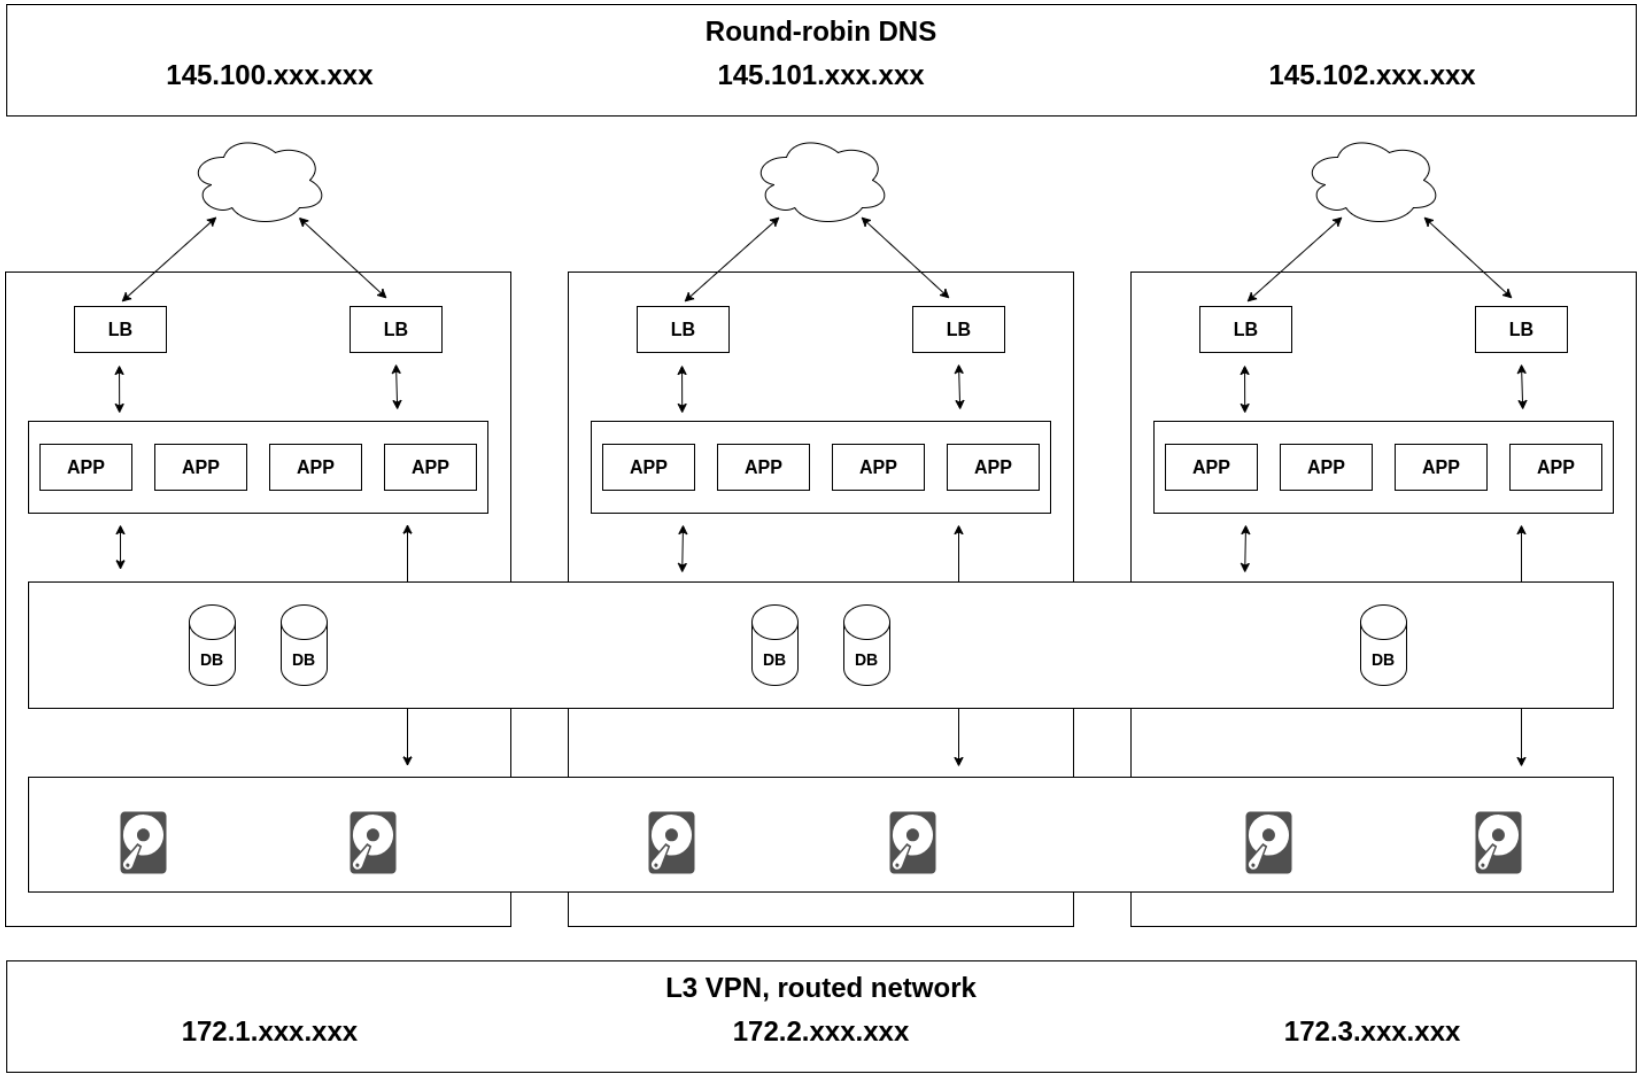

The big move

The big move

- **It was all about planning**
  - We had never done this before
  - Minimize stress-induced risk
  - Avoid manual operations, script as much as possible
  - 4 to 5 weeks per datacenter
- **Partly easy**
  - Nodes may be stateless such as webservers
- **Partly hard!**
  - Storage was the most complex
  - Even vendors (Scality) had no real prior experience
- **Everything while in production**

SURF

We are done!

- **Or are we?**
  - The proof of the pudding…
- **It must be tested**
  - Some involuntary testing…
  - Followed by some structured testing
- **Yes, we are done!**
- **But future enhancements are needed**
  - Controlling flow of traffic
  - Network interruptions are disruptive
  - But hard to troubleshoot

Thank you

13 March 2024