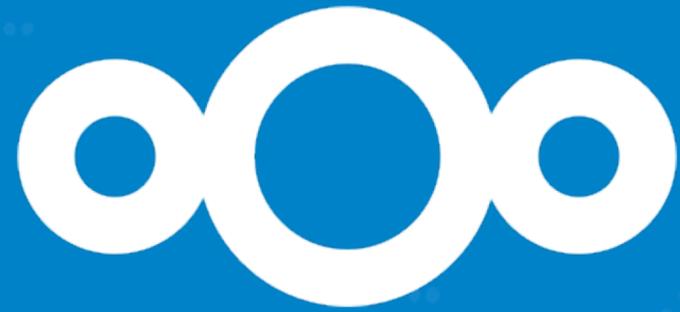


# An insiders look into scaling Nextcloud

Matthias Wobben





Matthias Wobben

[matthias@nextcloud.com](mailto:matthias@nextcloud.com)

- working in Berlin, Germany
- Senior Sales Engineer at Nextcloud



Before: 3<sup>rd</sup> level IT Engineer and Administrator at Systems Provider  
with focus on EFSS and collaboration portals

# Agenda

- I. Nextcloud in a nutshell
- II. Architecture
- III. Default Deployment
- IV. Example 1 | Education: Technical University Berlin
- V. Example 2 | Industry: Siemens
- VI. The idea behind Global Scale
- VII. Example 3 | Enterprise: Large-Scale Service provider



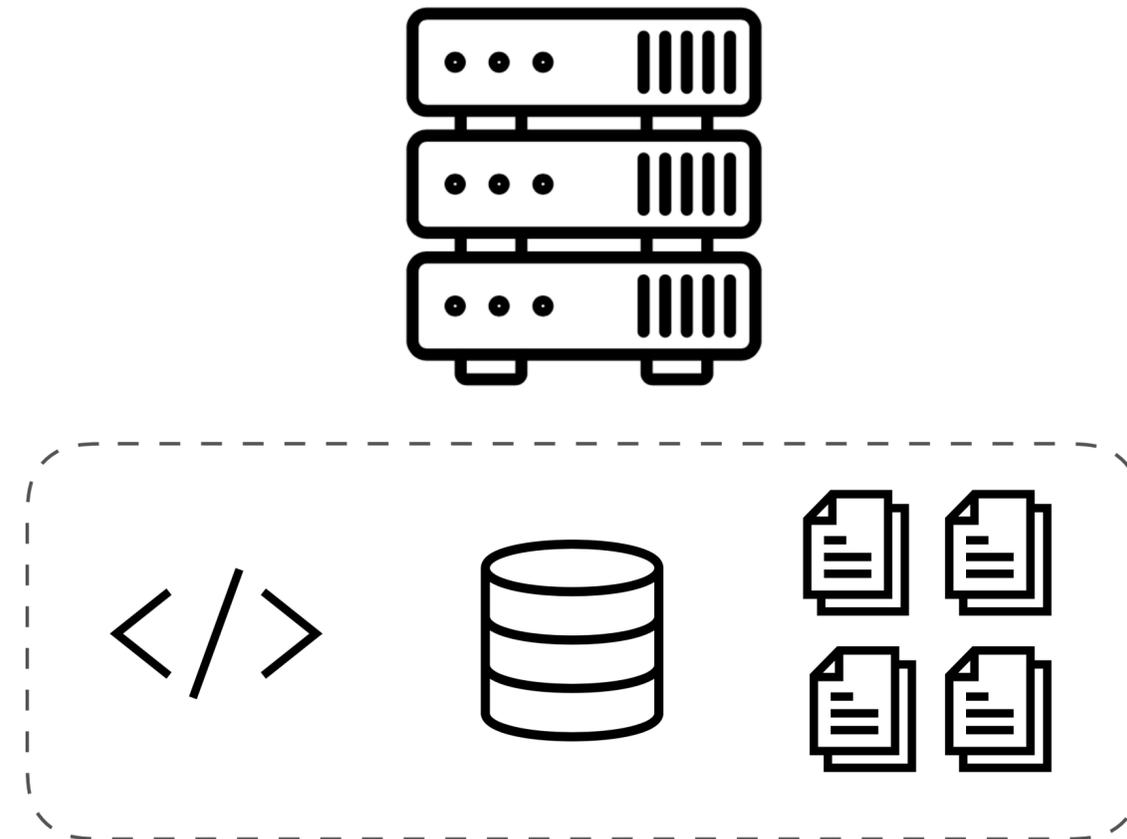
I. In a nutshell:



- secure, stable, flexible and scalable
- collaborative workspace and Enterprise File Sync and Share
- 100% open source

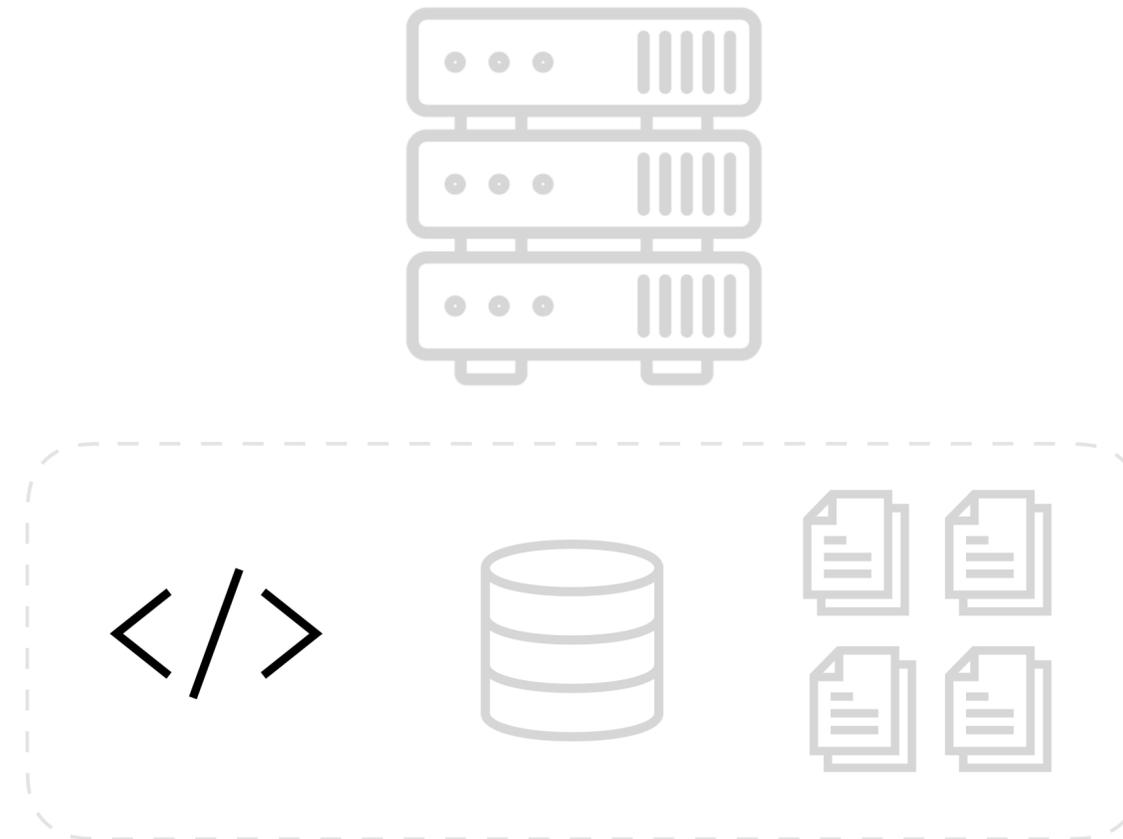


## II. Architecture



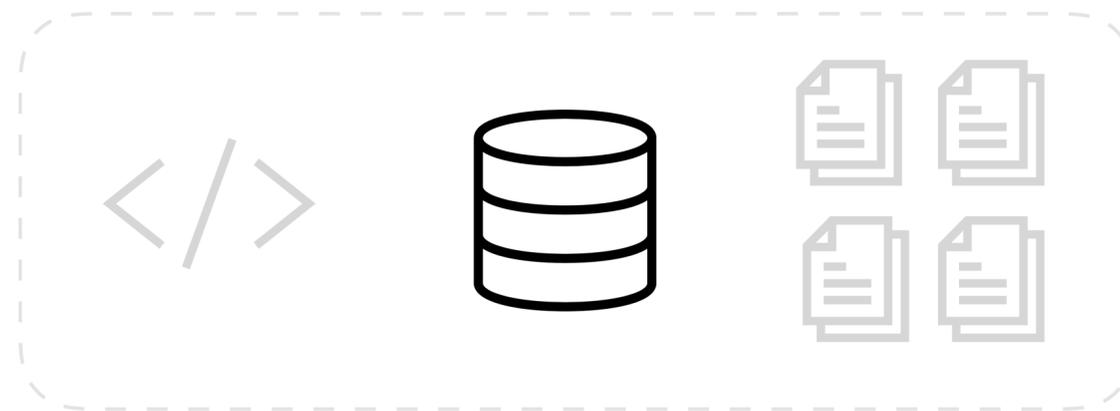
## II. Nextcloud

- server & app code
- PHP runtime
- web server



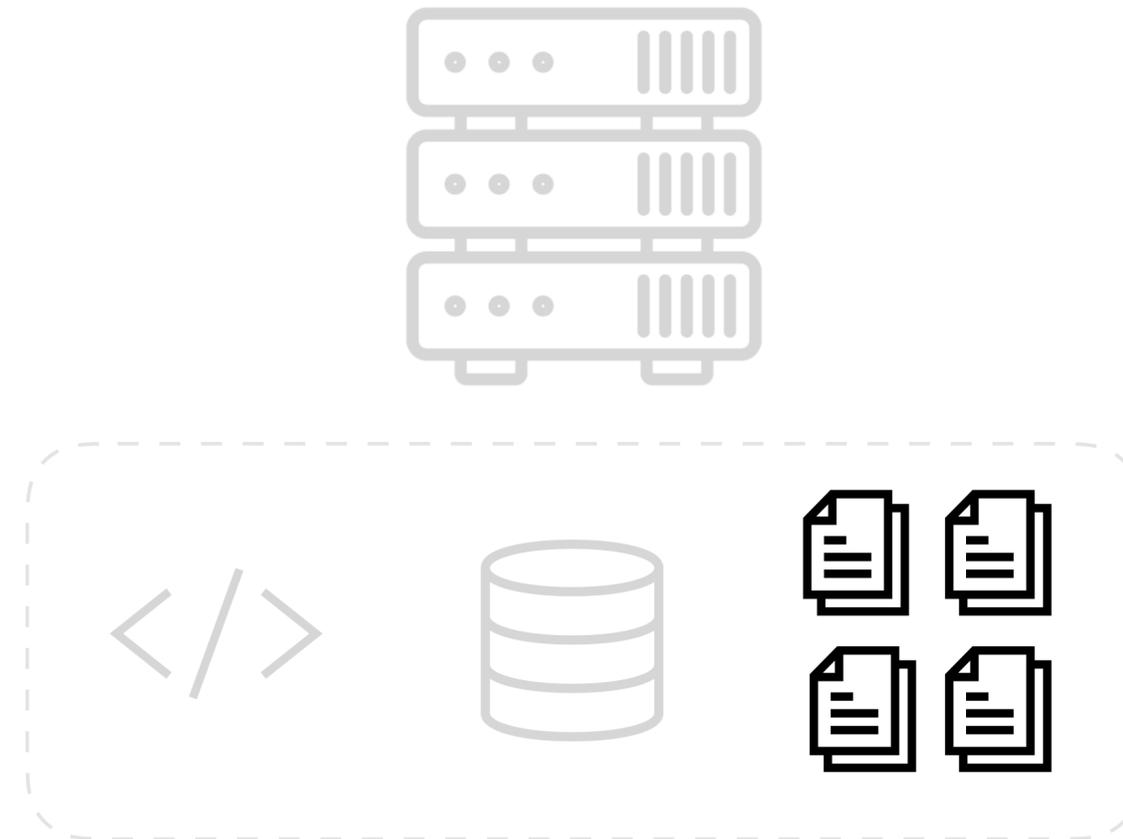
## II. Database

- MySQL/MariaDB or PostgreSQL
- stores everything except the binary files
- users, shares
- settings, ...

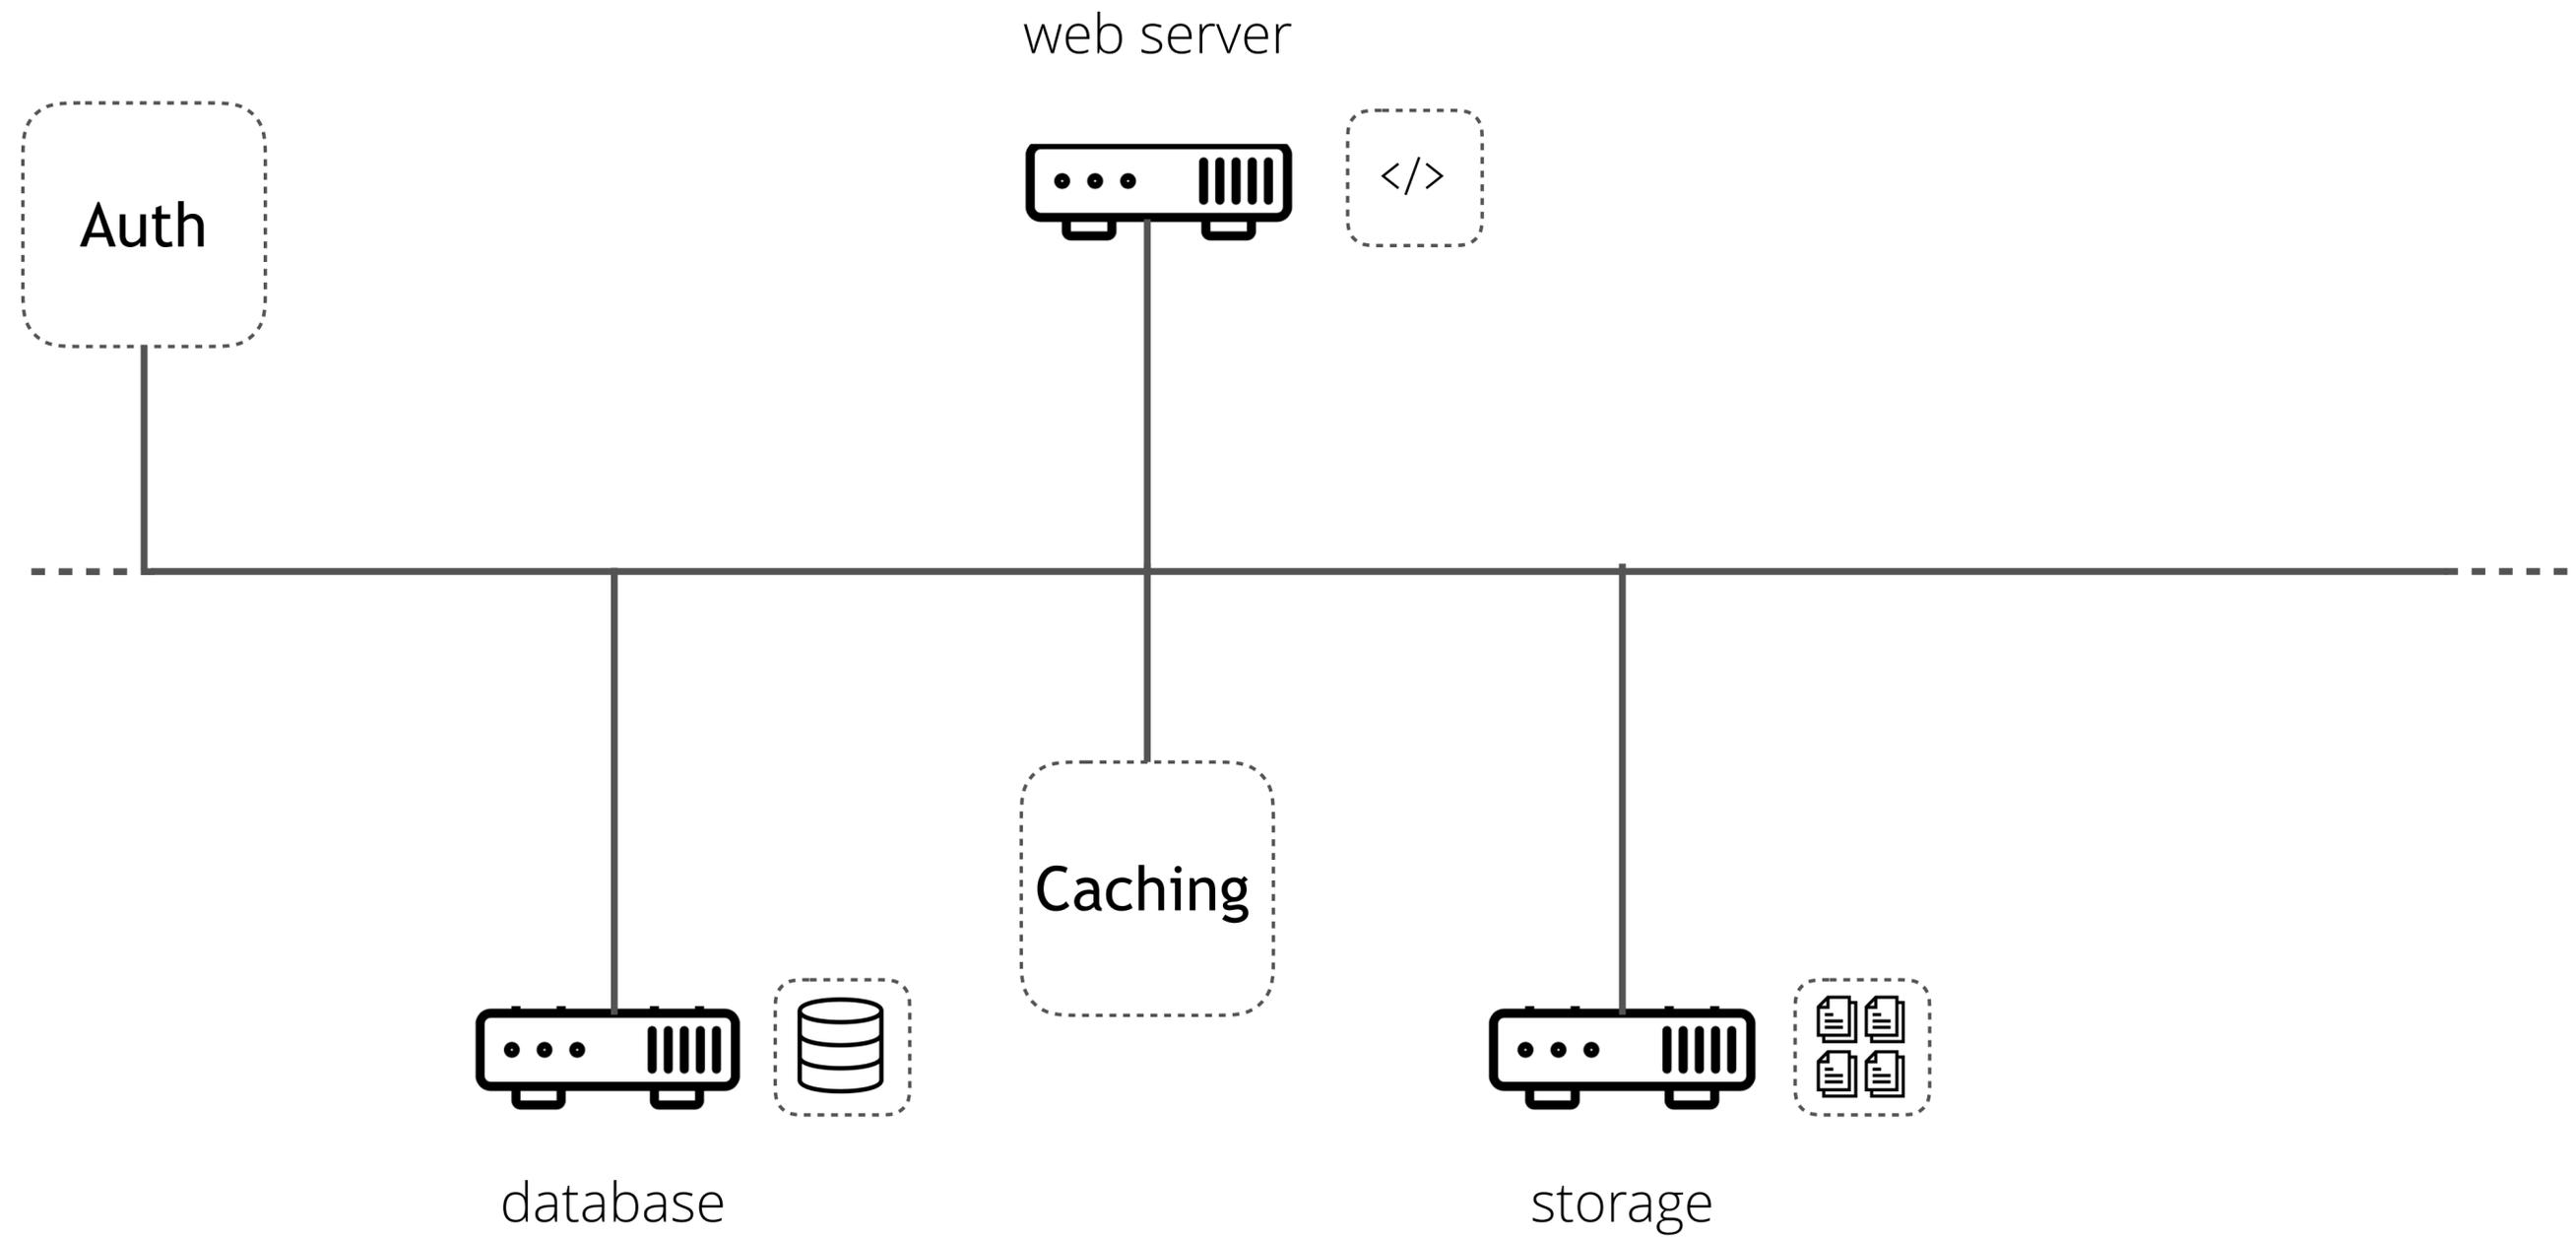


## II. Files

- user uploaded files
- versions
- trash bin
- previews
- ...

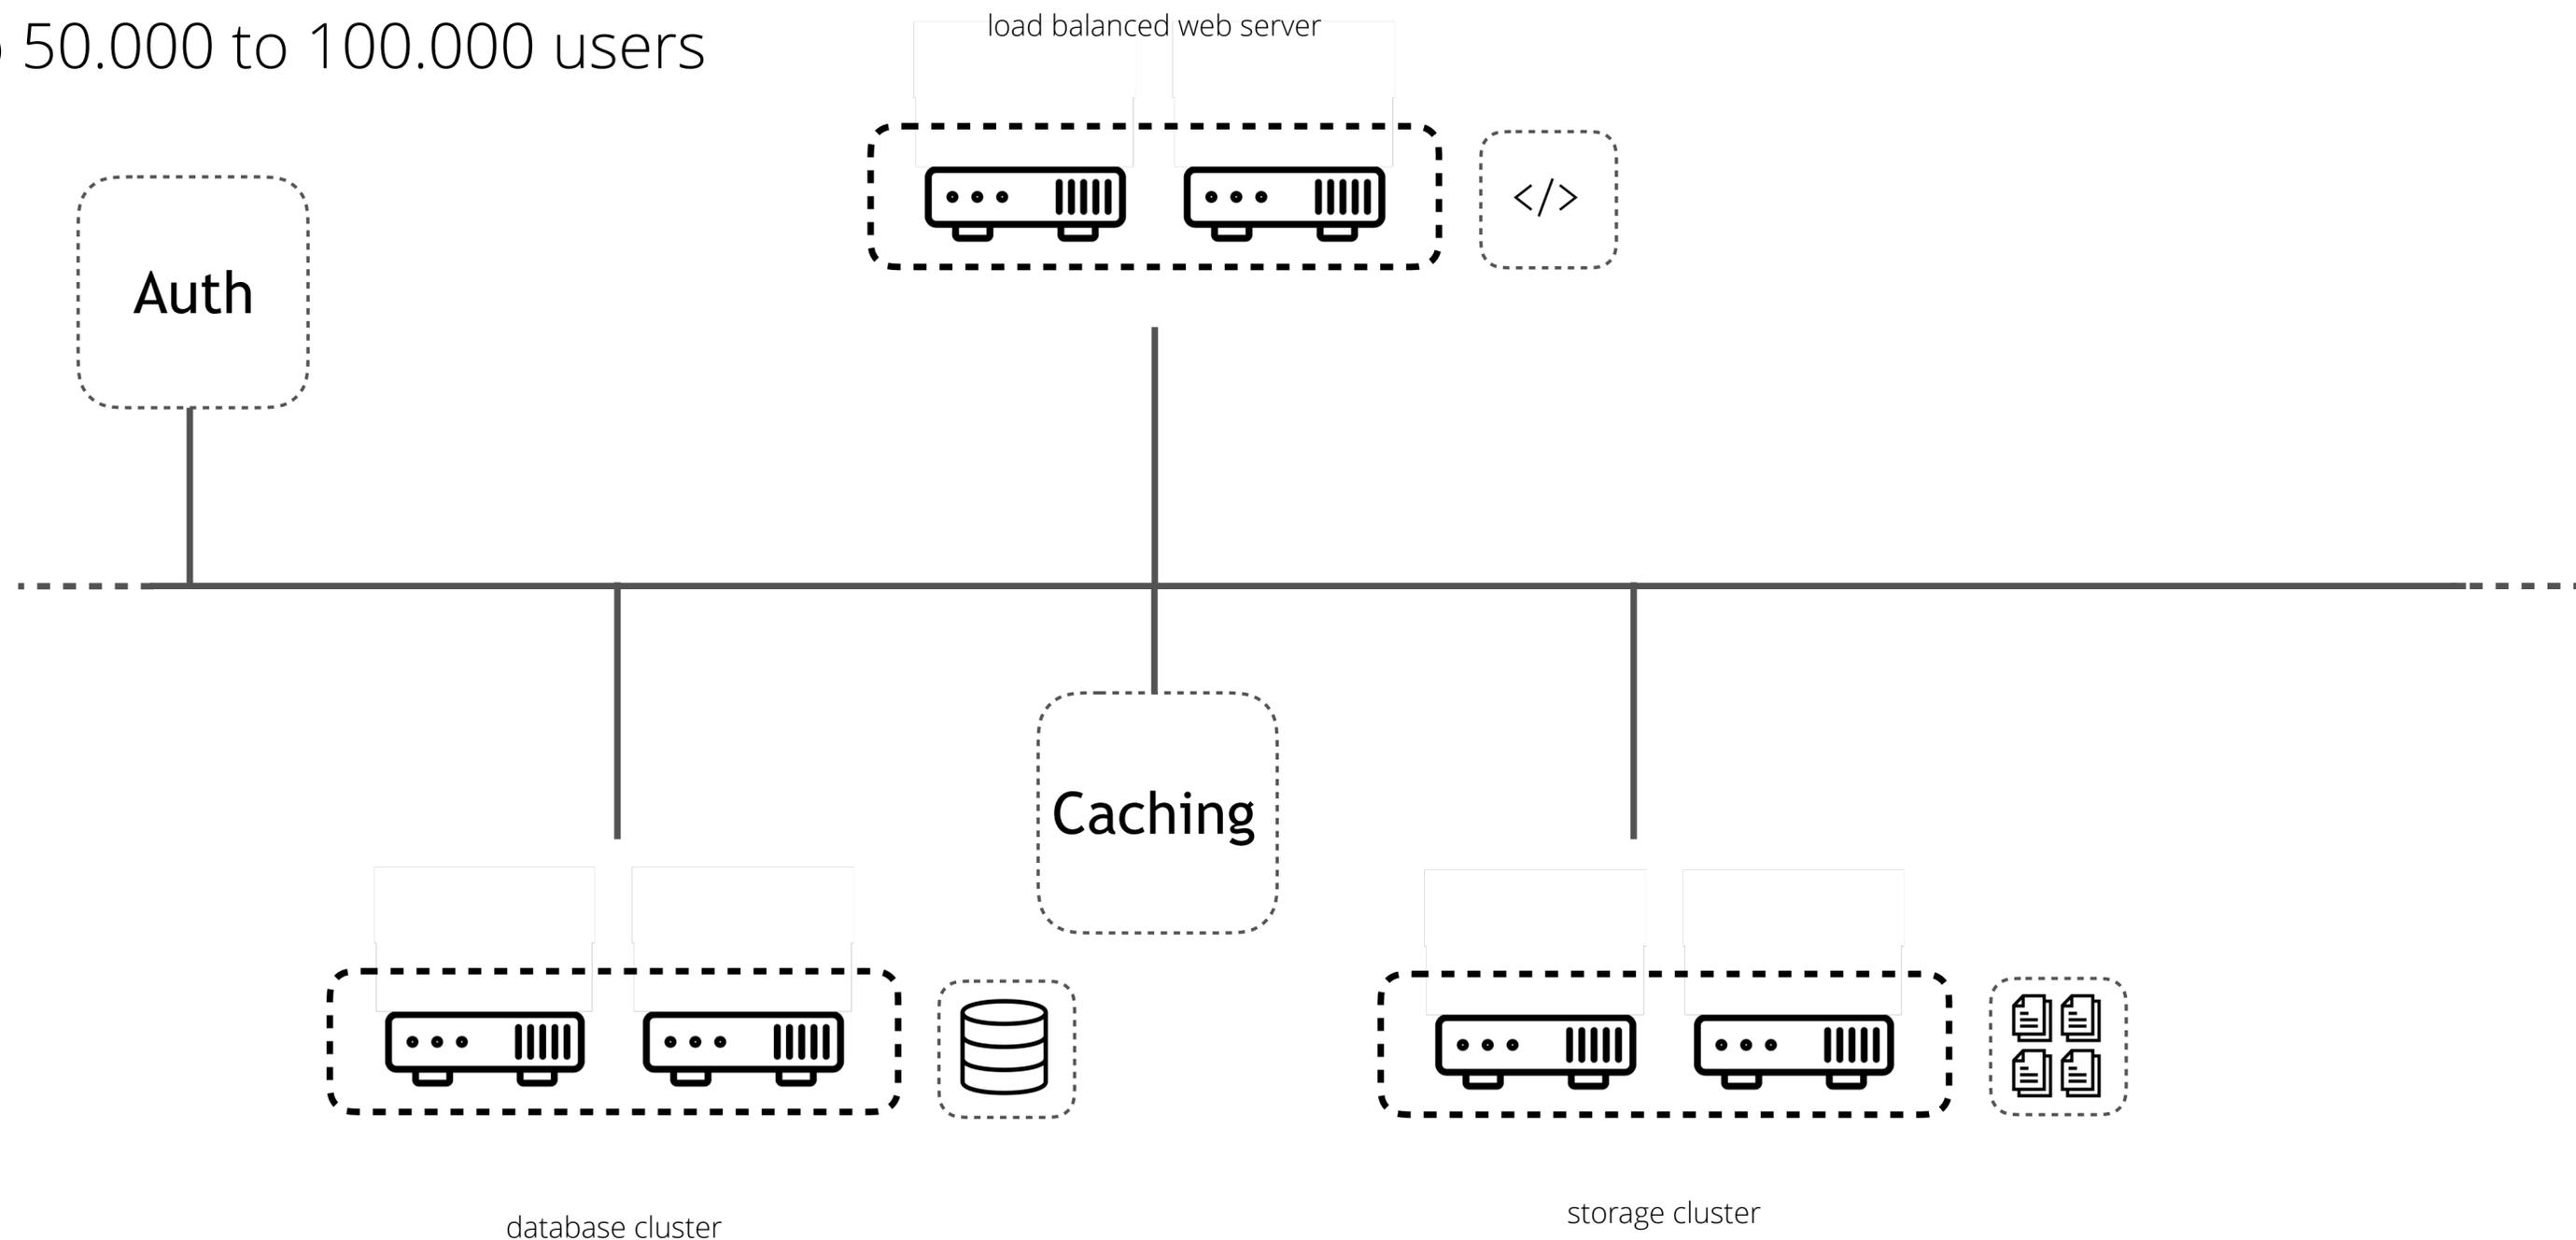


## II. Split it up

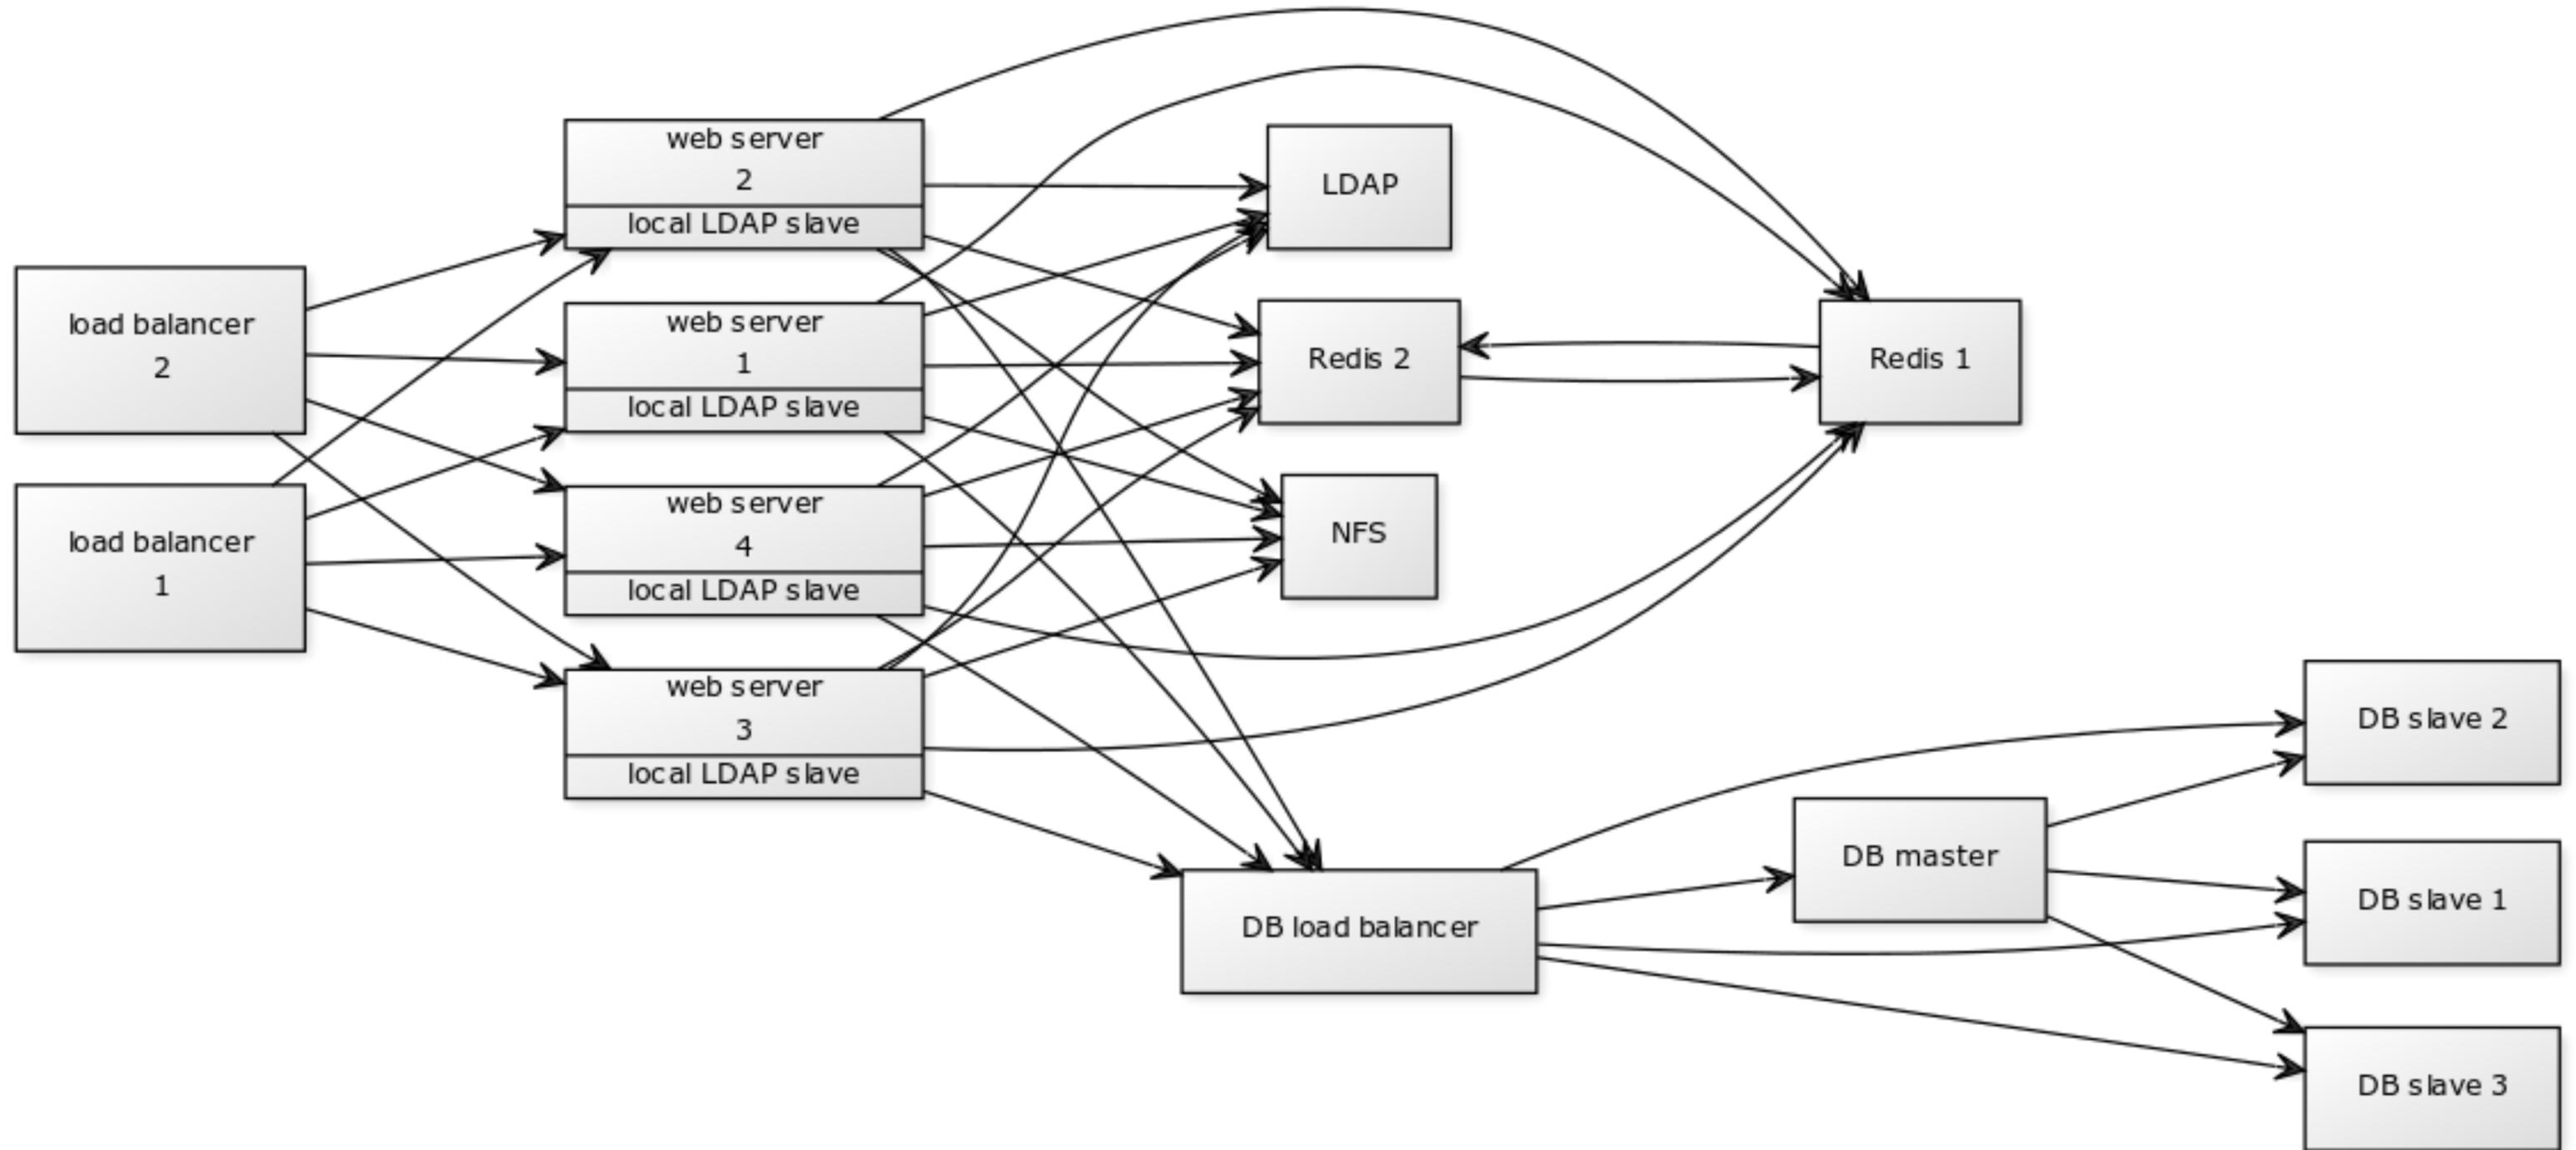


# II. Capacity

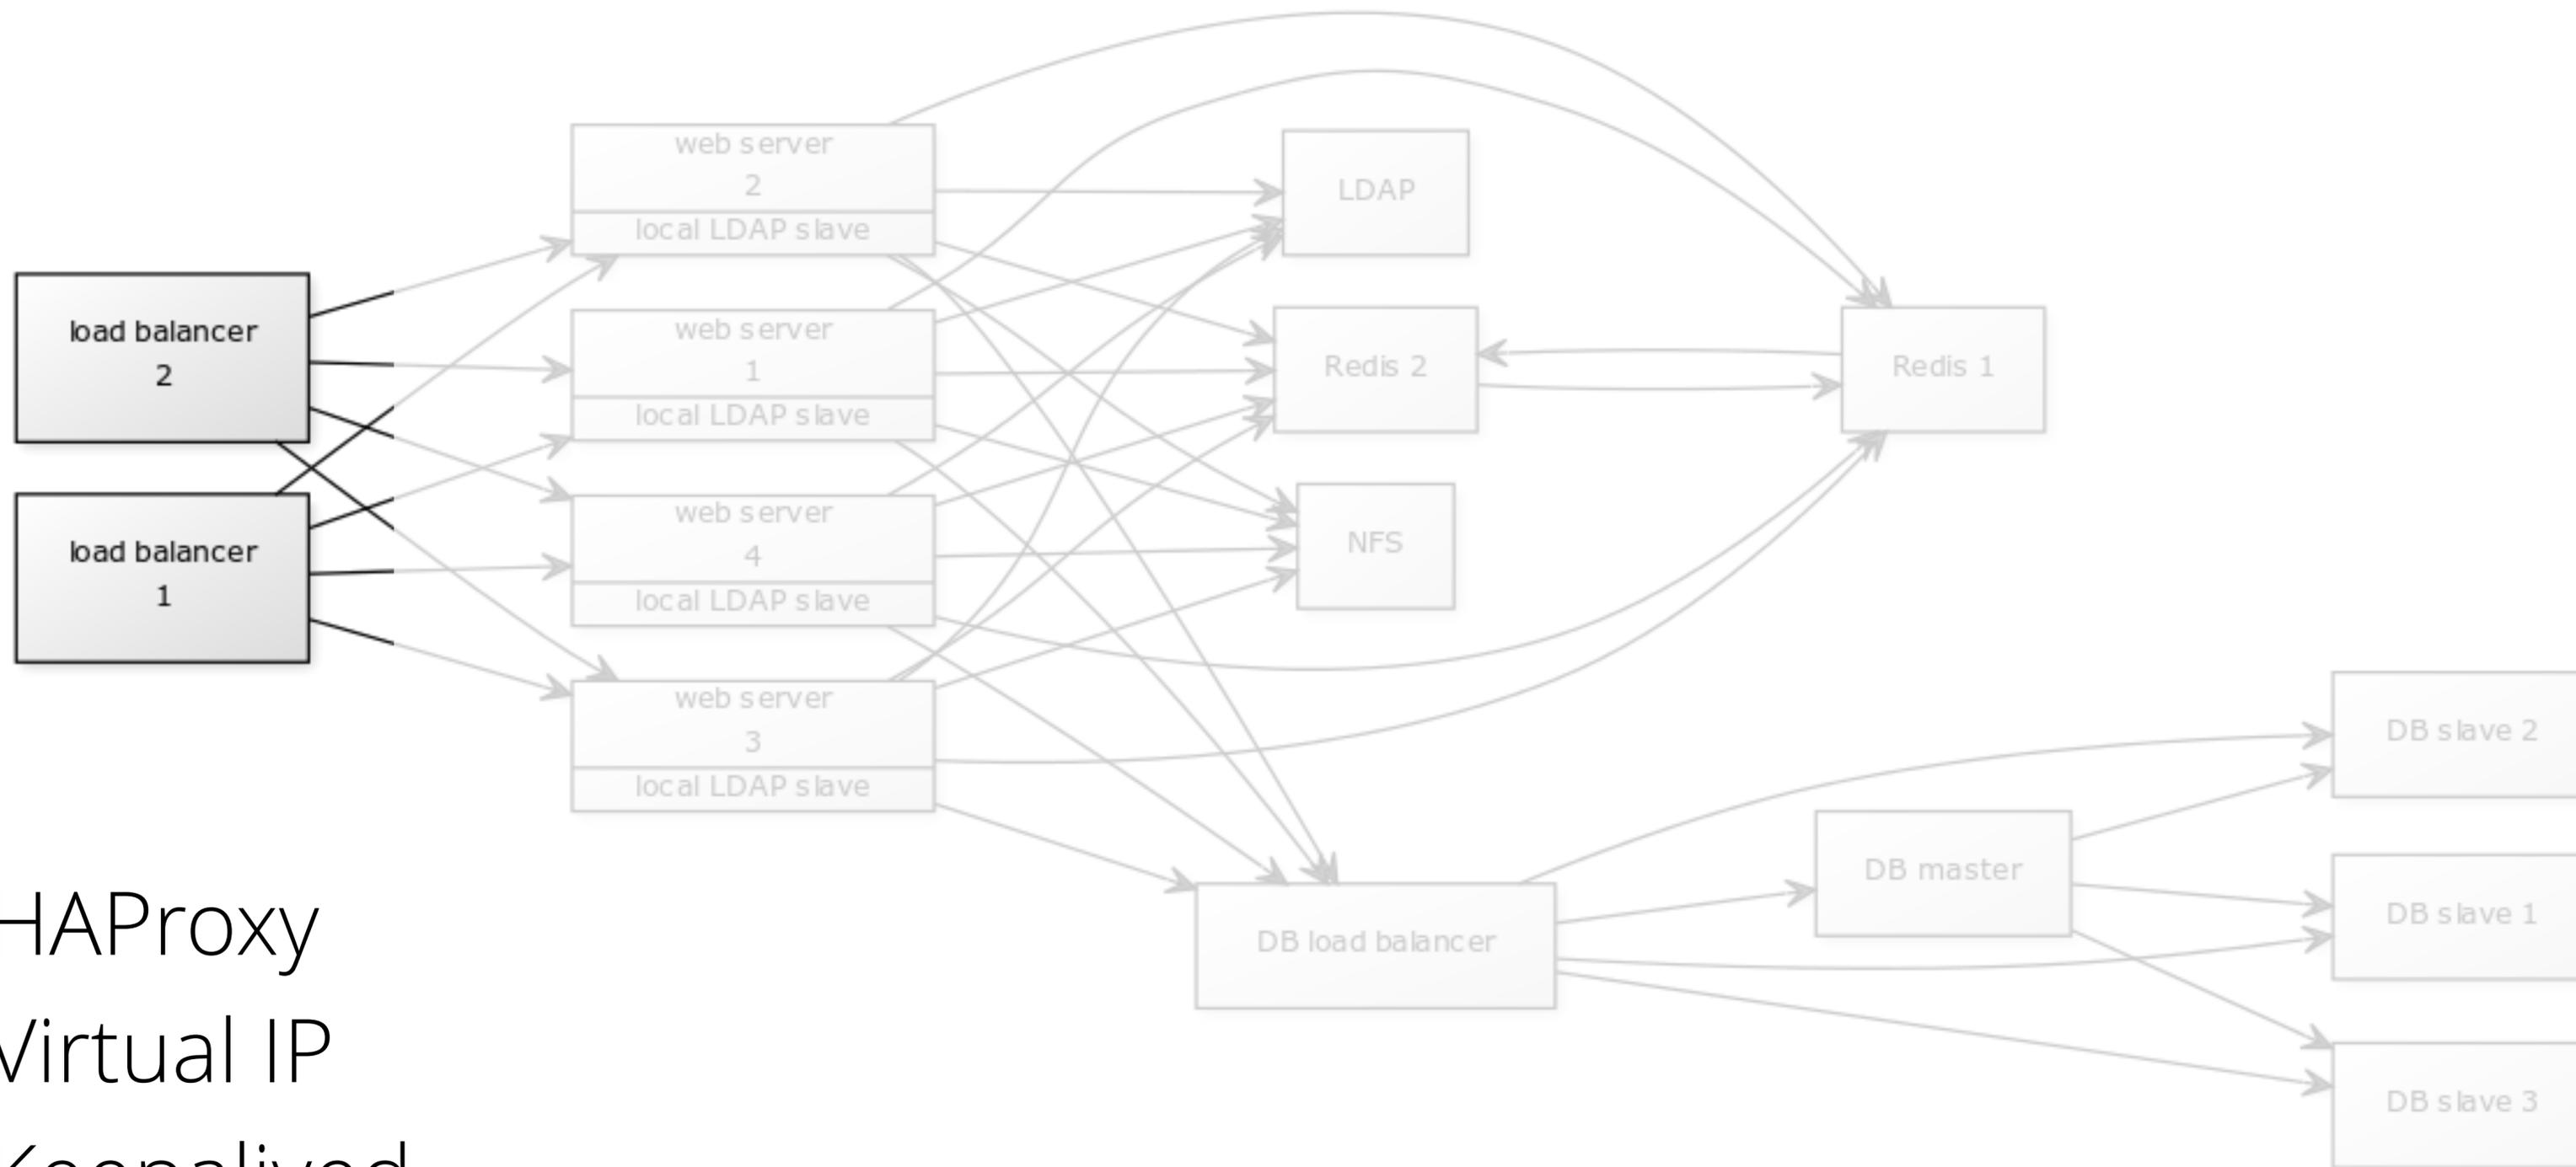
up to 50.000 to 100.000 users



# III. Default Deployment



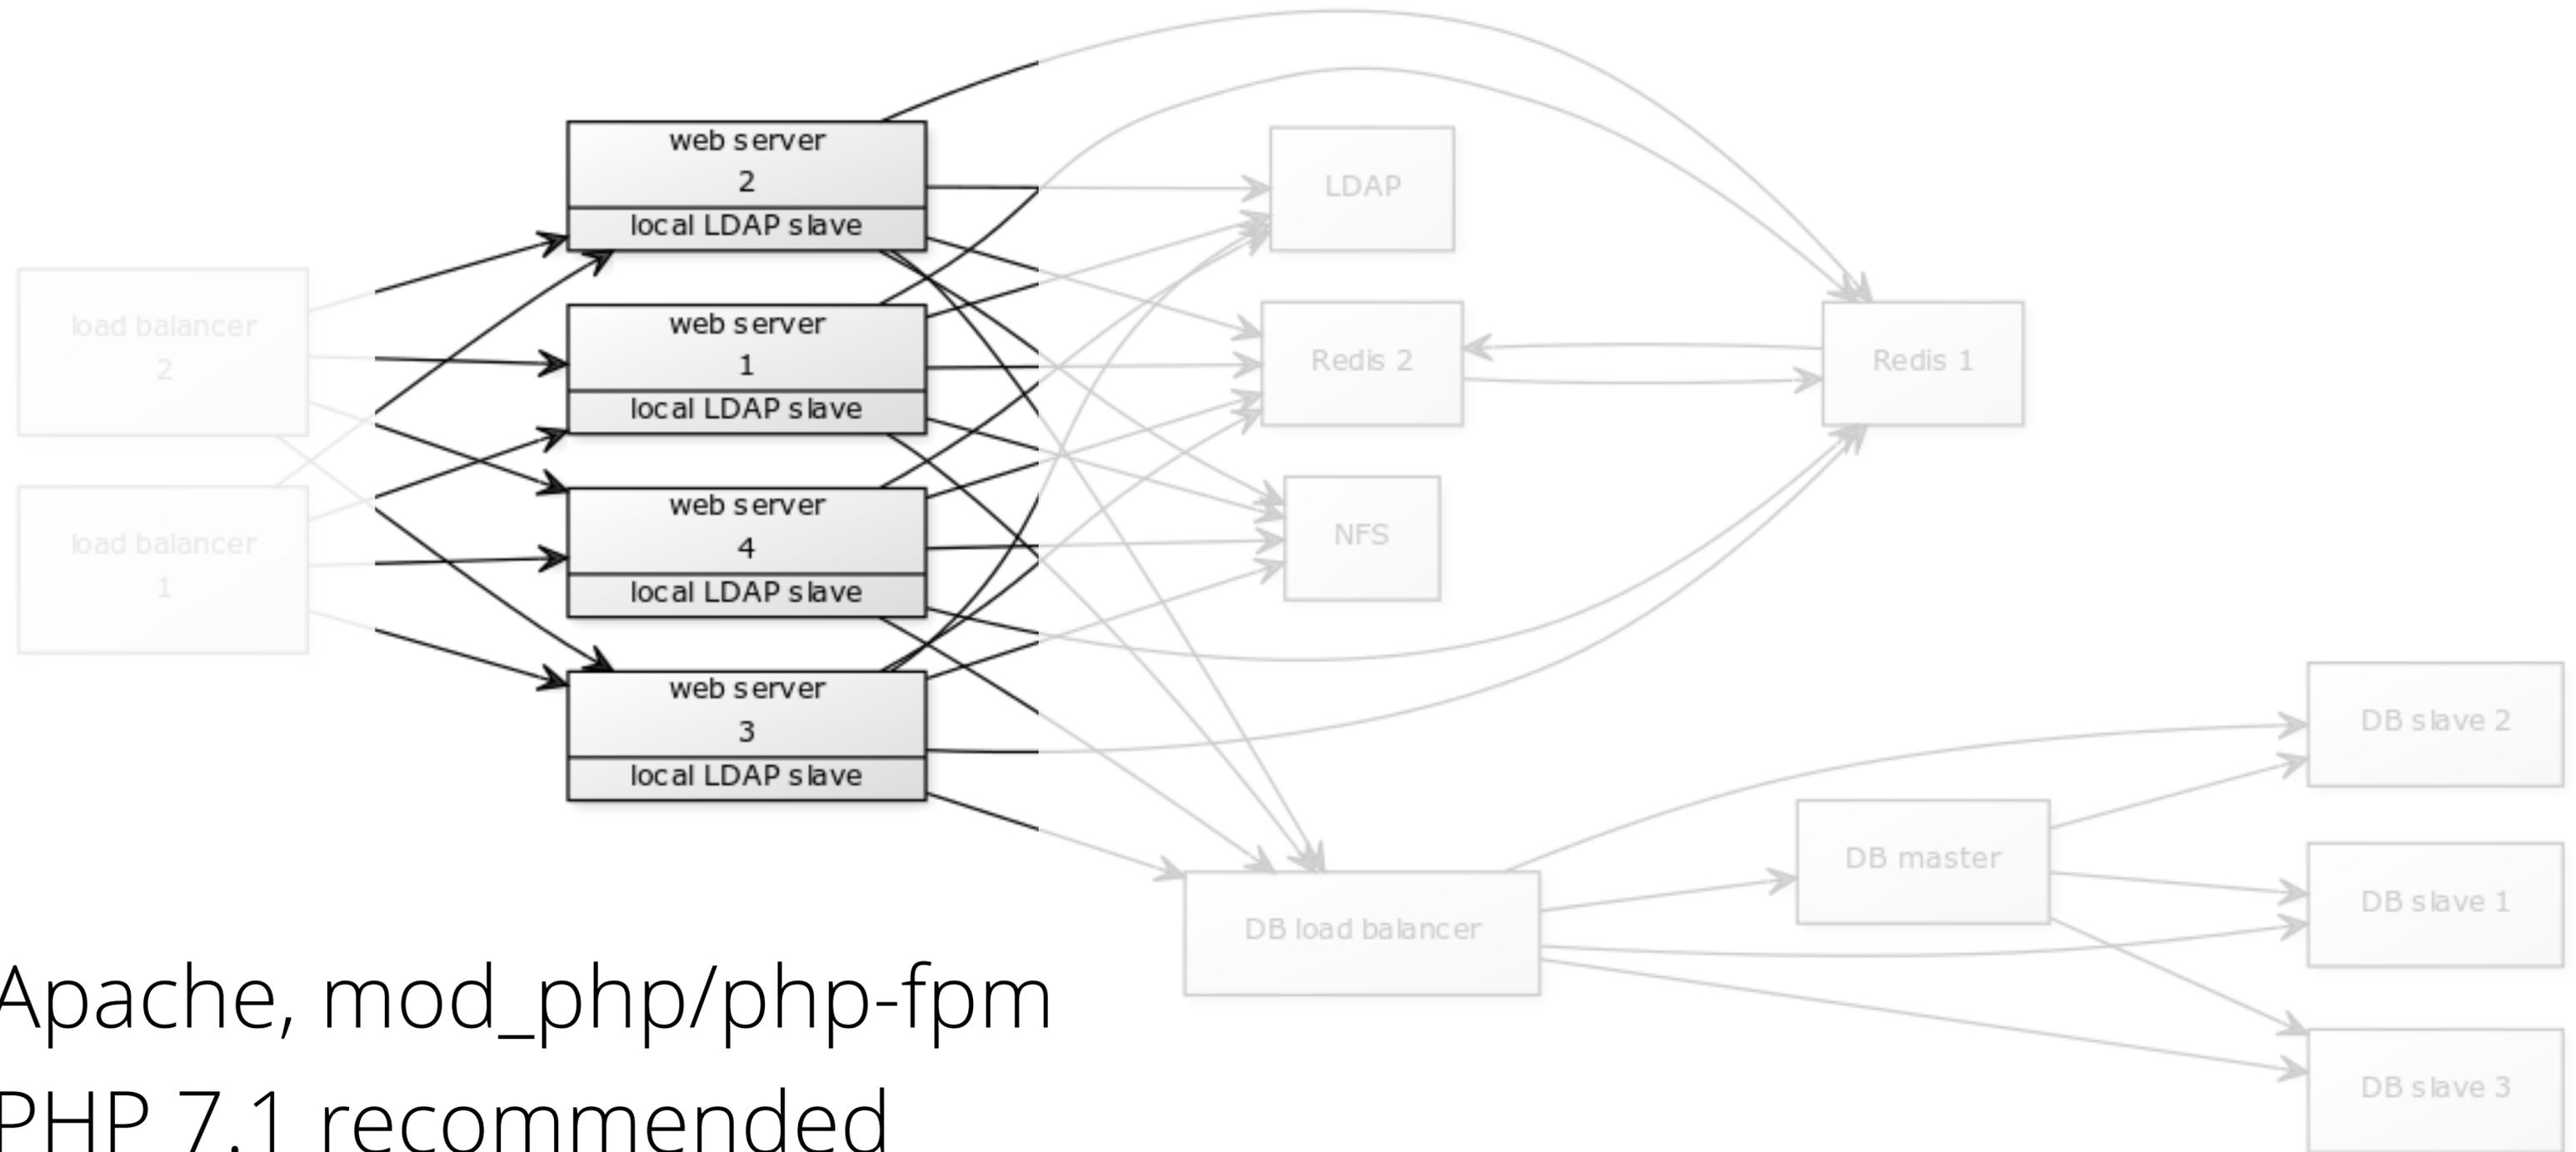
# III. Default Deployment



- HAProxy
- Virtual IP
- Keepalived

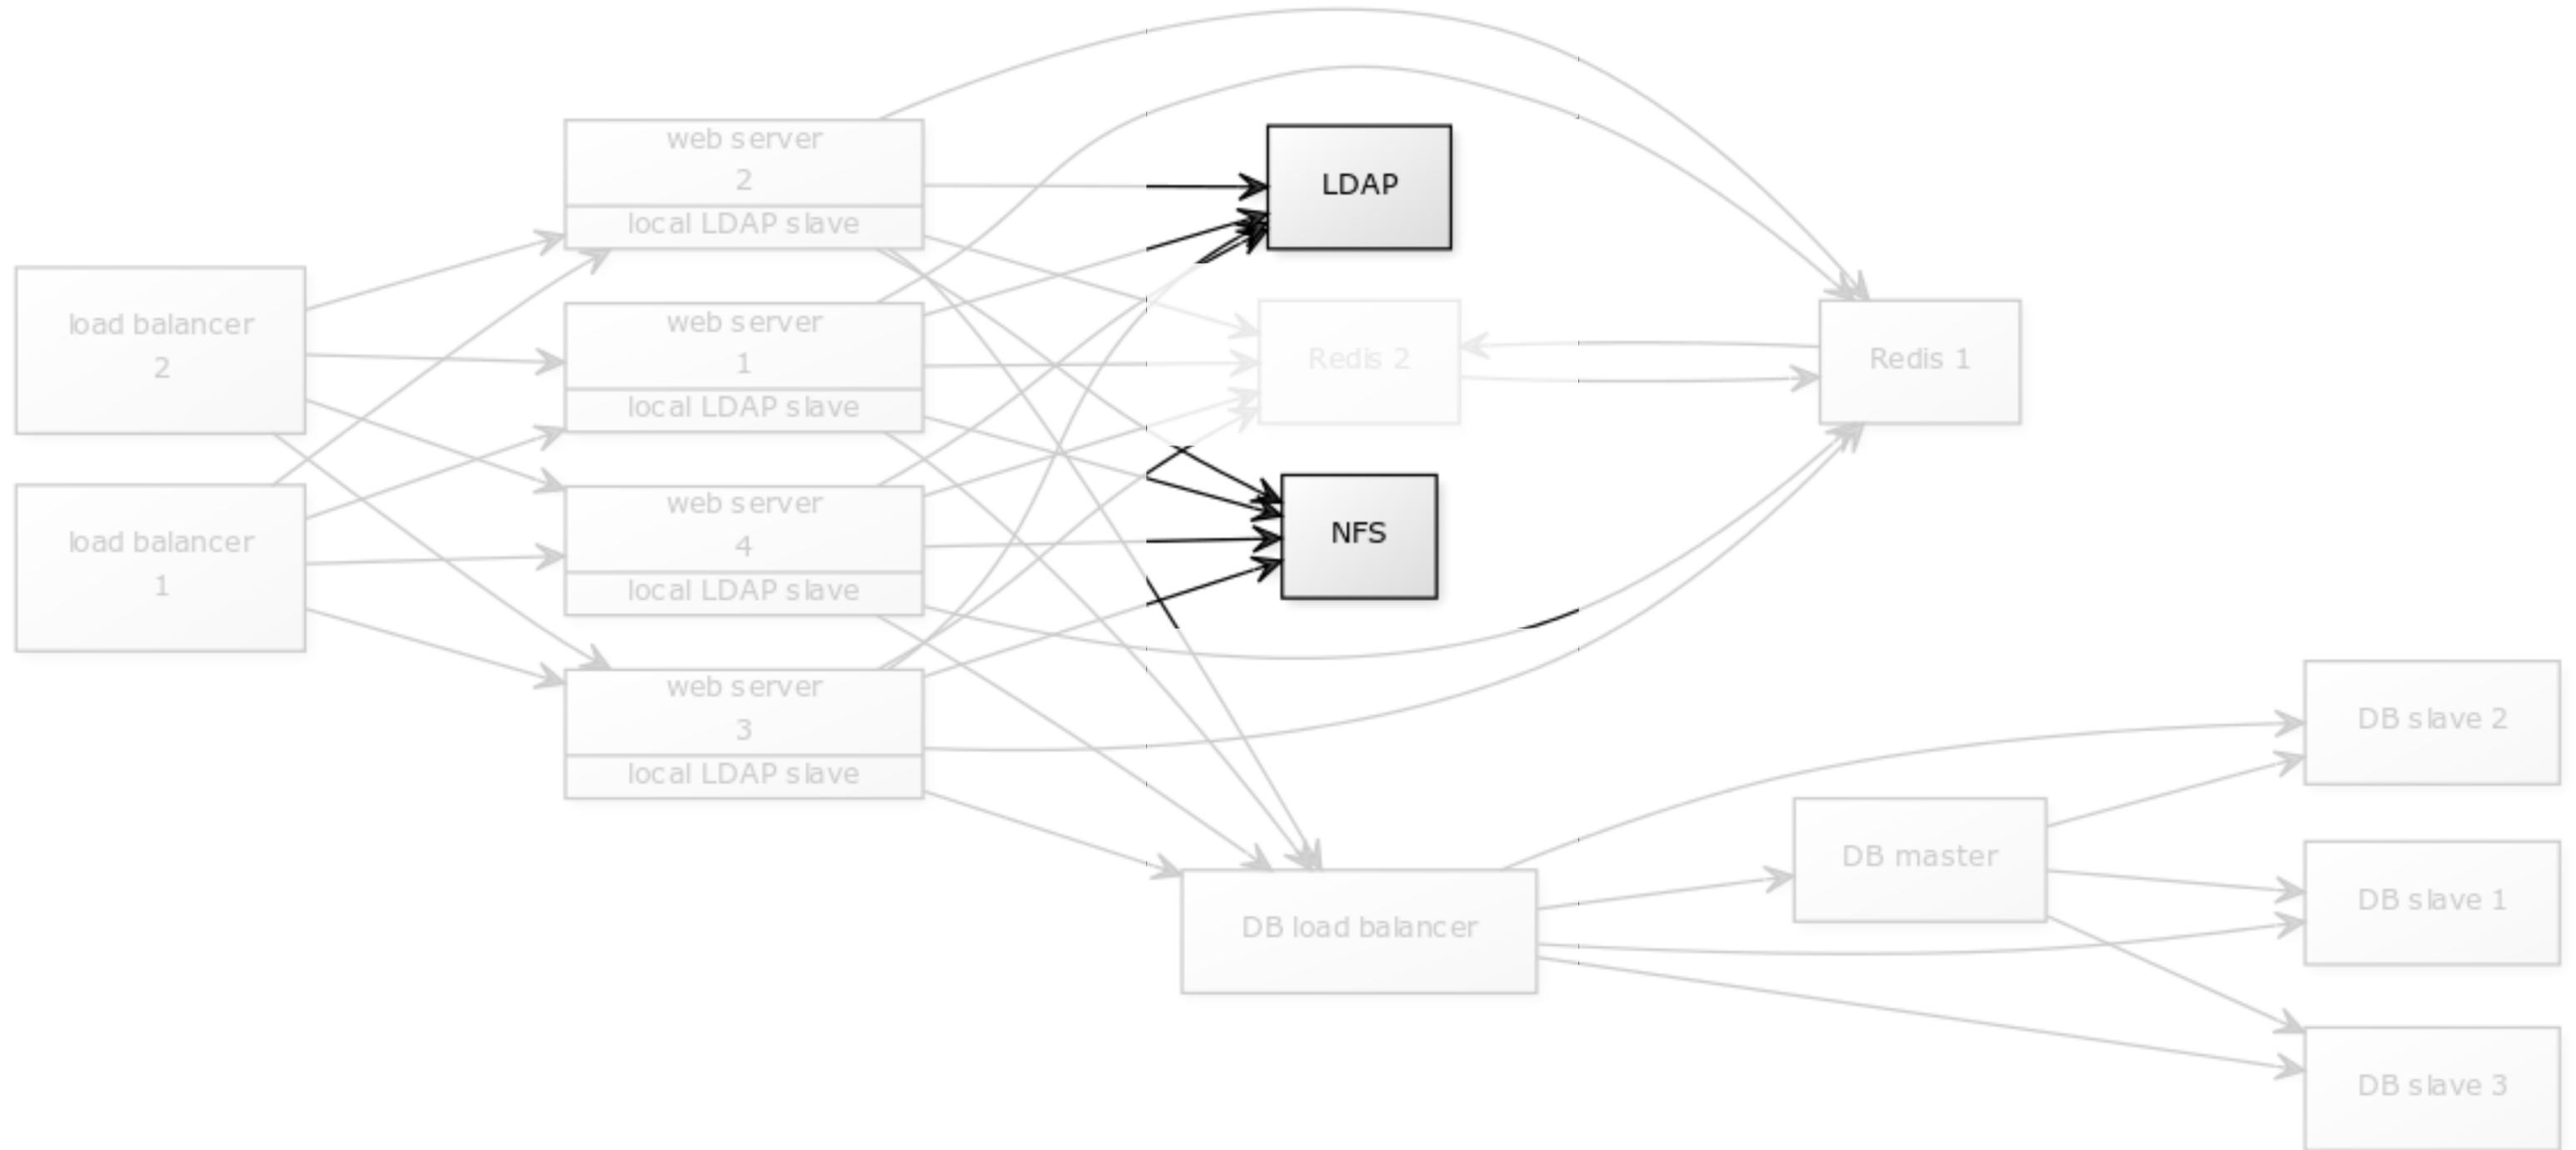


# III. Default Deployment

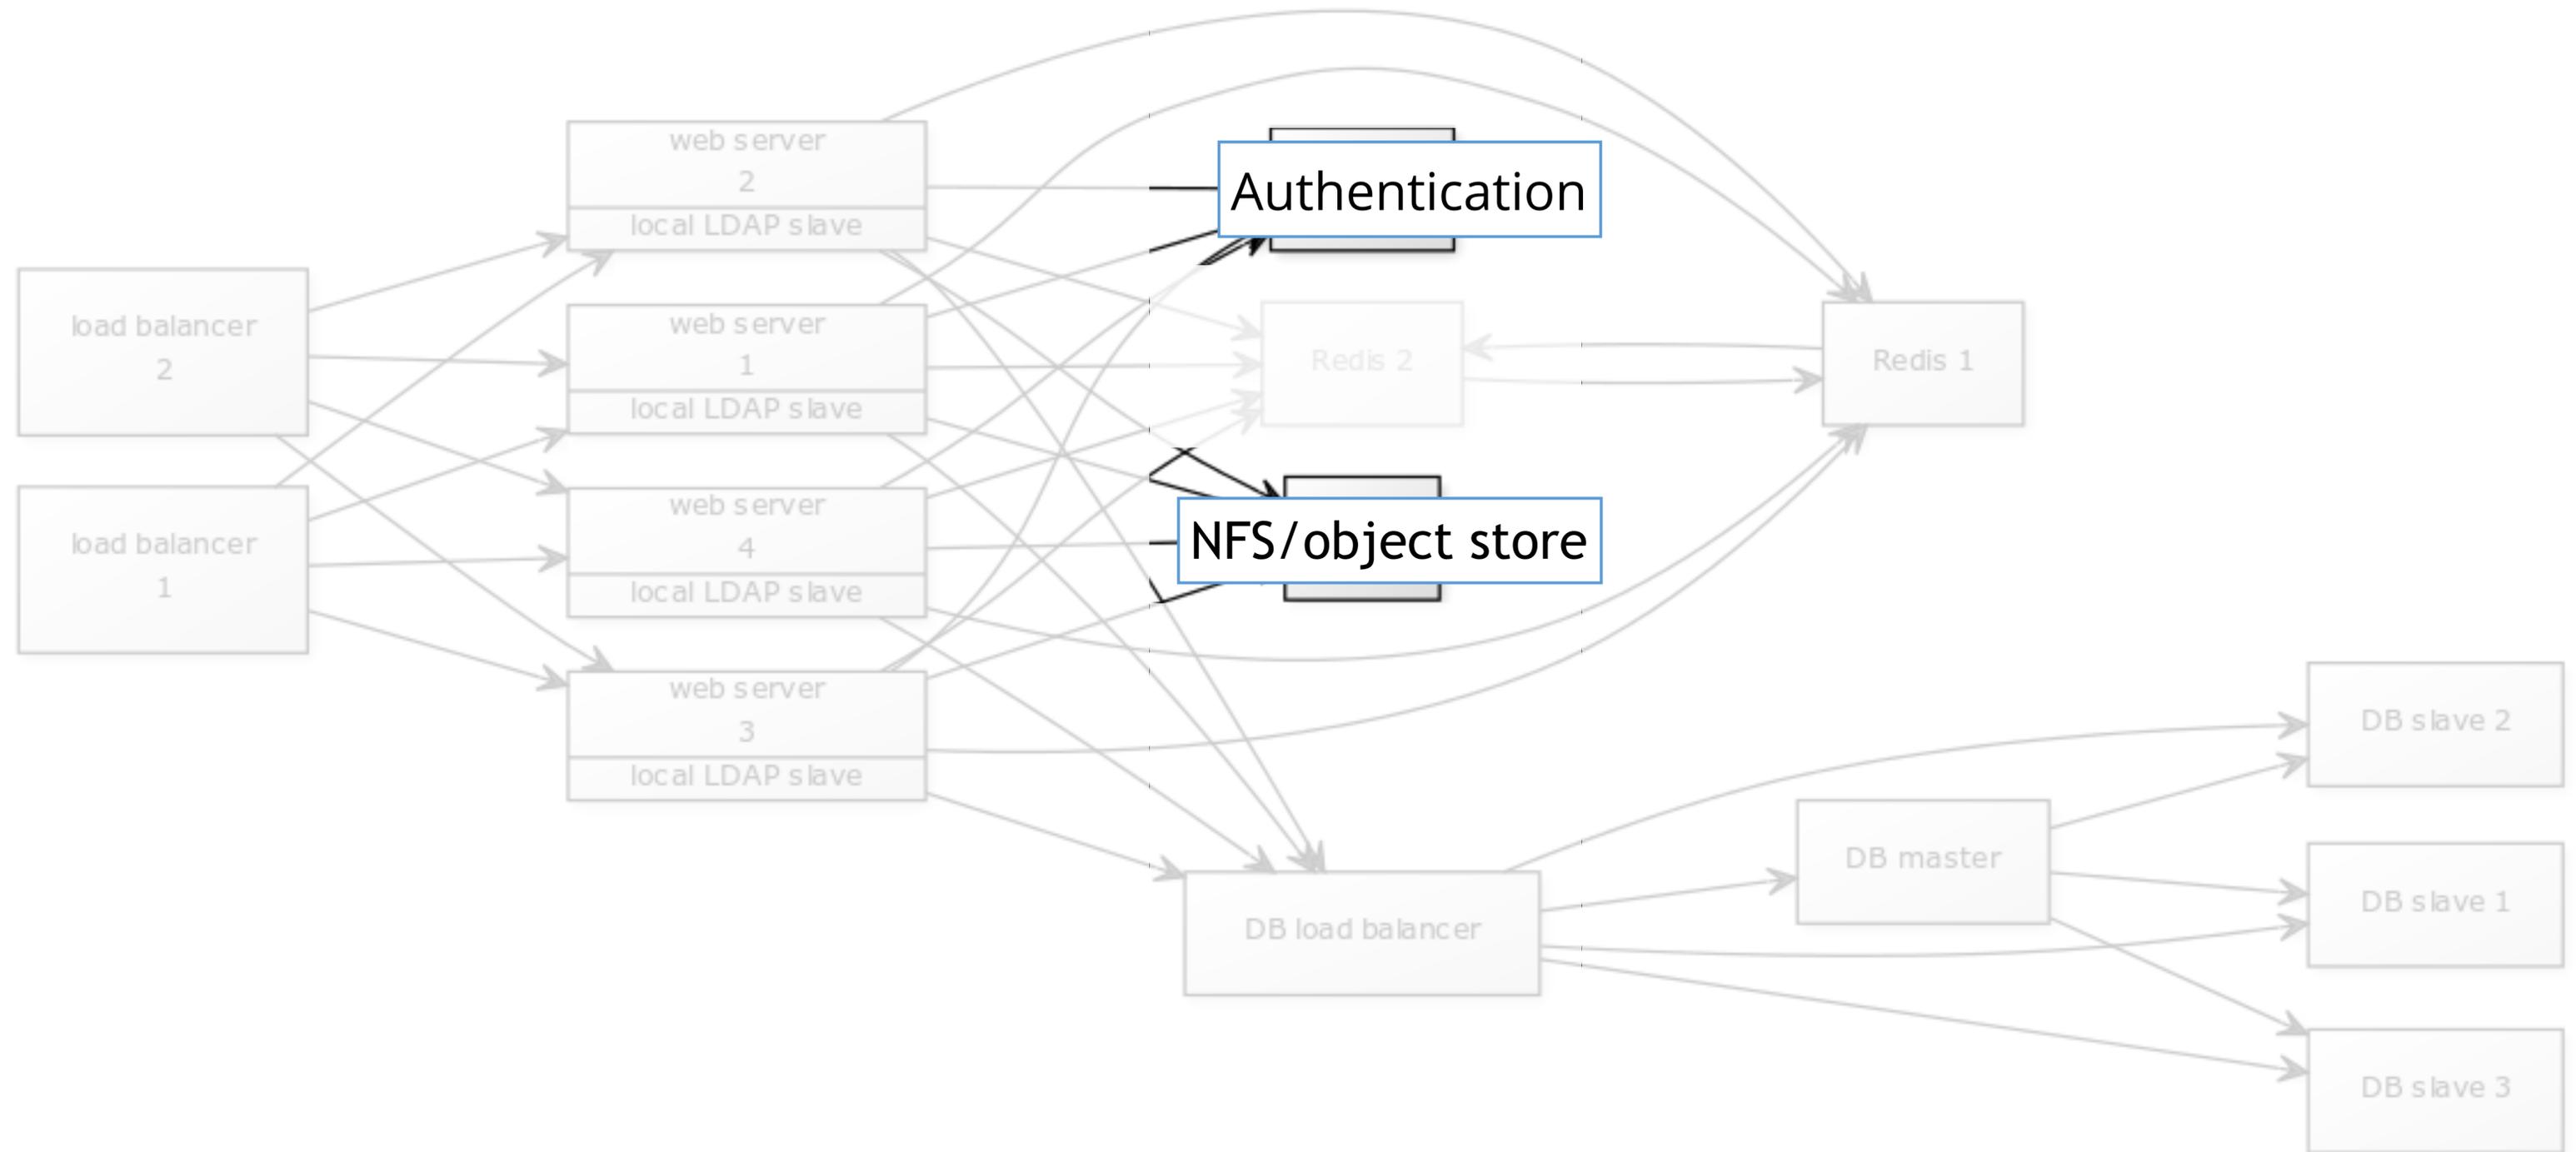


- Apache, mod\_php/php-fpm
- PHP 7.1 recommended

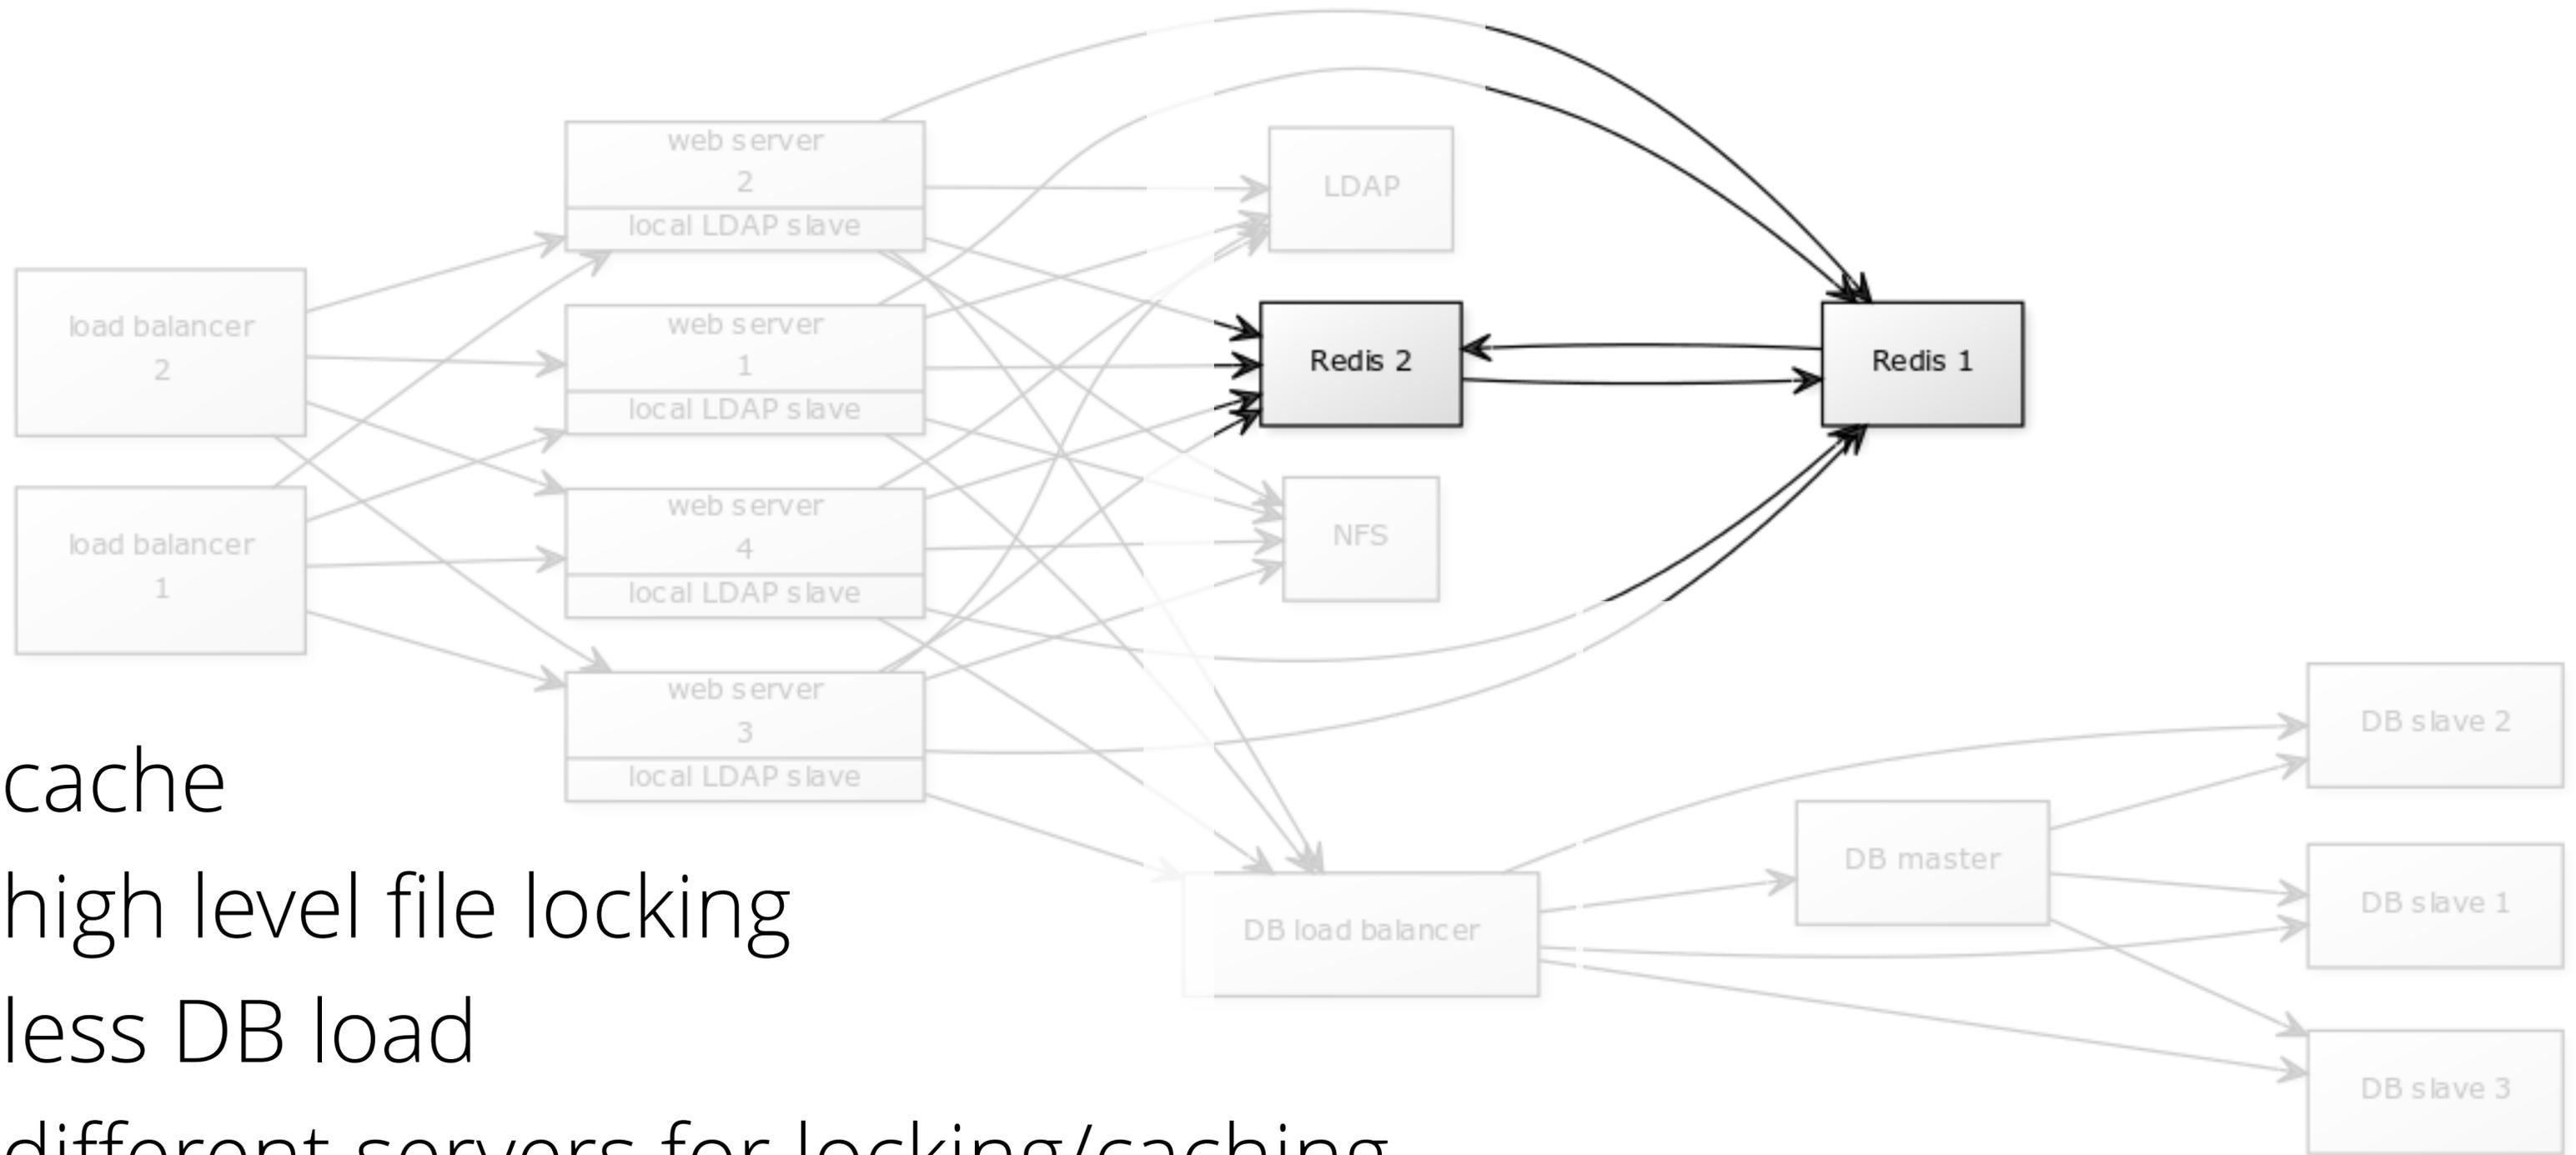
# III. Default Deployment



# III. Default Deployment

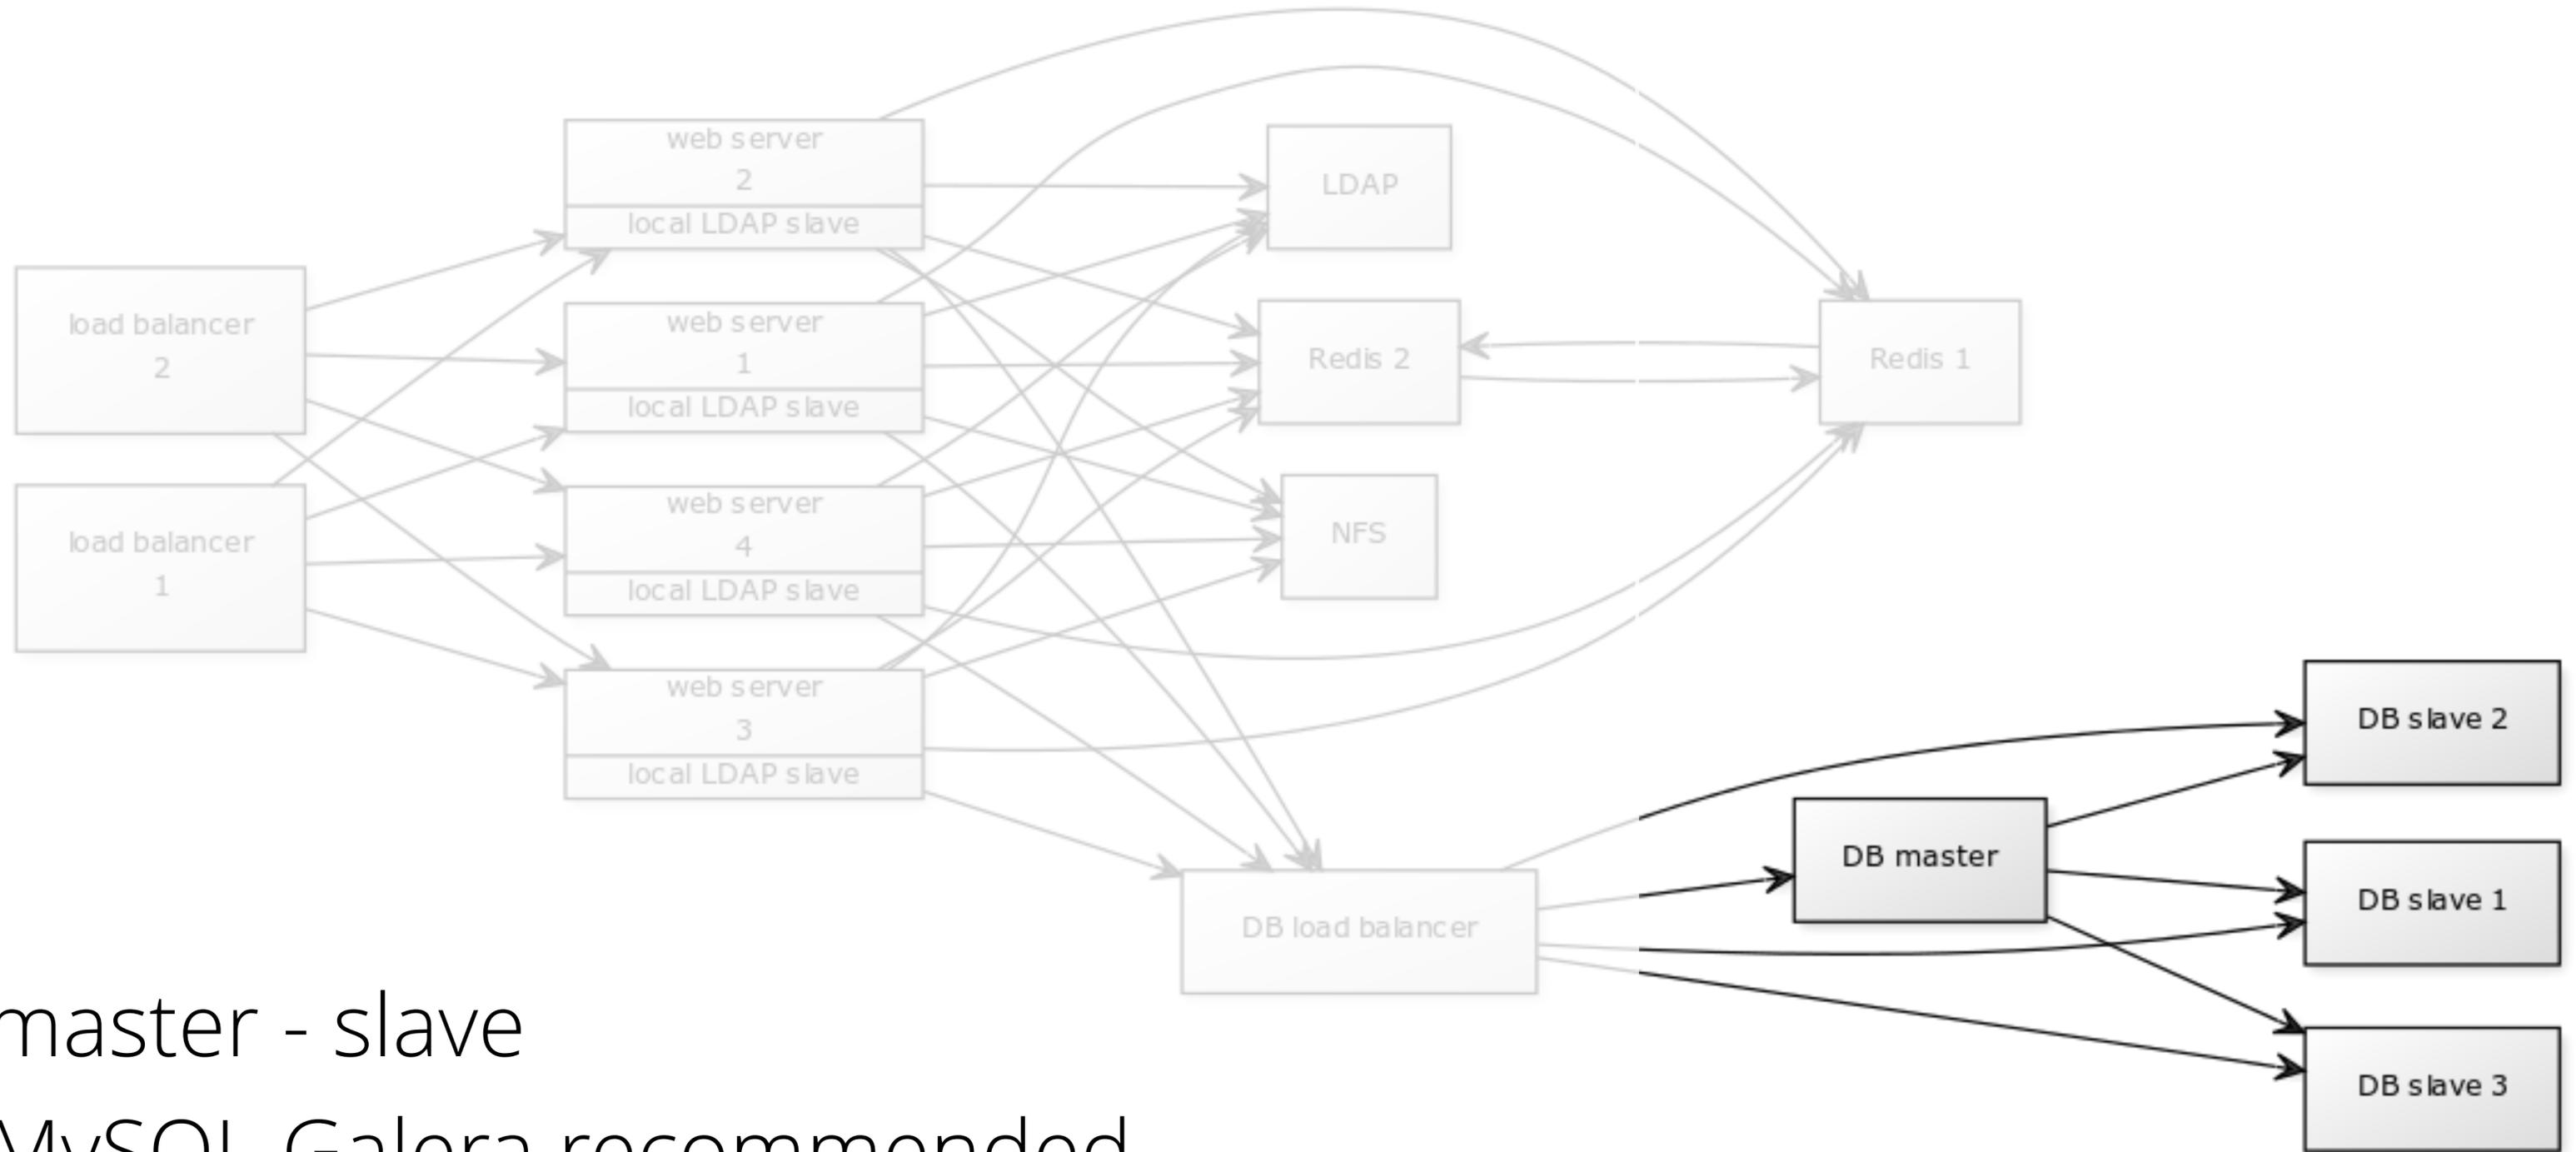


# III. Default Deployment



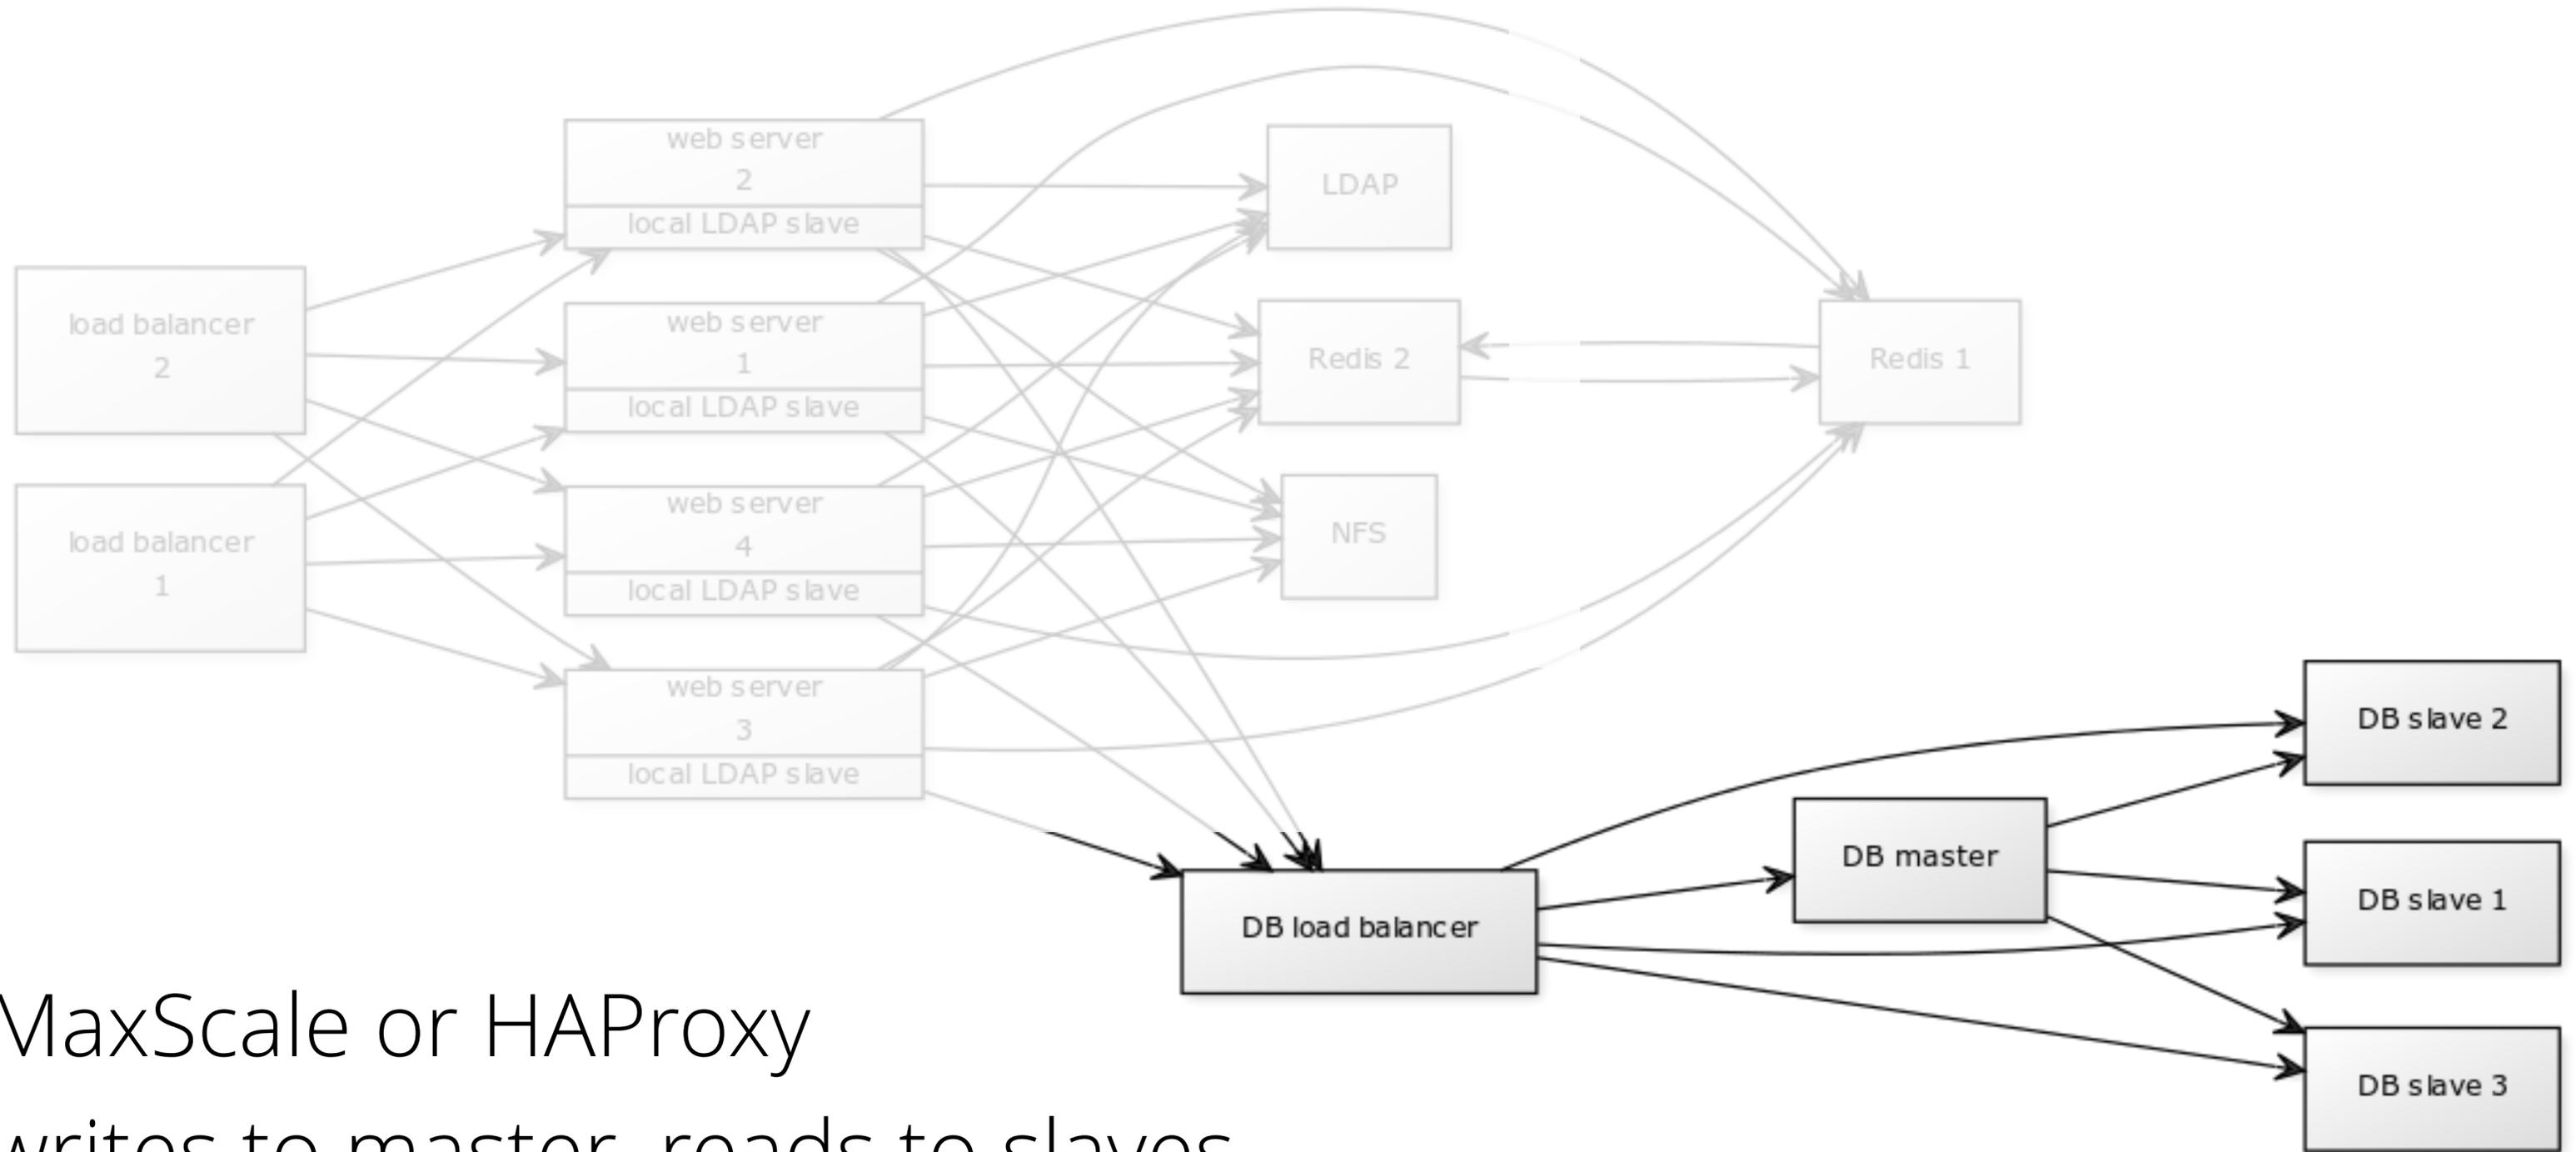
- cache
- high level file locking
- less DB load
- different servers for locking/caching

# III. Default Deployment



- master - slave
- MySQL Galera recommended

# III. Default Deployment



- MaxScale or HAProxy
- writes to master, reads to slaves

# IV. Example 1: Technical University Berlin



2012

First Review of Cloud Storage / Sync-n-Share solutions

2017

migration of 22.000 users to Nextcloud.

80.000 of 100.000.000 files are changed/moved on a daily basis

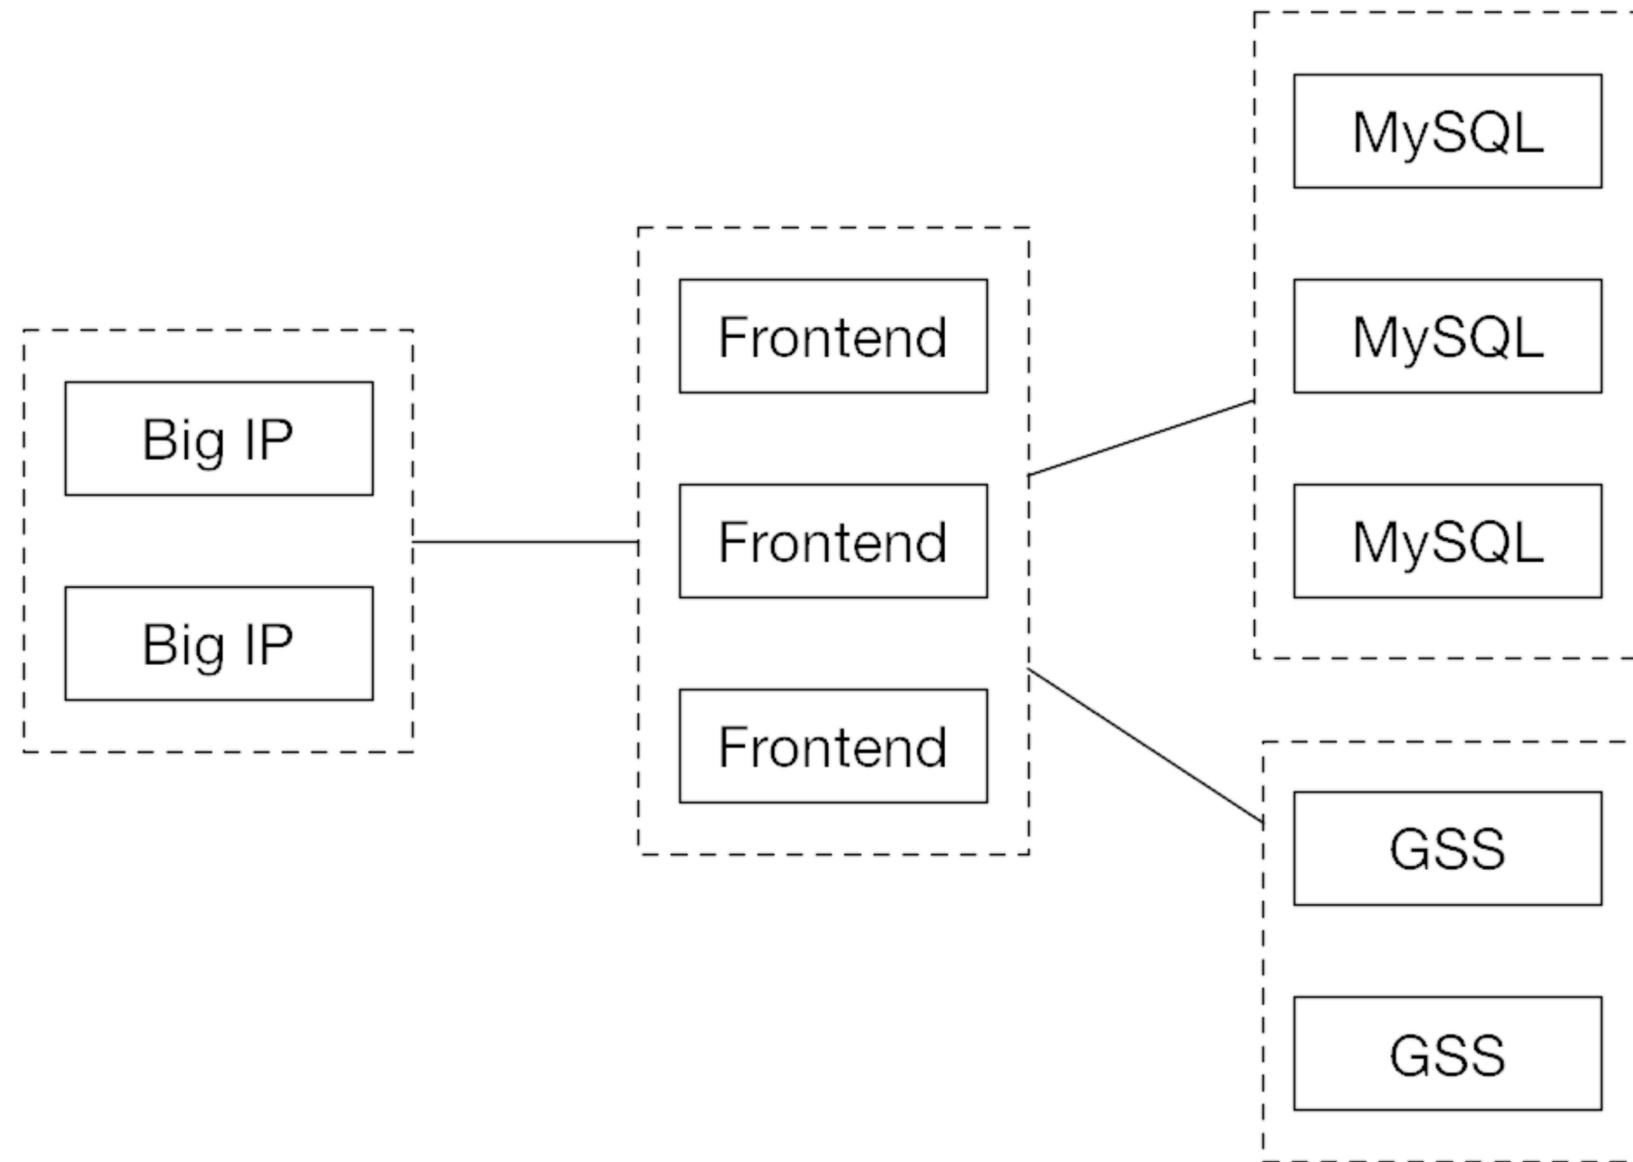
70 TB Storage

2018/2019

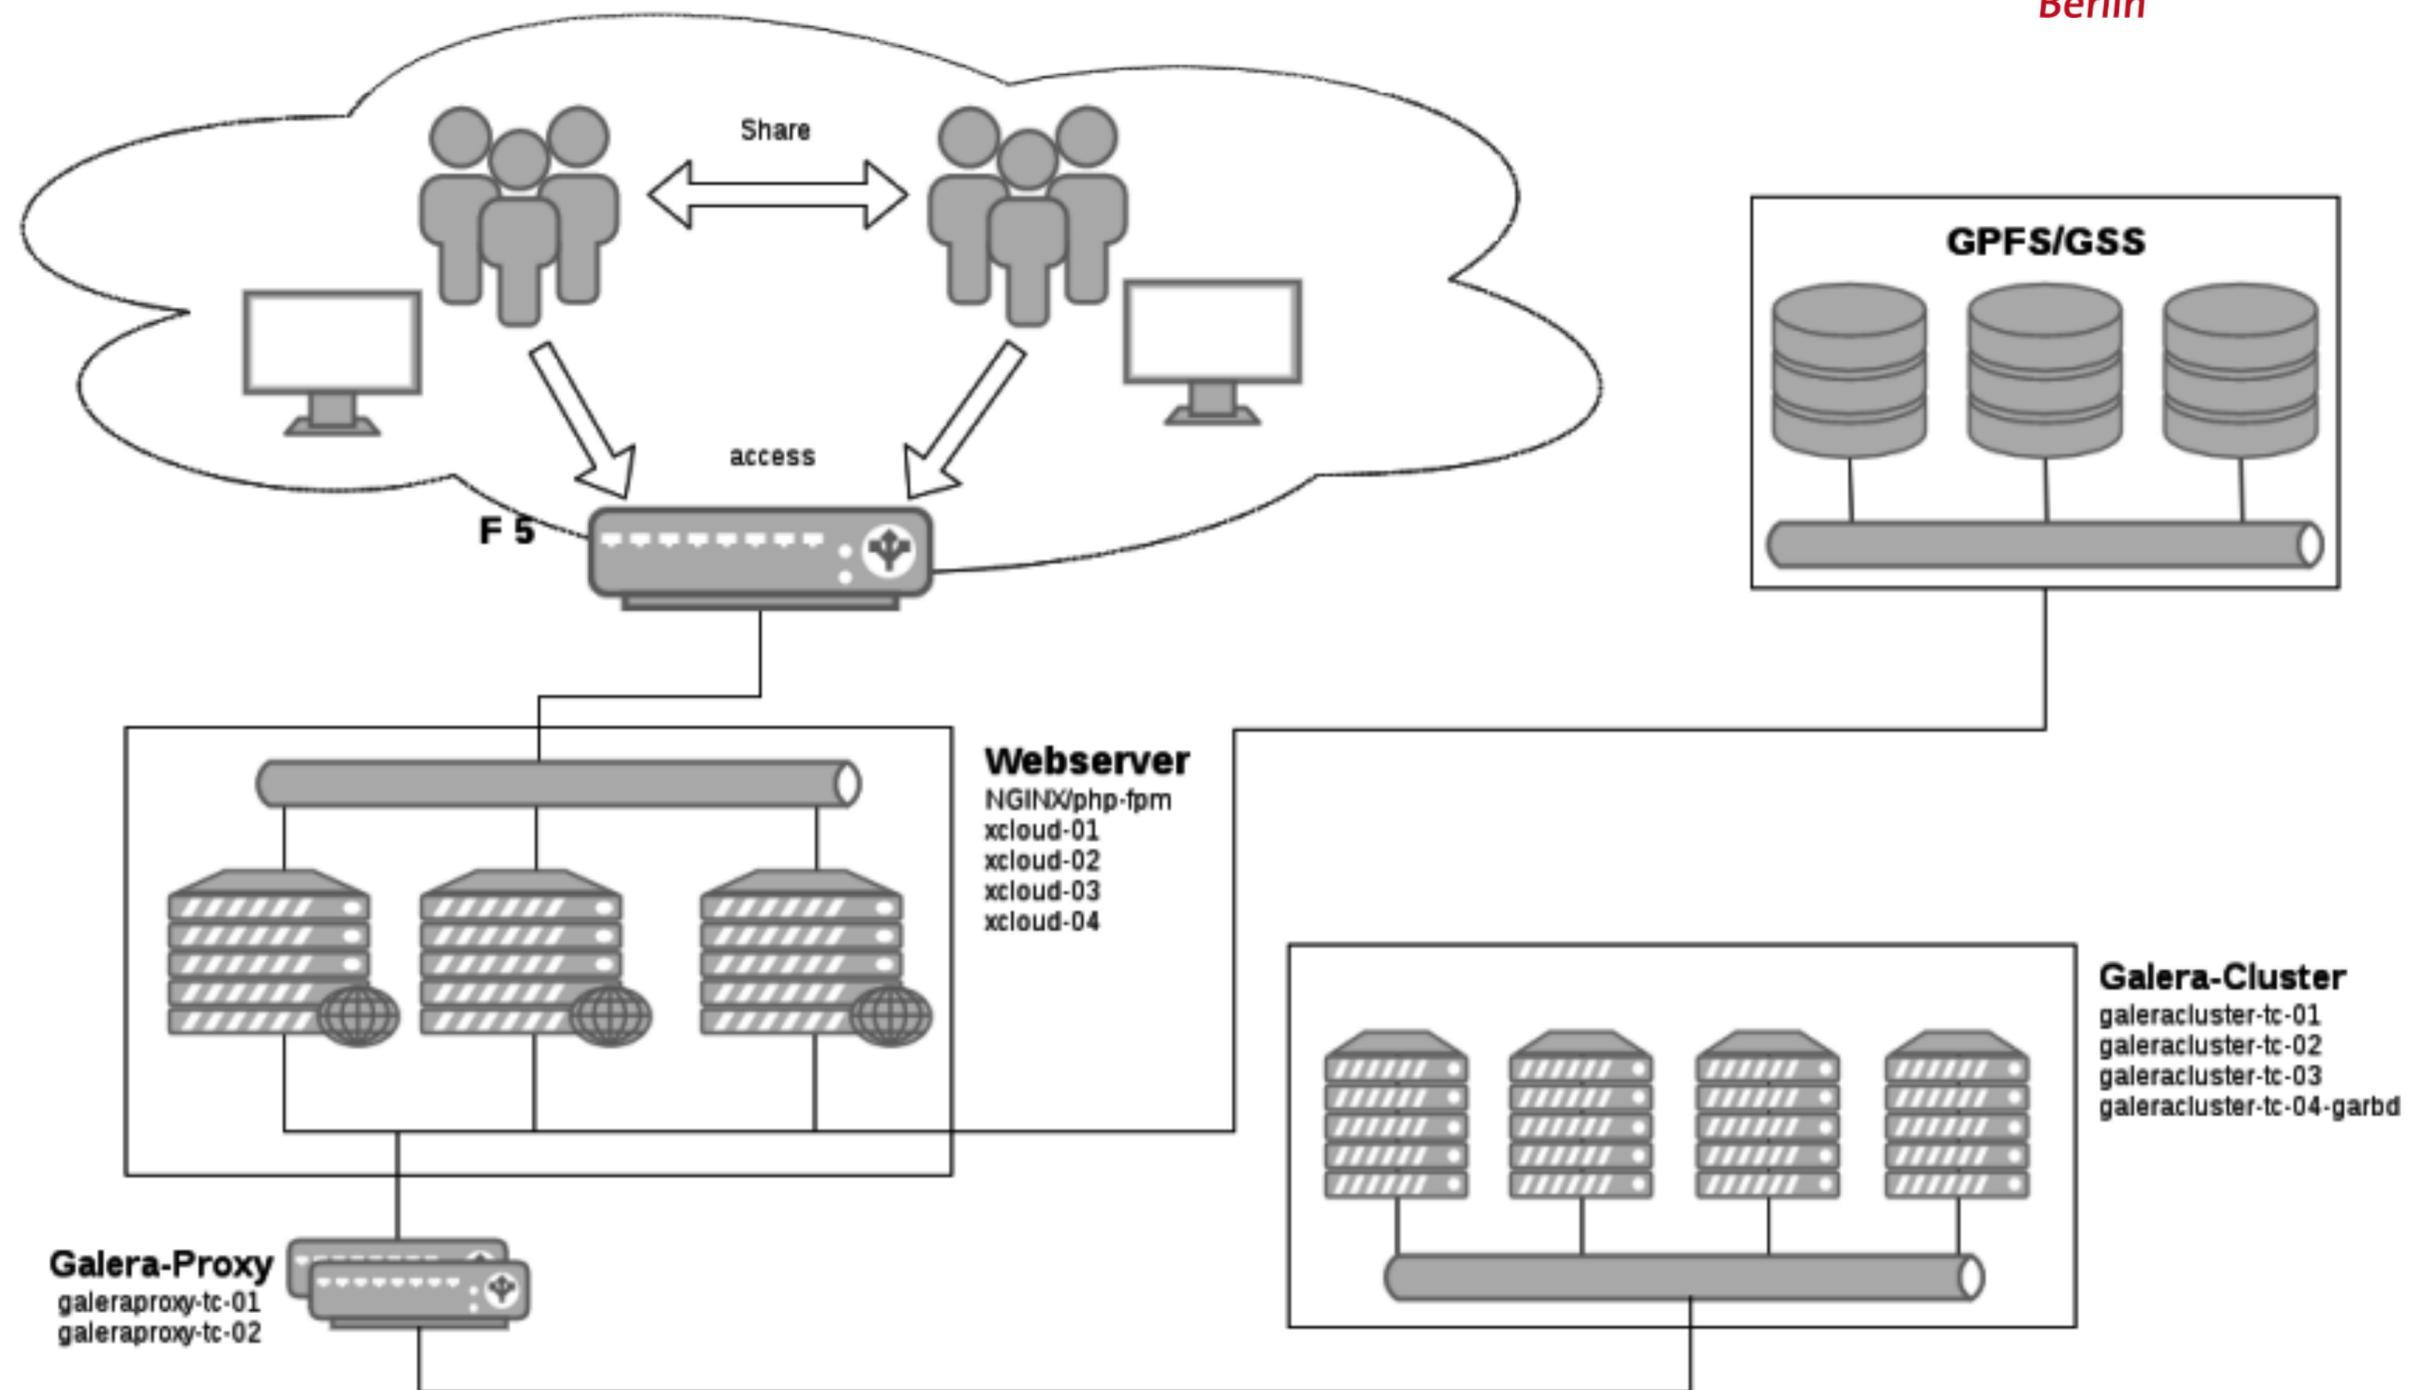
Collaboration Features, more partners and universities



# IV. Example 1: Technical University Berlin



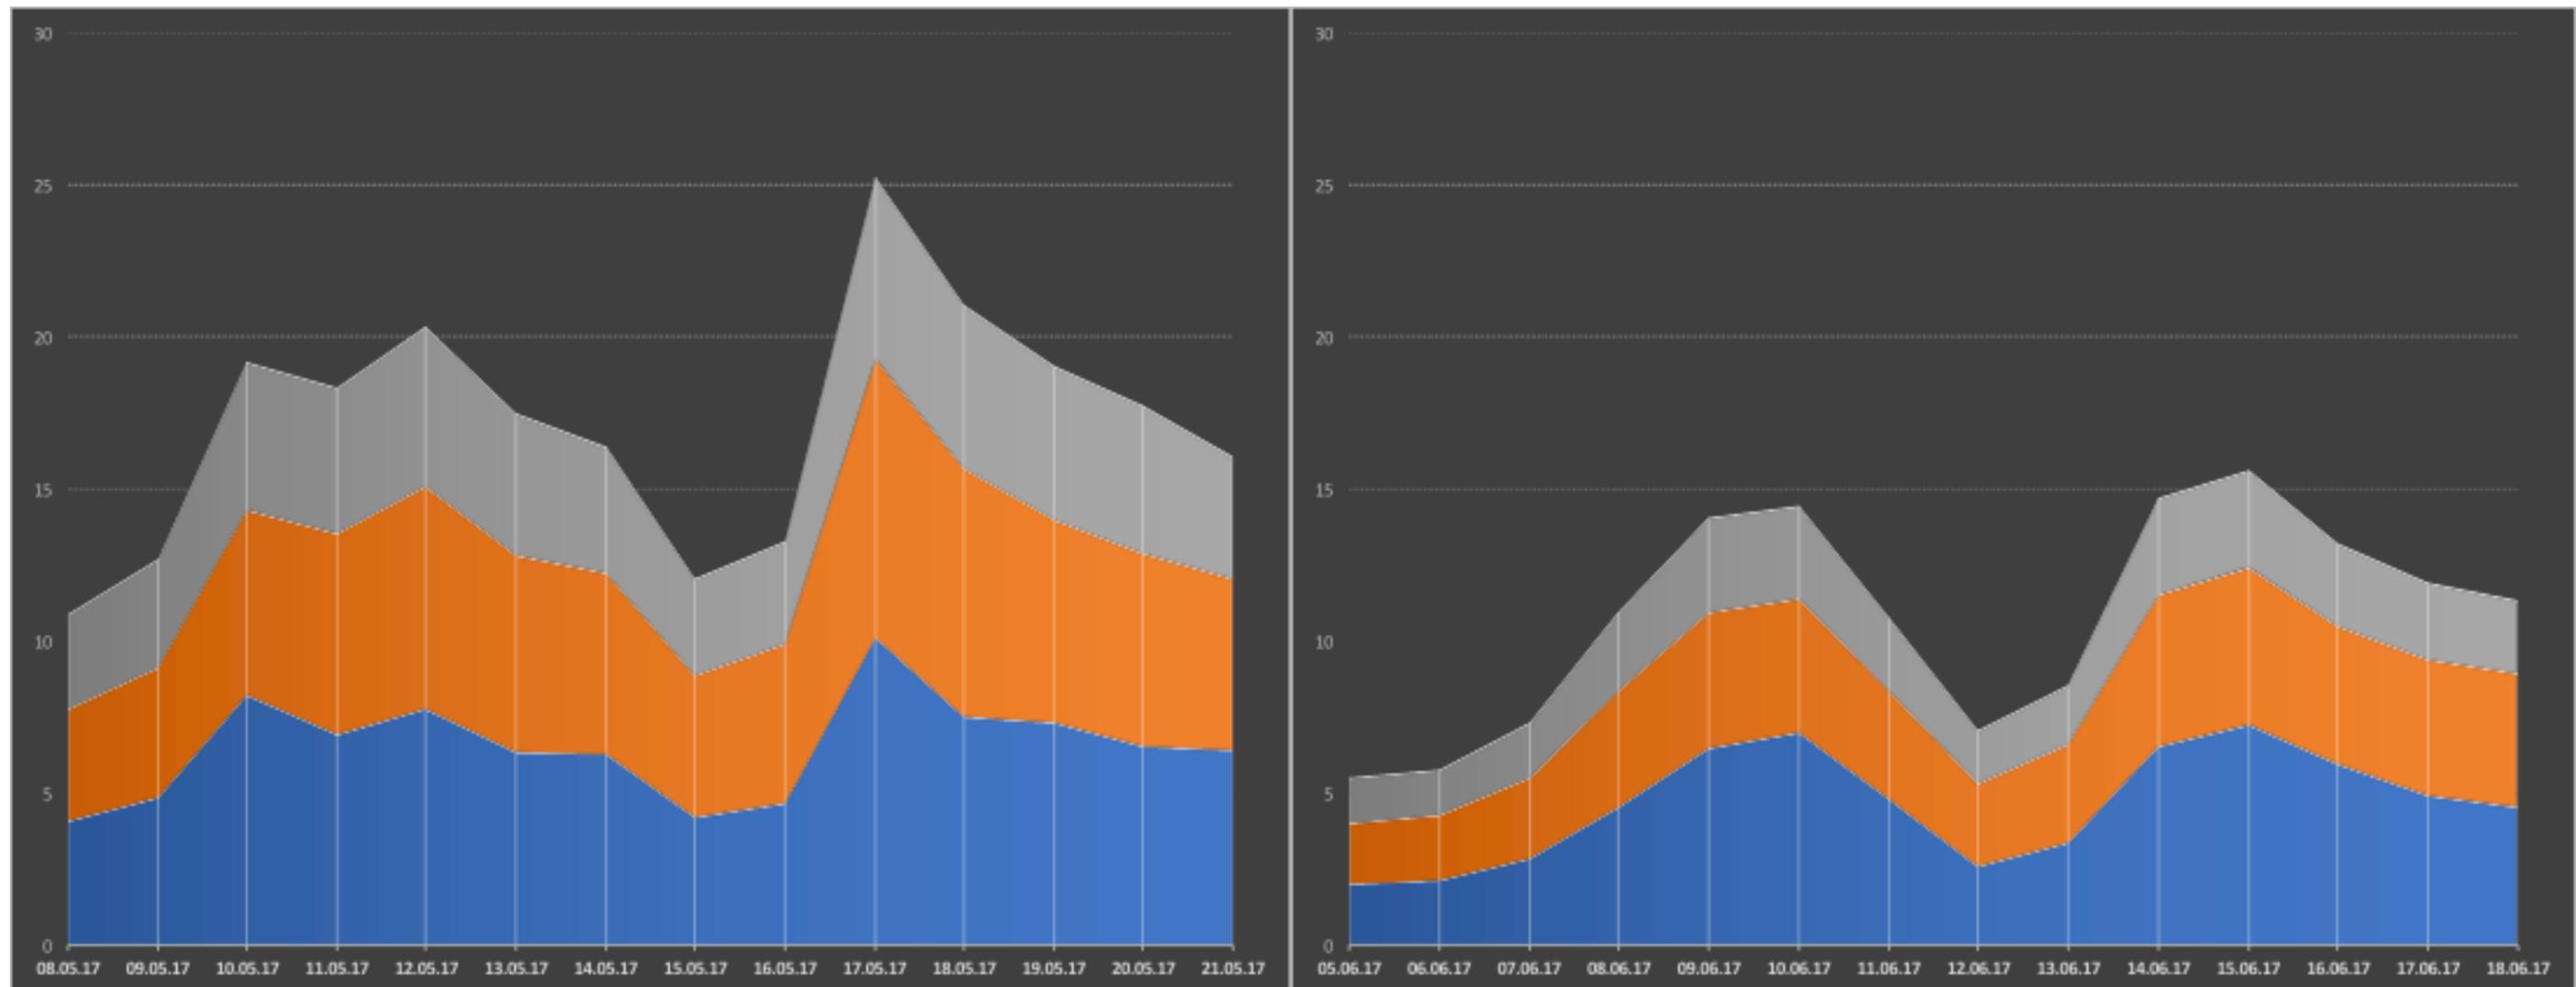
# IV. Example 1: Technical University Berlin



# IV. Example 1: Technical University Berlin



Real world test: DB cluster CPU load (2 weeks before/after migration)



ownCloud 9

Nextcloud 11



# IV. Example 1: Technical University Berlin



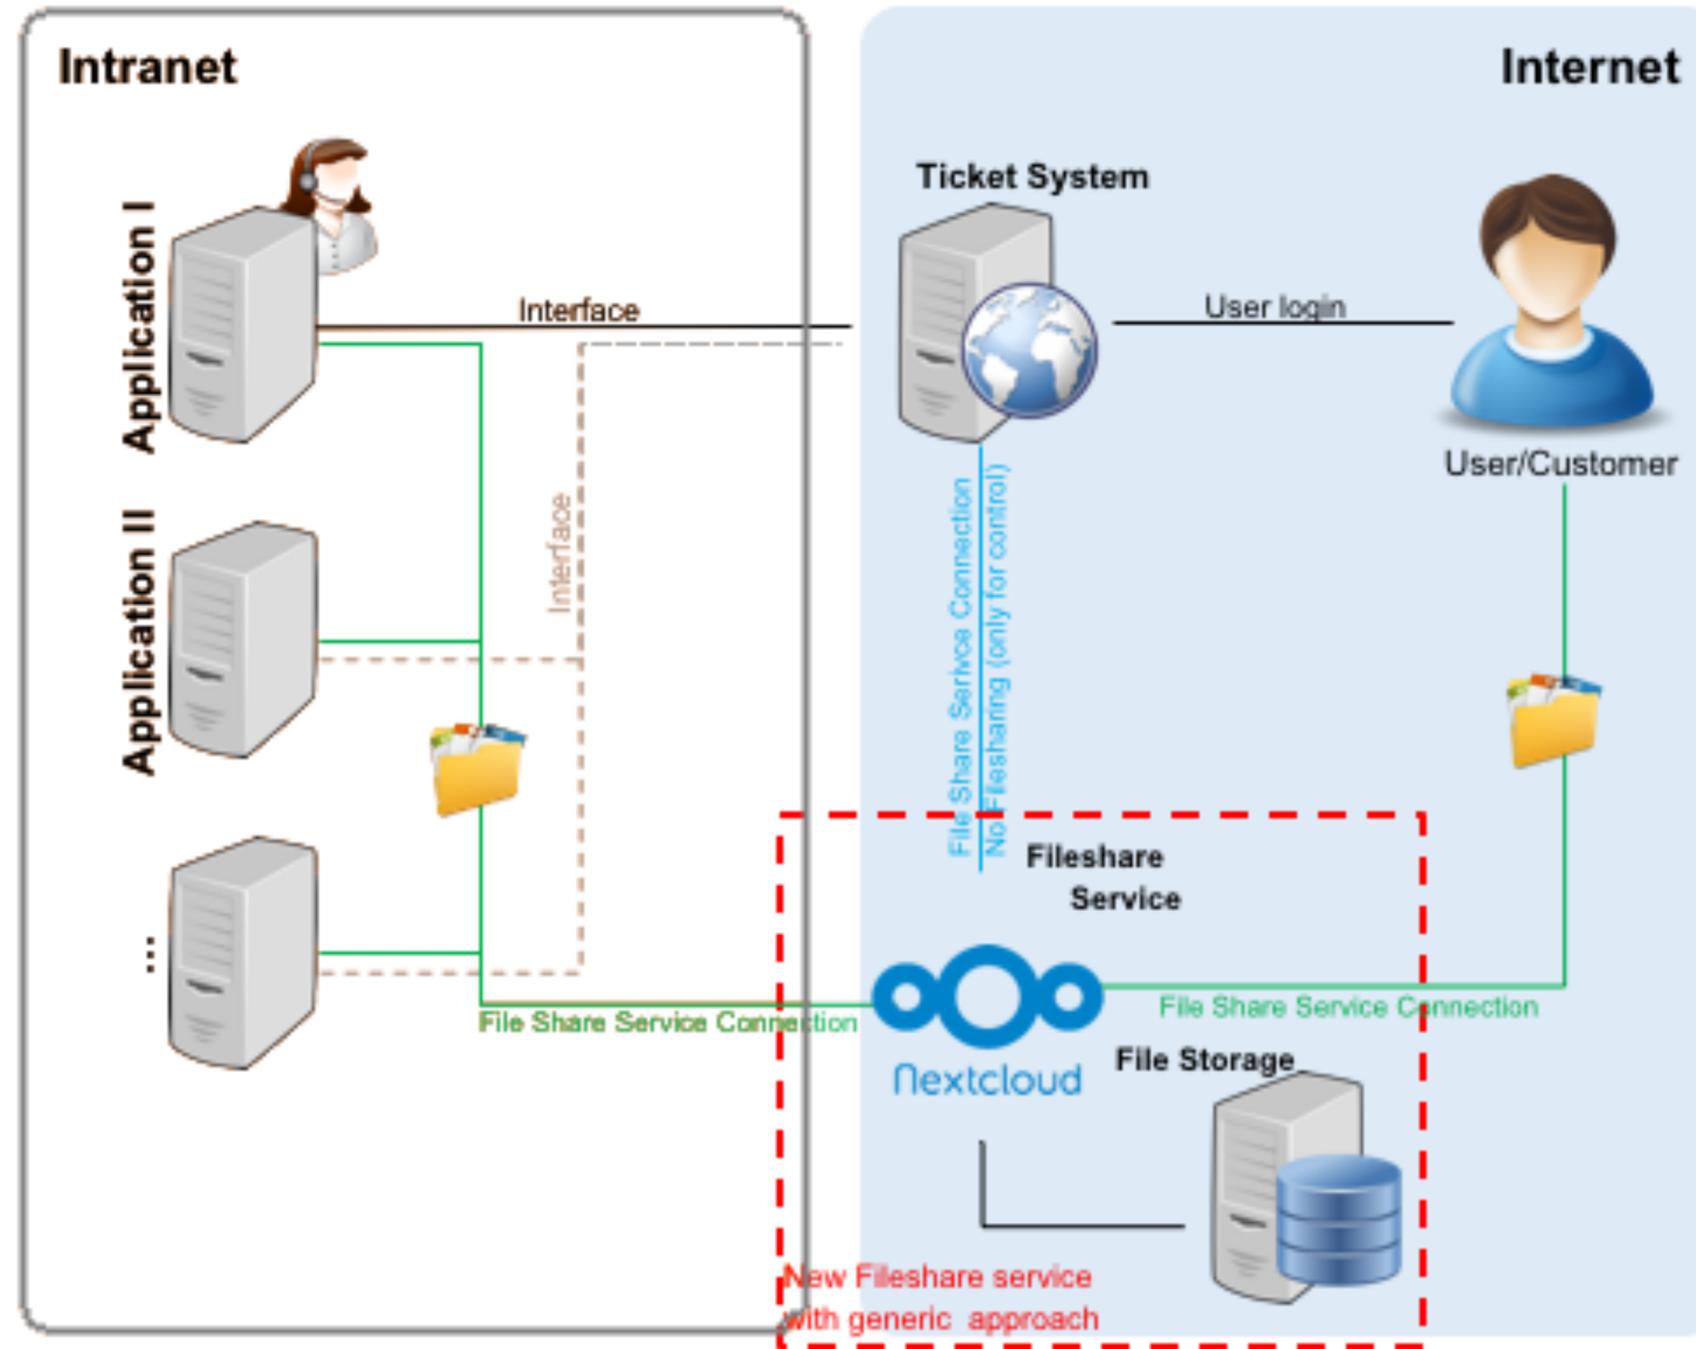
Performance improvements and features by Nextcloud:

- Response time improvements by 60%
- Reduced DB cluster CPU load by 40-50% in real world test.
- Multi bucket object store support
- Big LDAP performance improvements. Up to 85% faster.
- Significant performance improvements in external storage handling
- 60% faster propfinds
- 60% faster transfer of small files

## V. Example 2: **SIEMENS** Services Integration

1. Exchange of large files between customer and support
2. Customer provides files -> technician notified by ticket system
3. Technician can provide data to the customer as well

# V. Example 2: **SIEMENS** Services Integration



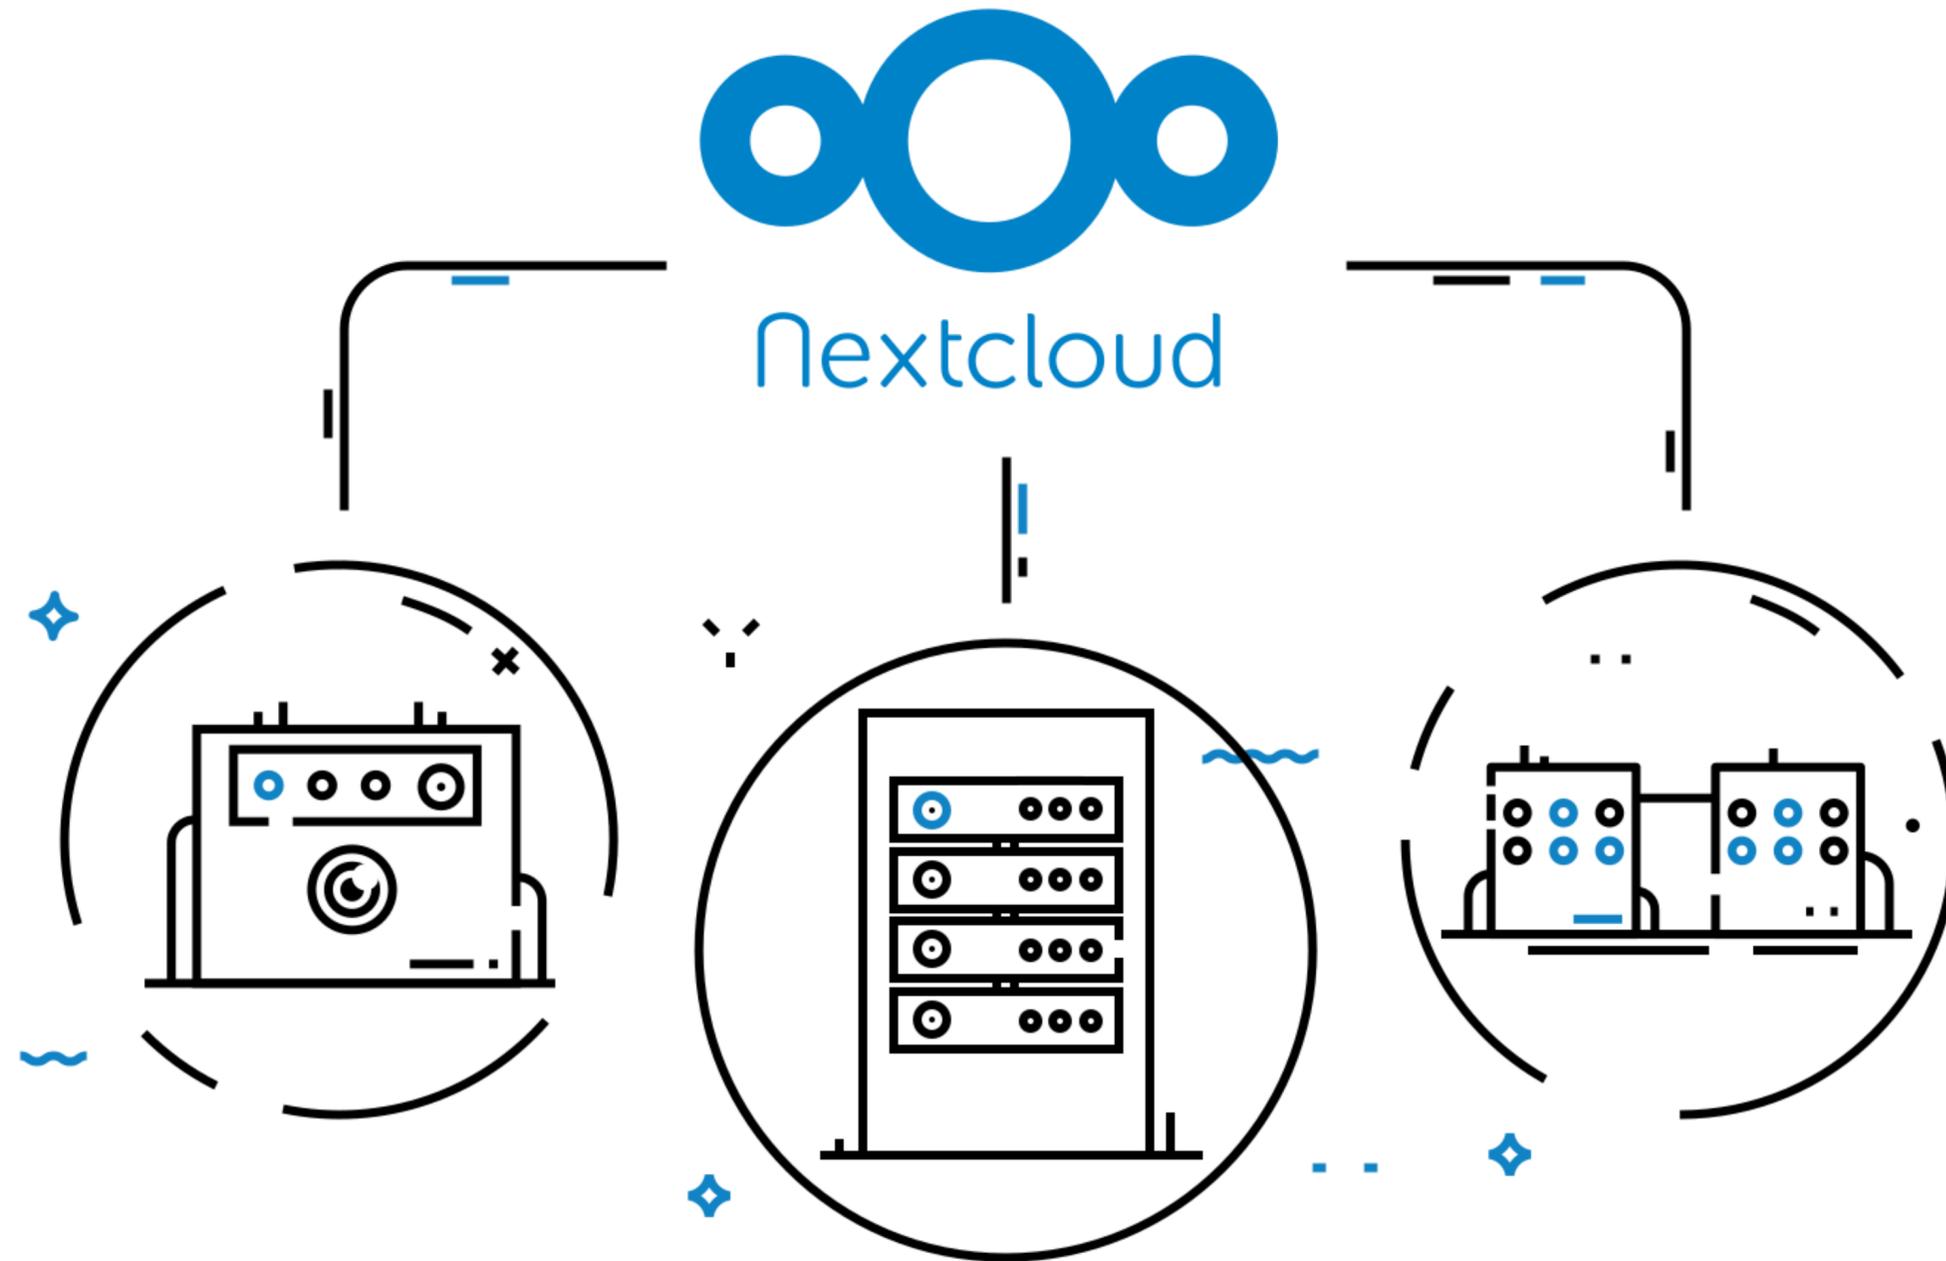
# VI. The idea behind Global Scale

## Pitfalls of the default deployment

- Components and WAN uplink will become bottlenecks
- Database particularly hard to scale beyond a 4 node Galera Cluster  
-> number of users and files are limited
- Scaling the storage becomes very expensive when dealing with PB of data



# VI. The idea behind Global Scale



# VI. The idea behind Global Scale

- Scalability > 500k users
- Global distribution
- Cost efficiency
  
- Independent nodes
- Federated sharing between nodes
- Lookup Server
- Global Site Selector
- Balancer
  
- First release in Nextcloud 12
- Refined user distribution logic in 13
- In production since 6 month for 10M users



# VI. The idea behind Global Scale

## Solution

- Many independent application servers
- No central DB, storage or caching instances
- Every Nextcloud node can use the default deployment



# VI. The idea behind Global Scale

## Solution

- A Node can be hosted in different hosting centers
- No fast interconnect between the sites is necessary
- User is local to a node, user data exists only in this local node
- Sharing via federation



# VII. Example 3: Large-Scale Service provider

## Request:

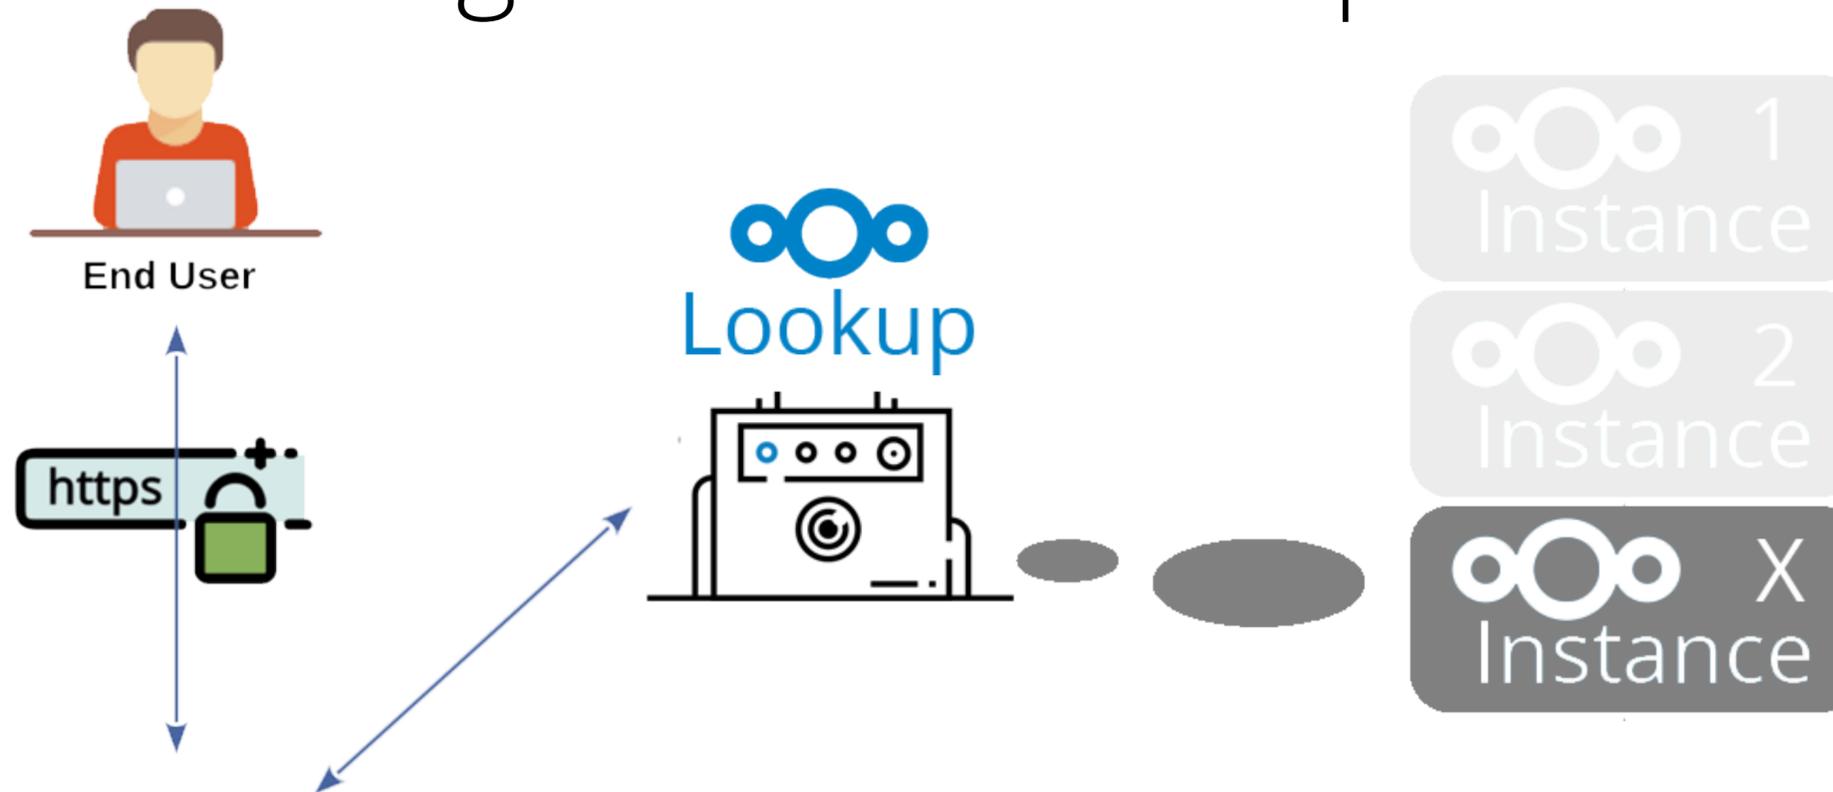
- Cloud storage for millions of users
- Multi-national distribution

## Requirements:

- Cost efficiency
- Independent nodes
- Single Interface



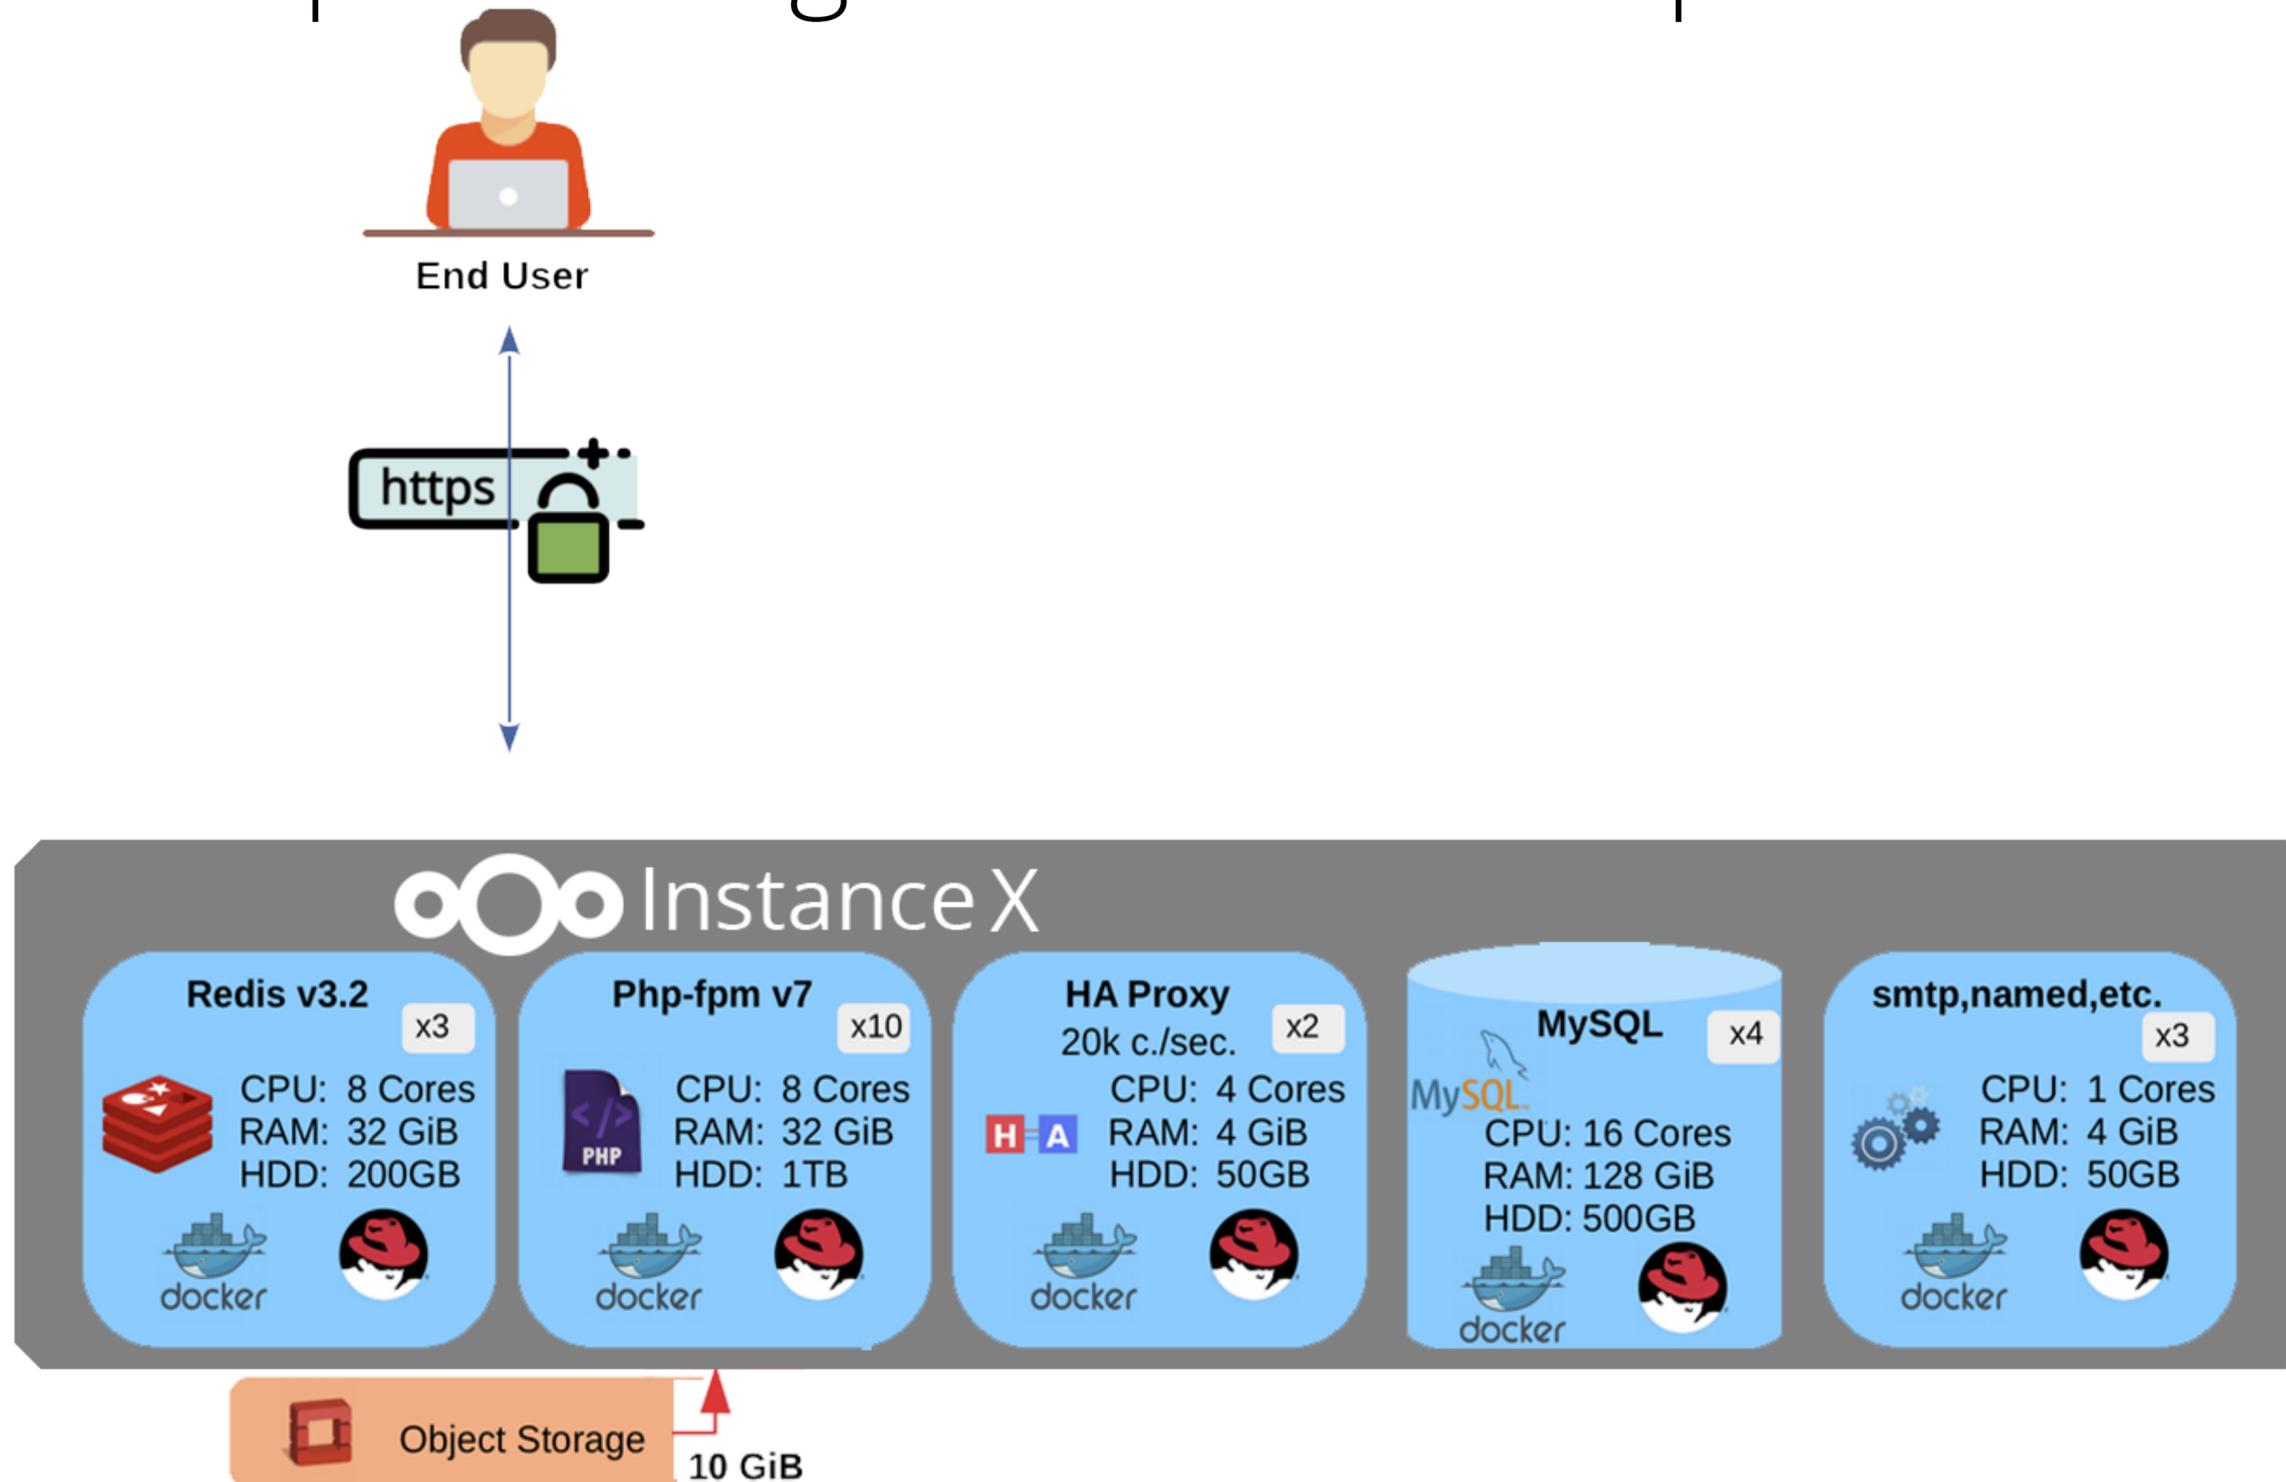
# VII. Example 3: Large-Scale Service provider



## Portal

Service	Count	CPU	RAM	HDD	Container
Redis v3.2	x3	8 Cores	32 GiB	200GB	docker
Php-fpm v7	x2	8 Cores	32 GiB	100GB	docker
HA Proxy	x2	4 Cores	4 GiB	50GB	docker
MySQL	x2	16 Cores	64 GiB	500GB	docker
smtp,named,etc.	x3	1 Cores	4 GiB	50GB	docker

# VII. Example 3: Large-Scale Service provider



Nextcloud supports you!

Scale with us

Contact us at

[sales@nextcloud.com](mailto:sales@nextcloud.com)



Your Questions



Thank you!

Nextcloud supports you!

Scale with us

Contact us at

[sales@nextcloud.com](mailto:sales@nextcloud.com)