Global Scale

The next level of scalability and federation

About Me

bjoern@nextcloud.com

schiessle

@schiessle

Björn Schießle

Senior Software Engineer at Nextcloud

5+ years experience with Nextcloud (and ownCloud) technology

Developed the first version and architecture of Federated Cloud Sharing

Agenda

1. Motivation & general idea
2. The next generation of Federated Sharing
3. The Global Scale components
4. Communication between the components
5. Re-balance Global Scale nodes
6. Summary

Traditional way to scale Nextcloud

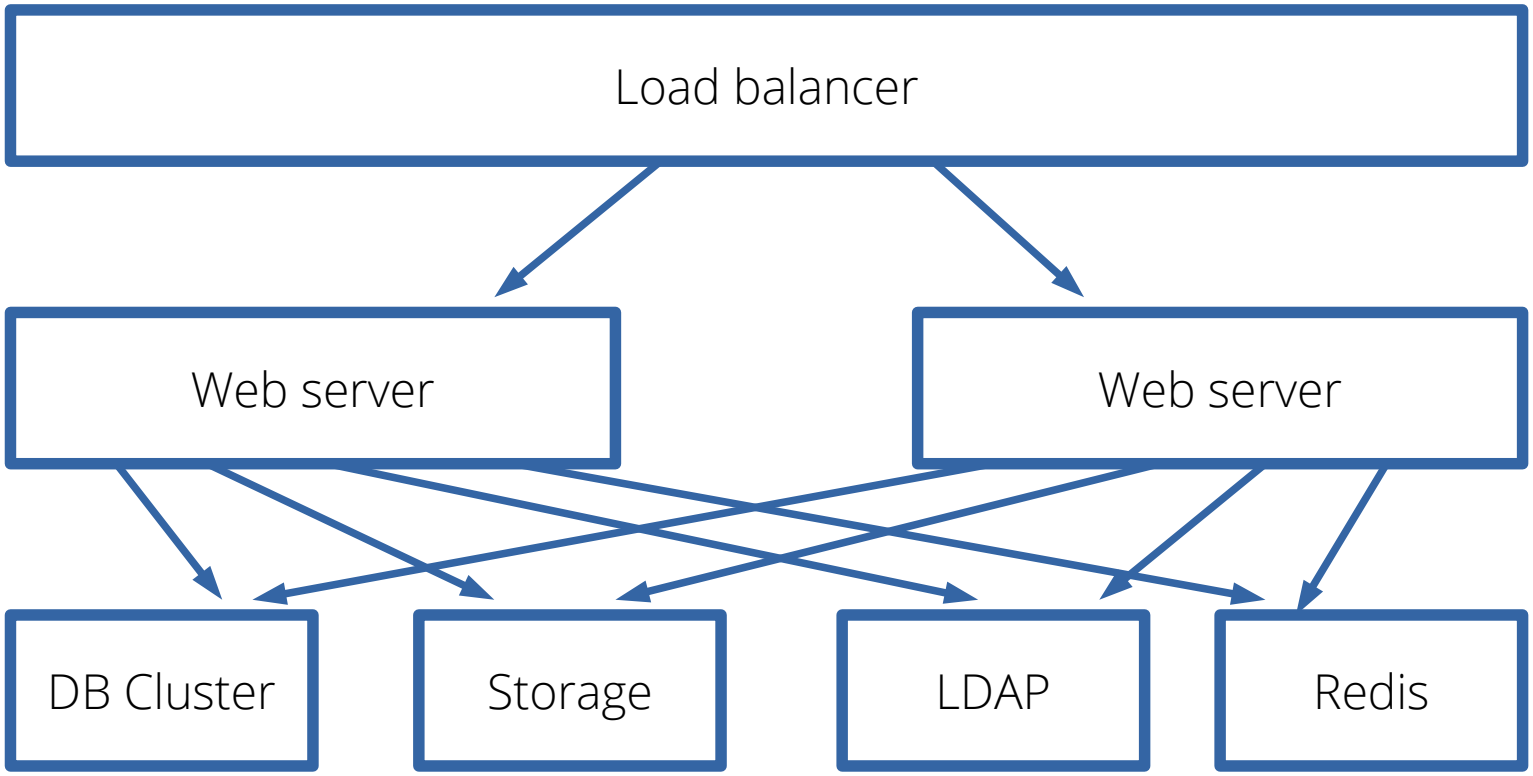

Limitations

- Database cluster can’t be scaled infinitely (typically 4 nodes is the maximum)
- At some point in time the database cluster will become a bottle neck
- Storage for a huge instances can become really expensive
- Maintaining huge instances can be expensive

What do we want to achieve?

- Scale to millions of users
- Improve performance by storing data next to the users
- Keep data at specific location to meet external and internal (legal) requirements
- Makes it easy to extend the network
- Improve cost efficiency by reducing hardware and maintenance costs

Globe Scale

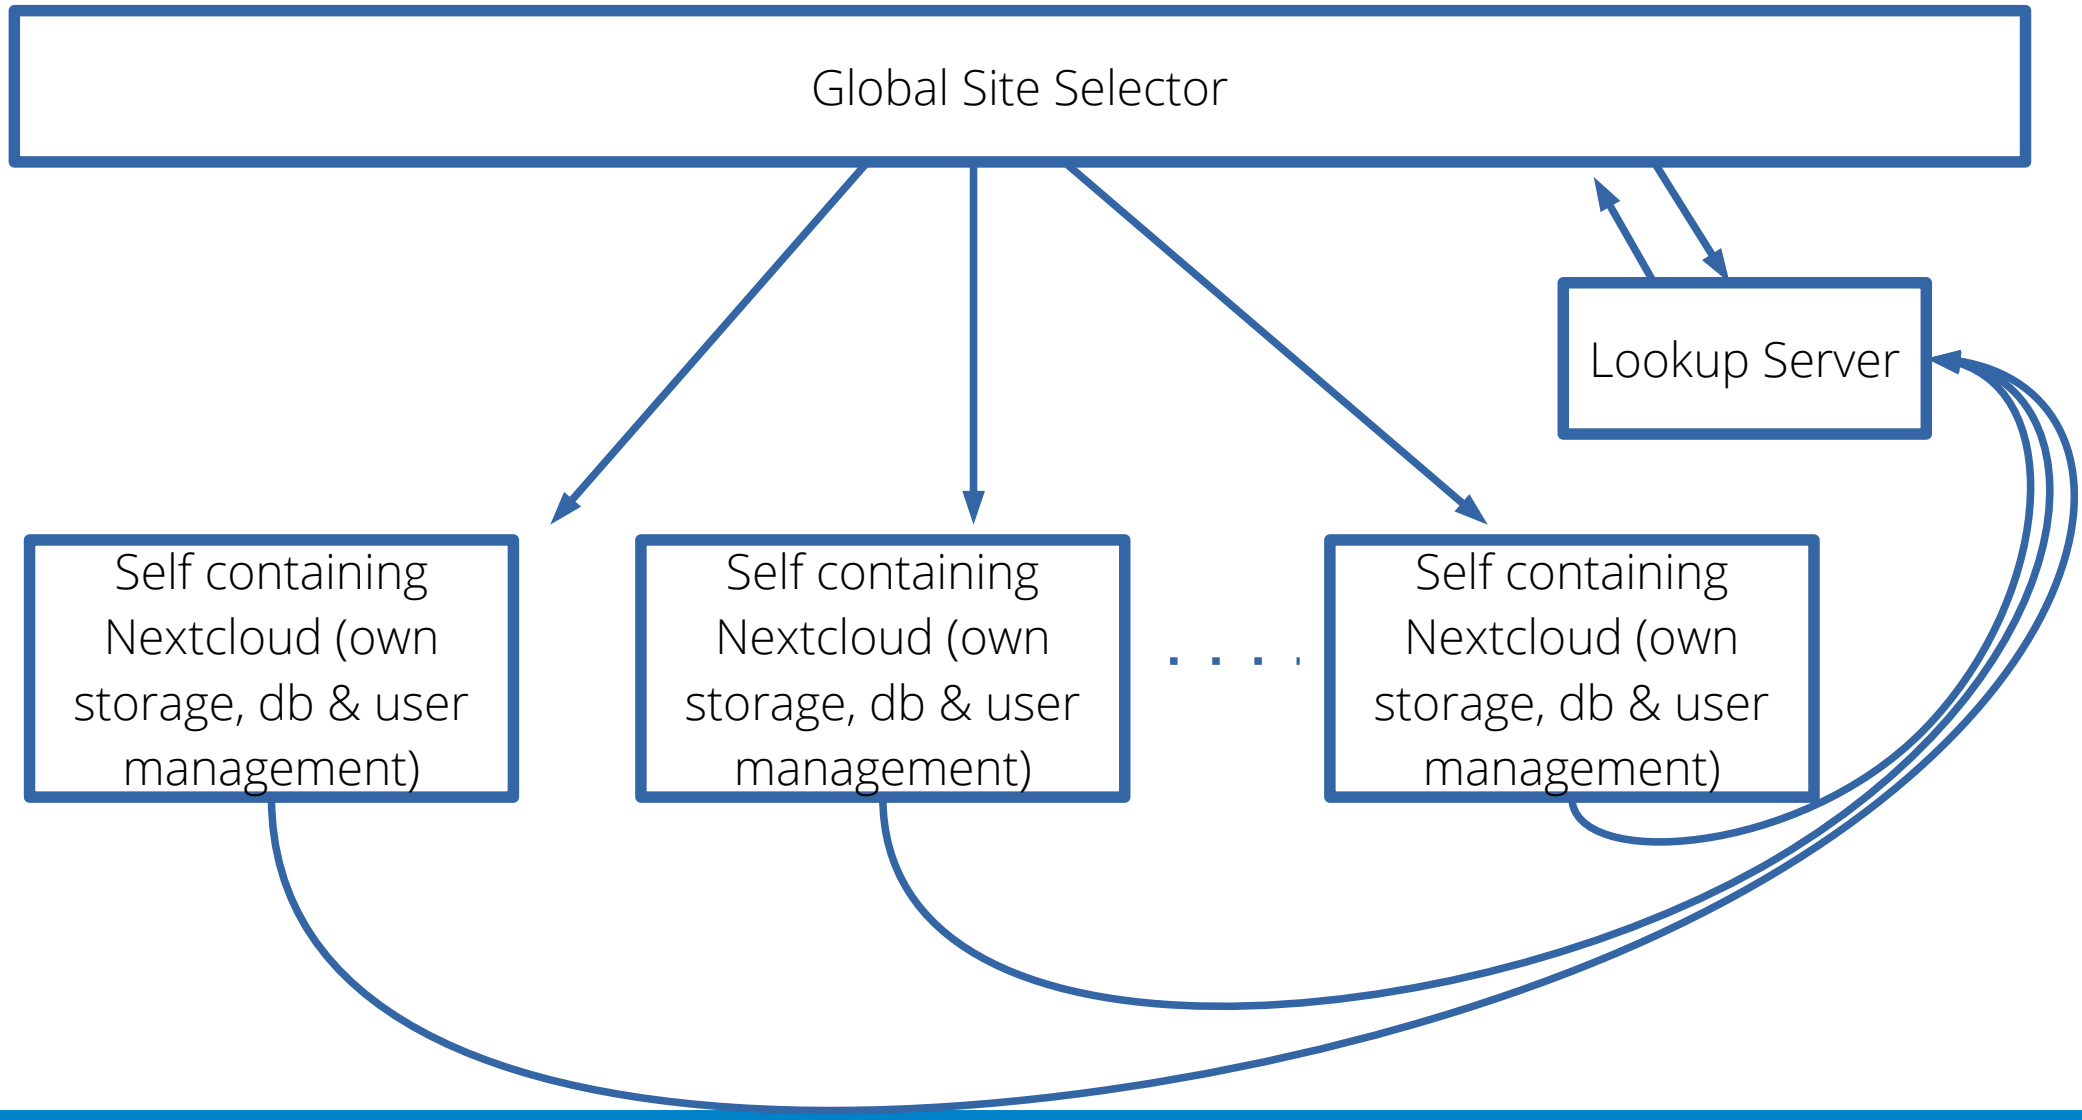
Global Site Selector
Lookup Server
Self containing Nextcloud (own storage, db & user management)
Self containing Nextcloud (own storage, db & user management)
Self containing Nextcloud (own storage, db & user management)

Federated Sharing

Federated Sharing

- User has a unique cloud ID user@server.tld
- The federated cloud ID can be used the same way as local user IDs
- Share gets send to the remote server
- User can accept/decline it
- If the user accept it we create a remote webdav mount
- Turn public links into full-featured federated shares

Federated Sharing 2.0

- Make it more generic/modular
- Federated group sharing
- Based on the Open Cloud Mesh Specification

Federated Sharing 2.0

Modularity

Federated Sharing 2.0 - Modularity

- Add new share types, such as...
  - Calendar
  - Contacts
  - Talk
- Enable other apps to add their own share types

Federated Sharing 2.0 - Modularity

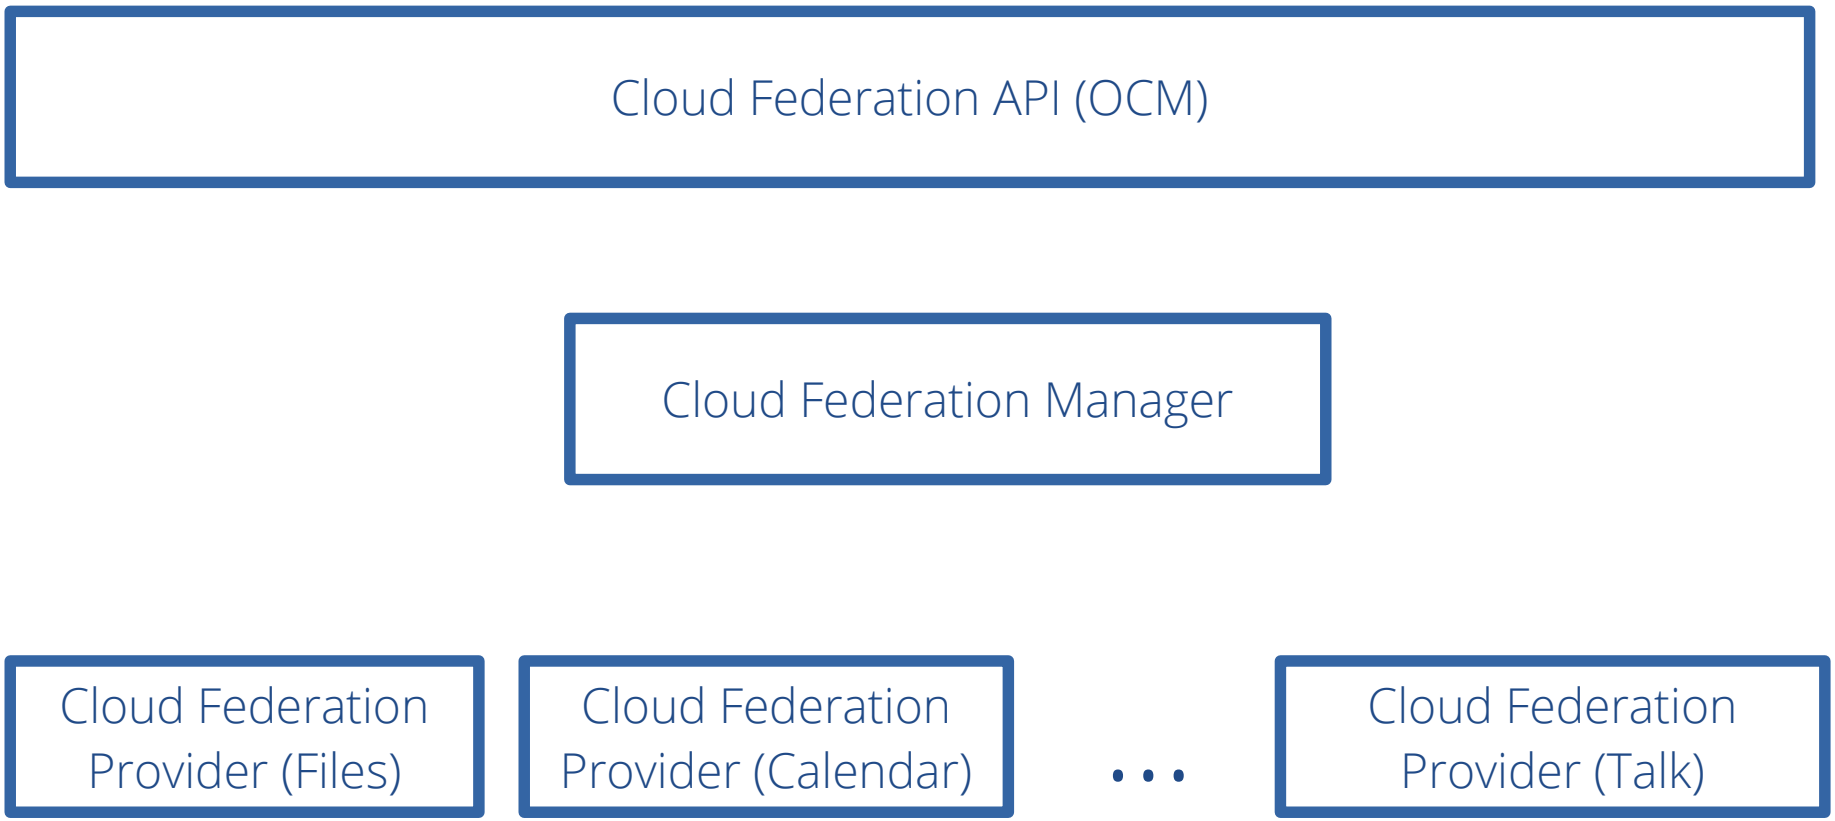

Federated Sharing 2.0 - Modularity

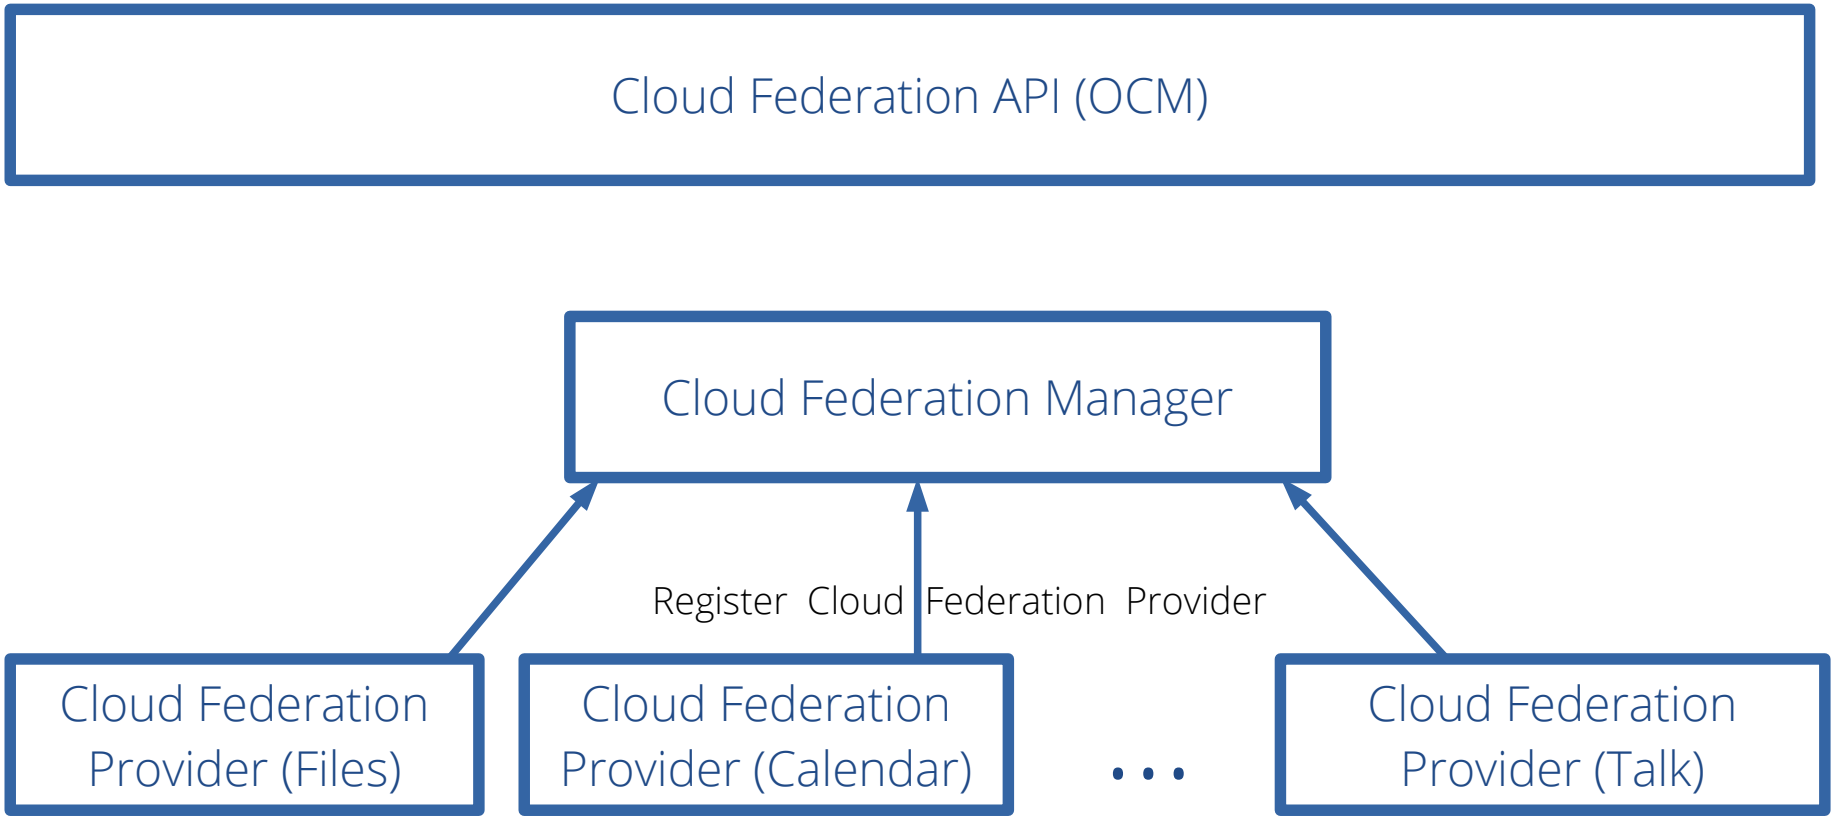

Federated Sharing 2.0 - Modularity

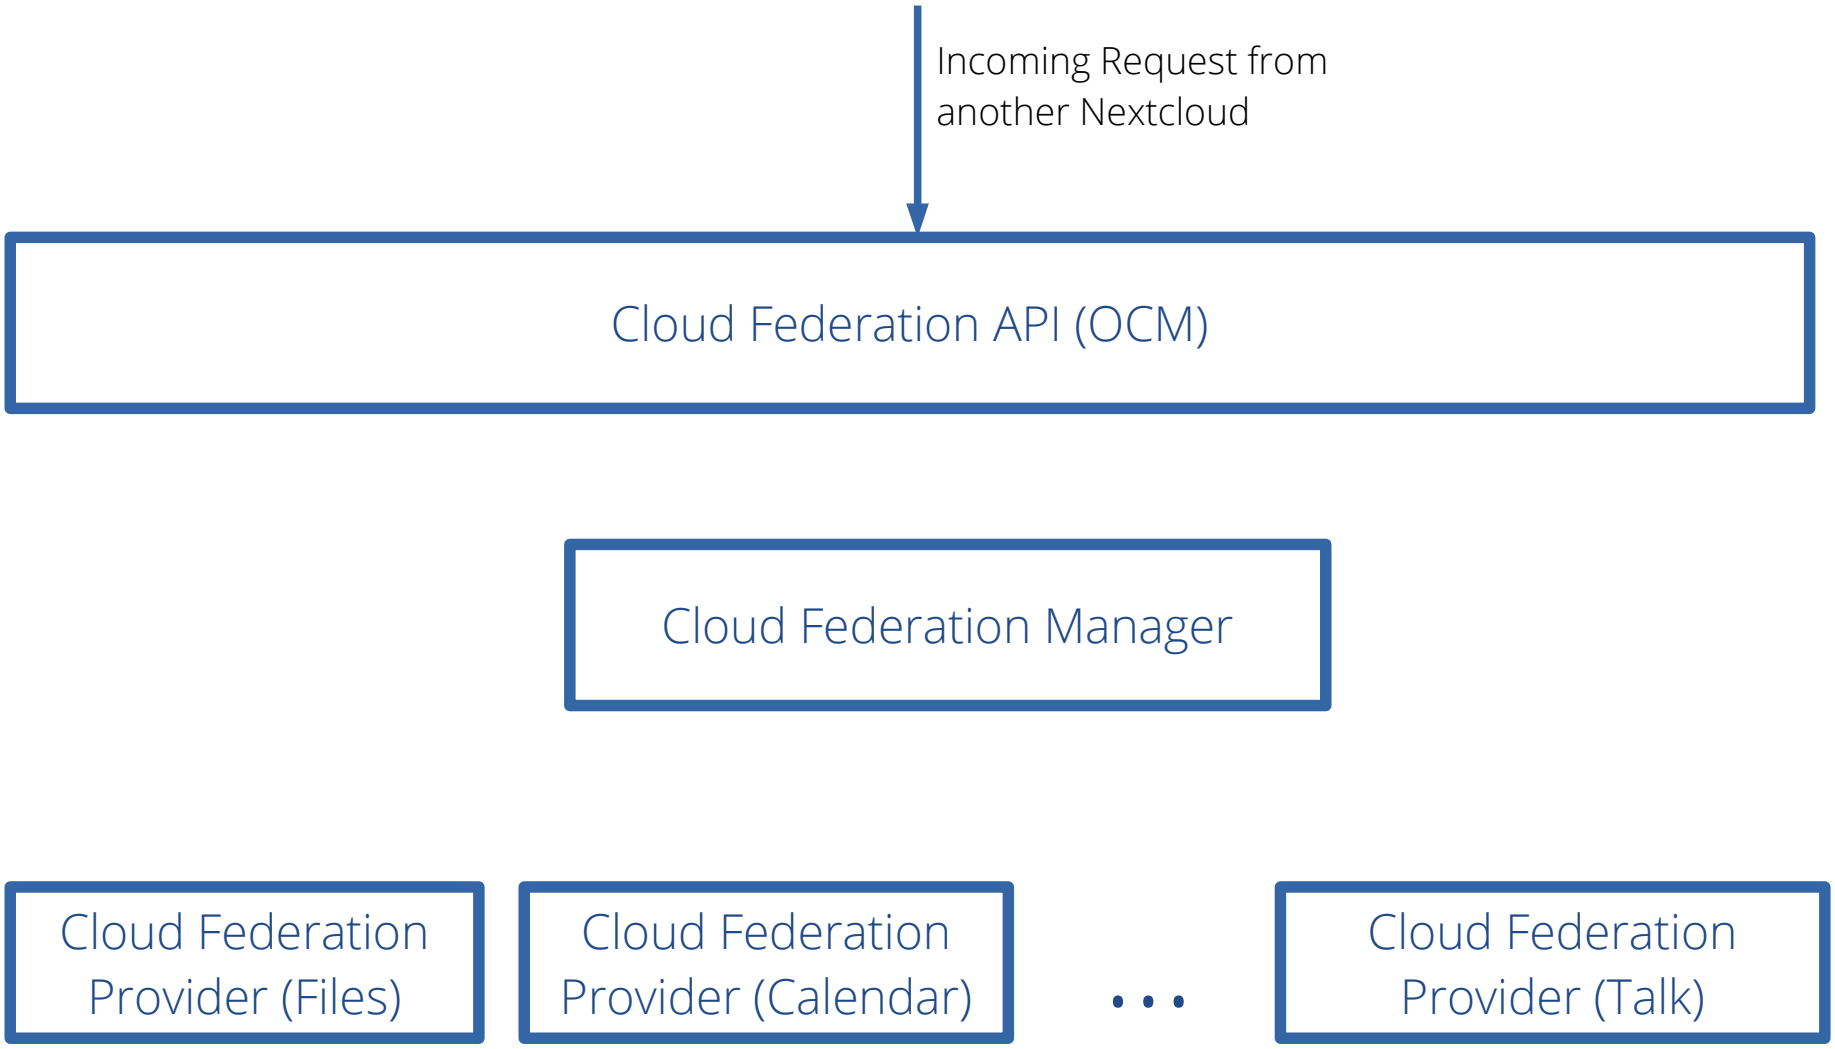

Federated Sharing 2.0 - Modularity

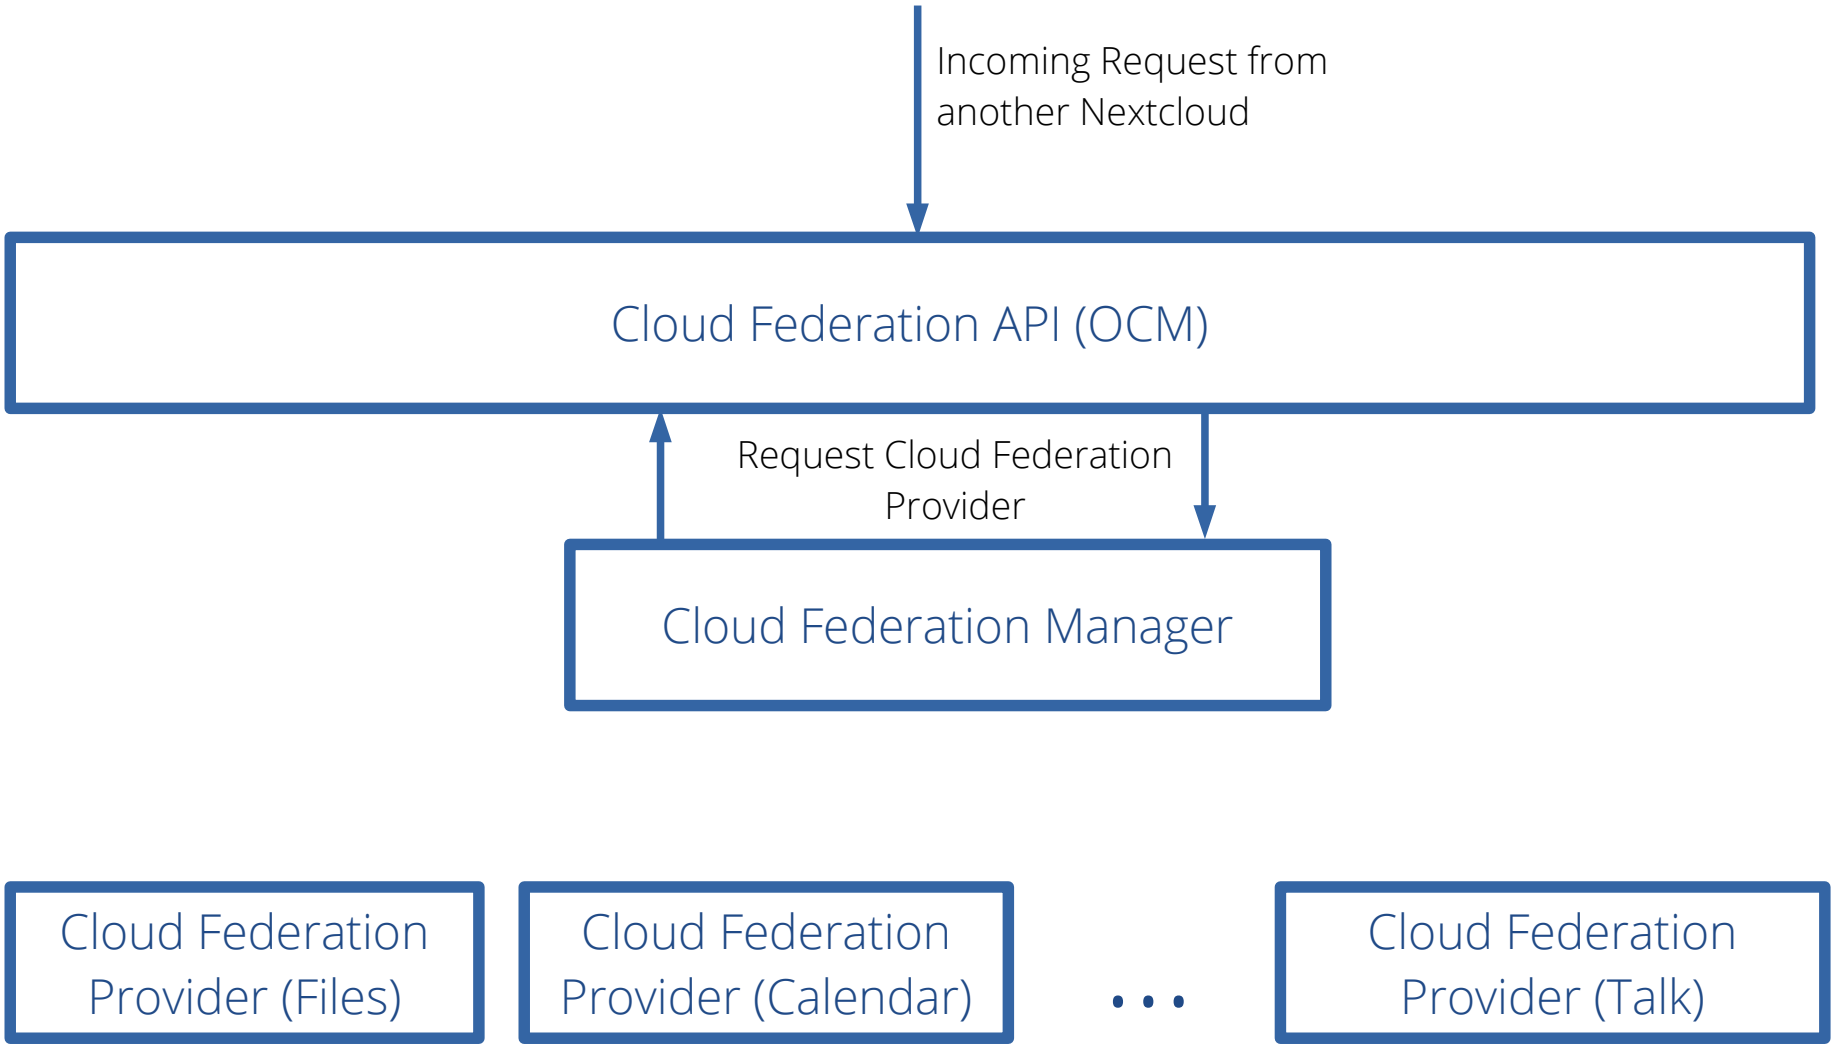

Federated Sharing 2.0 - Modularity

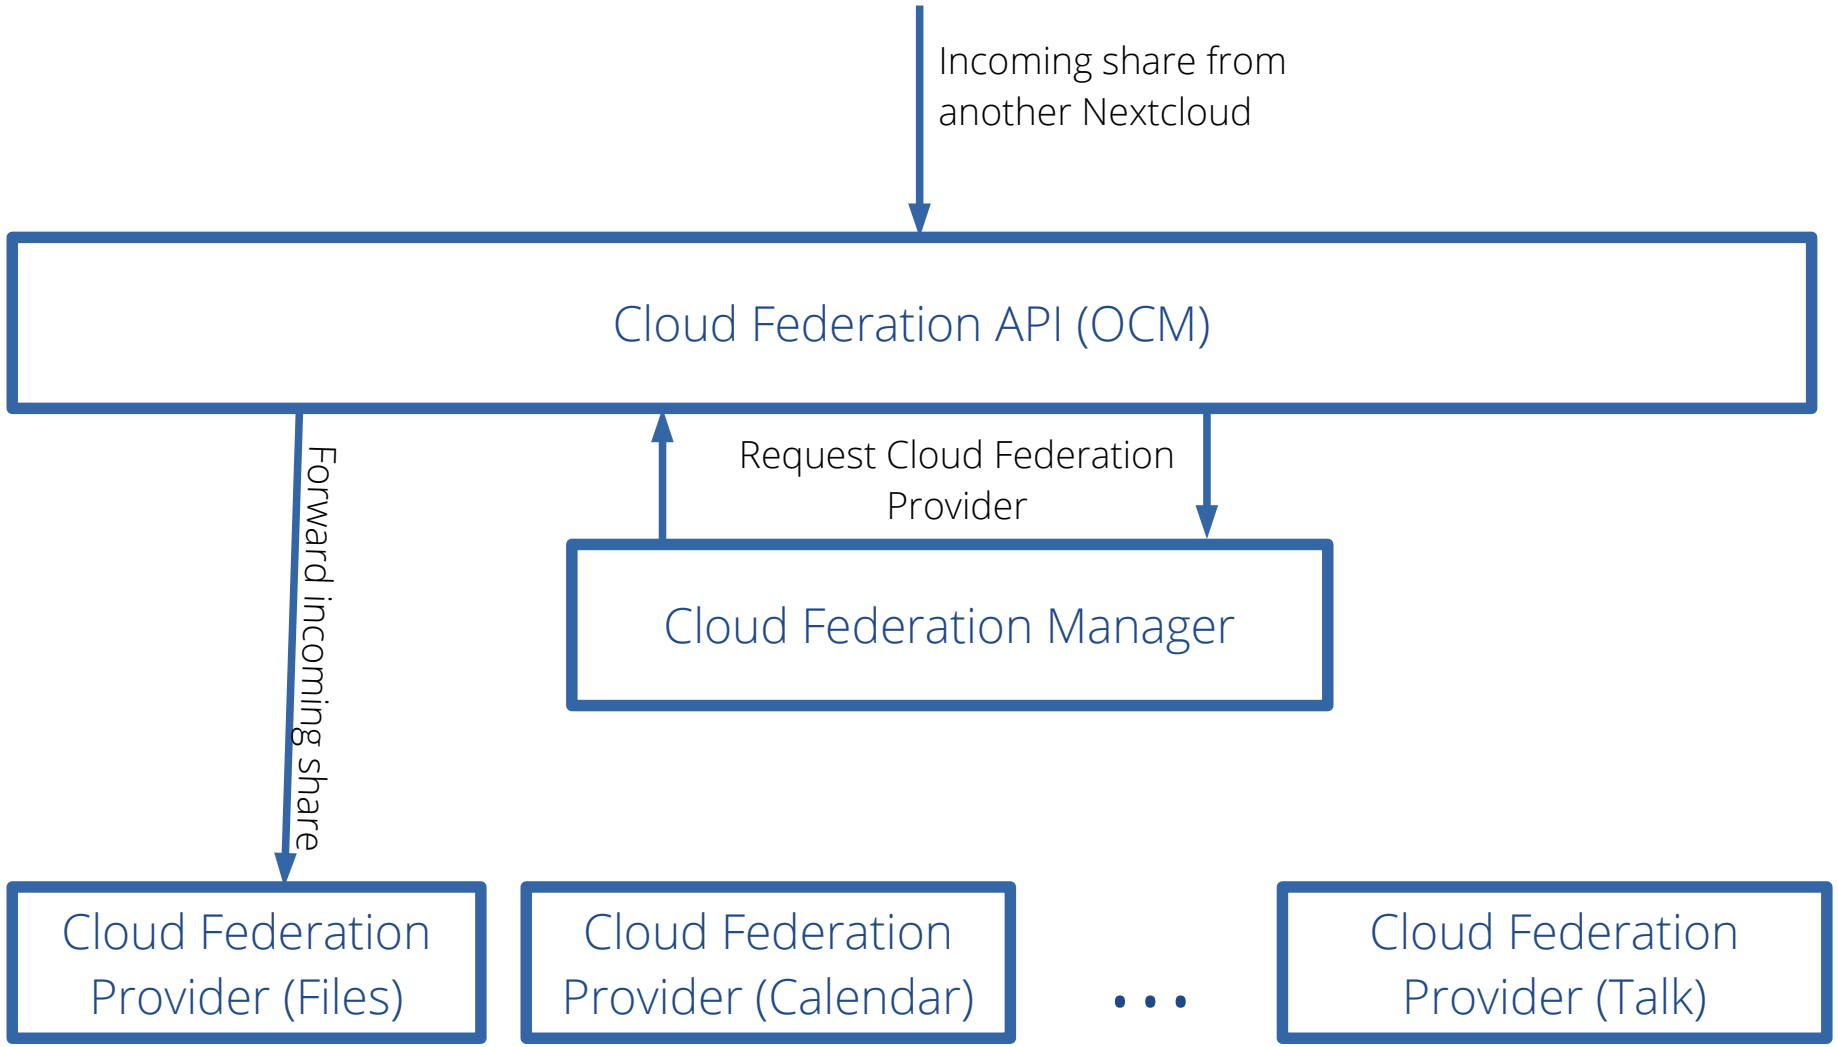

Federated Sharing 2.0

Groups

Federated Sharing 2.0 - Groups

Standard Setup

- Distinguish different share types: ‘group’ vs ‘user’
- Access groups the same way we access users: group@server.tld

Federated Sharing 2.0 - Groups

Global Scale Setup

- Assumption: All groups exist on all servers
- Broadcast each group share to all servers in the global scale setup

Global Scale

Use Cases

- Research networks with multiple institutes
- Large international companies with locations all over the world
  - Keep data next to the people
  - Meet local data-protection or export requirements
  - Meet internal requirements
- Companies with a large user base (literally millions of users)

General Architecture

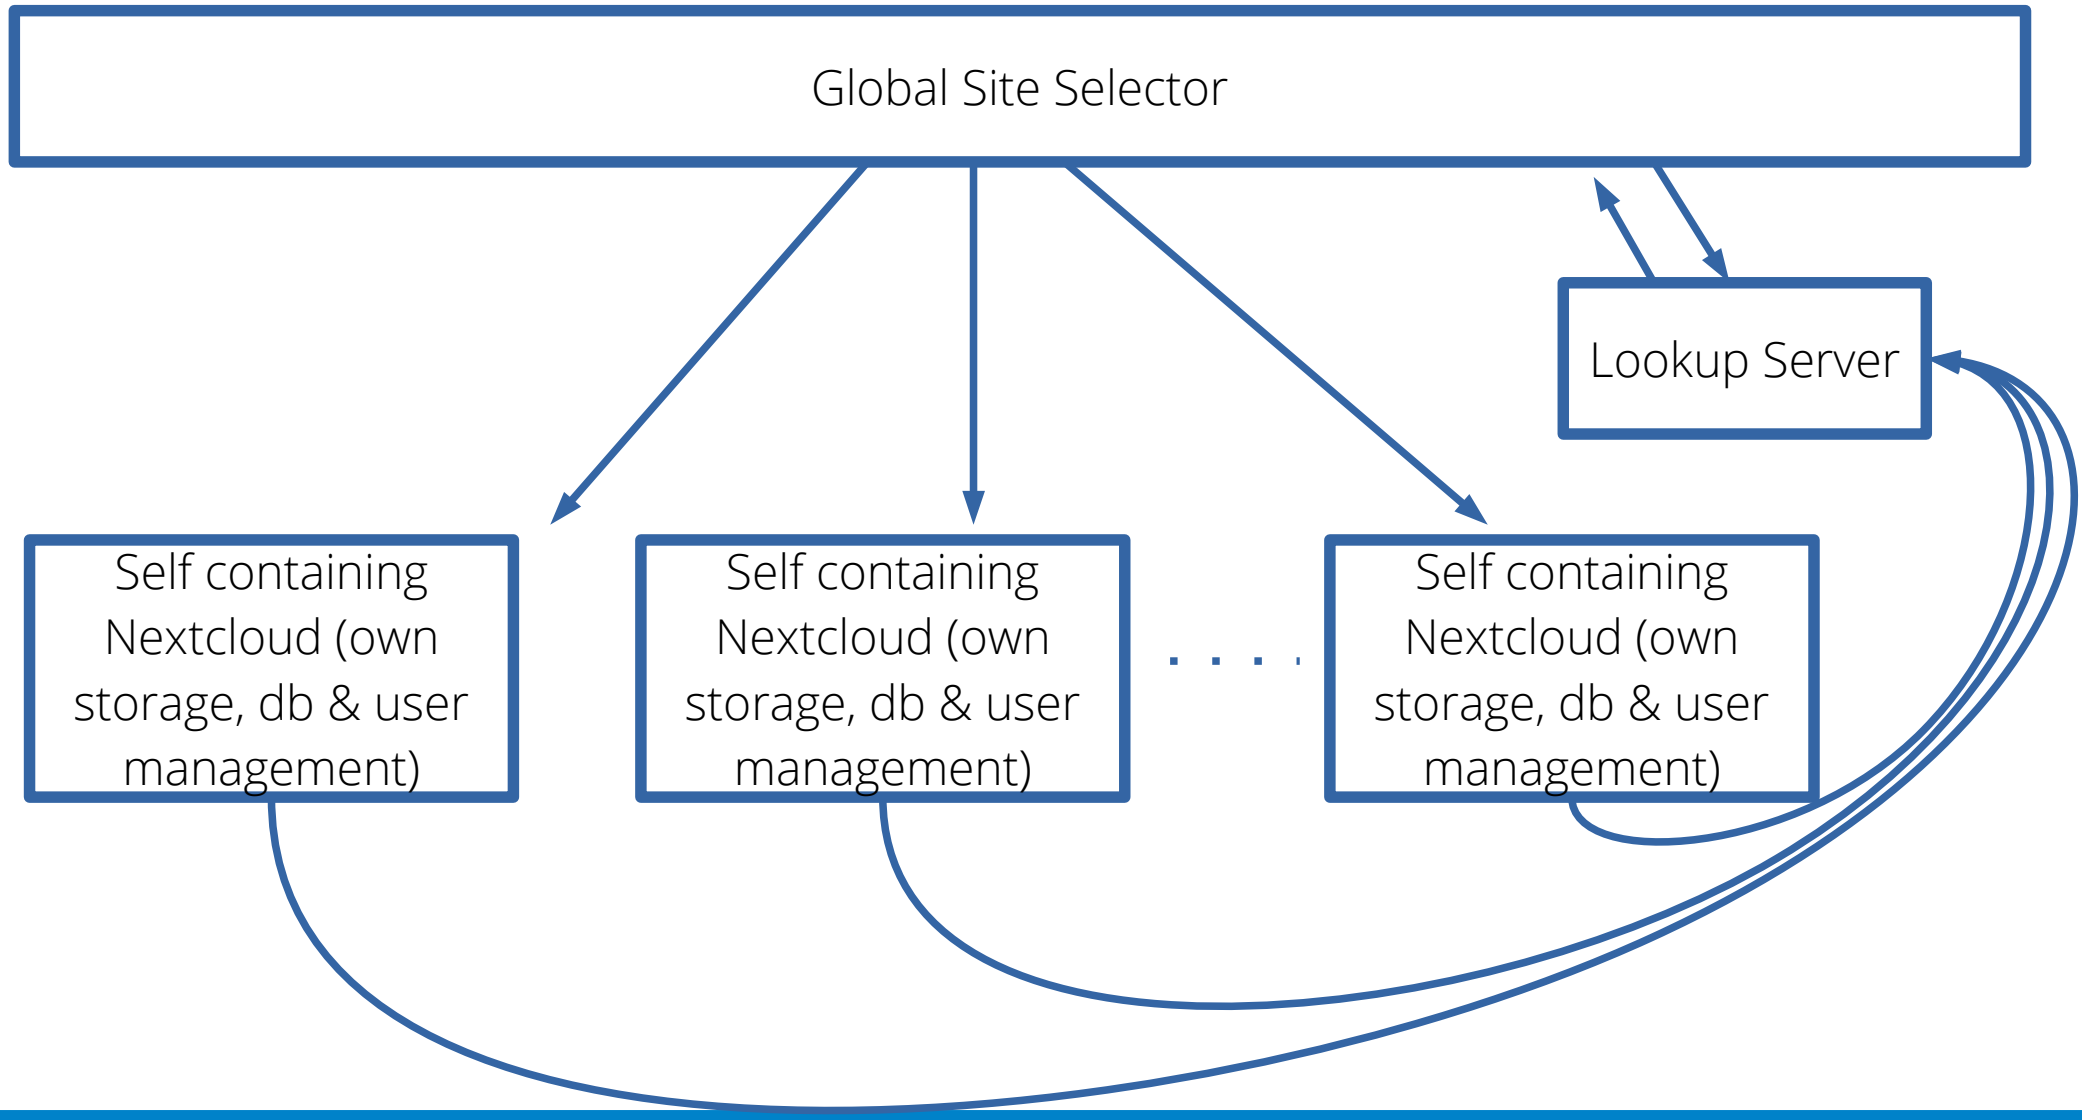

Global Scale

Components

The Global Site Selector (GSS)

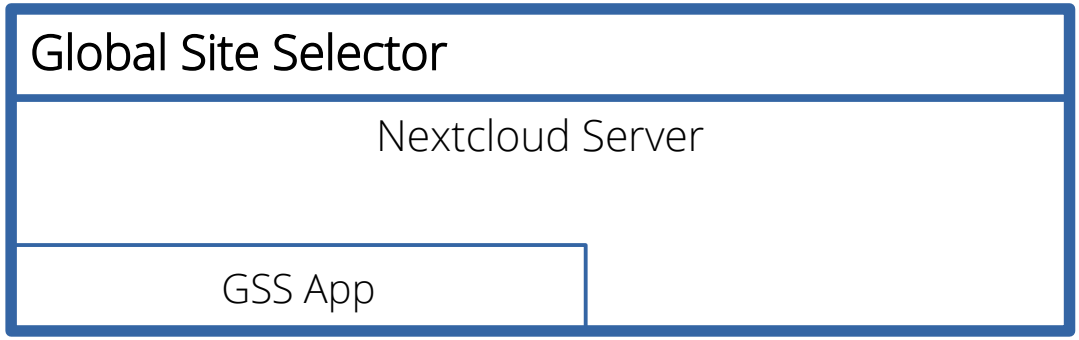

The Global Site Selector (GSS)

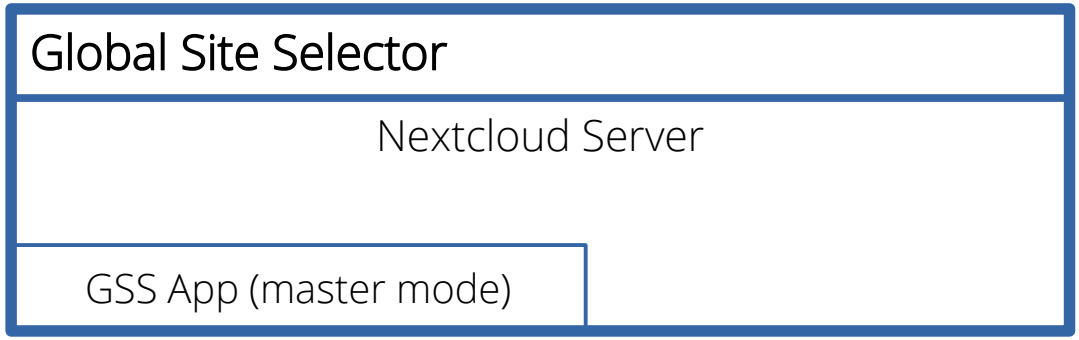

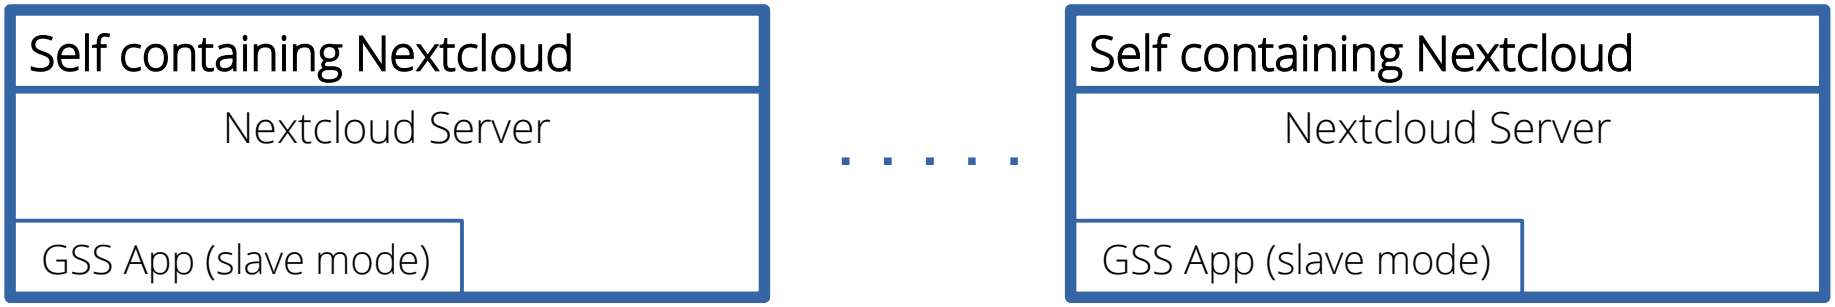

The Global Site Selector (GSS)

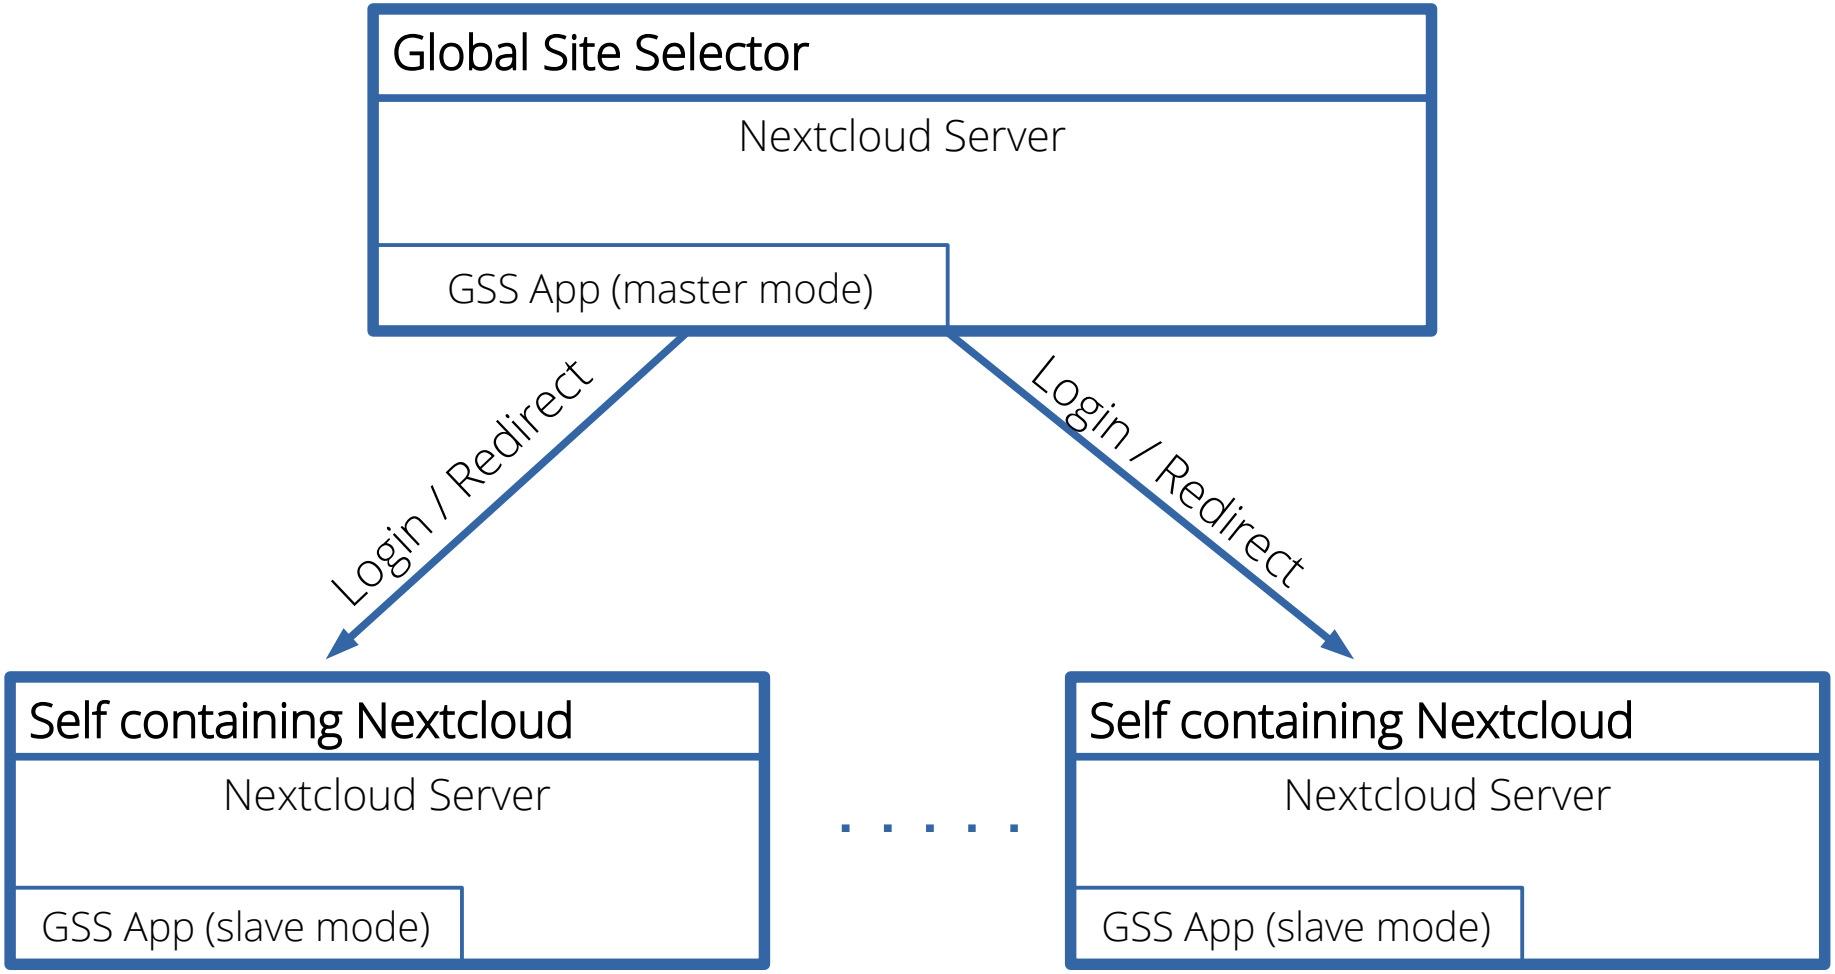

The Lookup Server

- Stand alone server written in PHP
- MySQL/MariaDB back-end
- Stores all the user records
- Used for sharing to search users and create a federated share to the right instance
- Tells the Global Site Selector on which instance the user is located on login

The individual Nextcloud instances

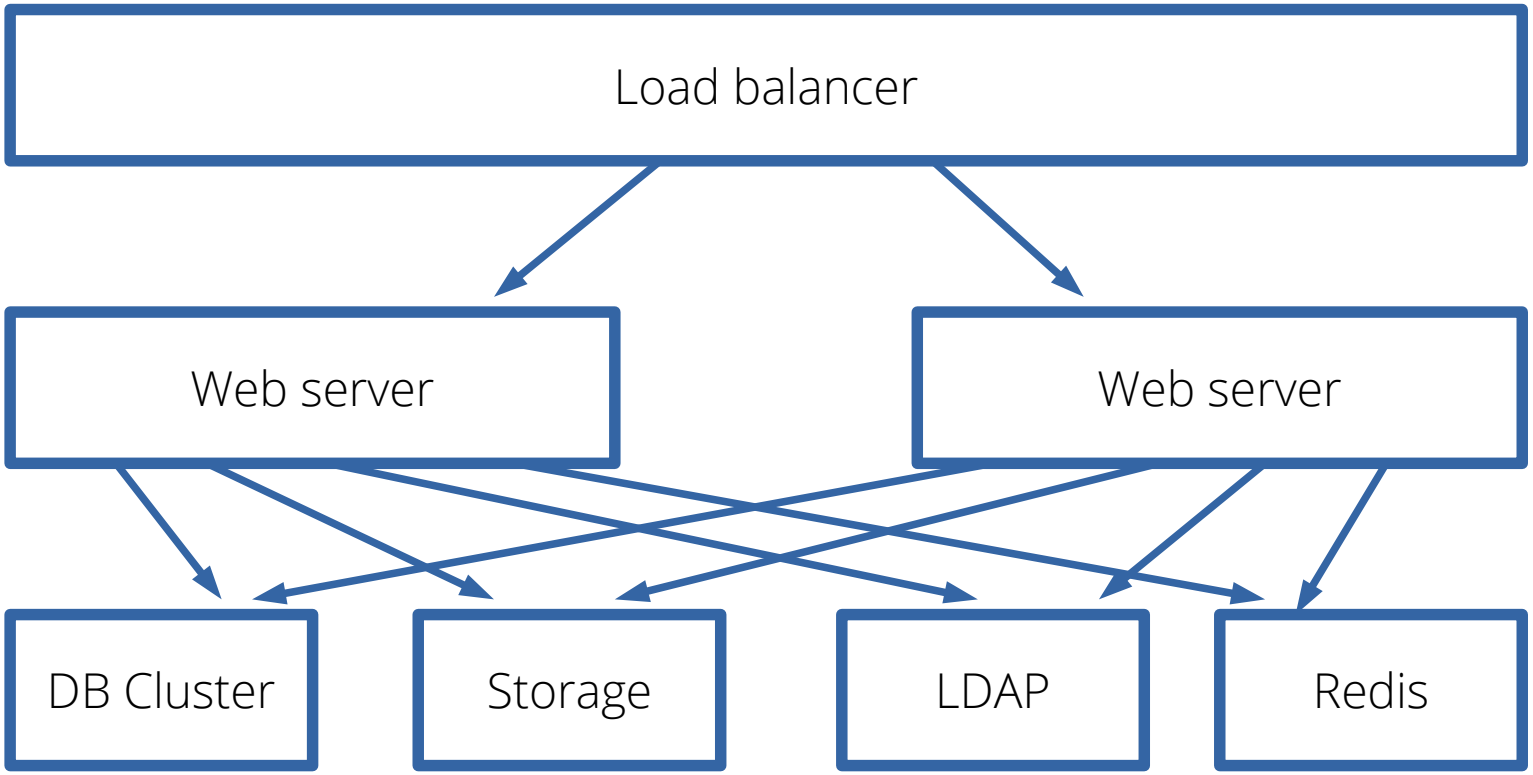

Small Nextcloud instances based on commodity hardware, following the traditional architecture for high availability

Global Scale

Communication

Login process

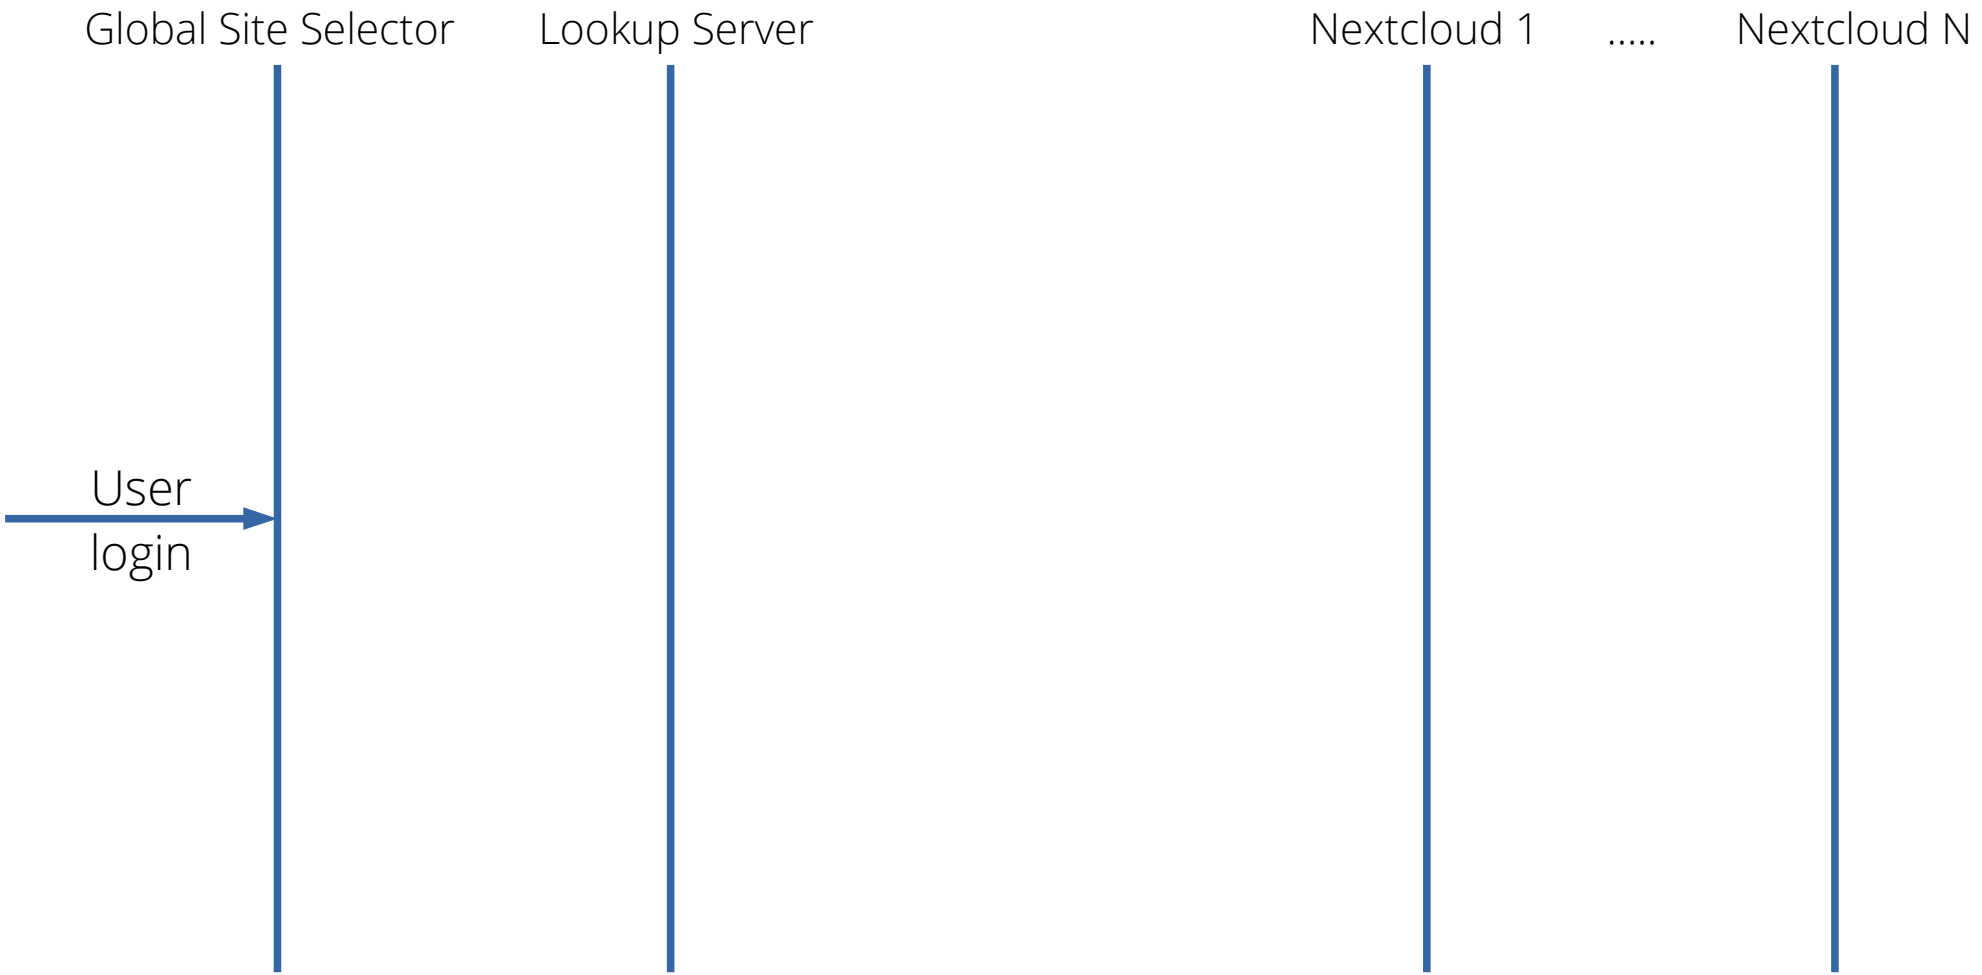

Login process

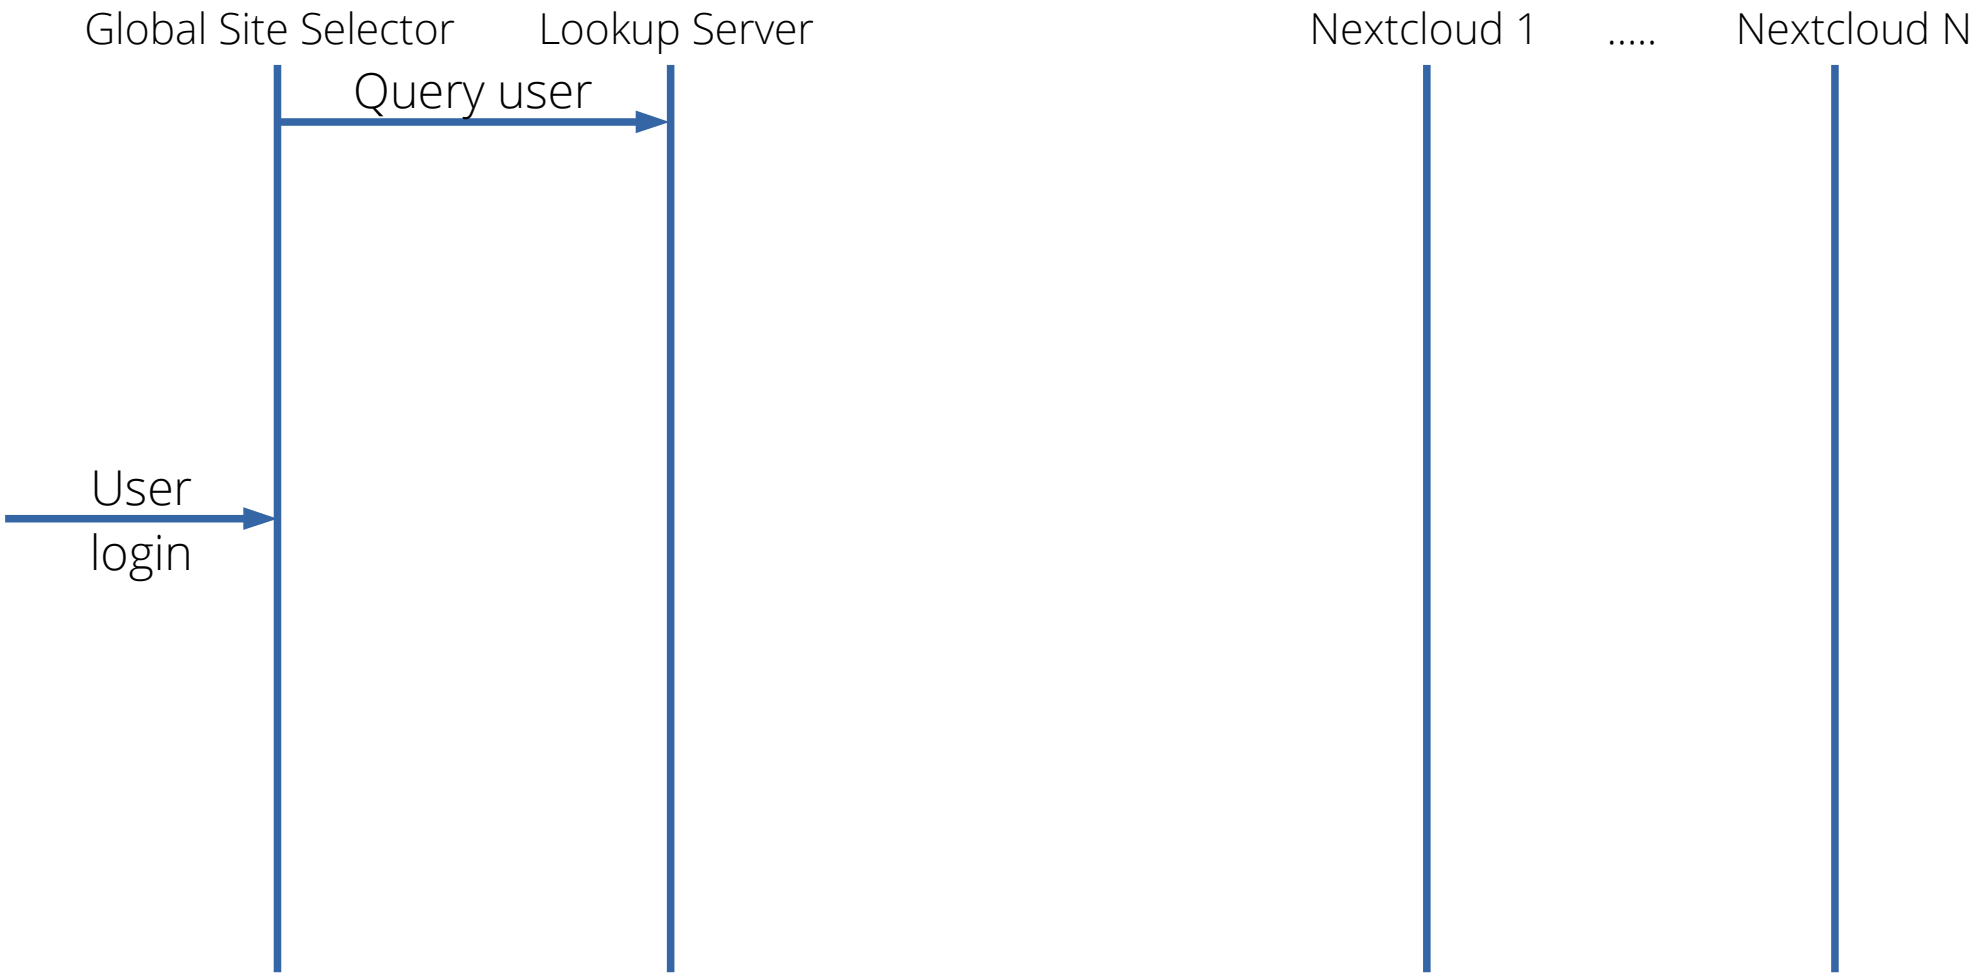

Login process

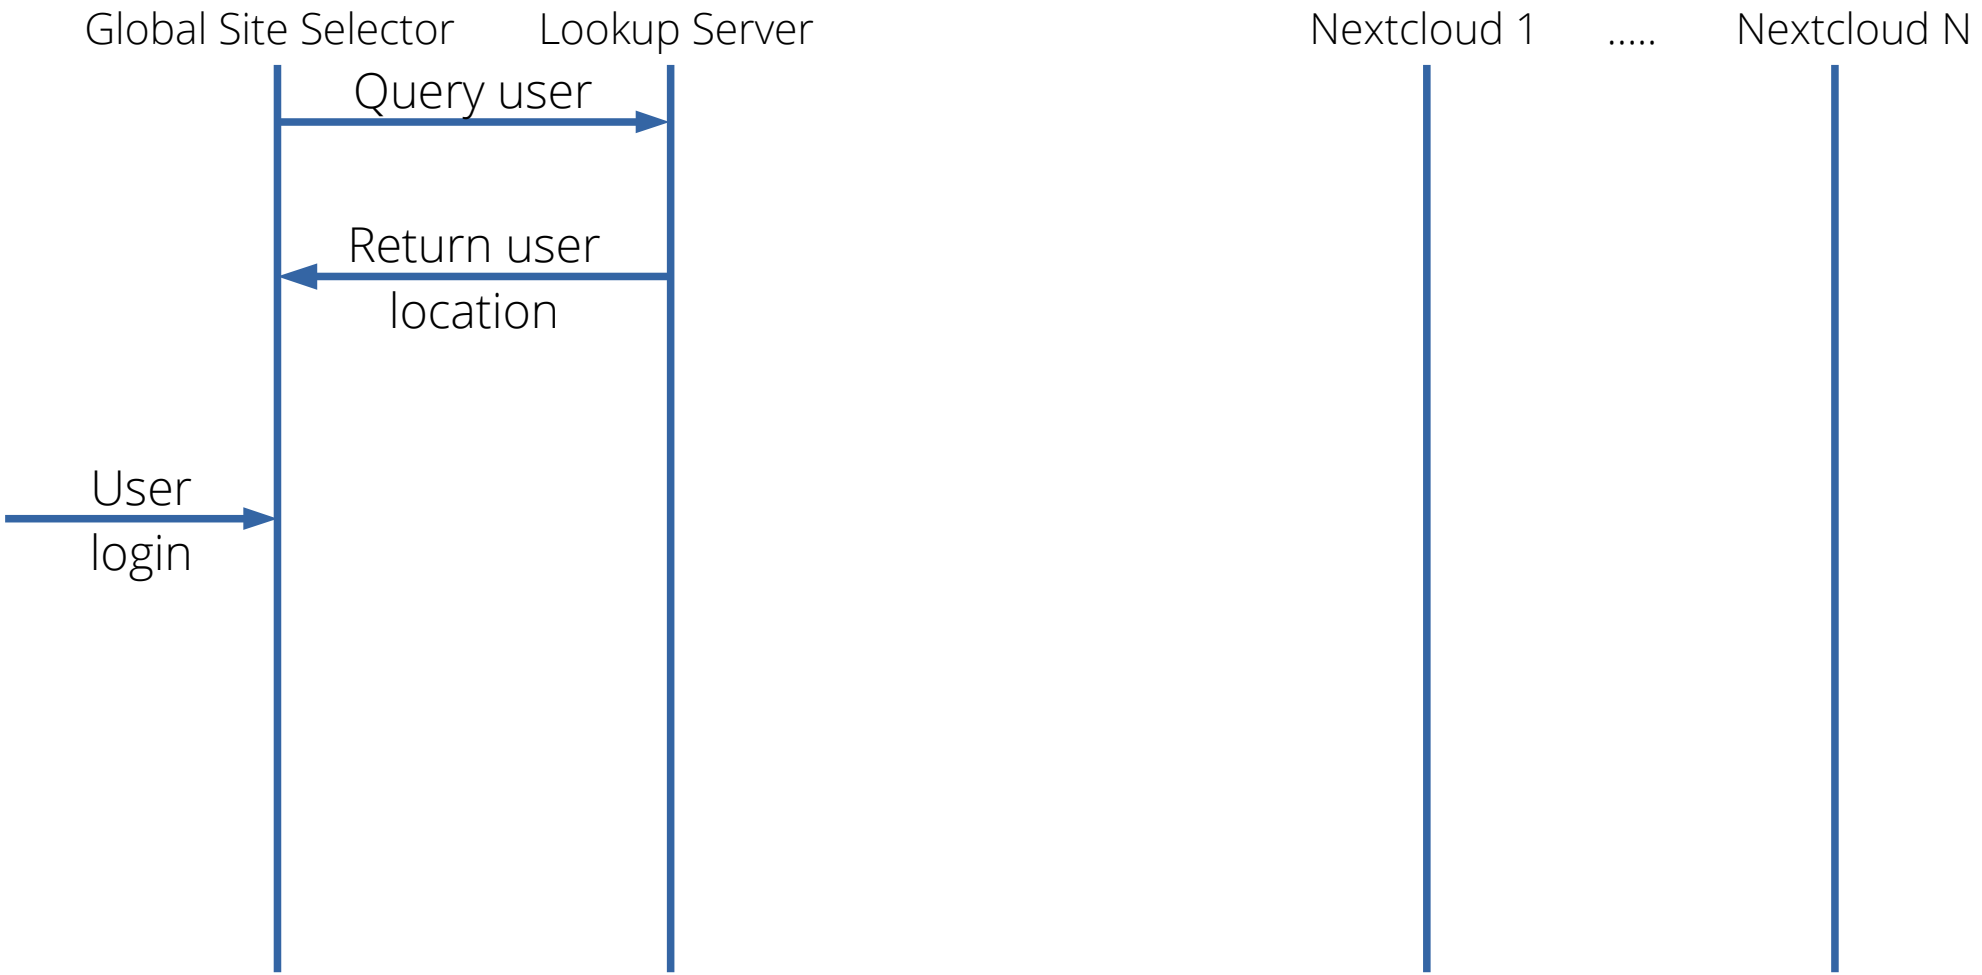

Login process

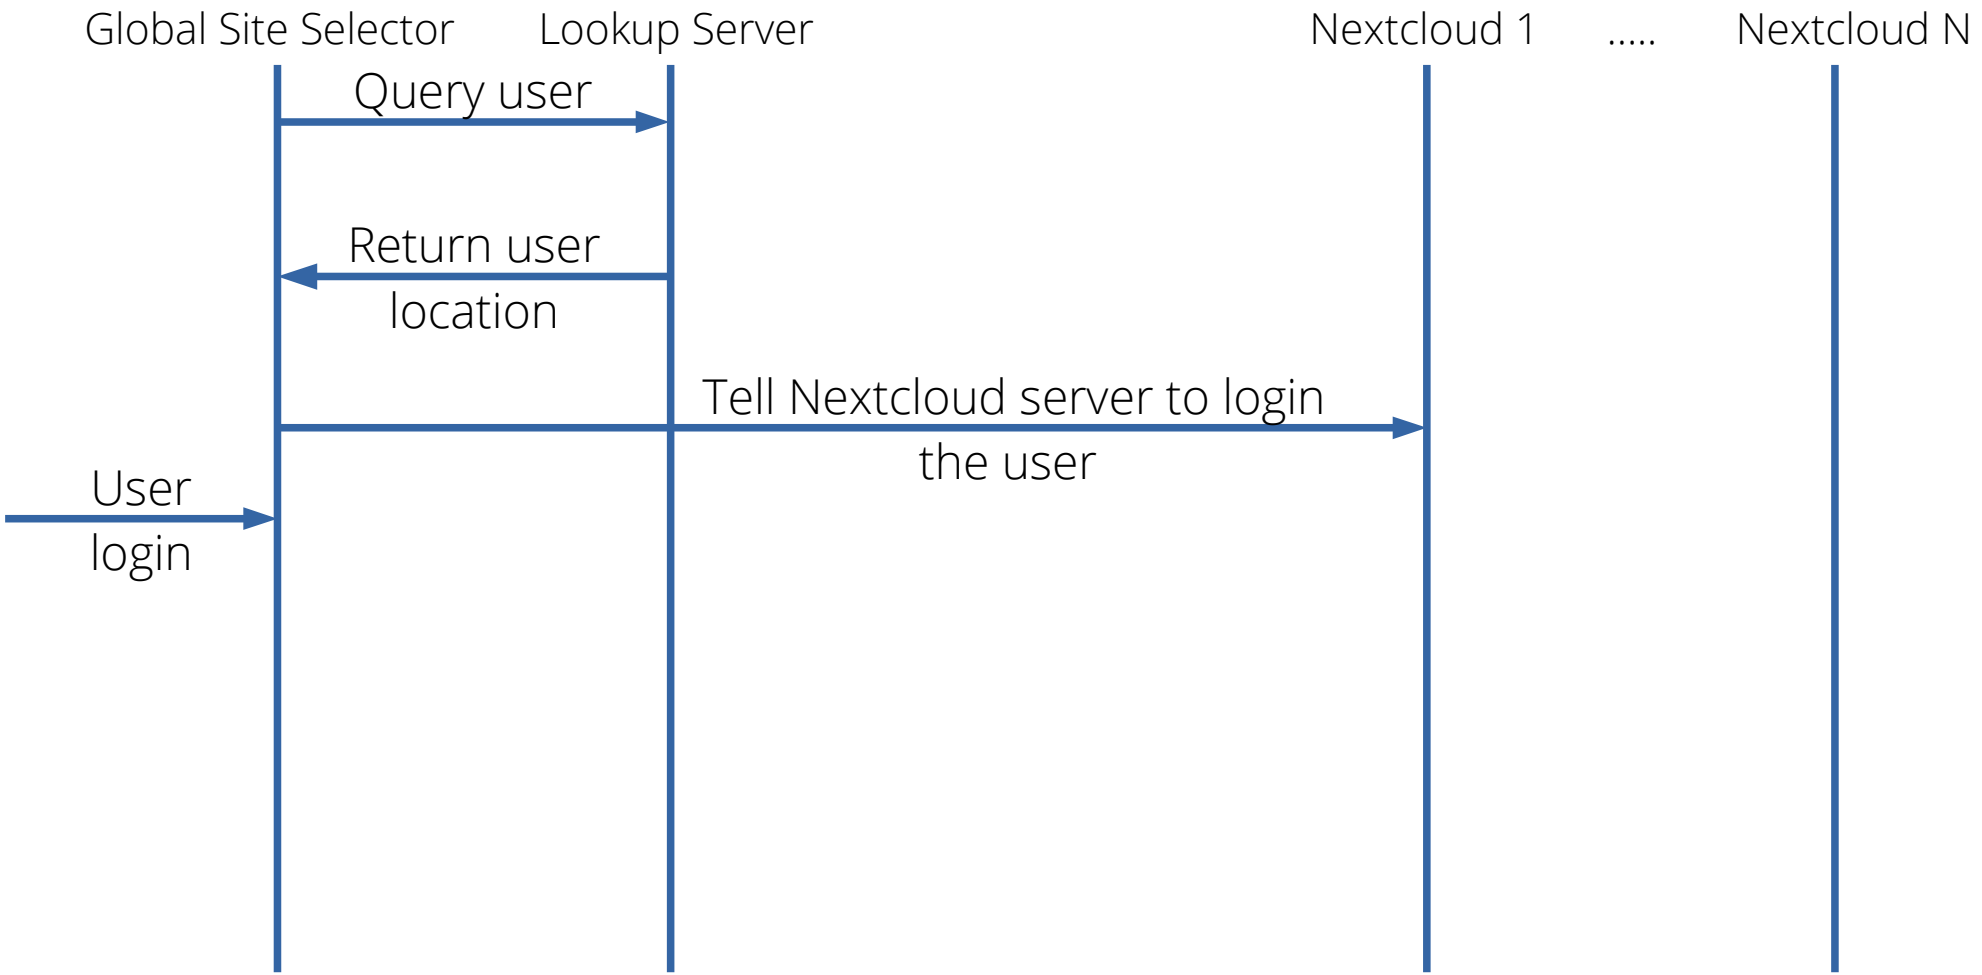

Login process

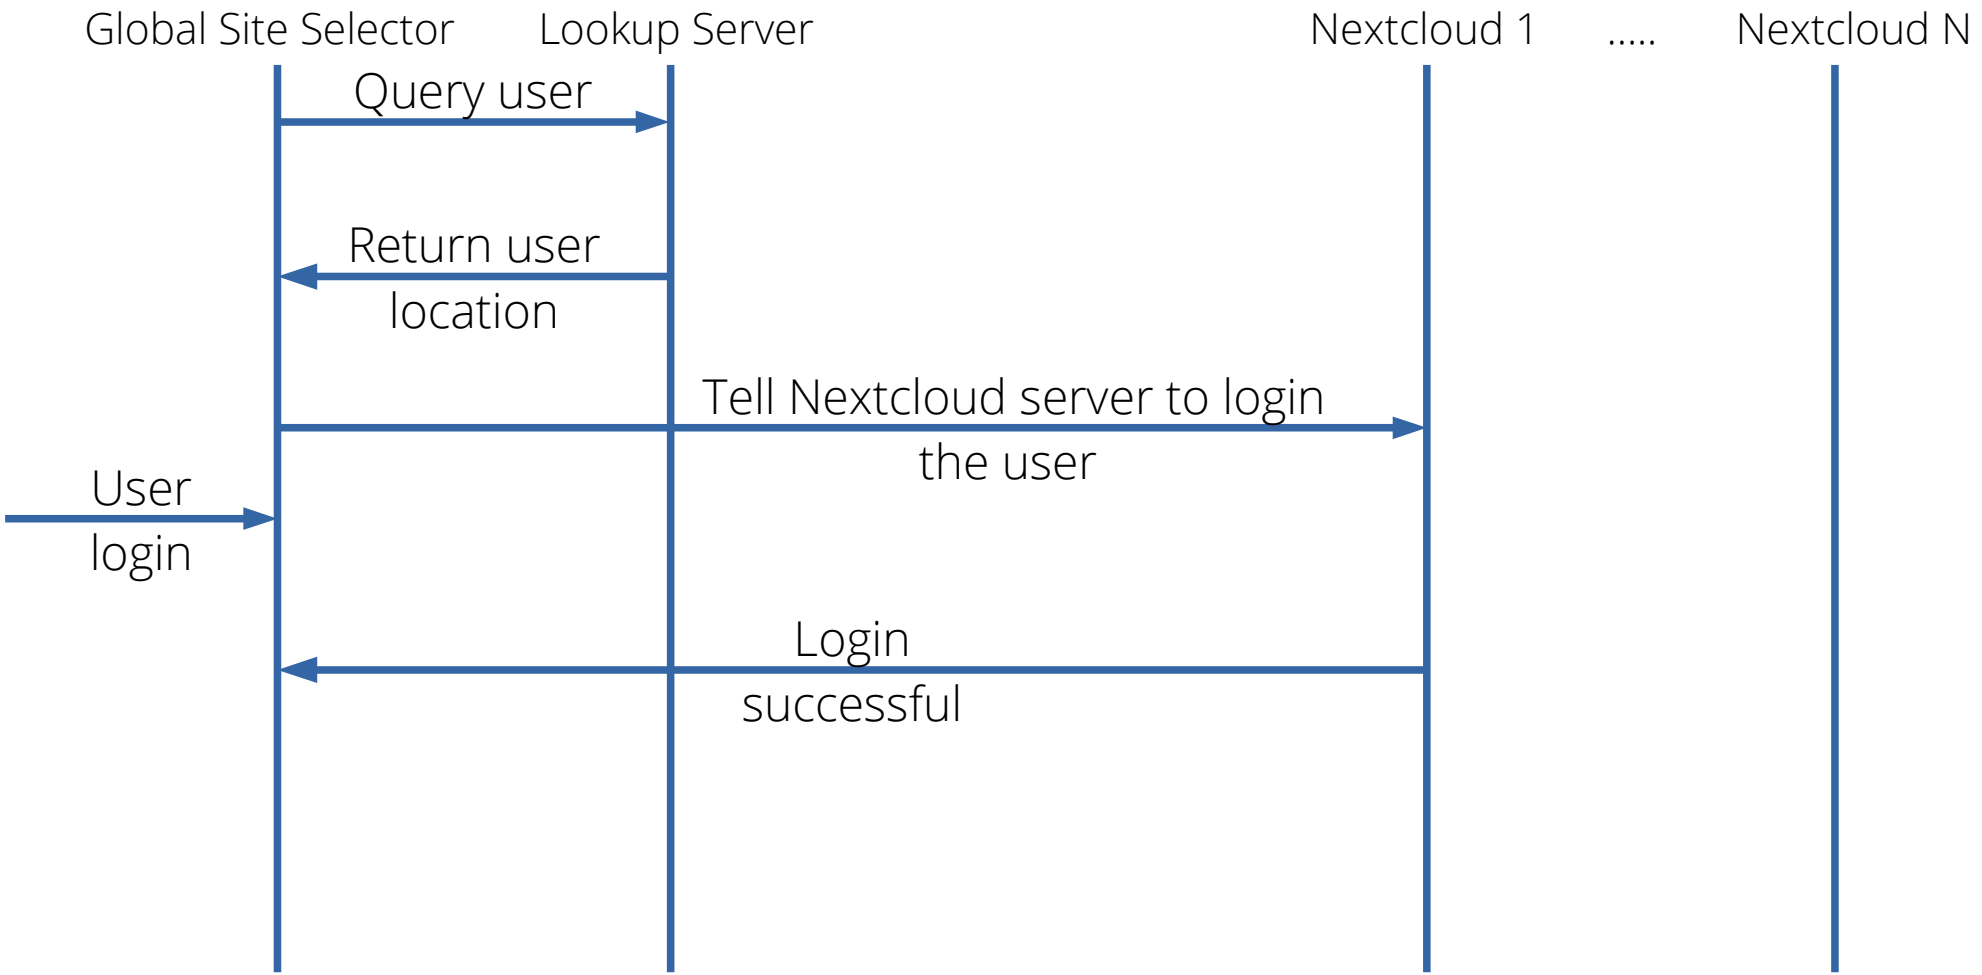

Login process

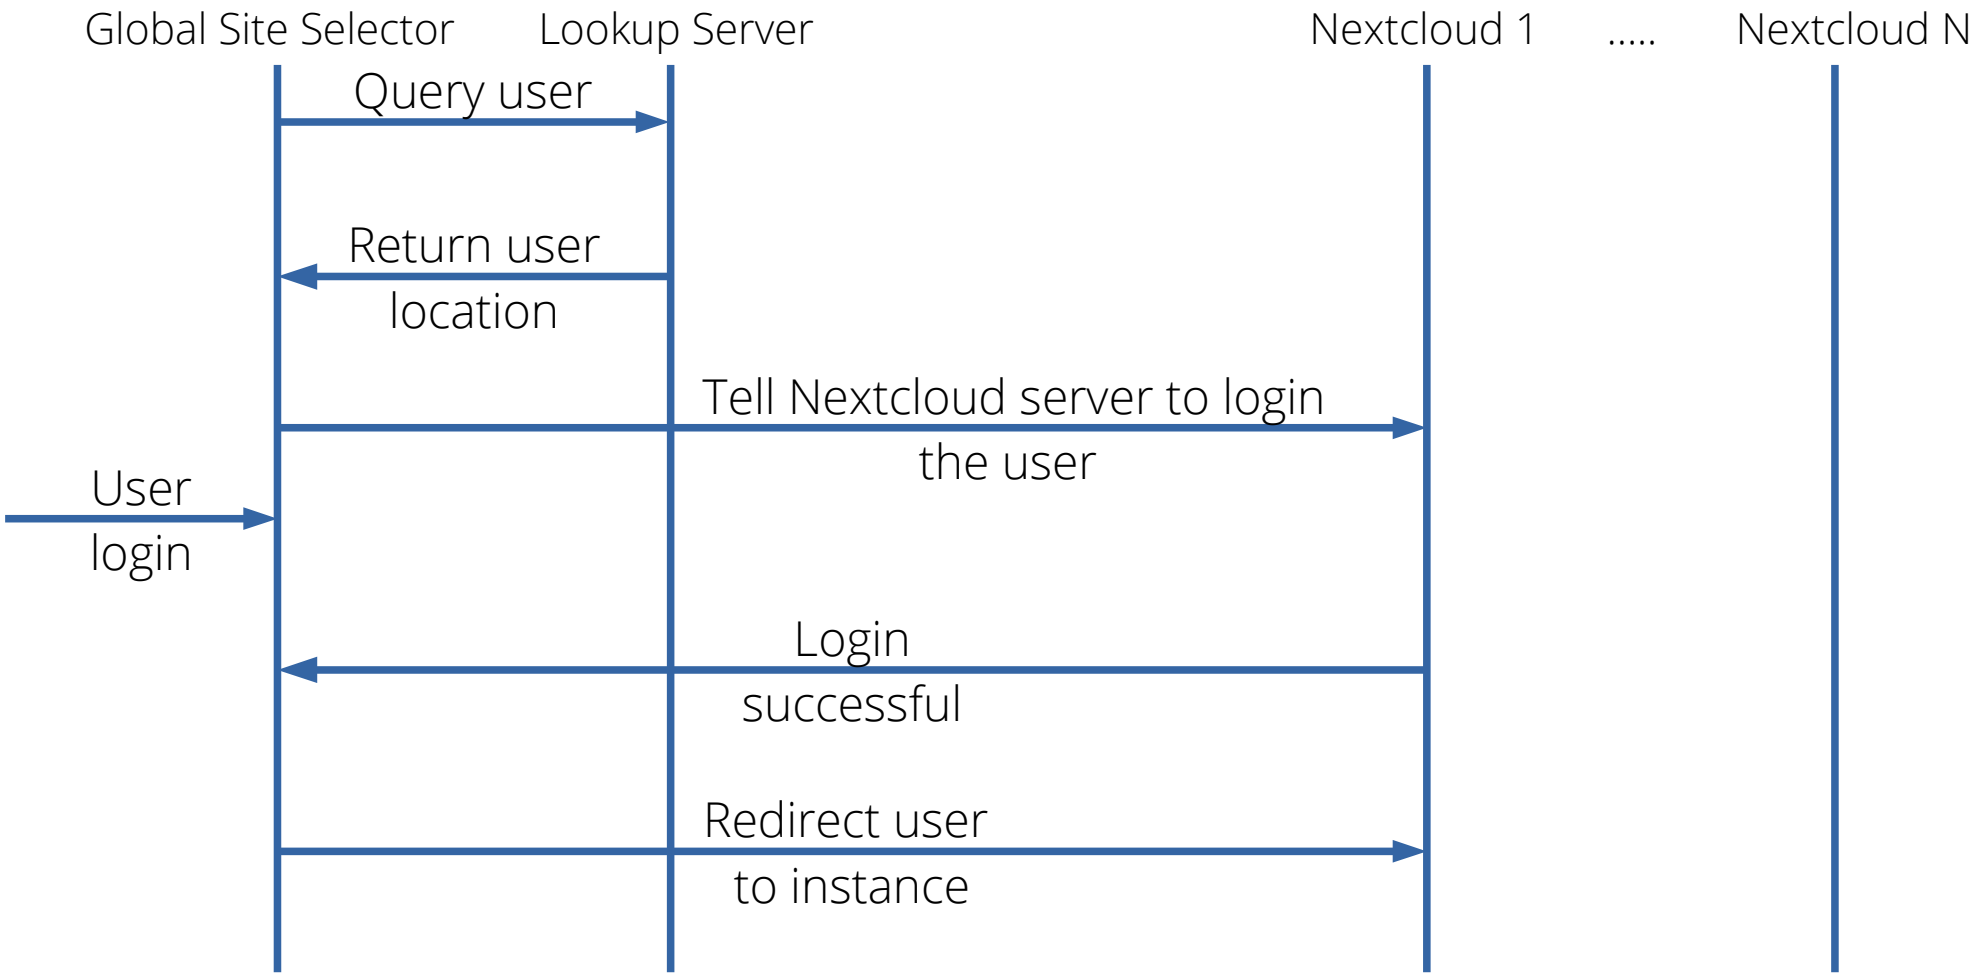

Global Scale

Re-balancing

User Migration

- Need a way to re-balance the instances
- User can be migrated from one instance to another
- Files will be transferred together with the share information and other important data
- Lookup Server will be notified
- Other Apps will be able to register a user migration process to migrate their data as well

Summary

- Enable Nextcloud to **scale** to **millions of users**
- **Cost efficient** because build on **commodity hardware**
- Maintaining many **small instances** is easier
- Keep data **next to the user** to improve **performance**
- Keep data in specific countries (meet data protection **requirements**)
- **Extend** network, e.g if you open a new office or if the organization grows
- **Re-balance** instances to **improve performance** or because **users move**
- Still **feels** like **one instance**
- Get **Cloud Federation** to the **next level**

Thank You!

Get in contact for a proof of concept

https://nextcloud.com/globalscale

Looking forward working with you to get OCM in production

https://github.com/schiessle/OCM-API