
# Digital vault

## High Security Laboratory

# 1

# Overview

## Generalities and problem statement

# High Security Laboratory

## Unique academic platform in France

### Objectives

- Expertise in computer security (audit, vulnerabilities assessment...)
- Pro-active defense against malwares and new threats
- Large scale experimentation and studies, publications
- Data collect and analysis
- Implementation and distribution of tools and softwa
- Validate and distribute research results
- http://lhs.loria.fr

# Physical Security

## Dedicated and isolated infrastructure

- A dedicated physical place (« bulletproof »)
- A full enclosure infrastructure (autonomous infra)
  - Almost self-sufficient (electricity, air conditionning)
- A dedicated network – RIPE – routing
  - Can simulate a virutal Internet
- DMZ for results dissemination and collaborations

## Enhanced security

- Different areas with different security levels
  - Office > Servers room > “Red room”
  - “Red room” completely isolated, meant to store and treat sensitive information
- Strengthened access control
  - Strong authentication  (entry pass + biometry)
  - Armoured doors and windows, alarms, airlock...

# Overview

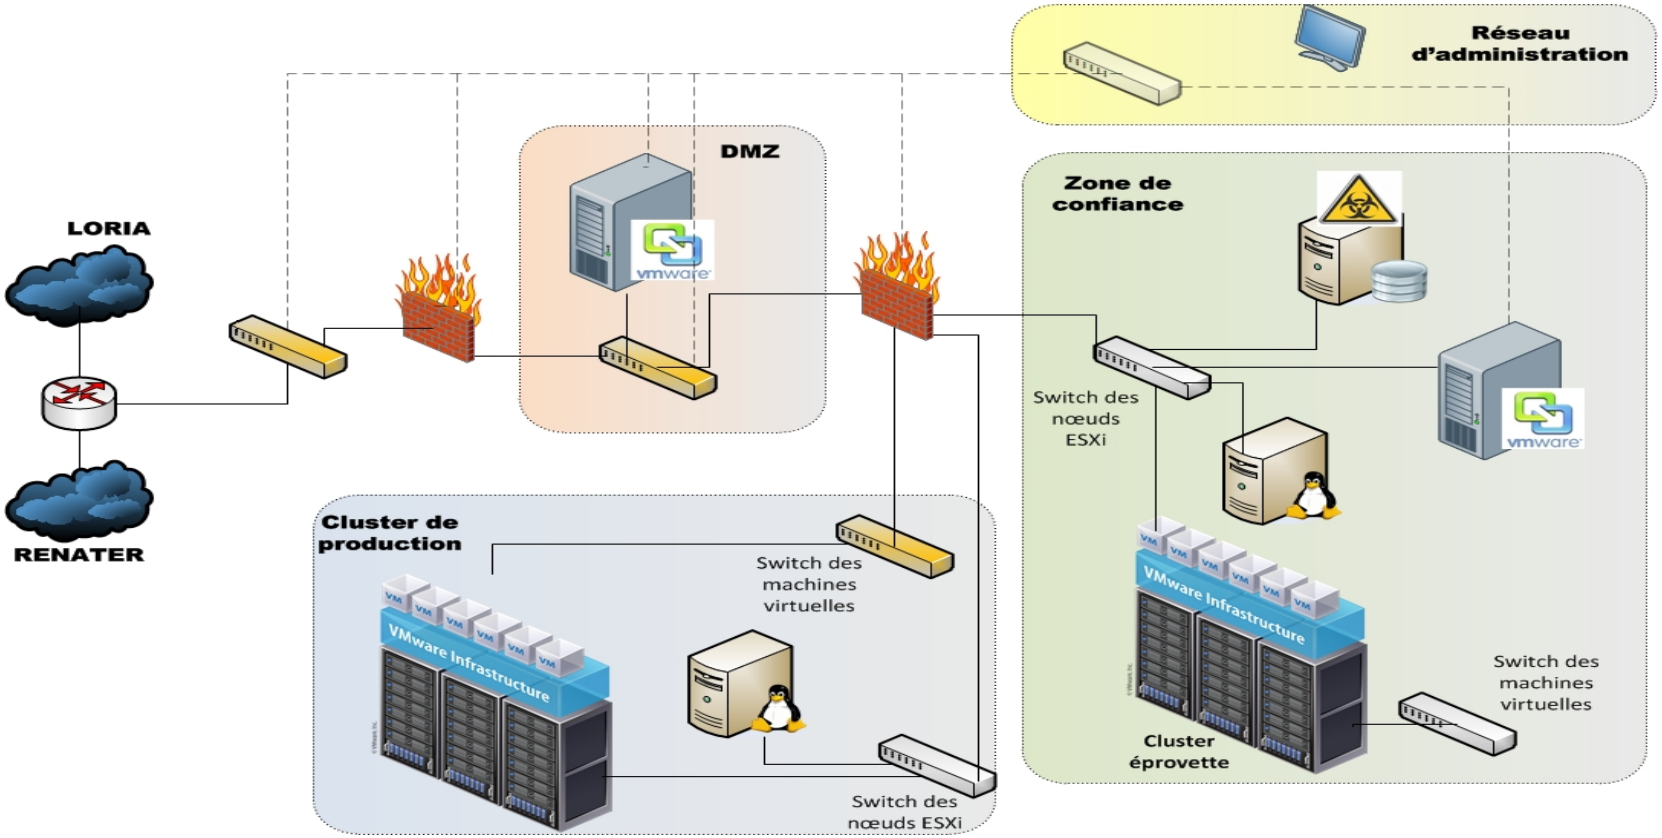
Réseau d'administration
DMZ
Zone de confiance
LORIA
RENATER
Switch des nœuds ESXi
Cluster de production
Switch des machines virtuelles
VMware Infrastructure
Switch des nœuds ESXi
Switch des machines virtuelles
Cluster éprovette

# An academic infrastructure for

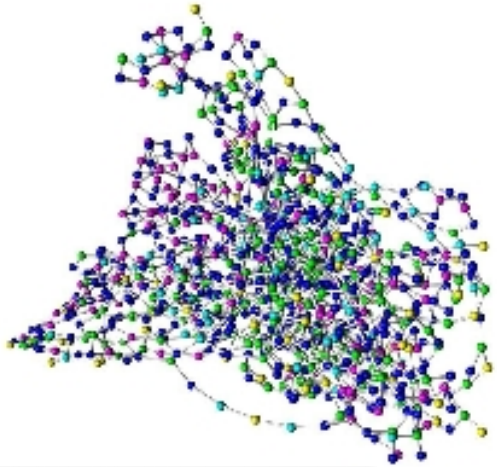

## Research in security

- Virology – Malware analysis
  - Packer analysis
  - Flow control graph ,dynamic execution in sandboxing
- Malware capture & attack recording – Global Internet security comprehension
  - Honeypots networks (servers and evolution to LIHC)
  - Darknet real-time and analysis (collaboration with NICT)
- Internet scanning campaigns
- SCADA industrial systems
- Hosting research projects
  - ADT SEA (Android security)
  - ADTs APISense / Crowdlab
- …

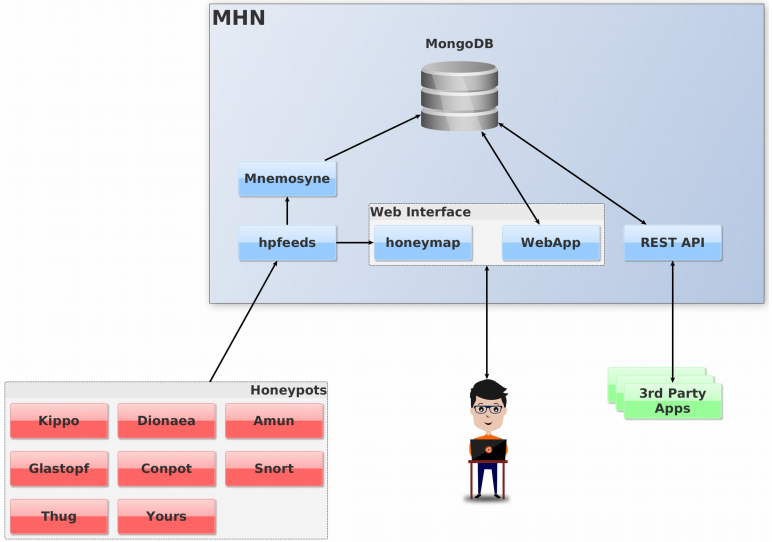


**Research in security ⟹ Sensitive data**

# Problem statement

## Hosting research projects

- Teams working on sensitive data
  - Crowd-sourcing, privacy/personal data, collected information
- Industrial collaborations
- An argument for projects proposals
- ZRR, “local sensible” (?)

## How can a delocated team work together on sensitive data ?

- Accounts, access rights
  - C Class outside the Inria / LORIA, dedicated RIPE Register, DNS, NTP, APT, ...
- Delocated, i.e. geographically distributed, means access via the (unsafe) Internet
- Serve as many projects as possible
  - Documents, code, but not only

# 2

# Digital Vault

## What is it ?

# Objectives

## Why “safes” ?

### Working on sensitive data

- **A group of persons** wants to collaborate on the same subject, work together on shared data, the same software, **in a secure way**.
- Examples
  - A development team working on a important code.
  - Writing a sensitive document.
  - Testing an algorithm on sensitive data.
  - A research team, collaboration, with NDA
  - ...

- « Reasonable trust » in the IT team

# Choices - Strategy

## Similar to a physical vault

- Data are in the vault, and never get out of it
  - A user CAN extract data if he want, but he is responsible,
  - Avoiding data leaks.
- The user “goes in the vault” and work.
- All the users have access to the whole data contained
- One and only one vault per physical host
  - But we can have several hosts in a vault
- The safe is closed when nobody works in it. Data are encrypted

## Consequences

- **Everything needed to work MUST be in the safe**
- Safes are auto-sufficient

# What it is not (only) – Related work

## A safe for digital data

- Not only a solution for storing and archiving sensitive data
  - Banks, electronic payslip...
- More secure than encrypted cloud storage
  - But encryption is not on the client side
- We want to work on the data as well in a secure and robust way

## Software forge

- Permits to work as a team on projects
  - Gitlab, github, Inria or Renater forge...
- What about the security, is it sufficient for highly sensitive projects ?
  - Projects isolation, fine grain authentication and access rights

## Trust

- Existing solutions mean you trust them completely
  - What can the administrators of these solutions really do (or forced to do) ?

# Constraints and limits

## Constraints

- One sole project per vault
  - Only accessible for a small community of users
- Not publicly announced, users must know its existence
- Administration of the vault delegated to the users
  - Based on the project, solution administrators may help
  - Physical access to the servers may be an issue

## Limits

- Number of users
  - At the moment, limited by LUKS / cryptsetup implementation
- Capacity and performances
  - Not an HPC cluster
  - However, storage can be outsourced safely
- Dedicated hardware
  - SCADA, GPU
- ...

# 3

# How it works

# Architecture

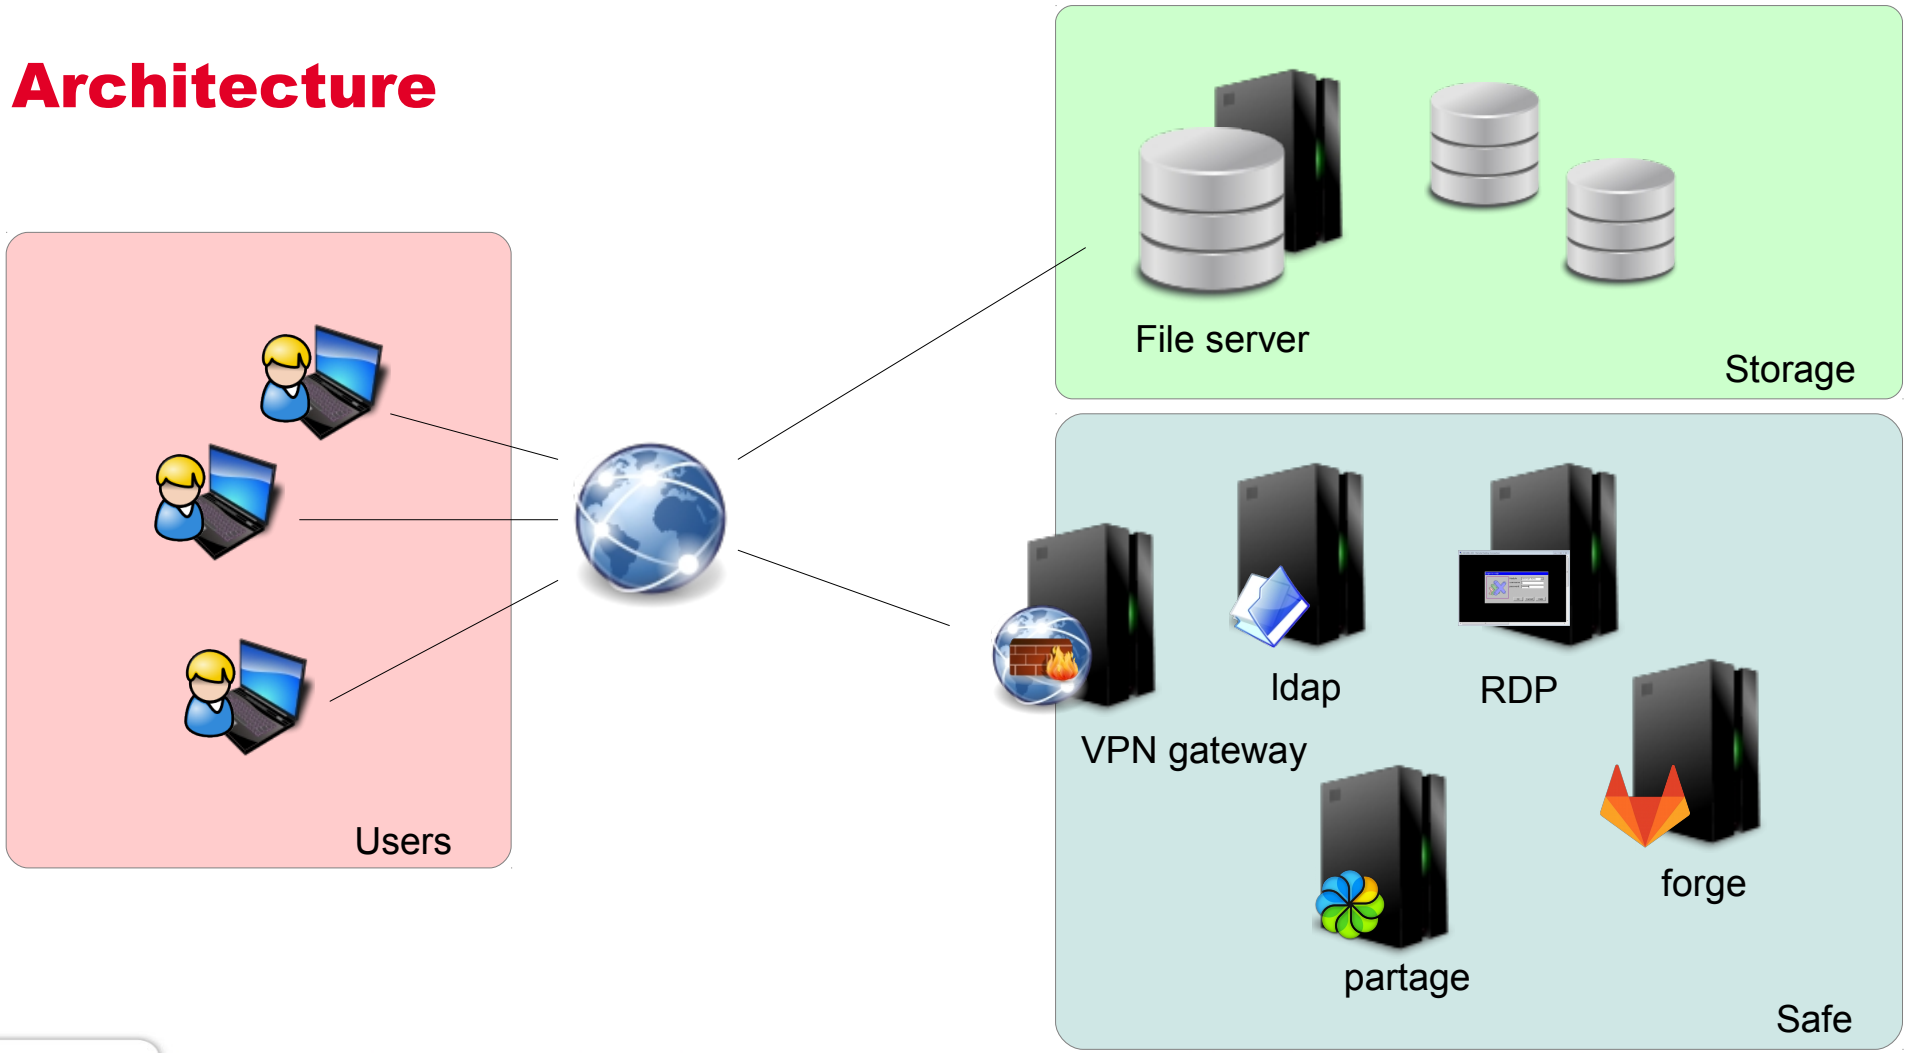

# Architecture

## Actors

- Client
  - Anywhere on the Internet, connect through a VPN connection
- Storage
  - Any entity capable of exposing crypted (AES 256) iSCSI volumes
- Vault itself
  - One or several hypervisors and associated VMs / services
  - VPN gateway, LDAP, forge, RDP, monitoring...

## Modes

- Fast
  - VMs are paused / unpaused → fast
  - Sate stored in RAM, not persistent
  - Cannot ensure confidentiality / integrity
- Safe / secure
  - VMs are suspended / restored → slow
  - State stored in the crypted volume, persistent
  - Same level of security than the data itself

# States

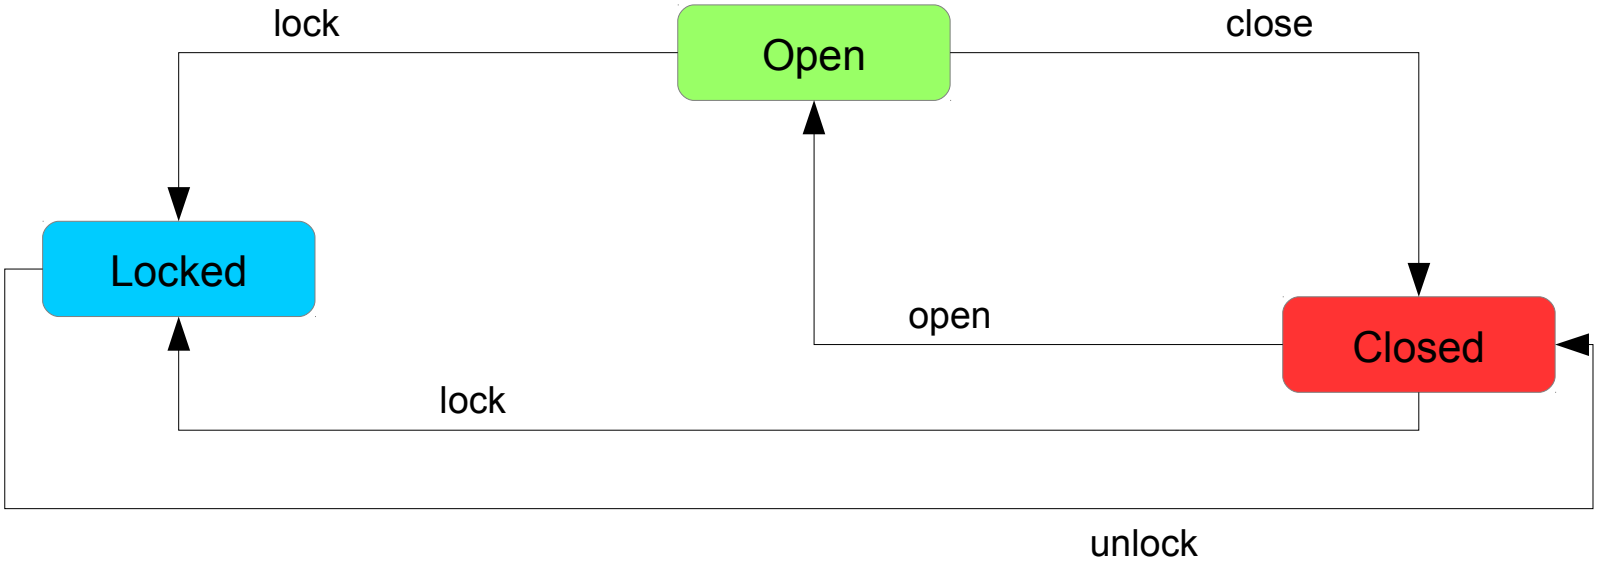


## Tasks

- Vault must be open to work on the data == user connected
- Tasks permit long treatments and batch processing
  - User programs a task with timer
  - User can disconnect from the vault which will remain open
  - When task finishes or timeout is reached, vault is closed

# Implementation

## Technical choices

- OpenVPN + OpenLDAP with client isolation
  - Once connected, users can use SSH or RDP to access the vault
- One or more Xen Hypervisors
  - Permits to use Dom0 to manage the vault
  - Dom0 not accessible via Internet
- LUKS / Cryptsetup
  - AES 256 block ciphering of an iSCSI volume
  - Only encrypted blocks between the hypervisor and the storage
  - Storage can be anywhere, e.g. cloud solutions
- Communication and synchronization via RabbitMQ
  - AMQP protocol
  - Several queues based on the operations and roles

# How it works

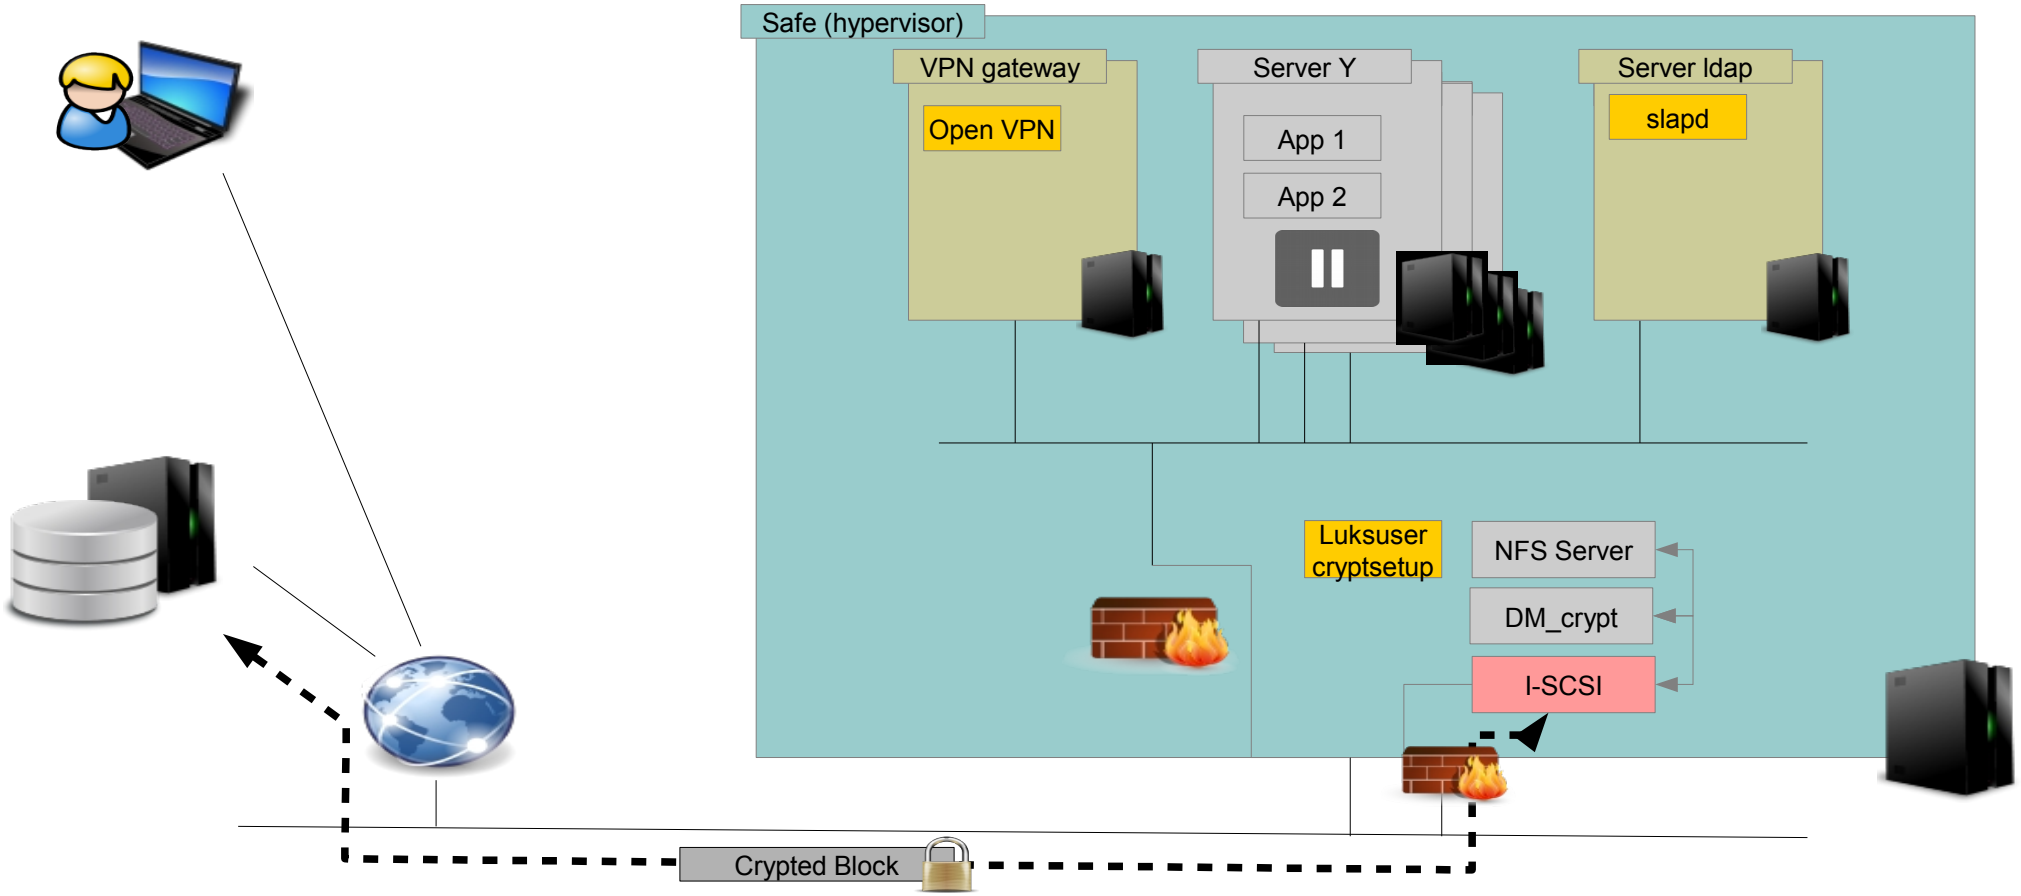
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
Luksuser
cryptsetup
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

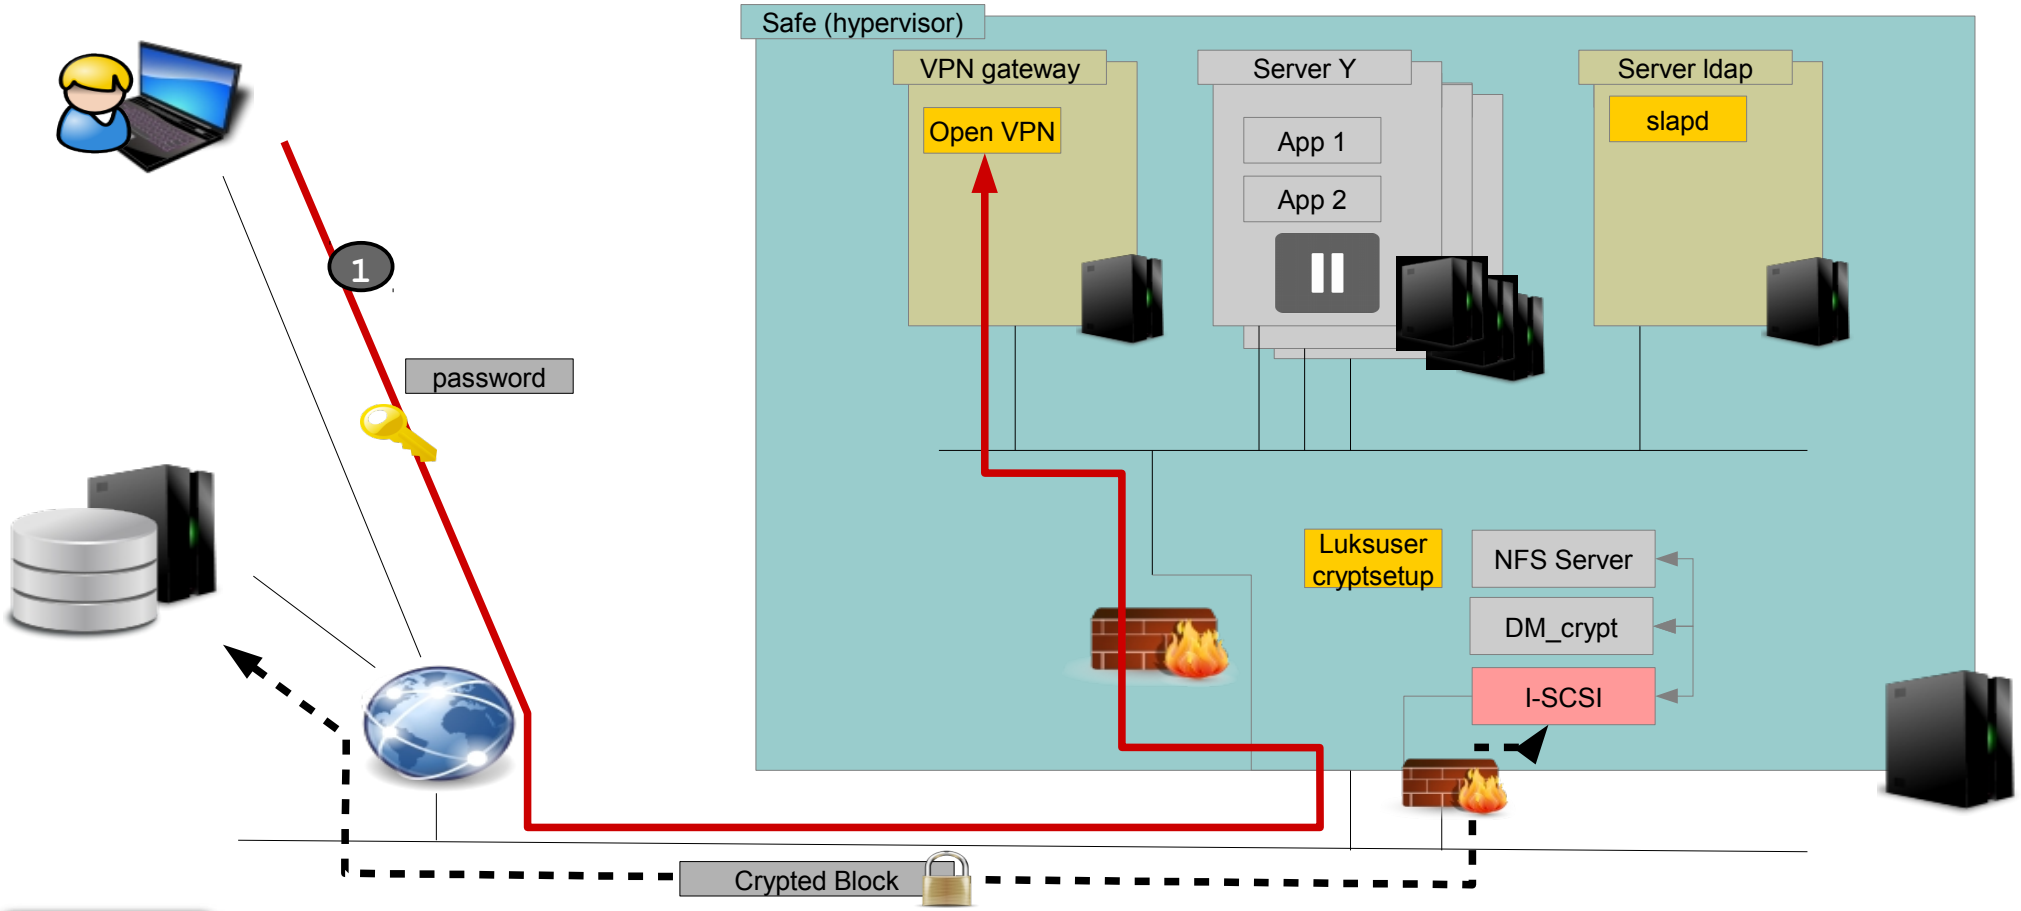
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
1
password
Luksuser
cryptsetup
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

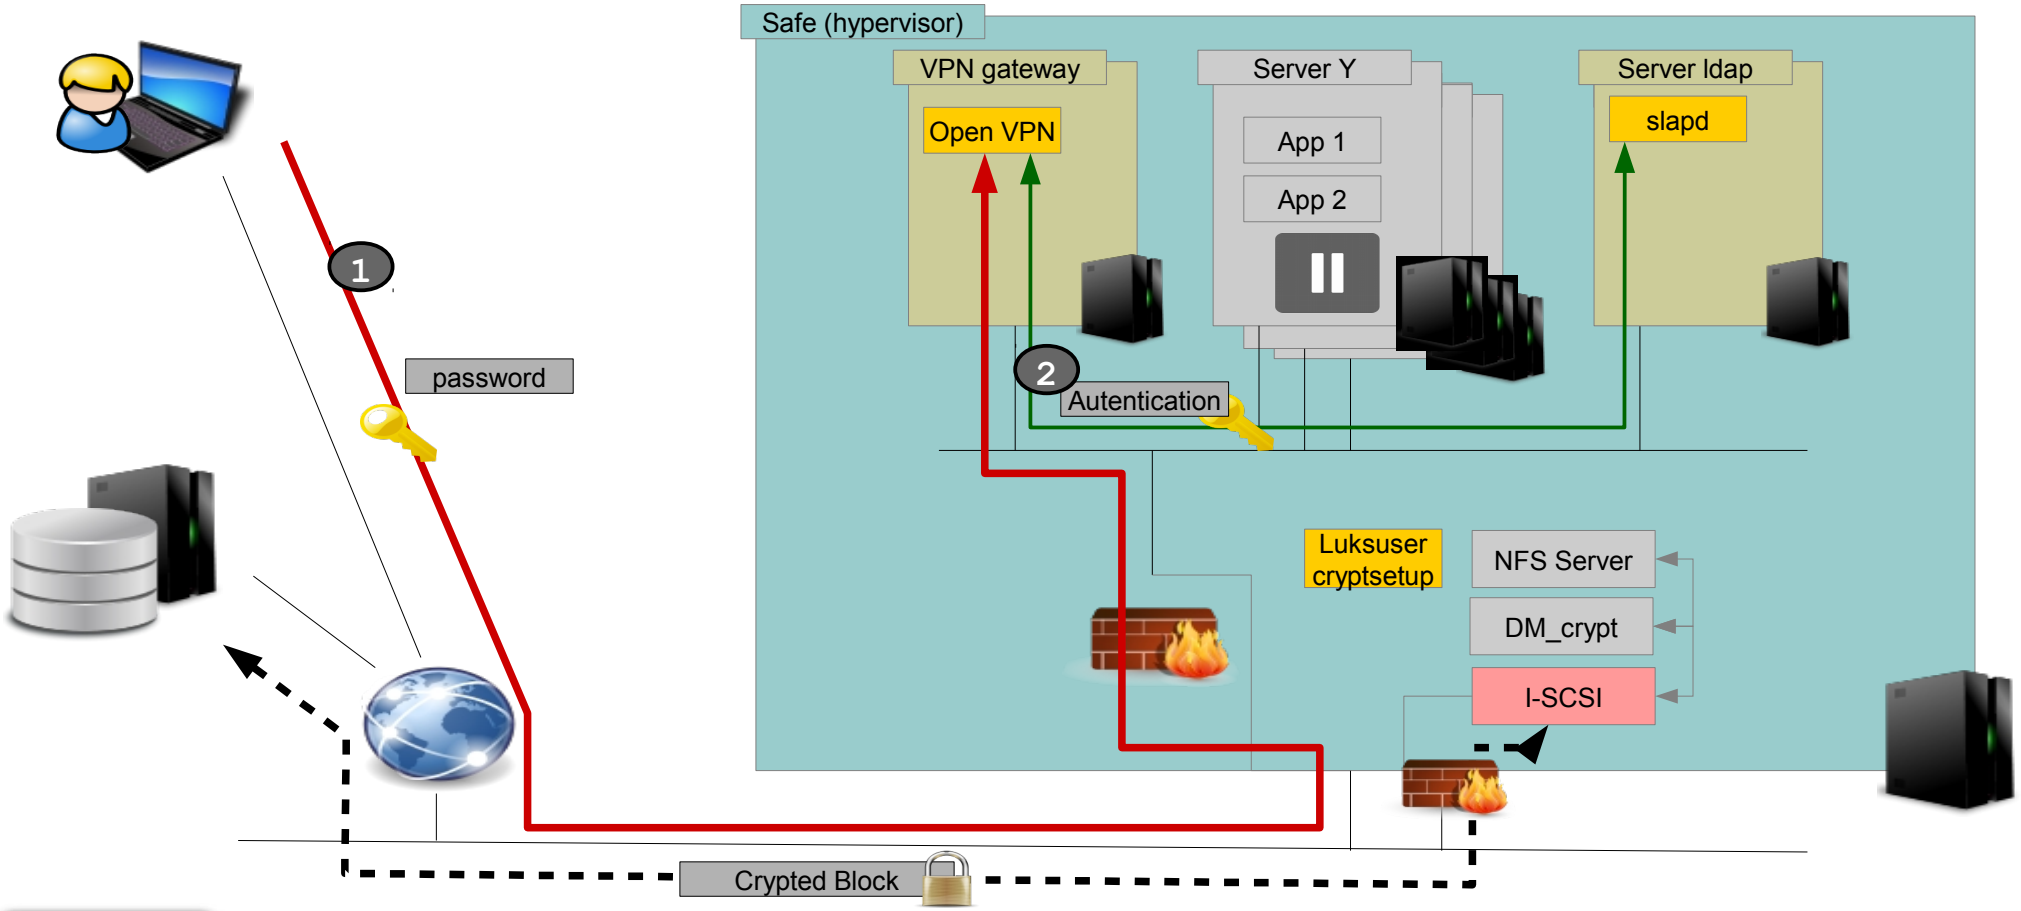
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
1
password
2
Autentication
Luksuser
cryptsetup
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

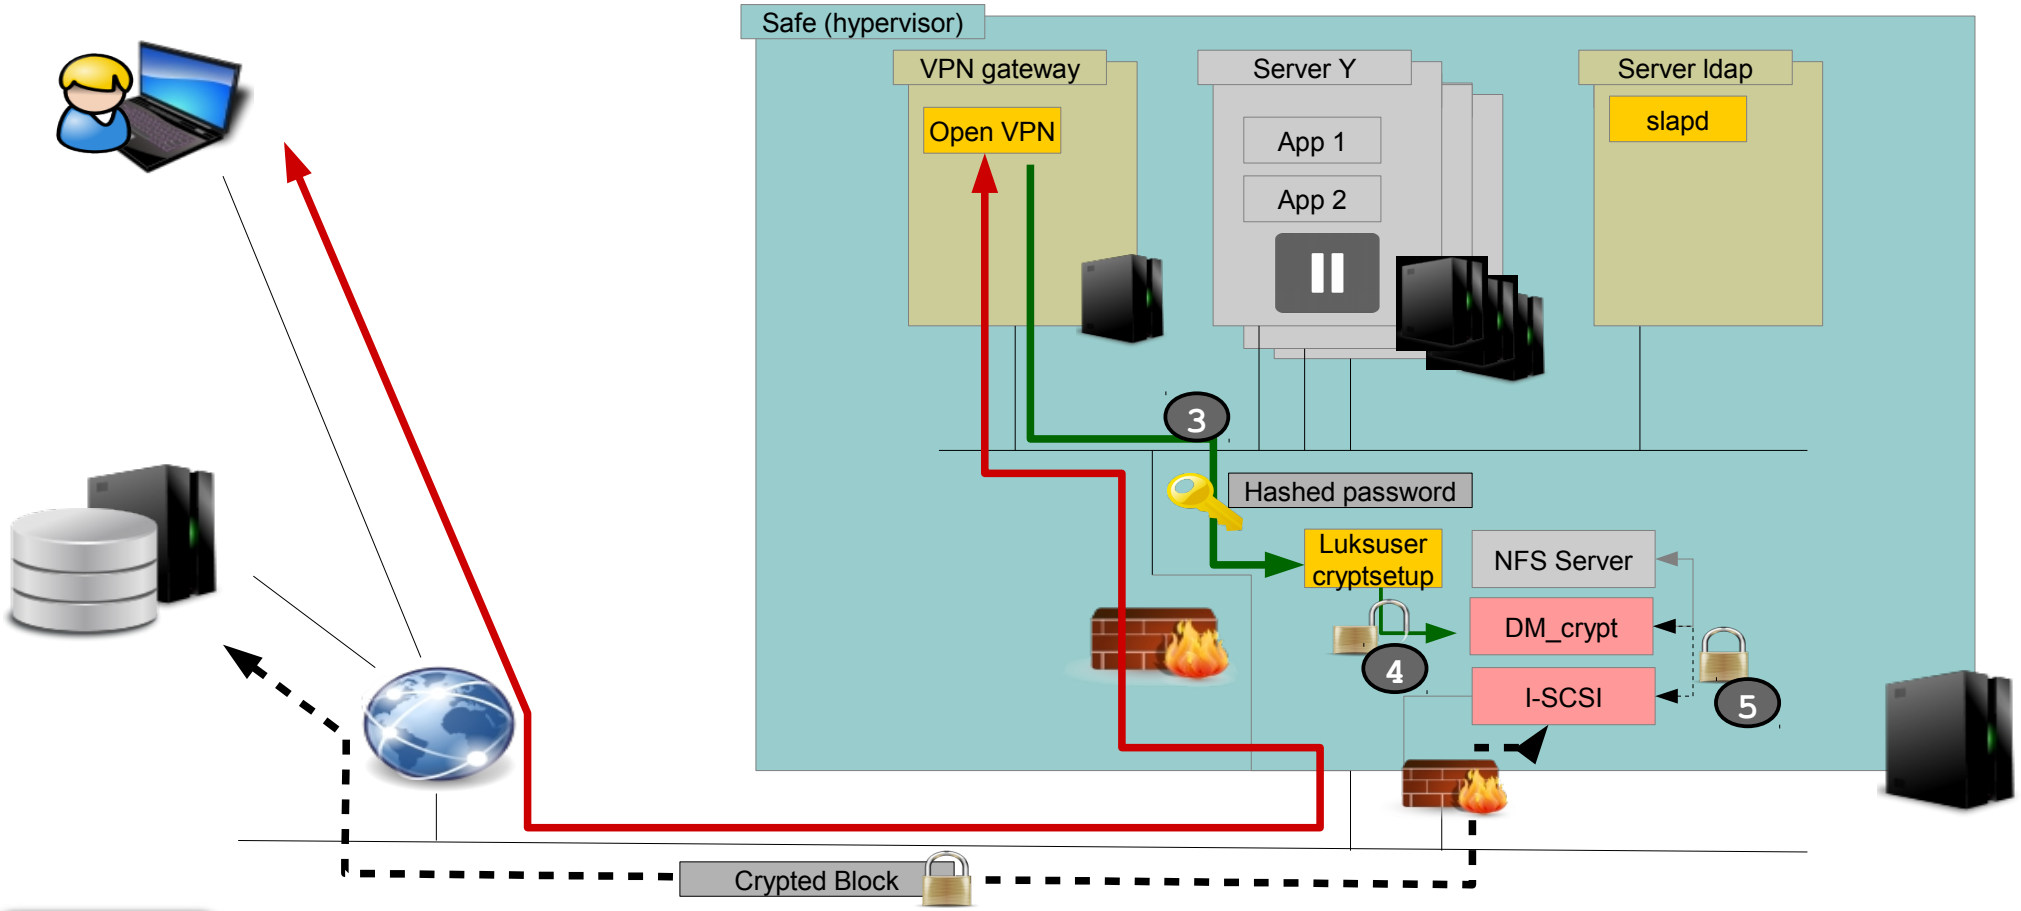
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
3
Hashed password
Luksuser
cryptsetup
NFS Server
DM_crypt
I-SCSI
4
5
Crypted Block

# How it works

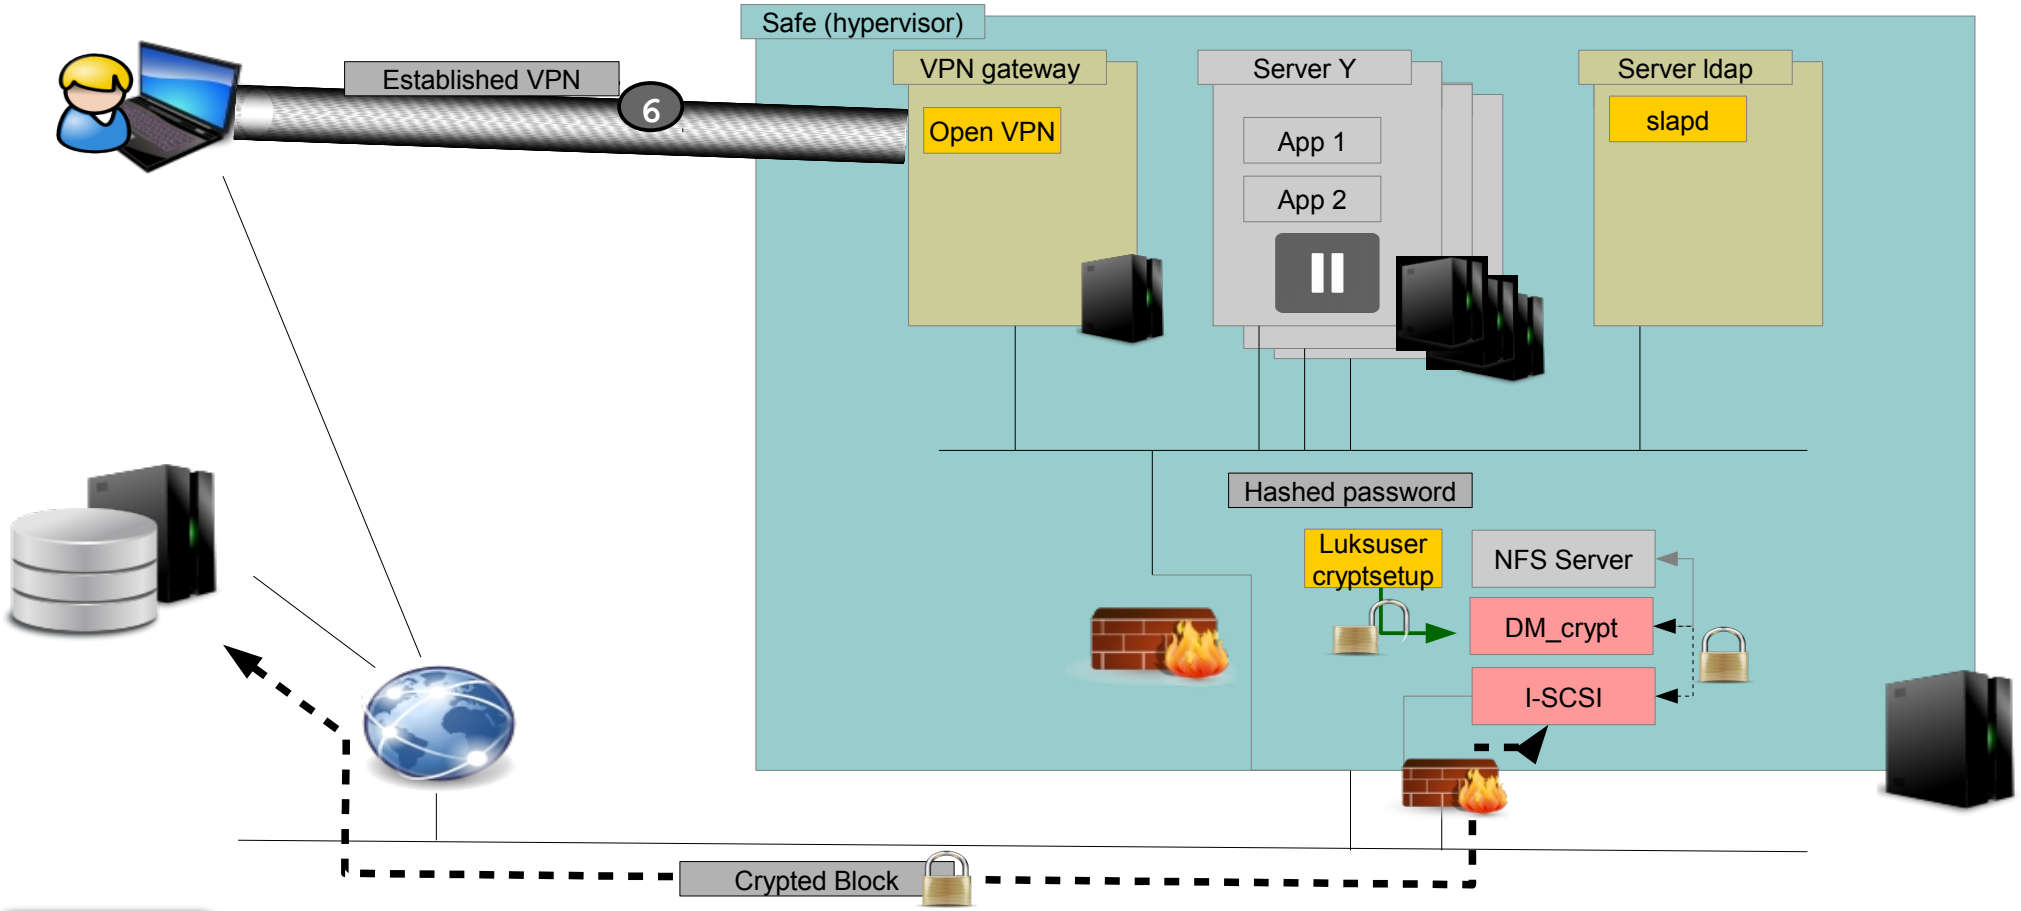
Safe (hypervisor)
Established VPN
6
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
Hashed password
Luksuser
cryptsetup
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

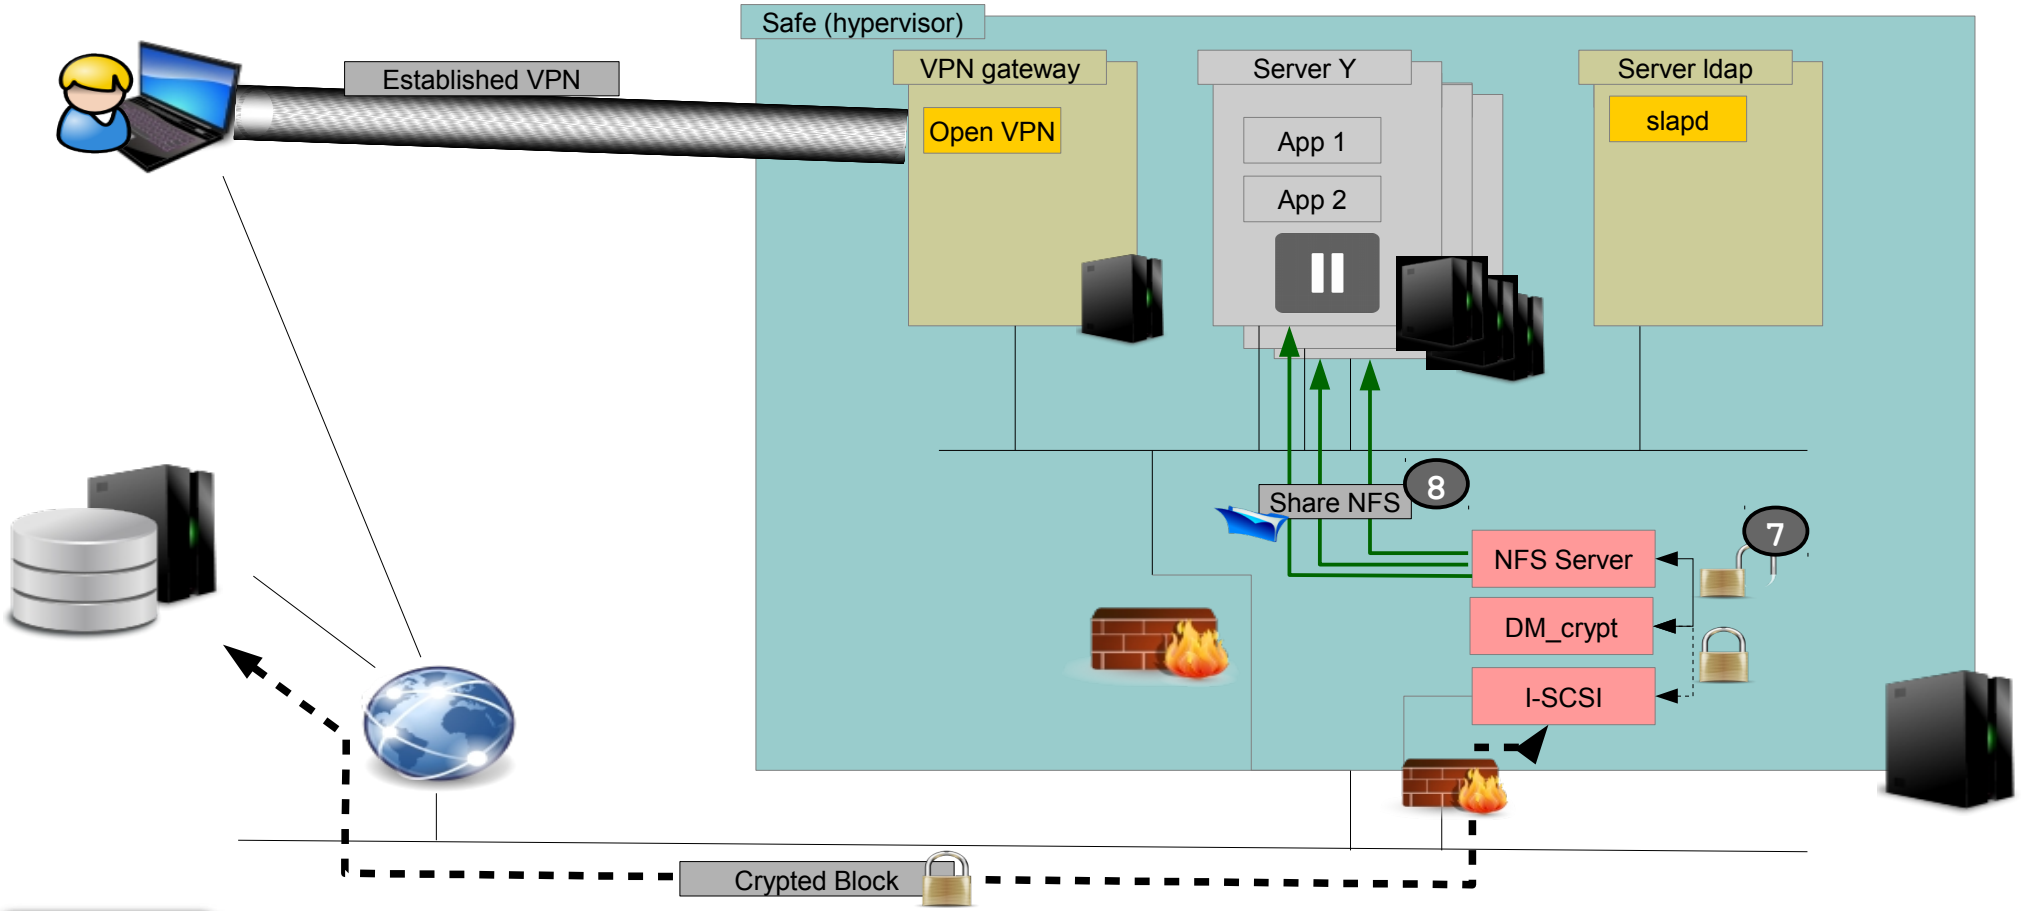
Safe (hypervisor)
Established VPN
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
Share NFS
8
7
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

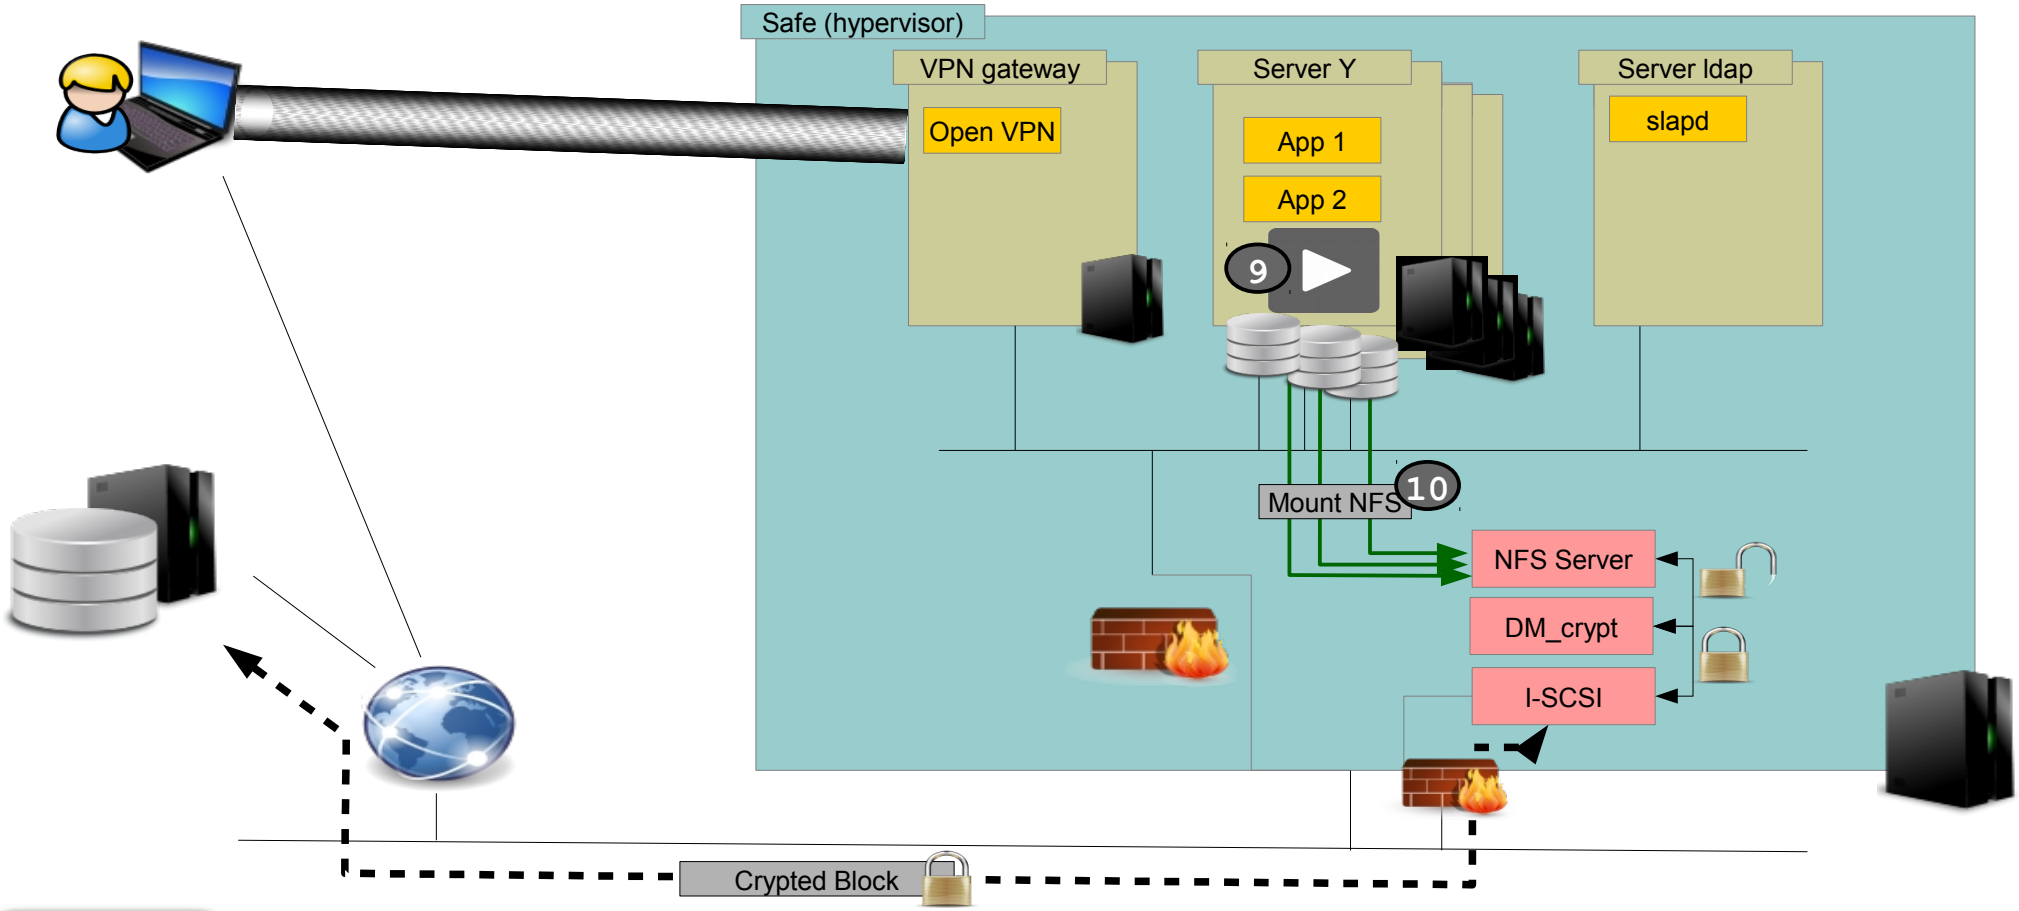
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
9
Server ldap
slapd
Mount NFS
10
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

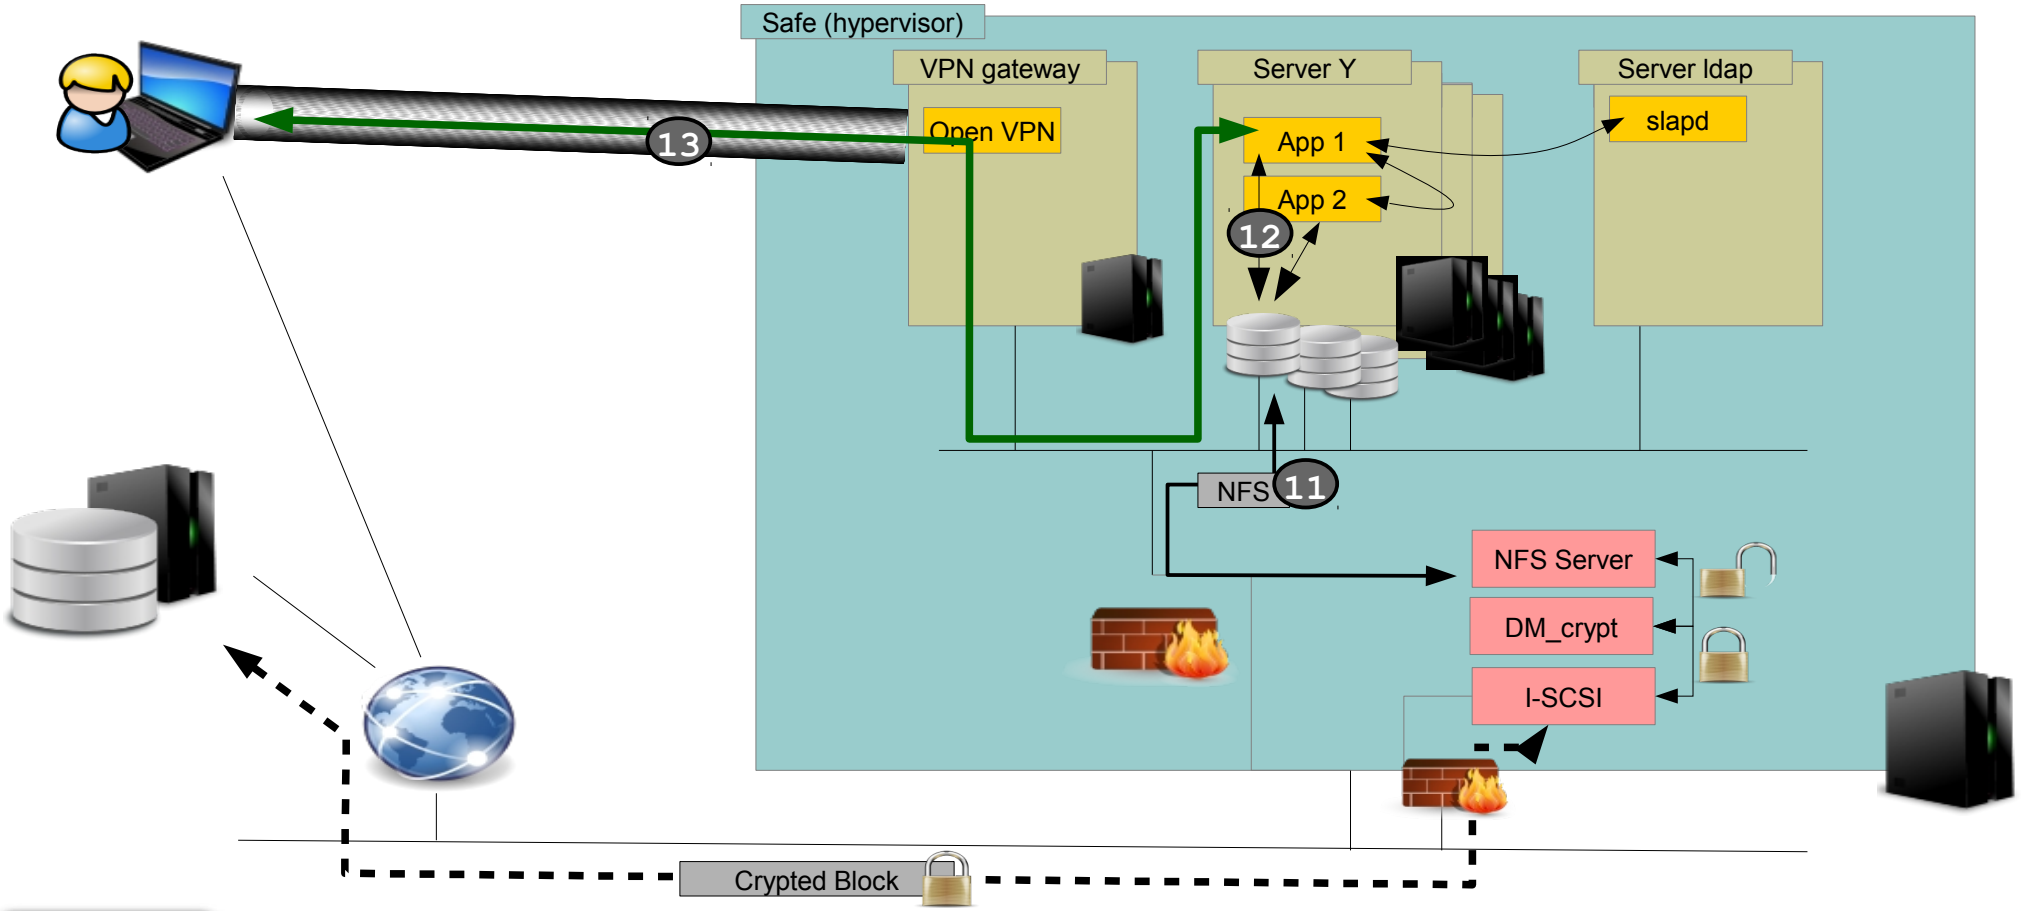
Safe (hypervisor)
VPN gateway
Open VPN
Server Y
App 1
App 2
Server ldap
slapd
13
12
11
NFS
NFS Server
DM_crypt
I-SCSI
Crypted Block

# How it works

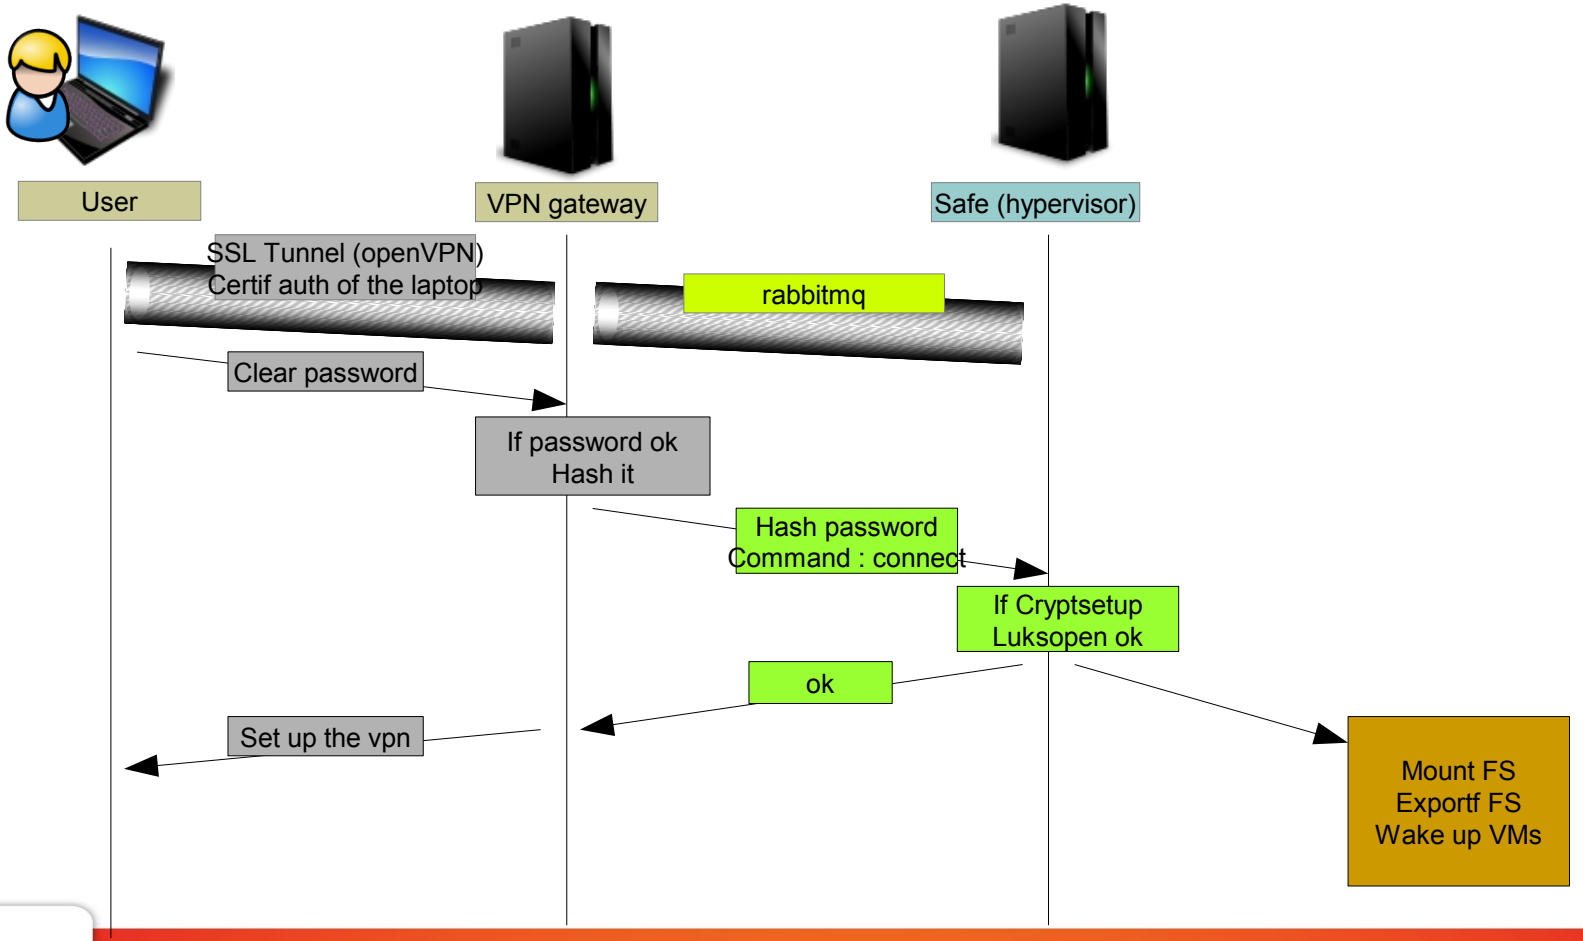
User
VPN gateway
Safe (hypervisor)
SSL Tunnel (openVPN)
Certif auth of the laptop
rabbitmq
Clear password
If password ok
Hash it
Hash password
Command : connect
If Cryptsetup
Luksopen ok
ok
Set up the vpn
Mount FS
Exportf FS
Wake up VMs

# Demo (... or not ...)

# 4

# Security considerations

# Security considerations

## Different levels (proposition)

| Access | Offer | Open/Close safe | Create user | sudo | sudo OpenVPN | Connection + sudo HyperVisor | Console HyperVisor | Create VM | Externalisation Logs |
|---|---|---|---|---|---|---|---|---|---|
| User | ZRR | Yes | No | Yes | No | No | No | No | Yes |
| | DR | Yes | No | Yes | No | No | No | No | Yes |
| | Paranoid | Yes | No | Yes | No | No | No | No | No |
| Admin | ZRR | Yes | Yes | Yes | No | No | No | Yes | Yes |
| | DR | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| | Paranoid | Yes | Yes | Yes | Yes | Yes | Yes | Yes | No |
| LHS Admin | ZRR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| | DR | No | No | No | No | No | Yes | No | Yes |
| | Paranoid | No | No | No | No | No | No | No | No |
| Puppet | ZRR | Yes automatically | | | | | | | |
| | DR | Yes manually | | | | | | | |
| | Paranoid | Yes manually | | | | | | | |

# Security considerations

## Hypervisor is considered safe

- Core of the solution, must be secured
- No network (or limited and isolated) access
- Ensure physical security

## At any time

- Only encrypted blocks on the network
- No key in memory on any VM.
  - Passwords are hashed on vpngateway, but not kept in memory.
- A compromised laptop can open the safe (after recording and replaying the user connection sequence and credentials, with keylogger for example)
- Firewalling on the network bridge of the hypervisor

# Security considerations

## When safe is open

- A hacked VM doesn't compromise the key,
  - But data leak is possible.
- A root on the hypervisor can catch LUKS key(s), and can see data.
- A user can extract data.
- ...

## When safe is closed

- No keys in memory or kernel, or on disk
- Only vpngateway can be attacked. Other VMs are suspended/paused.
- A root on the hypervisor can't mount FS and can't see data.
- ...

# Firewalling

## On the hypervisor

- Reverse Path Filtering
  - Avoid IP spoofing between the VMs
- Netfilter / iptables rules on the VMs bridge
  - Limit interactions between the VMs to the required minimum
  - Users will be able to add their own ones

## Virtual security appliance

- Firewall on a dedicated VM
- Permit to add mode functionalities
  - Traffic analysis, proxy...
- Many appliances or distributions (commercial or free)
  - Pfsense, shorewall, m0n0wall, smoothwall...

# 5

# Conclusion

## Future work

# Conclusion

## Not a simple vault

- Permits to **work** as a team on sensitive data
  - More than secure hosting / storage solutions
- Ensures data confidentiality and integrity
  - Virtually not limited in storage size
- Functional with all basic functionalities implemented
  - Looking for beta testers

# Future work

## Client side

- Dedicated client
  - Software, virtual appliance, dedicated hardware (raspberry)
  - More stuff for authentication in the client (OTP)
  - Some OpenVPN script modification to hash the password on the client side ?
- Site-to-site VPN ?
  - Would permit to bypass the hardware limitations
- Data leak Inspector on the vpngateway ?
  - Or on the virtual firewall ?
- Behavior analysis tools ?
- ...

## Authentication

- OTP with another device (e.g. cellphone and code sent by SMS)
  - Limited by the current version of OpenVPN that does not support dynamic challenges
- Do not use a simple hash of the password as LUKS key
- LDAP schema / cryptsetup modifications to use the user slot as authentication factor

# Future work

## Hypervisor side

- On-demand VMs
- Embedded HSM to separate keys outside the kernel ?
- More than one physical server
- More than one LUKS volumes
  - RW or RO
- RO user role
  - Open the vault to users than can not modify the data
- Backups and crash recovery
- Hardening the kernel
- Logs ex-filtration, log analysis tools

## Administration

- WEB interface
  - Vault state and statistics
  - Administration operations

# Thank you for your attention

# Contacts

## Technical staff

- Bertrand Wallrich – bertrand.wallrich@inria.fr
- Frédéric Beck – frederic.beck@inria.fr
- Loïc Rouch – loic.rouch@inria.fr
- http://lhs.loria.fr