RCauth Online CA service

Distributed operations and plans

🔗 rcauth.eu

Dissemination level: Public

@eoscfuture

EOSC-hub receives funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 777536.

Reminder - for n00bs and people with dodgy memories 1/2

- RCauth is an IGTF accredited IOTA (DOGWOOD class) CA
  - Online credential conversion
  - Connected to eduGAIN (R&S+Sirtfi) plus direct, e.g. EGI Check-in and eduTEAMS
- EOSC Hub and EOSC Future implementing a **High Availability setup across 3 sites**

Reminder - for n00bs and people with dodgy memories 2/2

- Private key is cloned and will be hosted in HSMs at each site
- Cloning is done by XORing key with random strings
- OTP randomness exchanged using different means (usually in-person)
- => key is 3-of-3 encrypted in transit
    - Any part, or any two of the three, will have *no information* about the key
    - Each part was transferred by different (trusted) means

Outline

- PMA and Ops membership
- Review of tasks
  - Key cloning
  - Deployment
  - HA Database (and network)
  - HA Networking
  - HA WAYF
  - Documentation
- Renaming the CA
- Site specific reports
- Q&A

PMA and Operations membership

RCauth Ops (alphabetical order):

- Will Furnell (STFC)
- Kyriakos Gkinis (GRNET)
- Jens Jensen (STFC)
- Nicolas Liampotis (GRNET)
- Mischa Sallé (Nikhef)

RCauth PMA

- Chaired by David Groep
- GRNET member need updating: Kostas replaces Nicolas

RCauth.eu

Activity overview (EOSC Hub/Future view)

- **Operational tooling**
  - Operator comms (205, 206)
  - Self audit (207)

- **High Availability setup** - run across NIKHEF, GRNET, STFC
  - Key cloning (201)
  - Deployment (202)
  - HA Database (203)
  - HA testing (204)
  - HAProxy frontends (229)
  - HA Network (232)

- **Operations**
  - Acceptance instance (228)
  - Service integration (208)
  - End user docs (209)
  - Monitoring docs (210)
  - Final PMA review (211)

Management of tasks

- JIRA dashboard

- Regular weekly ops calls for reviewing/planning

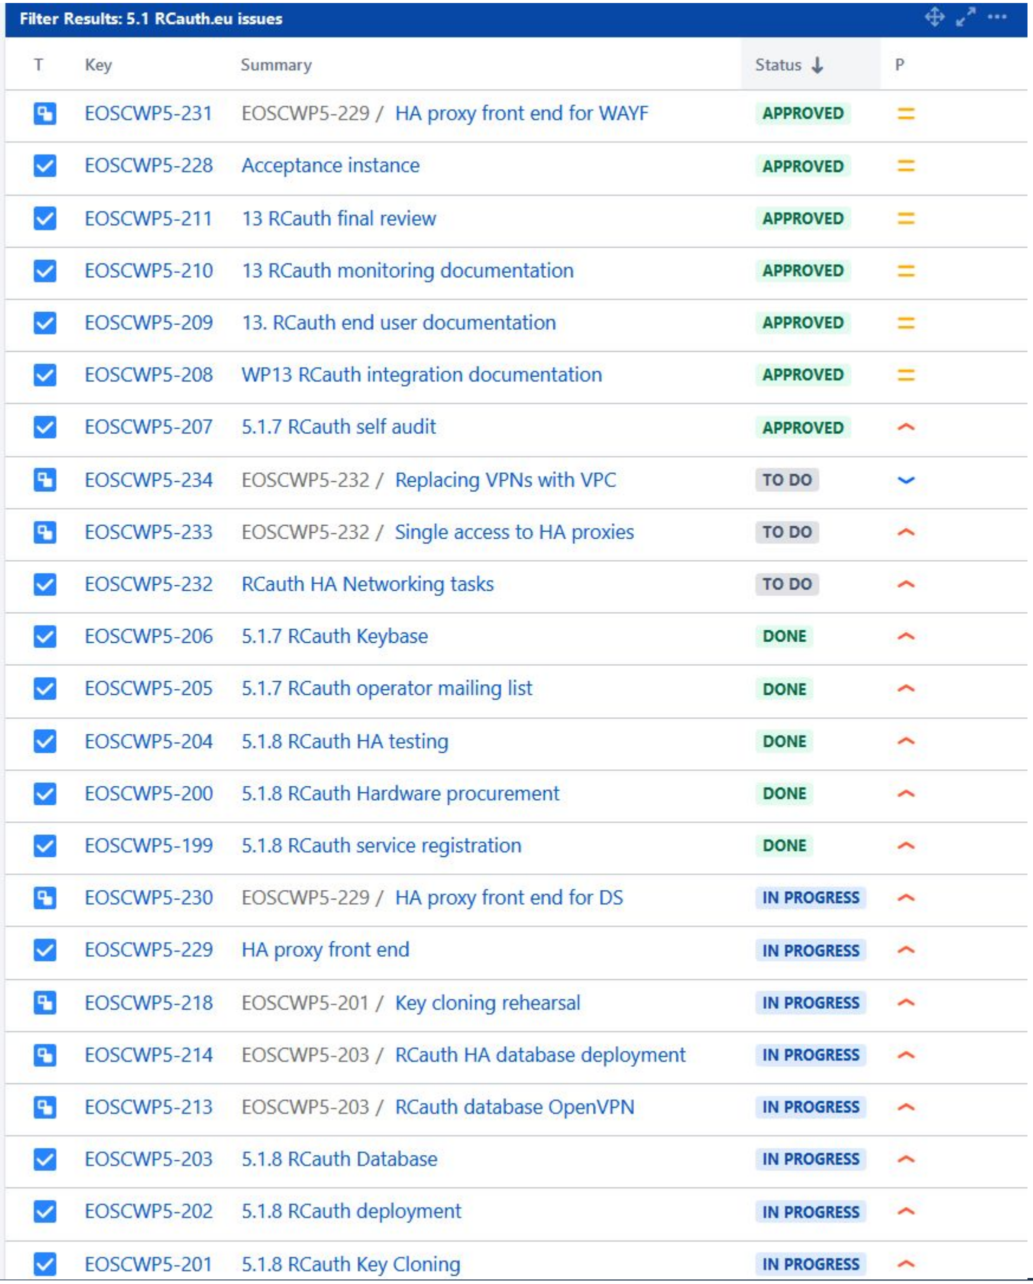

Filter Results: 5.1 RCauth.eu issues

T	Key	Summary	Status ↓	P
⧉	EOSCWP5-231	EOSCWP5-229 / HA proxy front end for WAYF	APPROVED	=
☑	EOSCWP5-228	Acceptance instance	APPROVED	=
☑	EOSCWP5-211	13 RCauth final review	APPROVED	=
☑	EOSCWP5-210	13 RCauth monitoring documentation	APPROVED	=
☑	EOSCWP5-209	13. RCauth end user documentation	APPROVED	=
☑	EOSCWP5-208	WP13 RCauth integration documentation	APPROVED	=
☑	EOSCWP5-207	5.1.7 RCauth self audit	APPROVED	⌃
⧉	EOSCWP5-234	EOSCWP5-232 / Replacing VPNs with VPC	TO DO	⌄
⧉	EOSCWP5-233	EOSCWP5-232 / Single access to HA proxies	TO DO	⌃
☑	EOSCWP5-232	RCauth HA Networking tasks	TO DO	⌃
☑	EOSCWP5-206	5.1.7 RCauth Keybase	DONE	⌃
☑	EOSCWP5-205	5.1.7 RCauth operator mailing list	DONE	⌃
☑	EOSCWP5-204	5.1.8 RCauth HA testing	DONE	⌃
☑	EOSCWP5-200	5.1.8 RCauth Hardware procurement	DONE	⌃
☑	EOSCWP5-199	5.1.8 RCauth service registration	DONE	⌃
⧉	EOSCWP5-230	EOSCWP5-229 / HA proxy front end for DS	IN PROGRESS	⌃
☑	EOSCWP5-229	HA proxy front end	IN PROGRESS	⌃
⧉	EOSCWP5-218	EOSCWP5-201 / Key cloning rehearsal	IN PROGRESS	⌃
⧉	EOSCWP5-214	EOSCWP5-203 / RCauth HA database deployment	IN PROGRESS	⌃
⧉	EOSCWP5-213	EOSCWP5-203 / RCauth database OpenVPN	IN PROGRESS	⌃
☑	EOSCWP5-203	5.1.8 RCauth Database	IN PROGRESS	⌃
☑	EOSCWP5-202	5.1.8 RCauth deployment	IN PROGRESS	⌃
☑	EOSCWP5-201	5.1.8 RCauth Key Cloning	IN PROGRESS	⌃

Review of tasks: Database (task 203)

Networking subtask

Interim solution - VPN

- Eventually should have dedicated VPC
- Protected with dedicated single-use PKI
- Databases accessible only over VPN
- Each site runs a VPN server & client

VPN setup works well:

- good stability and we've even seen automatic failover

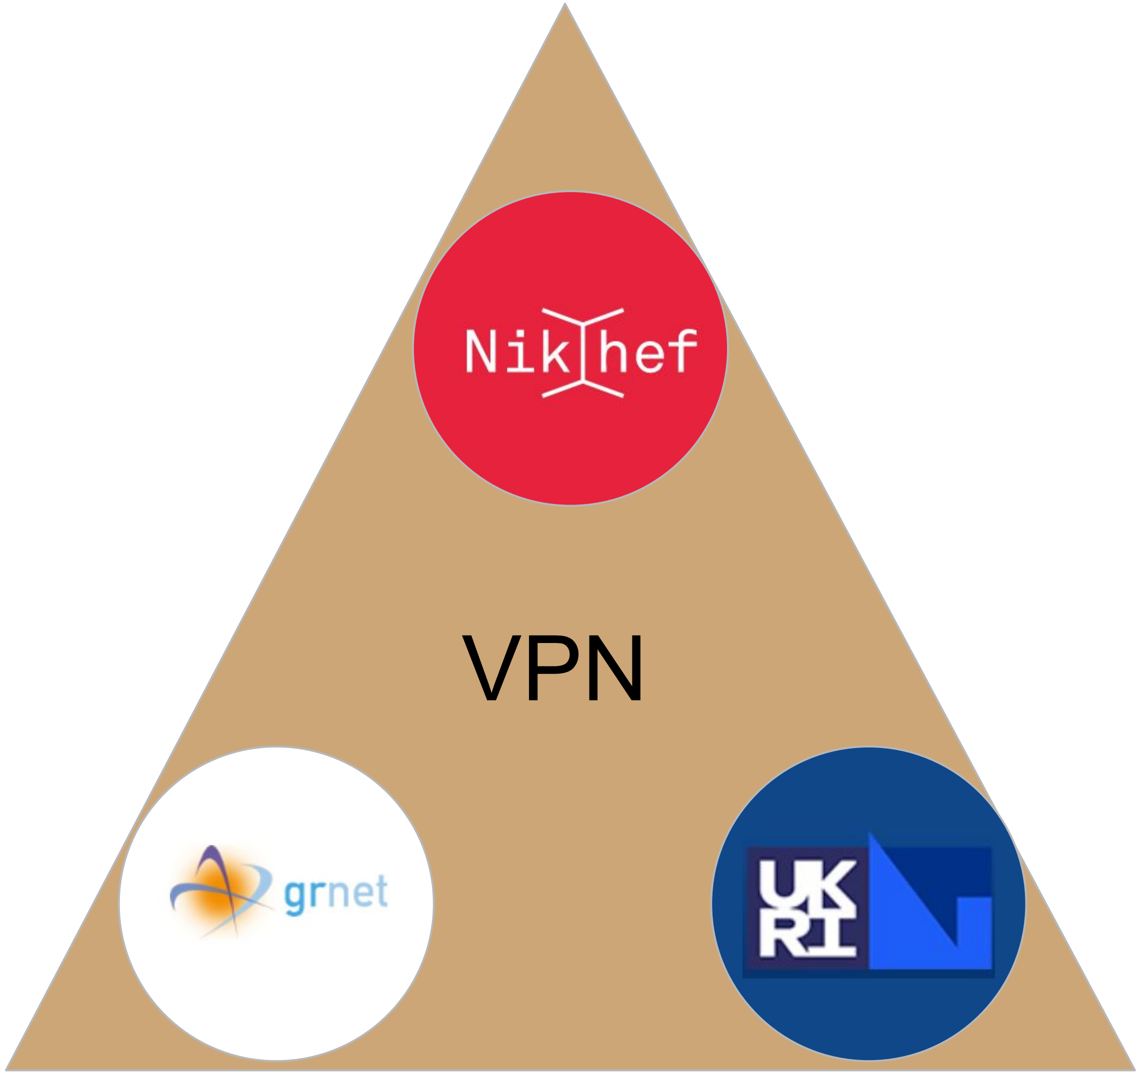

Review of tasks: Key cloning (task 201) - 1/4

Cloning steps 1/2

Note: This task is the one most affected by the current lockdown

- Agree plan with PMA [STFC, NIKHEF, GRNET] - **DONE**
- Develop software [STFC,NIKHEF] - **DONE**
- Generate secret A [STFC] - **DONE**
- Exchange A with NIKHEF [STFC] - **DONE**
- Share recipe for generating random numbers in HSM with GRNET [NIKHEF, STFC] - **DONE**
- Generate secret B [GRNET] - **DONE**
- Select additional methods for sharing keys - courier/snailmail, keybase or PGP email - **DONE**
- Exchange B with NIKHEF [GRNET] - **DONE**

...

Review of tasks: Key cloning (task 201) - 2/4

Cloning steps 2/2

- ...
- Generate C1 [NIKHEF]
- Exchange C1 with STFC [NIKHEF]
- Generate C2 [NIKHEF]
- Exchange C2 with GRNET [NIKHEF]
- Calculate S1 = S+A+C1 [NIKHEF]
- Exchange S1 with STFC [NIKHEF]
- Calculate S2 = S+B+C2 [NIKHEF]
- Exchange S2 with GRNET [NIKHEF]
- Calculate S from S1 [STFC]
- Install S in HSM [STFC]
- Calculate S from S2 [GRNET]
- Install S in HSM [GRNET]

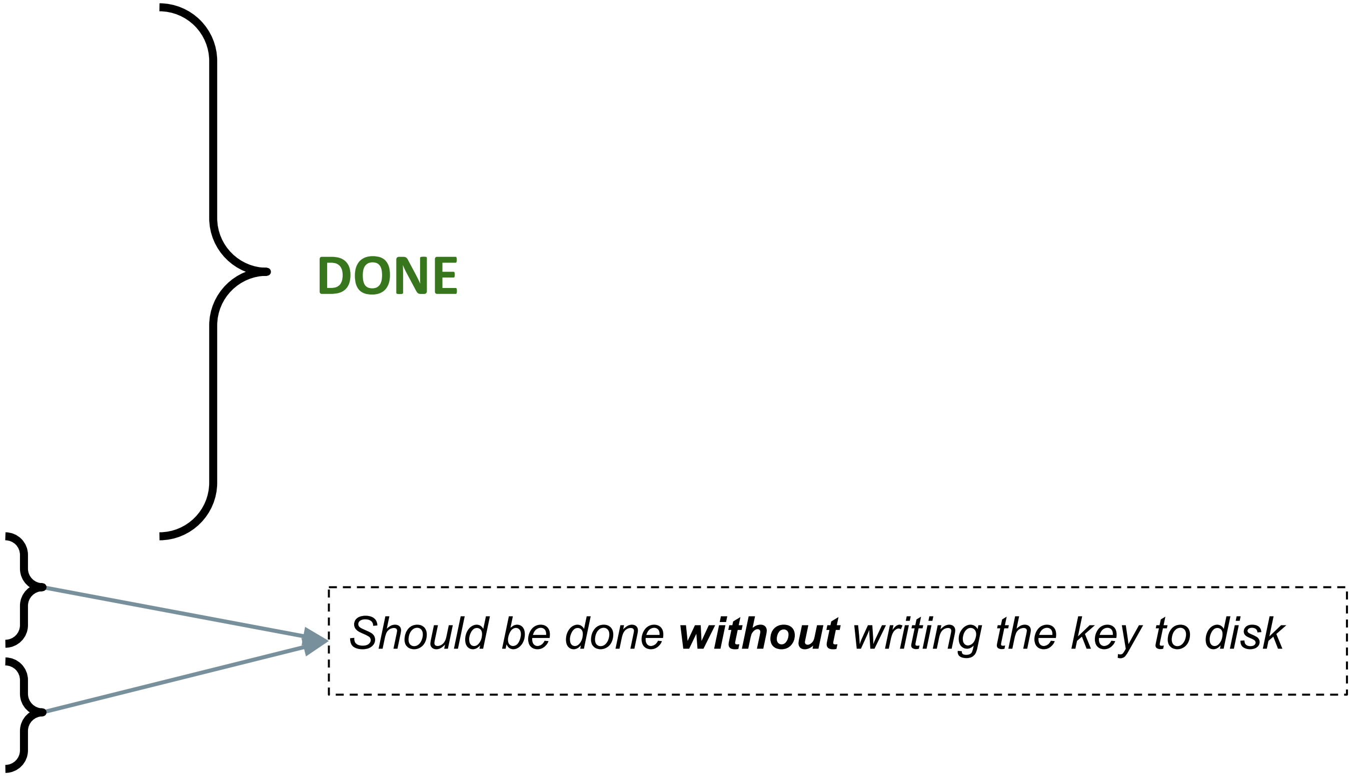

DONE

*Should be done **without** writing the key to disk*

Review of tasks: Key cloning (task 201) - 3/4

Methods used for random data exchange

- In person exchange of random data (pre-lockdown)
  - Written to portable and destructible media (CD, paper)
  - Paper is only machine readable with OCR…
- Sending random data via courier
  - GRNET sent its data to Mischa’s home during lockdown
- Keybase (self-destructing) exchange of dry run random data
- PGP-encrypted mail
  - Used for dry run
  - Used also for final secret

Review of tasks: Key cloning (task 201) - 4/4
Lessons learnt

- Hand-written secrets can be difficult
- Exchanging self-destructing messages over keybase
- Need python to de-/reconstruct keys in a portable way
- Python scripts written to support multiple versions:
  - Python is a very volatile language
  - Need to work with system default (particularly on offline systems)
  - Many features from python could not be used
- To keep things in memory have to be creative (e.g .p12 -> unencrypted RSA key input)

Review of tasks: Deployment (task 202)

1. Package/containerise software [NIKHEF] - **DONE**
2. Generate deployment recipe (ansible) [NIKHEF] - **DONE**
3. Set up infrastructure [STFC] - **DONE**
4. Set up infrastructure [GRNET] - **DONE**
5. Deploy delegation server [STFC]- **DONE**
6. Deploy delegation server [GRNET] - **DONE**
7. Access keybase git and deploy MyProxy/signing on infrastructure [STFC] - **DONE**
8. Access keybase git and deploy MyProxy/signing on infrastructure [GRNET] - **DONE**

Review of tasks: Database (task 203) 1/2

1. Generate OpenVPN recipe [STFC, NIKHEF, GRNET] - **DONE**
2. Set up VPN endpoint [STFC] - **DONE**
3. Set up VPN endpoint[GRNET] - **DONE**
4. Set up VPN endpoint [NIKHEF] - **DONE**
5. VPN functional tests [all] - **DONE**
6. VPN performance tests [all] - **DONE**
7. VPN monitoring [all] - **DONE**
8. Database deployment recipe [NIKHEF] - **DONE**
9. Database synchronisation configuration [NIKHEF] - **DONE**
10. Deploy database [STFC] - **DONE**
11. Deploy database [GRNET] - **DONE**
12. Database monitoring [STFC] - **DONE**
13. Database monitoring [GRNET] - **DONE**
14. Set up synchronisation [STFC] - **DONE**
15. Set up synchronisation [GRNET] - **DONE**
16. Database synchronisation testing [NIKHEF] - **DONE**

More on the next slide

Review of tasks: Database (task 203) 2/2

- Database connection done over OpenVPN
- Galera cluster over secure VPN
- Lots of resilience testing done:
  - jump between the 3 Delegation Servers in middle of OIDC flow

    (using e.g. https://github.com/msalle/test_oidc_client and via Round-Robin setup in the HAproxy)
  - Database sync is faster than client's HTTPS
  - Did NOT manage to break it (-:
- Also need an HA database for Shibboleth session:
  - reuse same HA MariaDB different DB
  - documentation not very clear and sometimes even wrong
  - seems to work now (some warnings though)

Final(ish) Architecture (1/3)

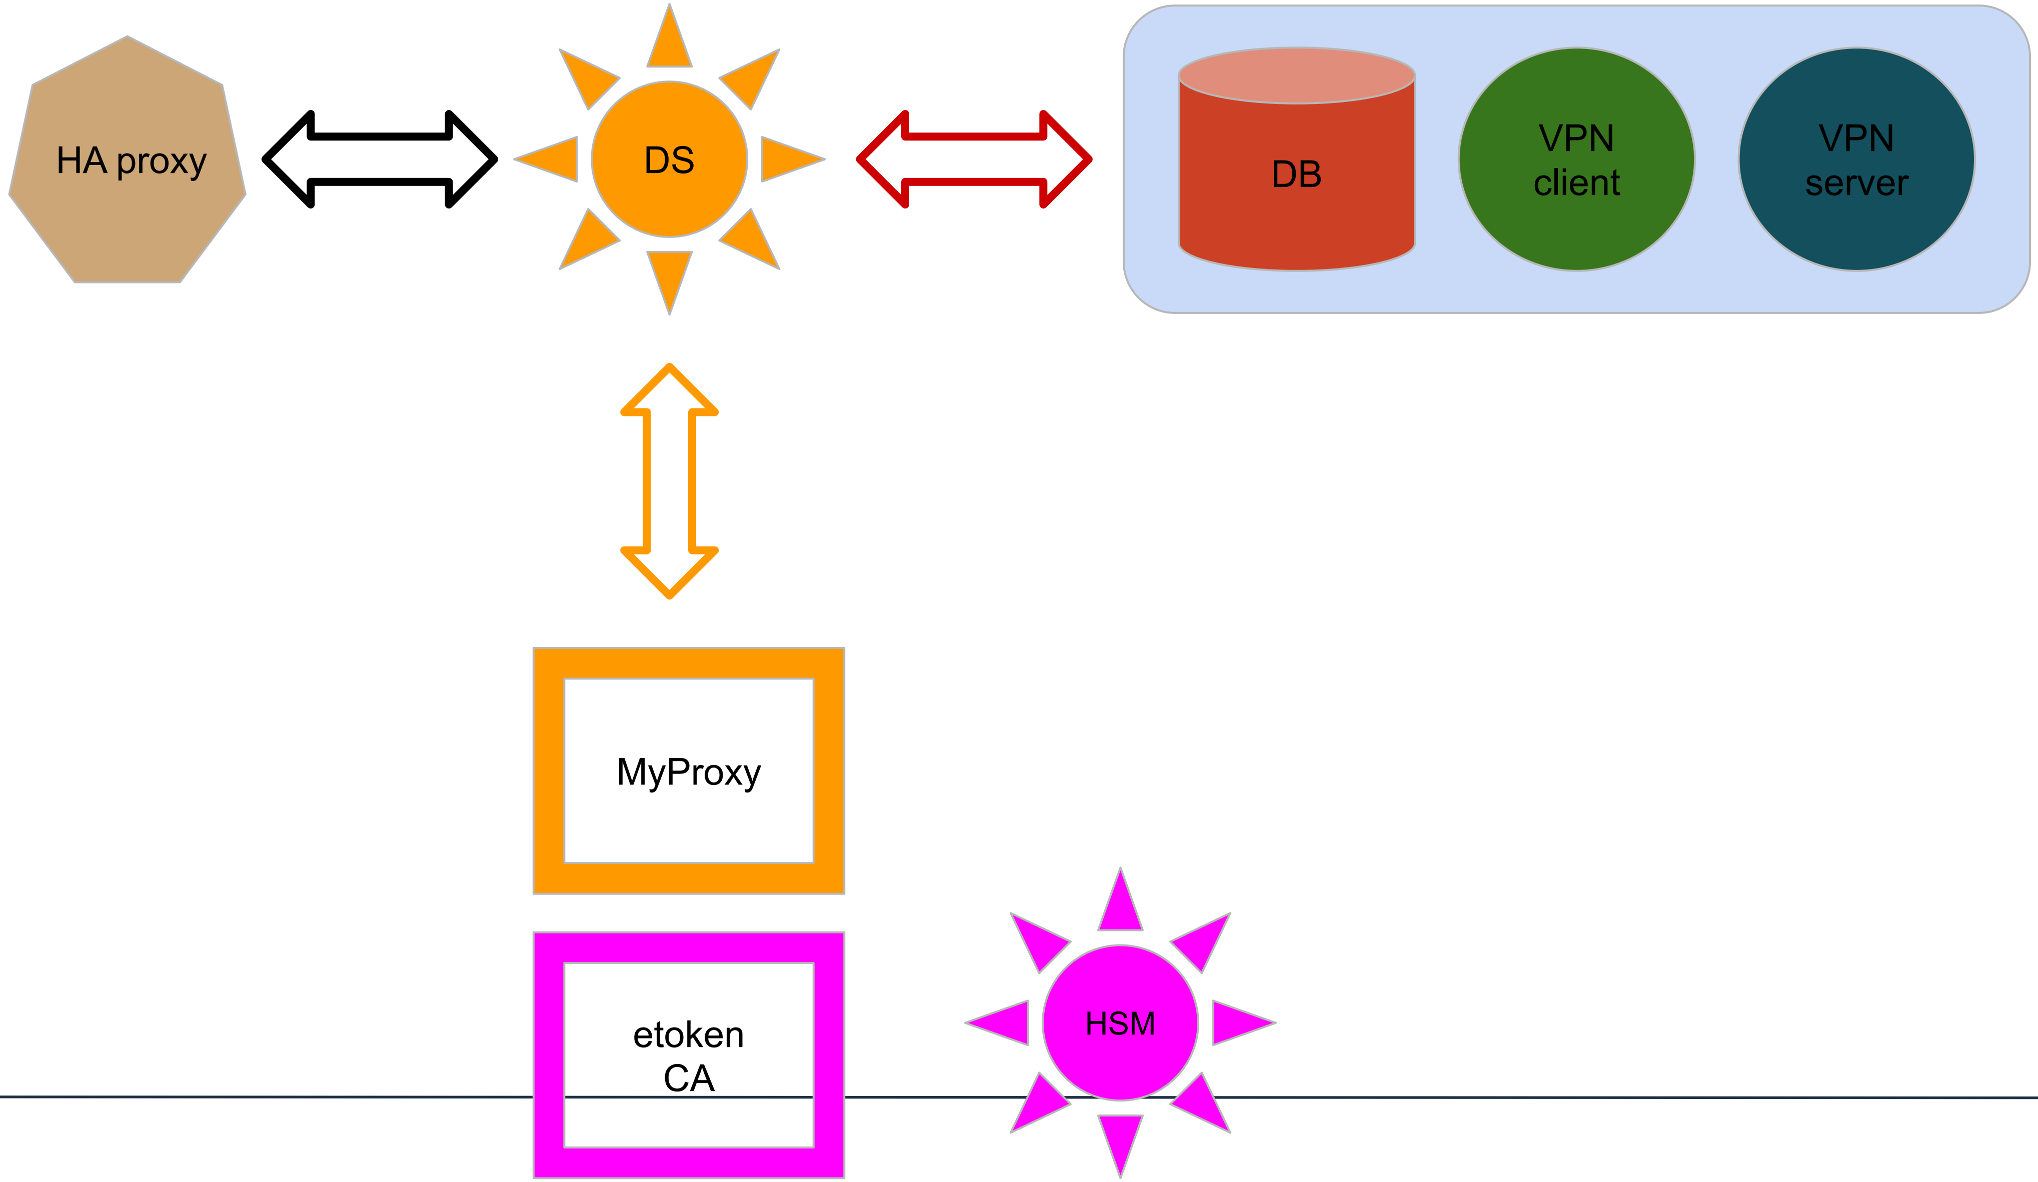

RCauth.eu

Final(ish) Architecture (2/3)

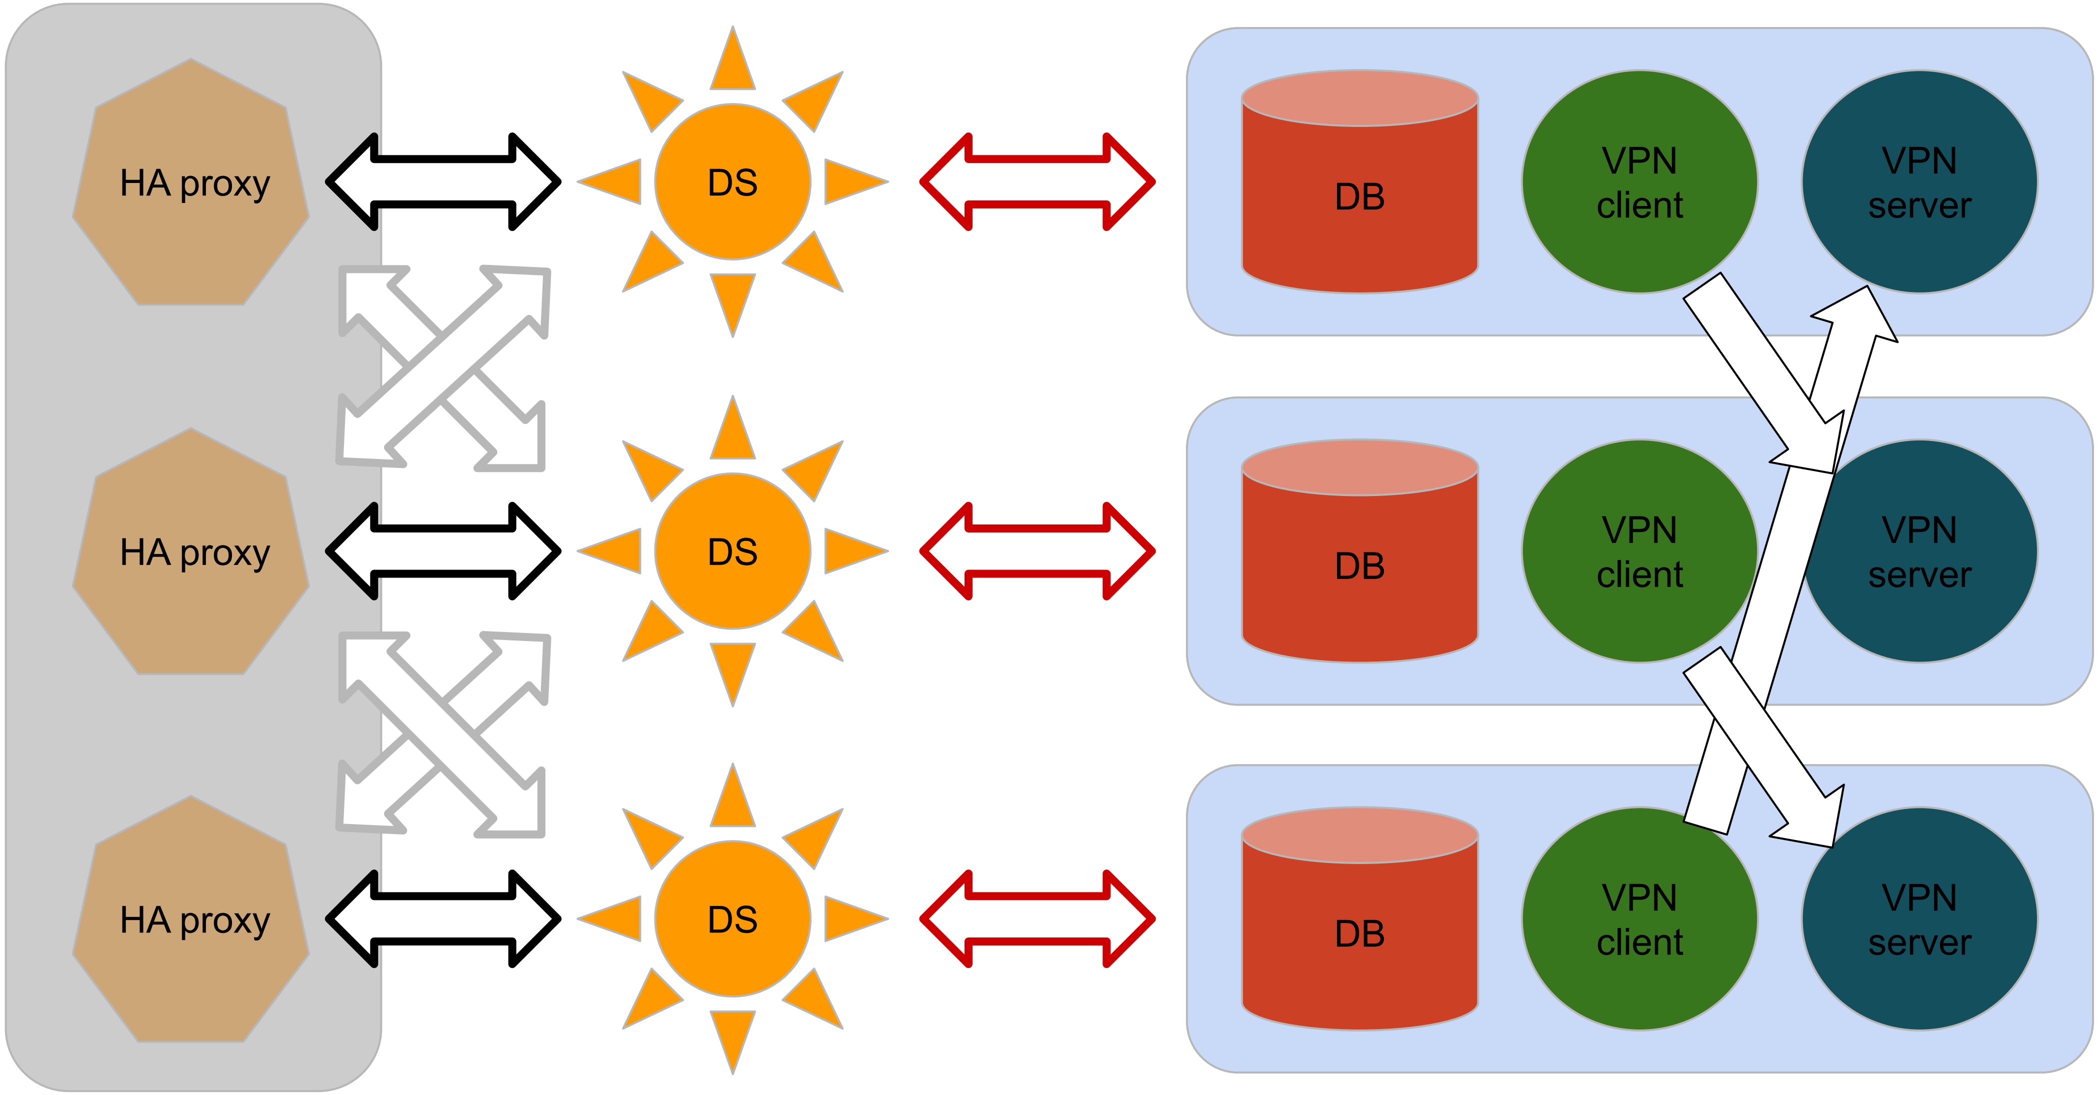

Final(ish) Architecture (3/3)

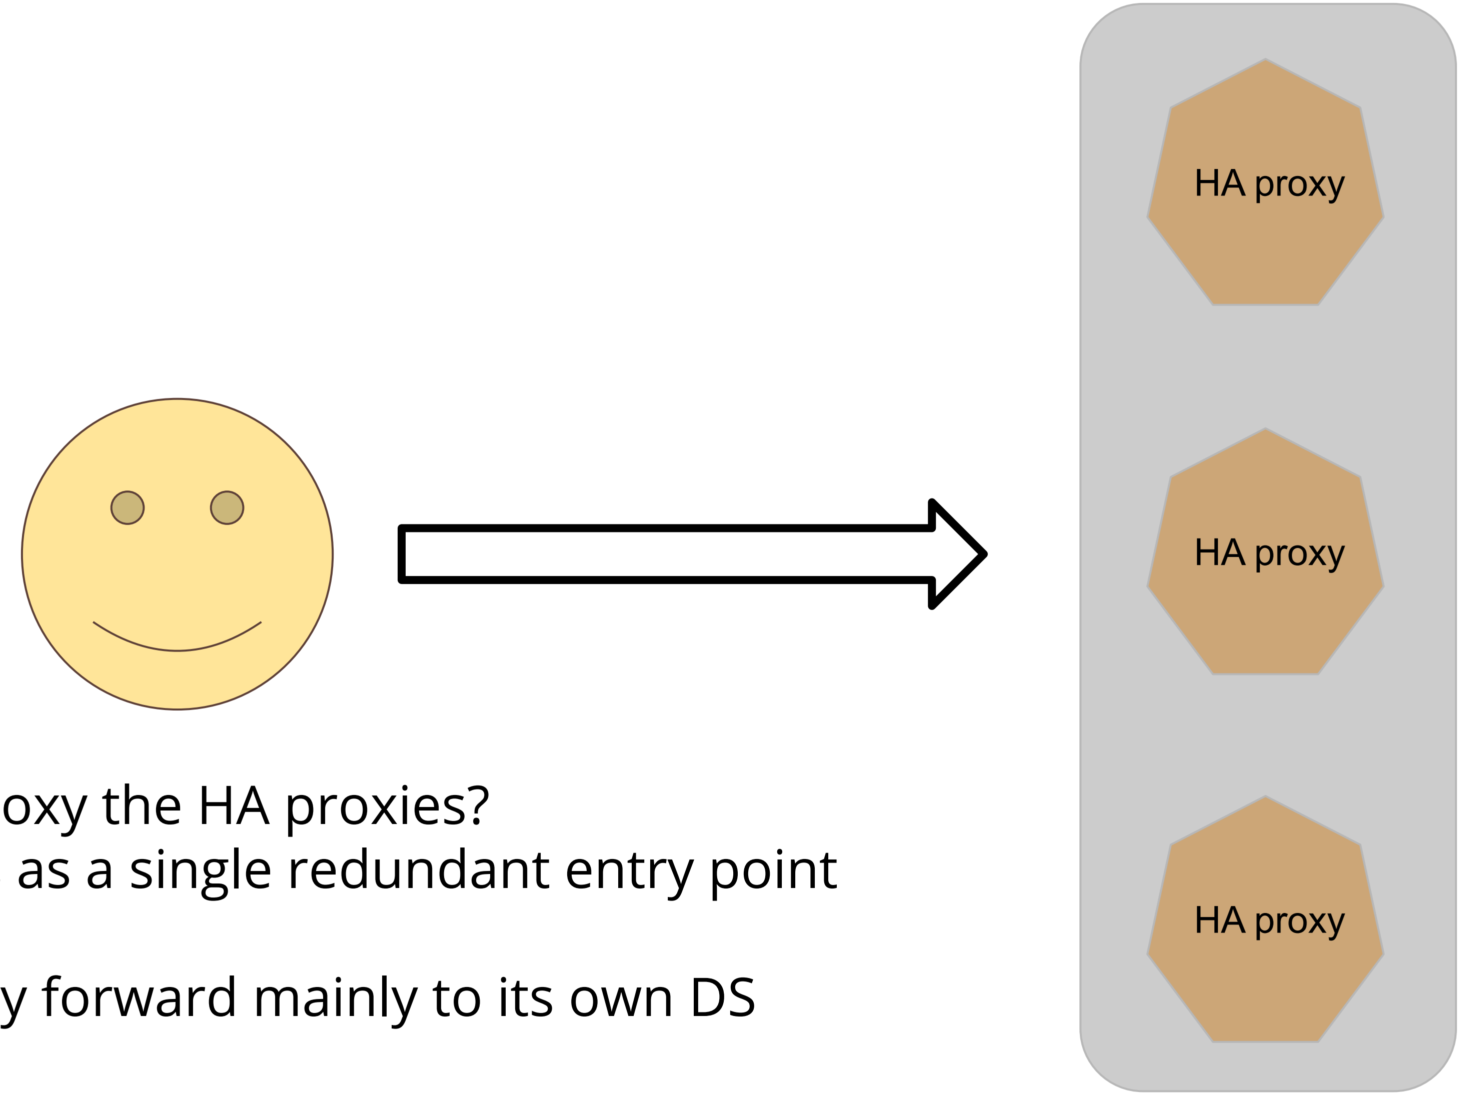

How to HA proxy the HA proxies?
I.e. it appears as a single redundant entry point

Each HA proxy forward mainly to its own DS

Renaming the CA 1/3

‘There are only two hard things in Computer Science: cache invalidation and naming things.’

- Phil Karlton

Renaming the CA 2/3

Current:

- Name in CP/CPS: “Research and Collaboration Authentication Pilot Issuing CA”
- DNS: `pilot-ca1.rcauth.eu`
- RCauth WAYF metadata in eduGAIN:
  - Name: “RCauth Pilot Online CA
  - Description: “RCauth Pilot Online CA for providing end-user proxy certificates to Science Gateways and other portals”

Production:

- Name in CP/CPS: “Research and Collaboration Authentication ~~Pilot~~ Issuing CA”
- DNS: ~~`pilot-`~~`ca1.rcauth.eu`
- RCauth WAYF metadata in eduGAIN:
  - Name: “RCauth ~~Pilot~~ Online CA
  - Description: “RCauth ~~Pilot~~ Online CA for providing end-user proxy certificates to Science Gateways and other portals”

Renaming the CA 3/3

Renaming side effects:

- Migration of MasterPortals:
  - Probably all need to be done simultaneously
- Update of RCauth WAYF metadata in eduGAIN
  - Could be done independently from the hostname update

RCauth in Production

Need to switch to production key & database all *at the same time*

- Databases must synchronise Nikhef’s production database:
  - Nikhef test DB needs to disconnect
  - STFC and GRNET need to kill their database
  - Nikhef prod DB needs to start cluster
  - STFC and GRNET rejoin and synchronize
- Revocation & CRL issuance needs to be extended to all sites
- Need single DNS to redirect to all HA proxies

Site specific reports

Site Specific Reports - Nikhef

Virtually no downtime

- One short unexpected network outage at Nikhef 19th May
- RCauth offline for about ½ hour

Local HA setup (see next slide) still very useful

Site Specific Reports - Nikhef: Local HA setup

- Created local HA setup to ease recovery from hardware failure:
  - duplicate backend nodes (i.e. 2 NUCs)
  - duplicate frontend nodes (i.e. 2 Delegation Servers)
  - all 4 on private LAN
  - automatic failover in case of failure of one of the backend nodes
  - 2nd frontend node probably hot spare for now
  - very useful for maintenance

- Could add both Delegation Servers to the European-wide HA setup

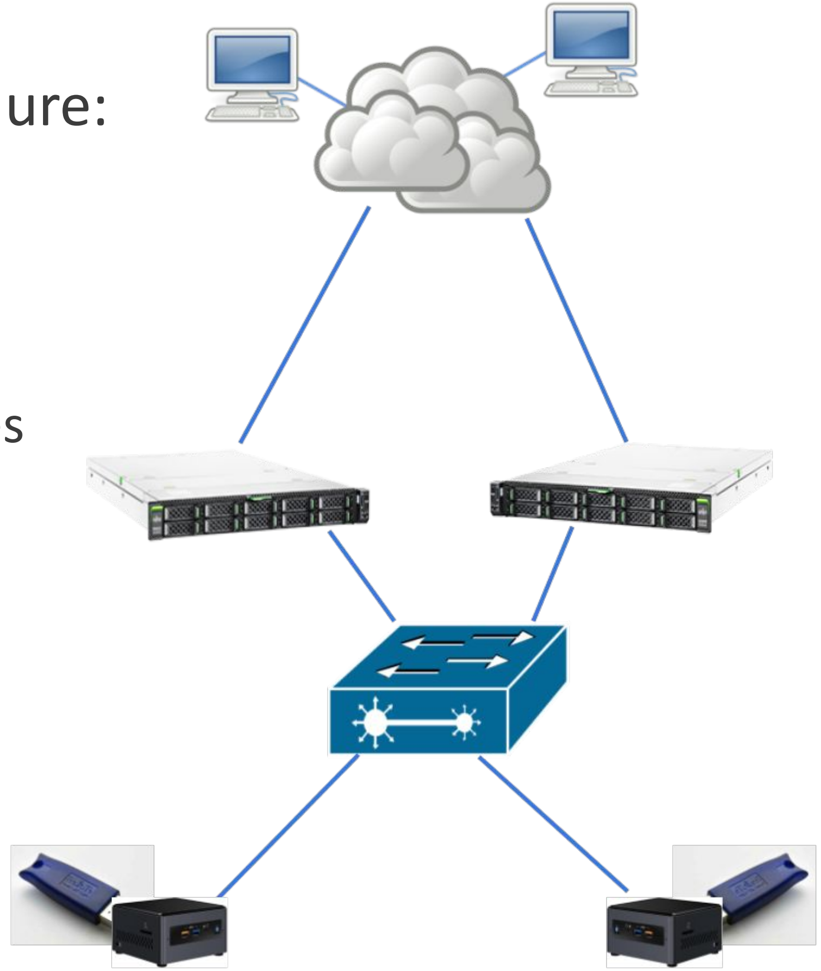

Site specific reports: GRNET

- Connected MyProxy service with the HSM device.
    - Used test certificate and private key stored in the HSM to sign certificates
- Installed and configured an HAProxy service in front of the delegation server
- Exchanged secret data with NIKHEF

Site Specific Reports - STFC

COVID related:

- Access to machine room still restricted (as of Jun 2021)

What’s good?

- Remote operations have gone well
- Will Furnell has picked up sysadmin of UK eScience and RCauth infra incl HSMs

What’s bad?

- Somewhat temperamental site firewall? (recent upgrade was not 100% smooth)

Thank you for your attention!

Questions?

Contact

RCauth Operations team
`ops-management(AT)rcauth.eu`

🔗 rcauth.eu 🐦 @eoscfuture