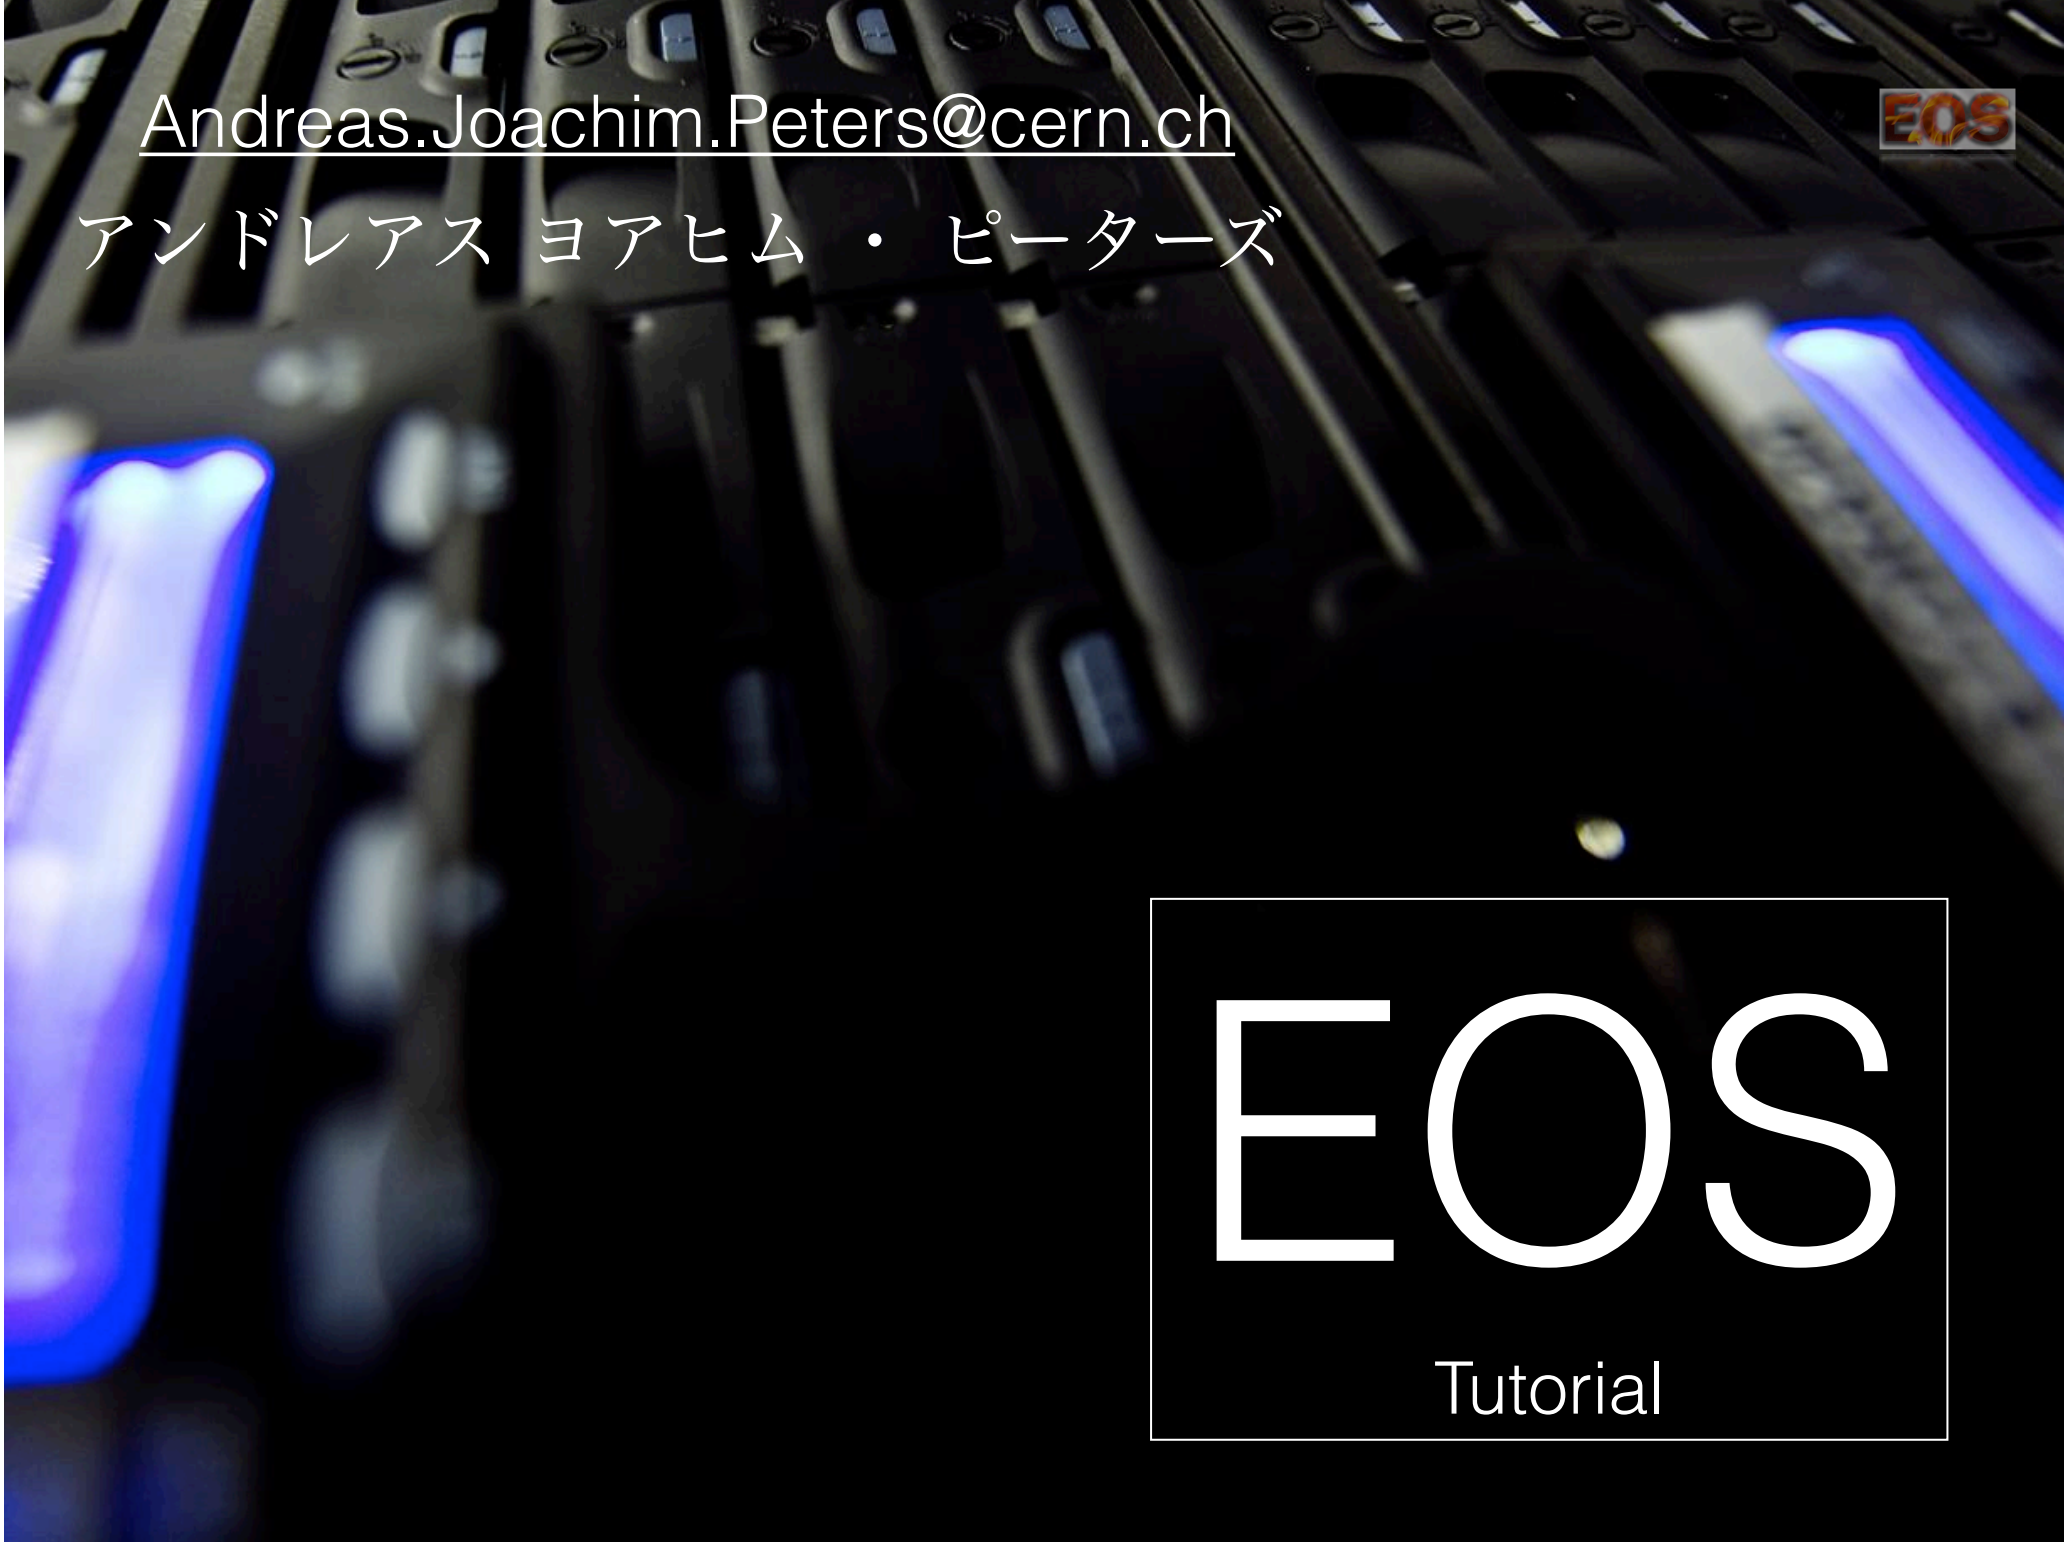

Andreas.Joachim.Peters@cern.ch

アンドレアス ヨアヒム ・ ピーターズ

EOS

Tutorial

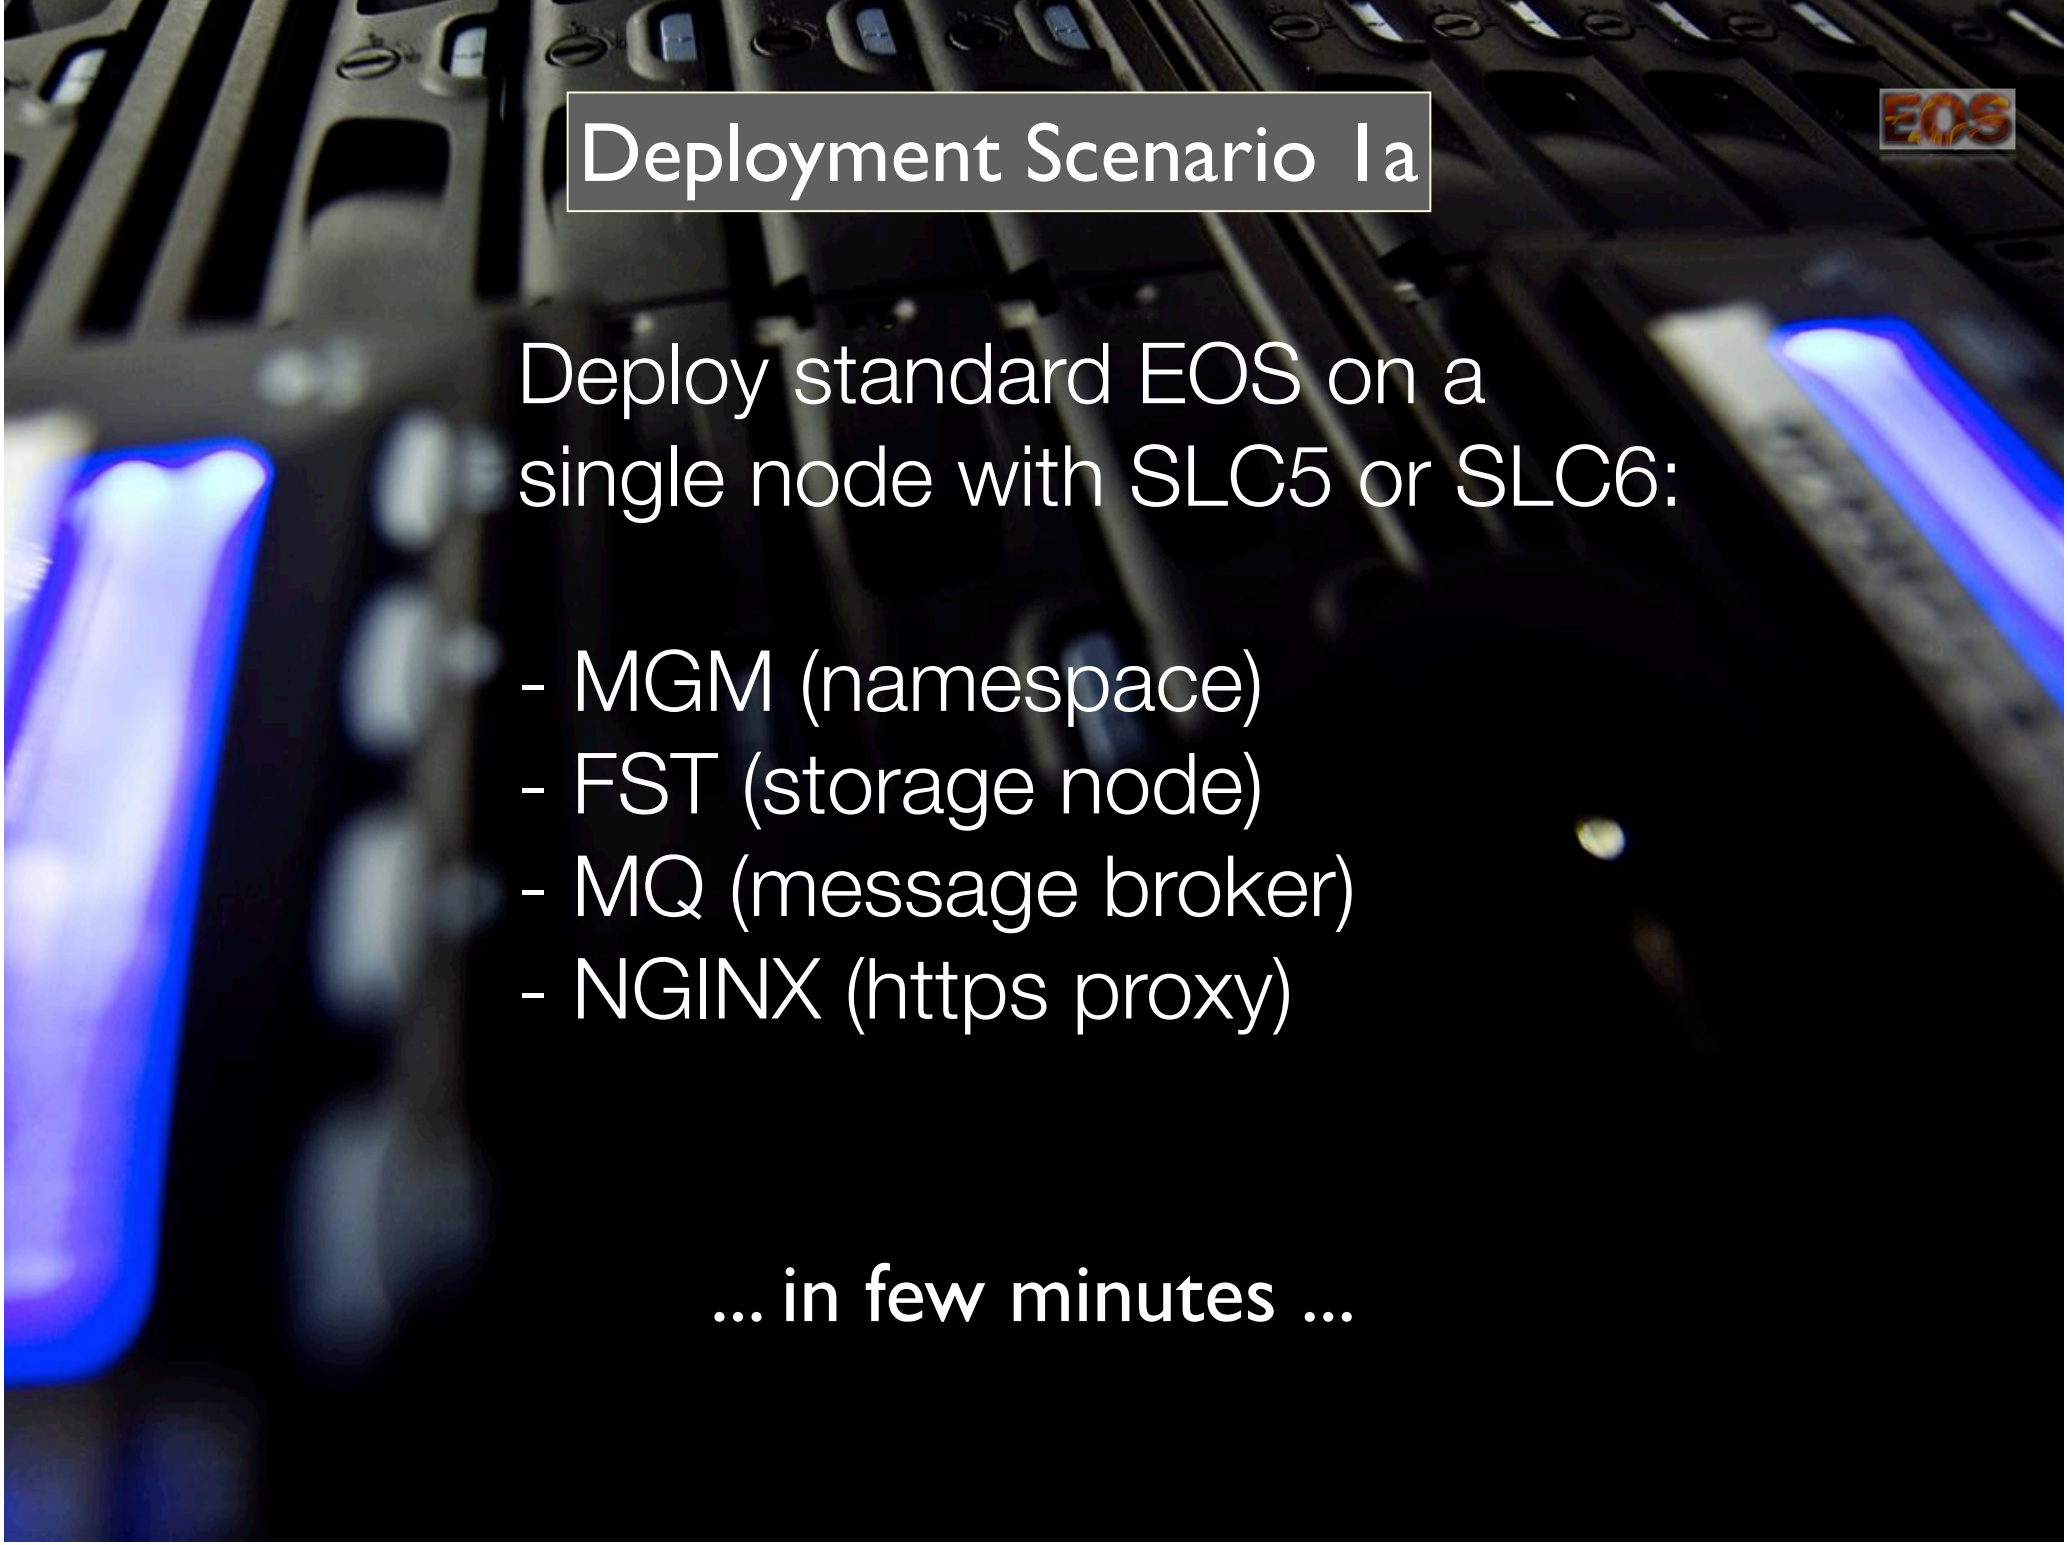
Deployment Scenario 1a
Deploy standard EOS on a
single node with SLC5 or SLC6:
- MGM (namespace)
- FST (storage node)
- MQ (message broker)
- NGINX (https proxy)
... in few minutes ...

Preparation

Download EOS-Deploy
http://eos.cern.ch/rpms/eos-deploy

```
wget http://eos.cern.ch/rpms/eos-deploy
chmod u+x eos-deploy
```

If you want to use other storage partions than / create for each of them a
symbolic link:

```
/var/eos/fs/0 => /data0
/var/eos/fs/1 => /data1
/var/eos/fs/2 => /data2
/var/eos/fs/3 => /data3
aso.
```

You can also create these links afterwards ...

... in few minutes ...

Installation 1a

Run the tool ...

```
./eos-deploy
```

Answer few question:

```
instance name     : test.foo.bar
2nd mgm name      : <return>
email             : yourmail@foo.bar
SLC flavour       : 6
#filesystems      : 4
storage servers   : <return>
```

... just wait ...

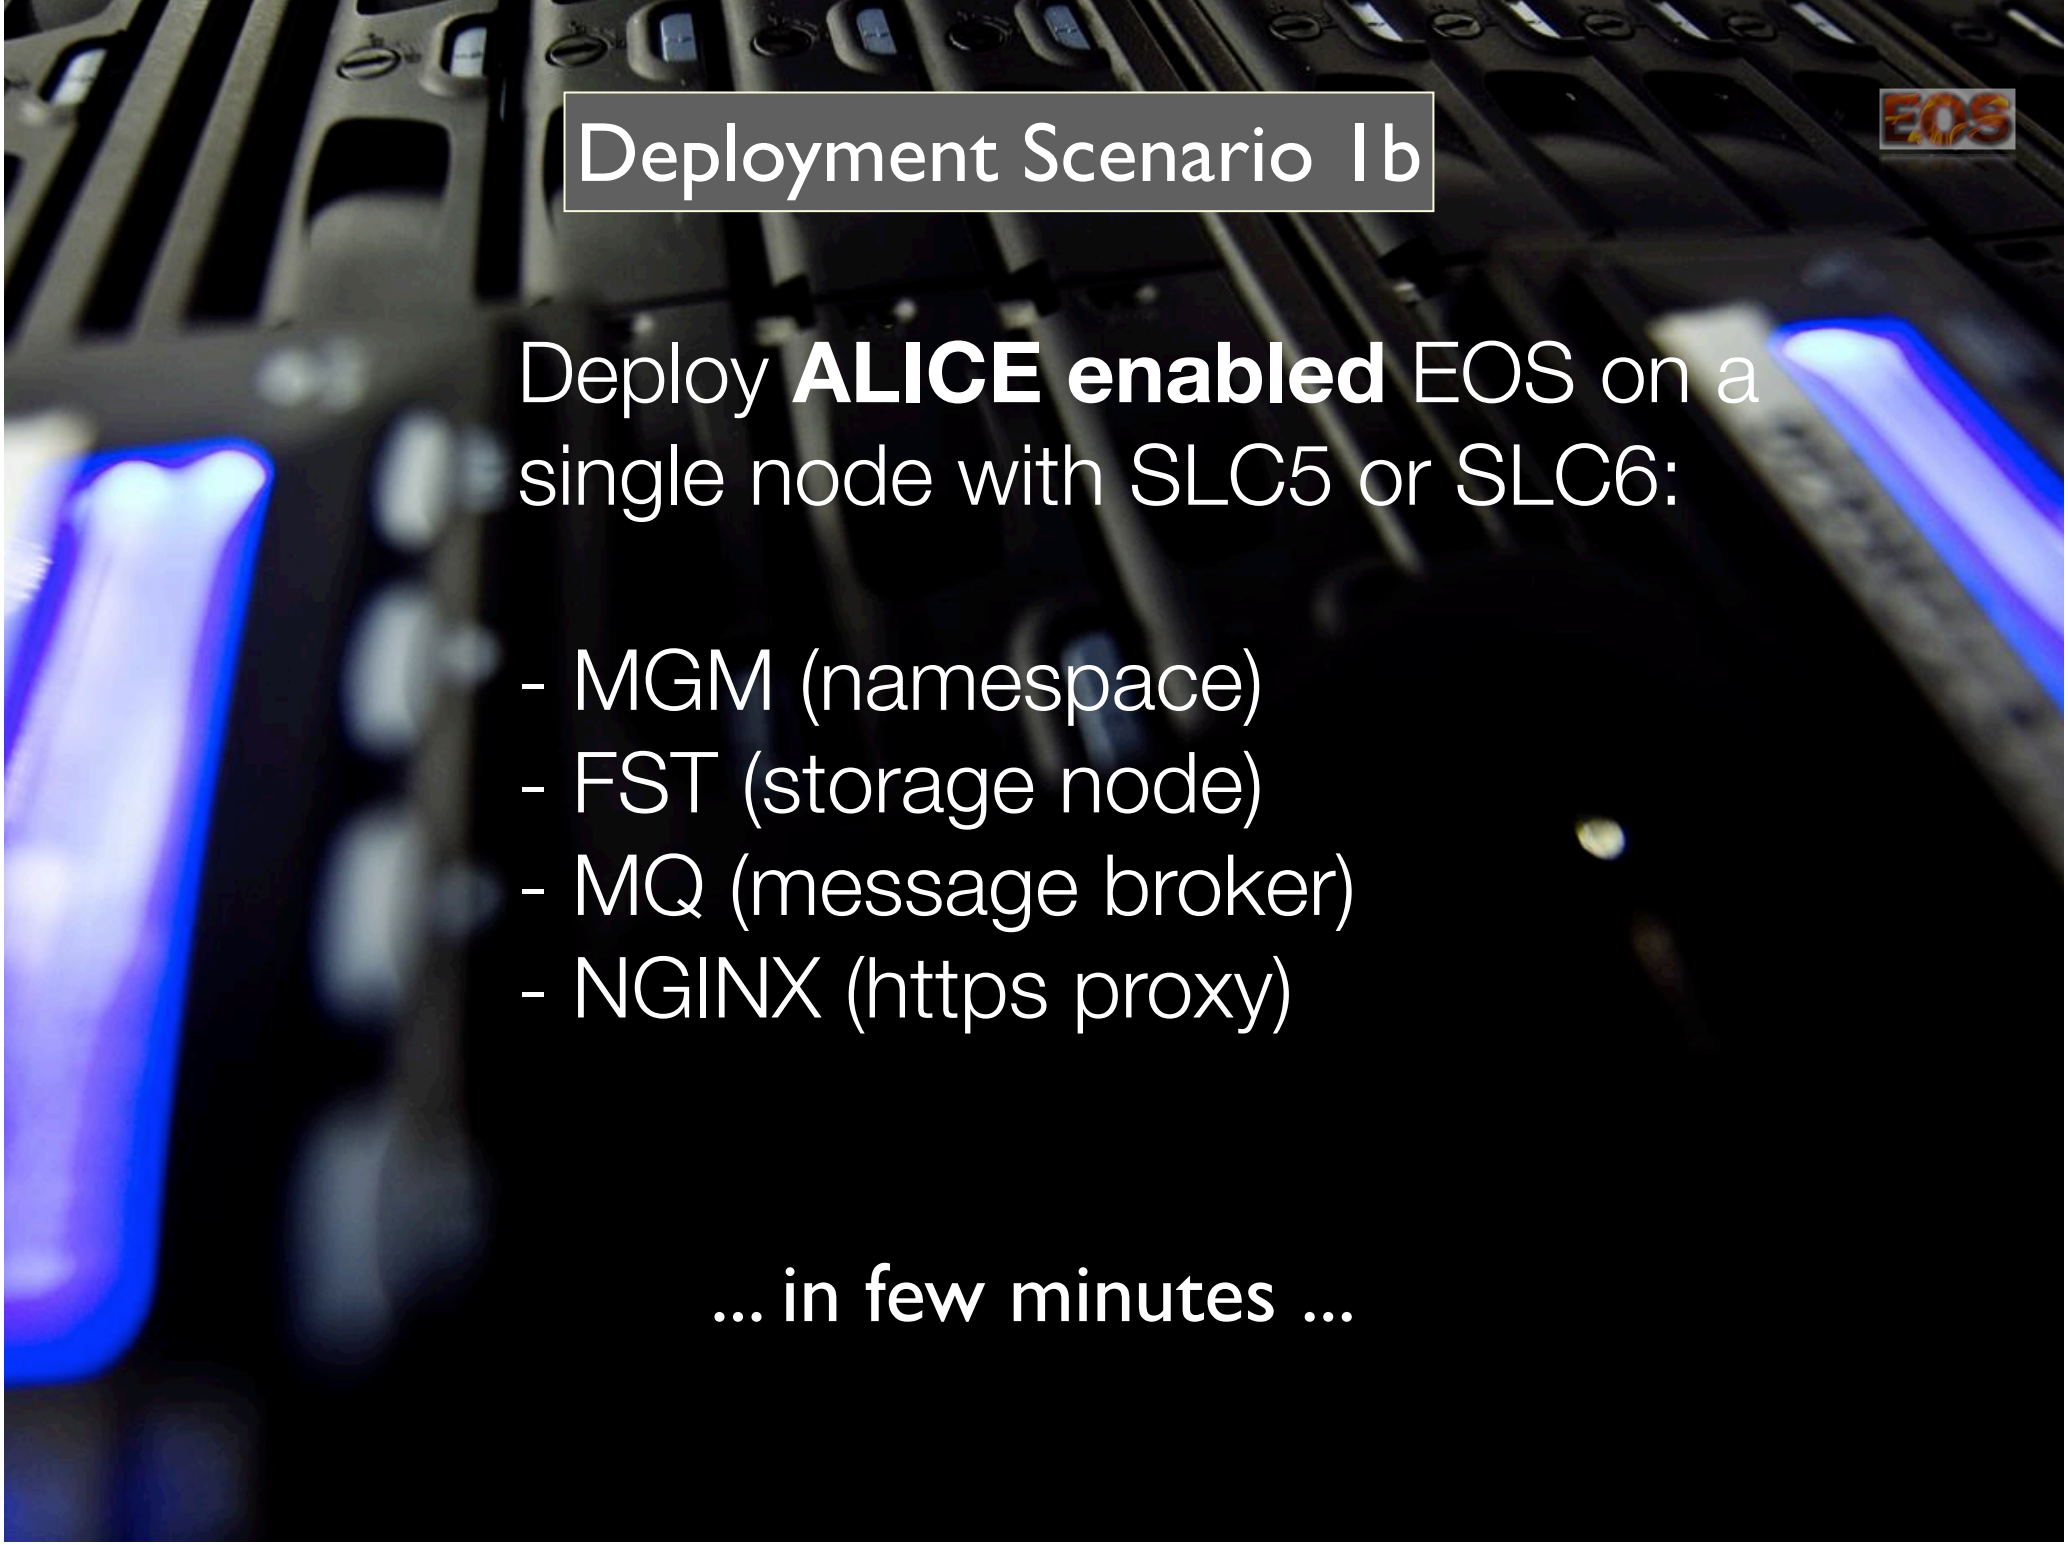
Deployment Scenario 1b
Deploy **ALICE enabled** EOS on a single node with SLC5 or SLC6:
- MGM (namespace)
- FST (storage node)
- MQ (message broker)
- NGINX (https proxy)
... in few minutes ...

Installation 1b

Run the tool ... but start the instance name with **alice.**

```
./eos-deploy
```

Answer few question:

```
instance name      : alice.foo.bar
alice SE name      : ALICE::FOO::BAR
2nd mgm name       : <return>
email              : yourmail@foo.bar
SLC flavour        : 6
#filesystems       : 4
storage servers    : <return>
```

... just wait ...

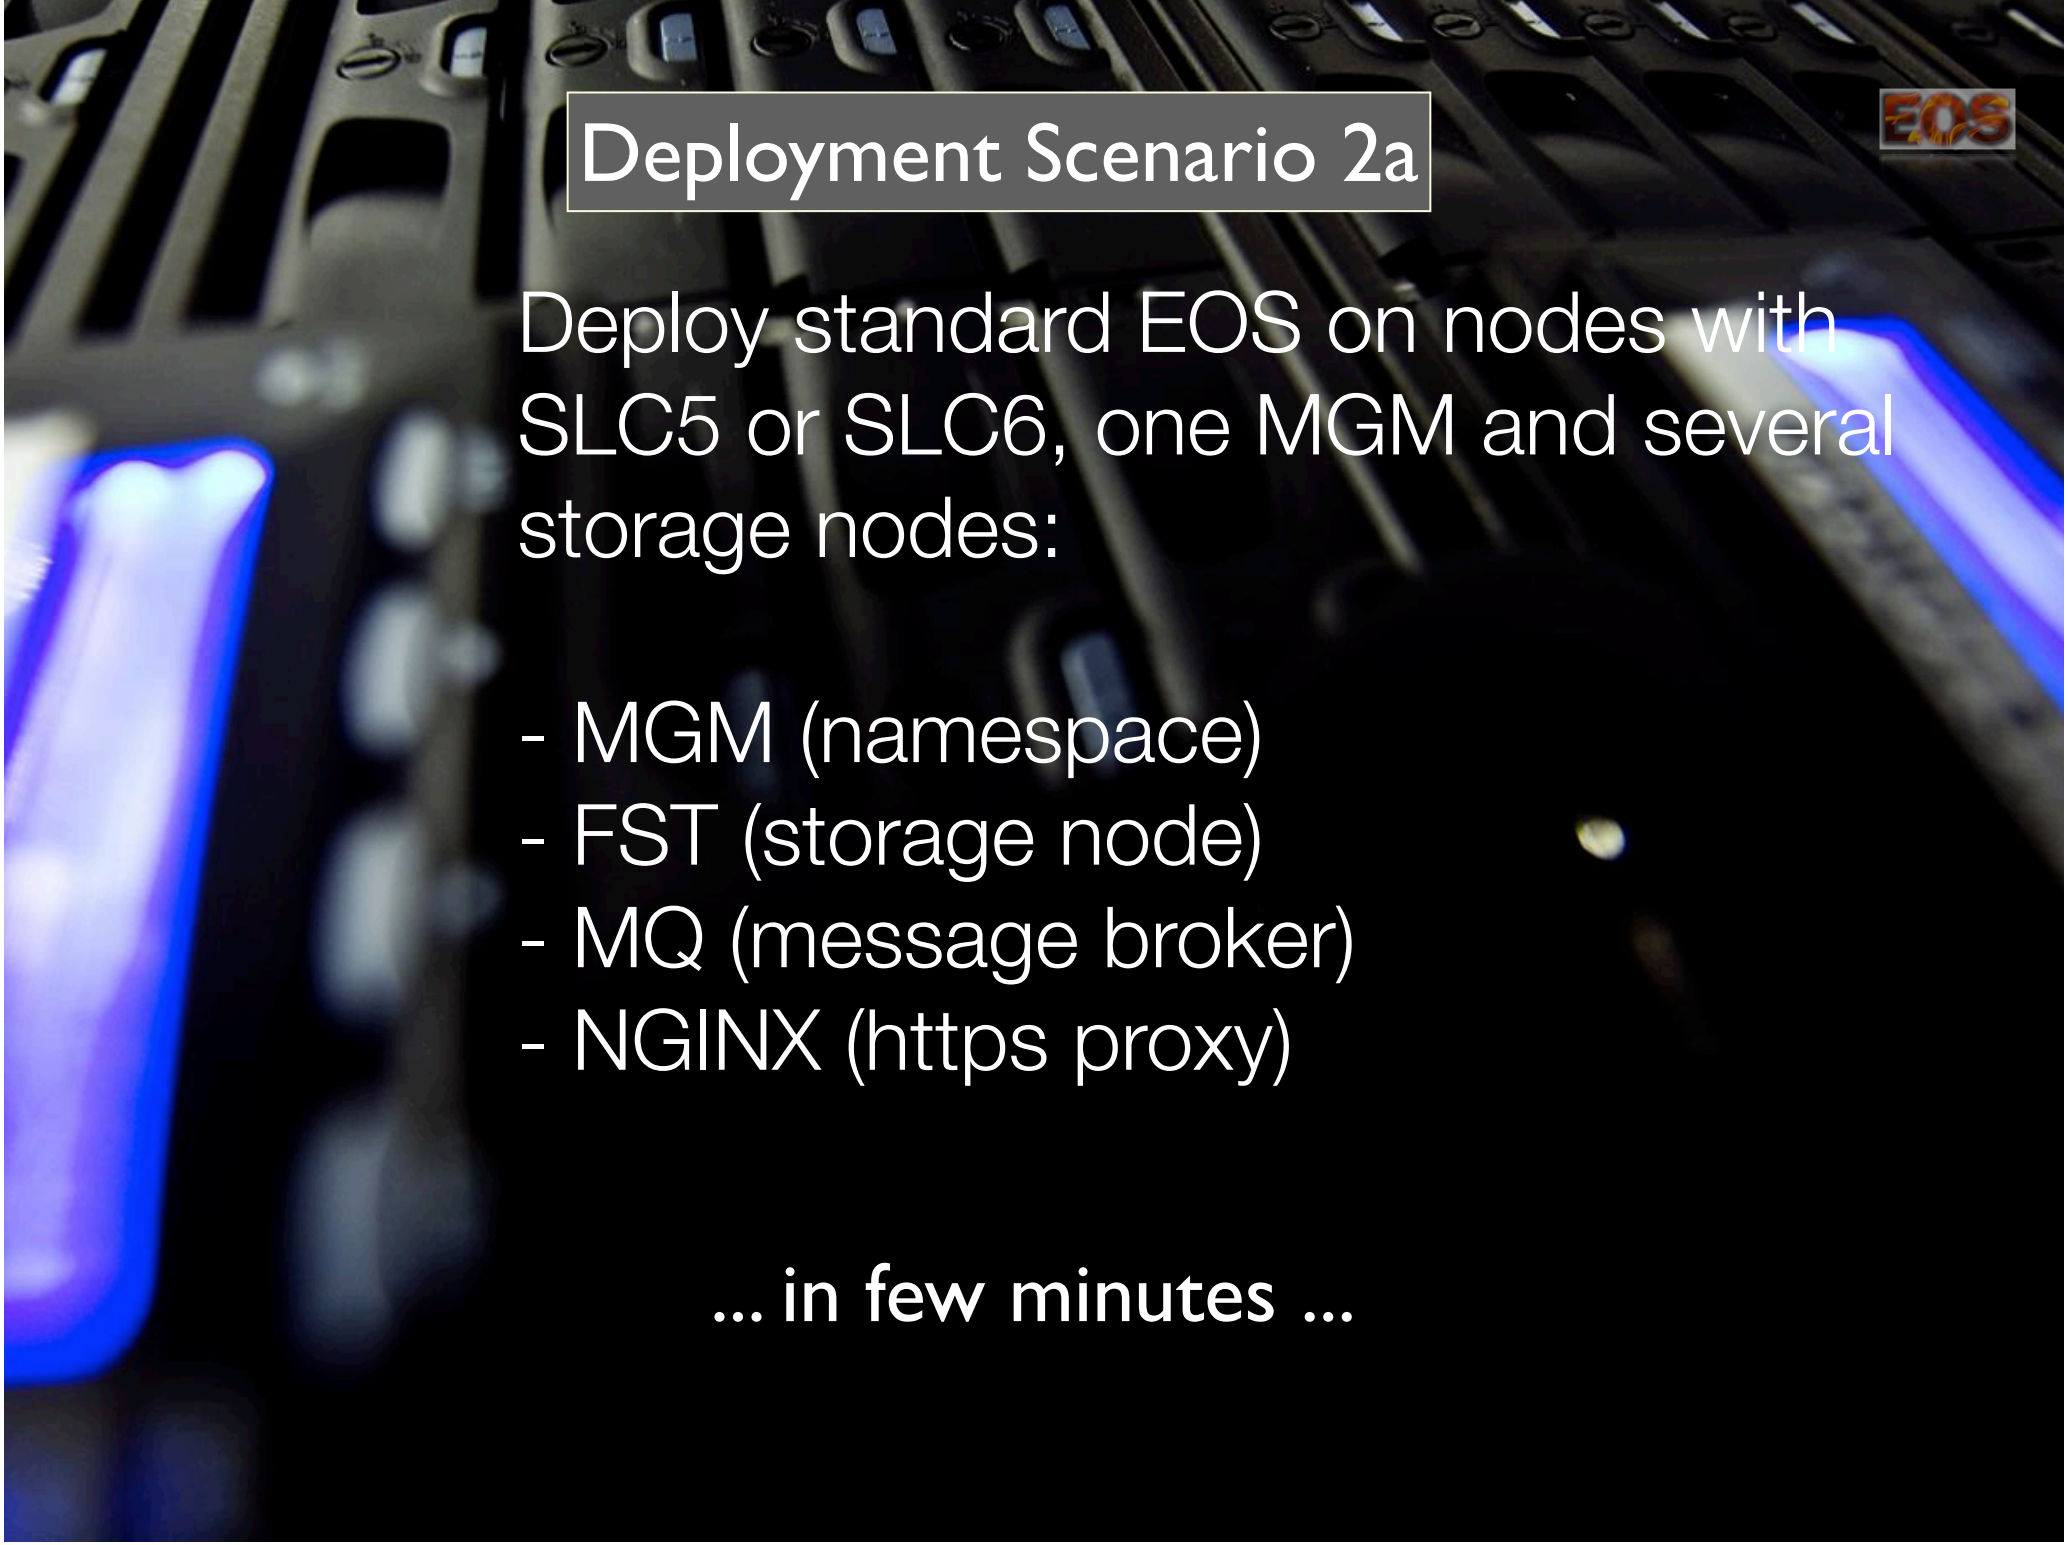

Deployment Scenario 2a

Deploy standard EOS on nodes with SLC5 or SLC6, one MGM and several storage nodes:

- MGM (namespace)
- FST (storage node)
- MQ (message broker)
- NGINX (https proxy)

... in few minutes ...

Installation 2a

Run the tool ...

```
./eos-deploy
```

Answer few question:

```
instance name     : test.foo.bar
2nd mgm name      : <return>
email             : yourmail@foo.bar
SLC flavour       : 6
#filesystems      : 4
storage servers   : fst1.foo.bar fst2.foo.bar
```

... just wait ...

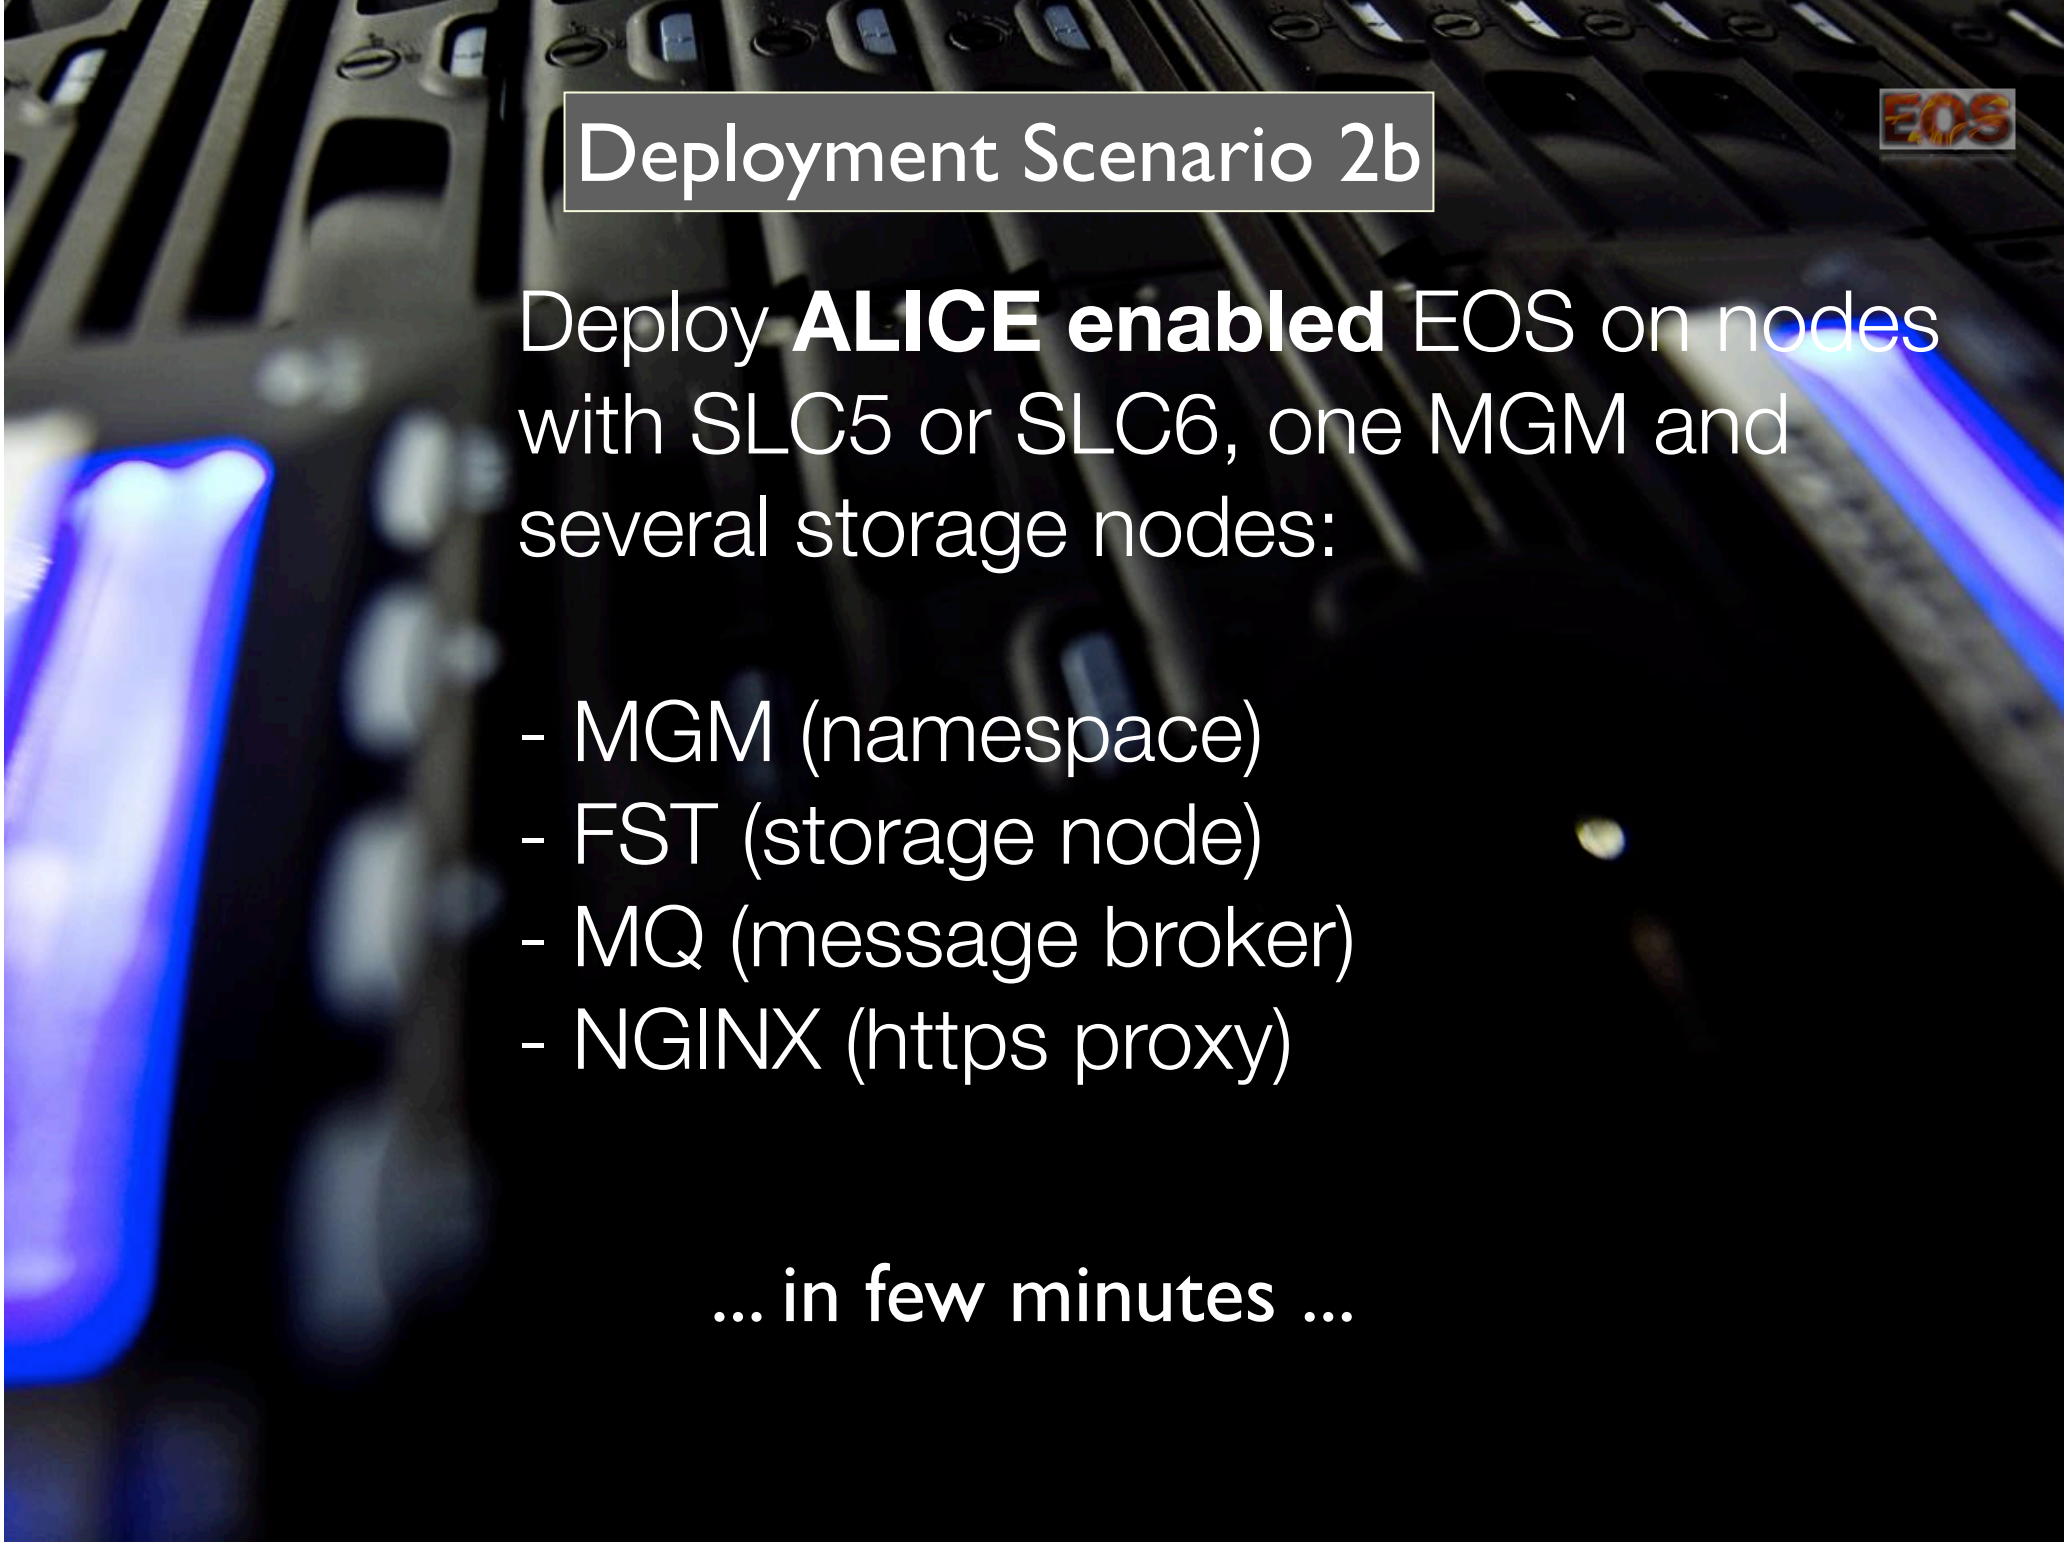

Deployment Scenario 2b

Deploy **ALICE enabled** EOS on nodes with SLC5 or SLC6, one MGM and several storage nodes:

- MGM (namespace)
- FST (storage node)
- MQ (message broker)
- NGINX (https proxy)

... in few minutes ...

Installation 2b

Run the tool ... but start the instance name with **alice.**

```
./eos-deploy
```

Answer few question:

```
instance name      : alice.foo.bar
alice SE name      : ALICE::FOO::BAR
2nd mgm name       : <return>
email              : yourmail@foo.bar
SLC flavour        : 6
#filesystems       : 4
storage servers    : fst1.cern.ch fst2.cern.ch
```

... just wait ...

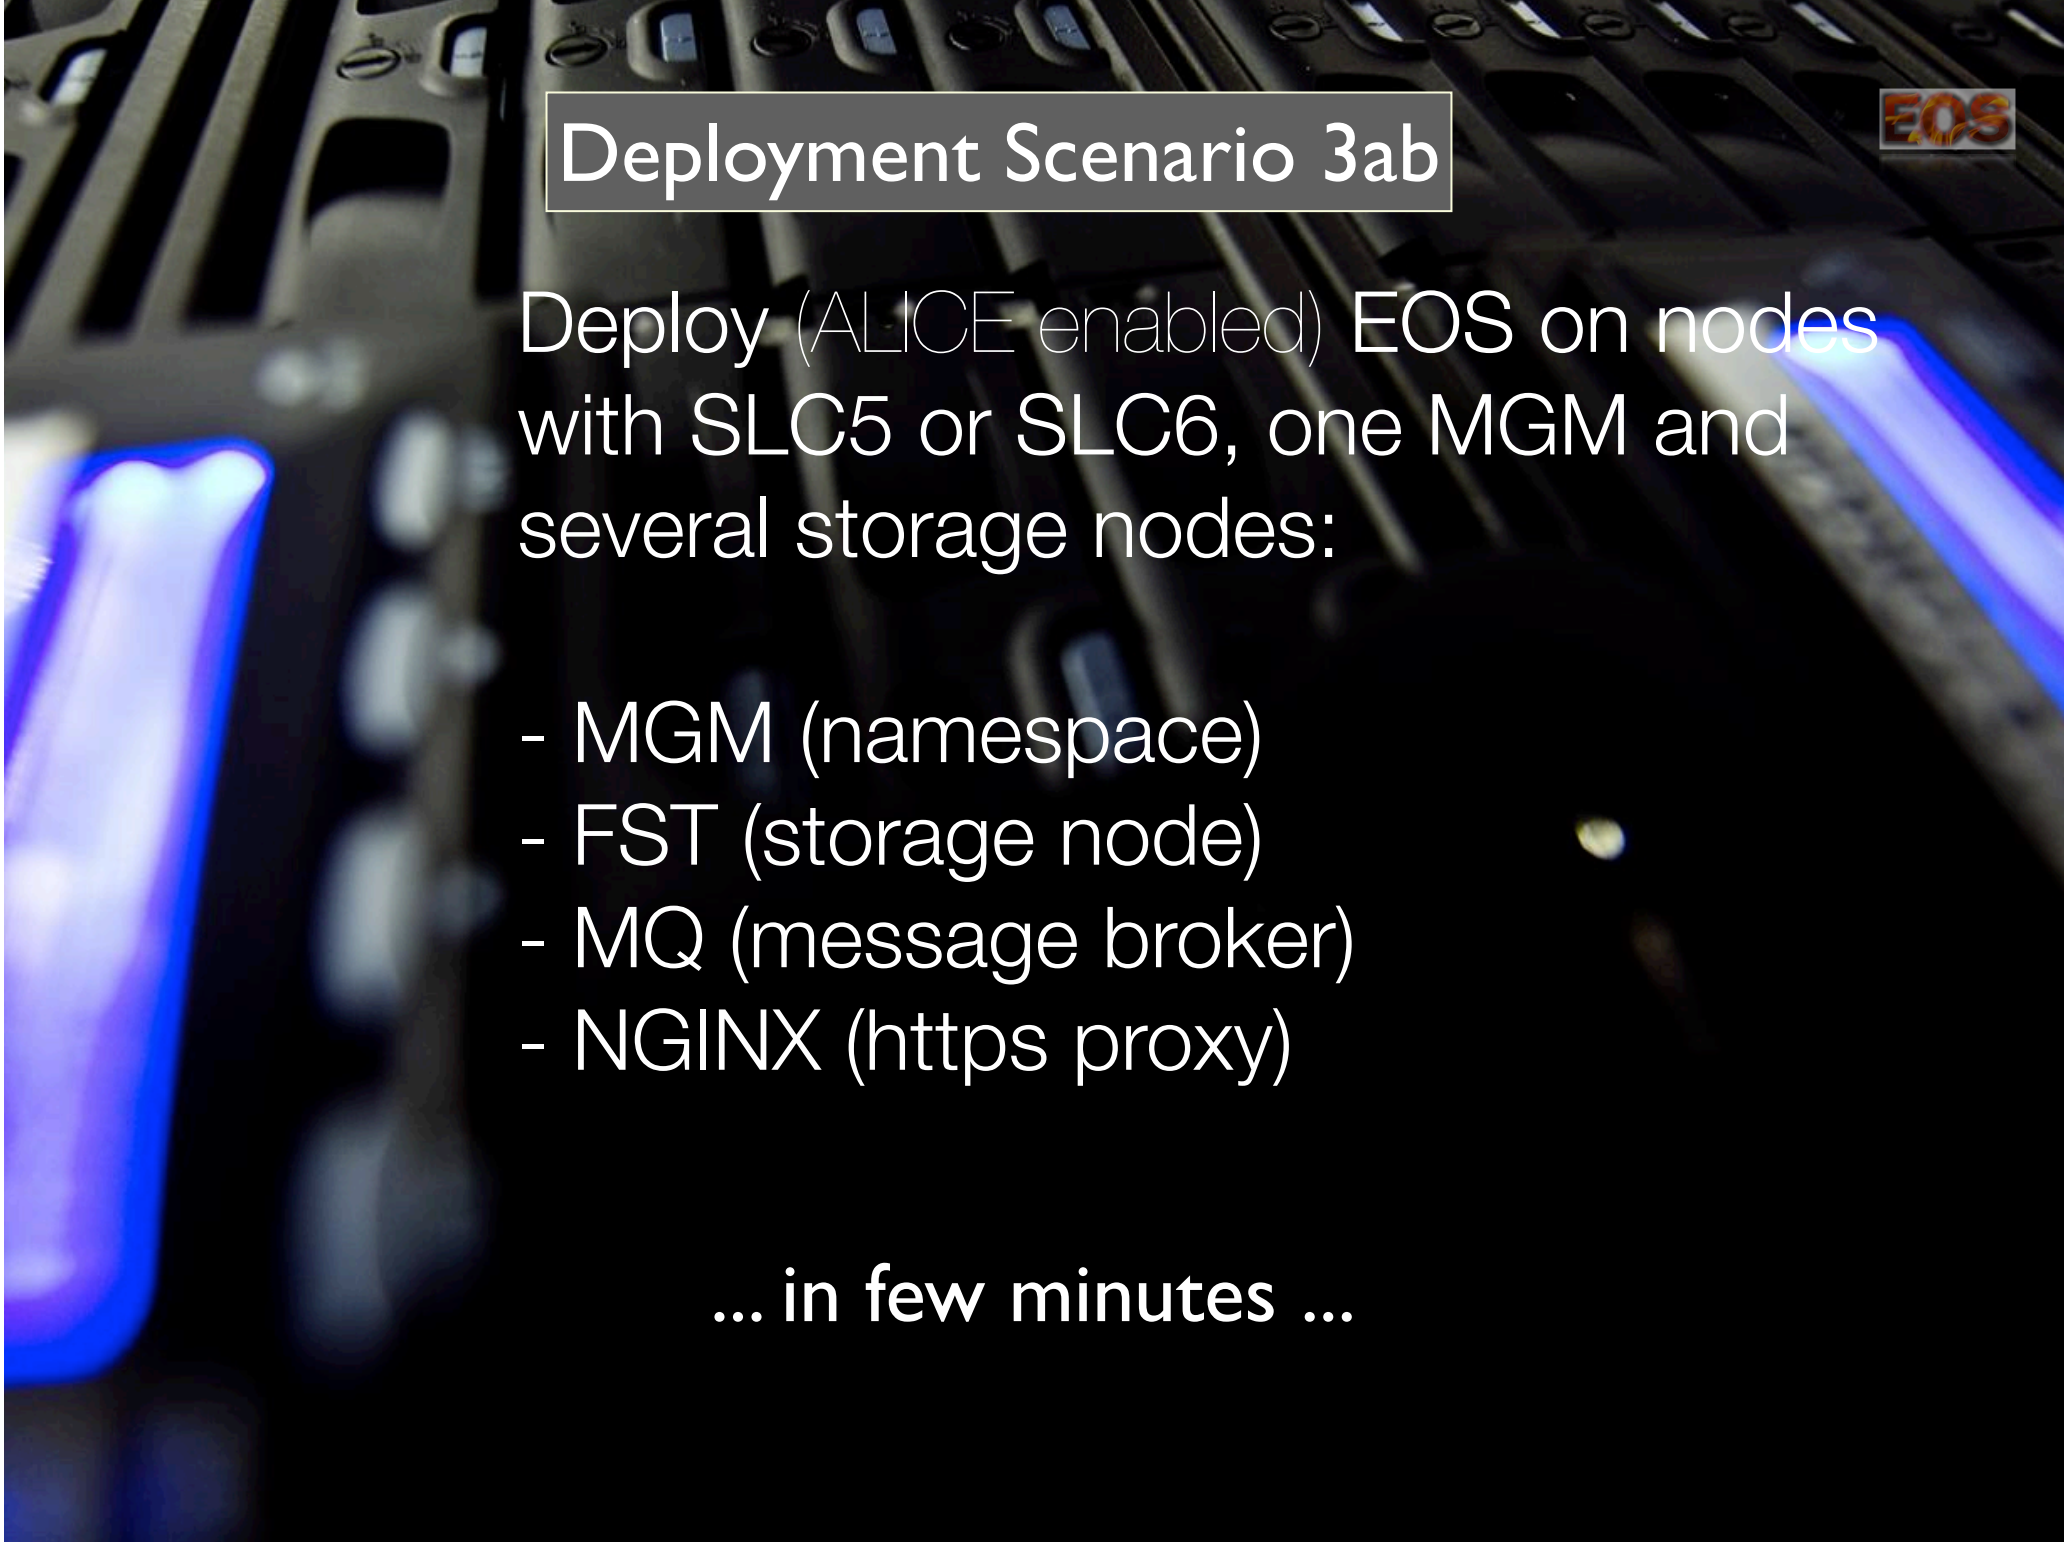
Deployment Scenario 3ab
EOS
Deploy (ALICE enabled) EOS on nodes with SLC5 or SLC6, one MGM and several storage nodes:
- MGM (namespace)
- FST (storage node)
- MQ (message broker)
- NGINX (https proxy)
... in few minutes ...

Installation 3ab

Run the tool ...

```
./eos-deploy
```

Answer few question:

```
instance name      : test|alice.foo.bar
alice SE name      : ALICE::FOO::BAR
2nd mgm name       : mgm2.foo.bar
email              : yourmail@foo.bar
SLC flavour        : 6
#filesystems       : 4
storage servers    : fst1.cern.ch fst2.cern.ch
```

... just wait ...

What happens?

Configuration steps `./eos-deploy`

MGM Step 1

- first checks if you have a host certificate, if not it creates a self-signed certificate and stores it under /etc/grid-security/[daemon]/hostcert.pem | hostkey.pm
=> this is used by the HTTPS server and presented to client browsers - you should get a real host certificate to avoid security questions by browsers

- downloads LCG CA files, creates empty GRID map file and configures the EOS repository. It masks the XRootD package from the EPEL repository and download all required RPMS from the YUM repository

- creates the master EOS configuration file /etc/sysconfig/eos and evt. modifies the default MGM XRootD configuration file /etc/xrd.cf.mgm to configure ALICE authorization

- opens required ports in the MGM firewall: 1094,1096,1097,8000,443,8443

- configures this machine to be the MGM master and start's EOS, NGINX and FUSE services

```
service eos start|stop|status ...
service nginx start|stop|status ...
service eosd start|stop|status ...
```

- basic EOS configuration: define space,group & authentication mappings


What happens?
Configuration steps ./eos-deploy
FST
- define EOS repository, exclude XRootD from EPEL, disable firewall on port 1095,8001
- register <n> filesystems under /var/eos/fs/<#> to the MGM
- start FST service
service eos start|stop|status fst

What happens?

Configuration steps `./eos-deploy`

MGM Step 2

- define space/filesystem configuration settings: scaninterval, autorepair, graceperiod, drainperiod, disk headroom
- enable kerberos5 authentication

- evt. create the ALICE homedirectory /eos/<instance>/grid and own it exclusivly by ID(aliprod@CERN)
- evt. map all UNIX access to ID(aliprod@CERN)

If you have given a second MGM host, it configures the second MGM host identical to the first one with the exception that the second MGM is started as a slave (ro) MGM. See the documentation reference later to get the details about the Master/Slave configuration.

How can I see it works?

EOS Shell

```
eos -b space ls
eos -b group ls
eos -b node ls
eos -b fs ls
eos -b whoami
eos -b find /eos/
xrdcp /etc/passwd root://localhost//eos/<instance>/testfile
```

UNIX Shell

```
bash> df /eos/
```

HTTP Browser

`https://<instance>:8443` => provide KRB5 user+passwd
`https://<instance>:443` => have client cert in the browser, otherwise you are nobody
`http://<instance>:8000` => you are nobody

When the browser works with kerberos or with a certificate (you have to add you DN to the grid-map file!) you can add yourself to be a SUDOer

```
eos -b vid set membership <your-uid> +sudo
```

If you reload the web page you can now use the administrative tabs to see the EOS configuration

What is still missing?

Firewall

Make the firewall settings persistent:
1094: XRootD MGM port (only on MGMs)
1095: XRootD FST port (only on FSTs)
1096: XRootD SYNC port (only on MGMs)
1097: XRootD MQ port (only on MGMs)
443: https X509 port (only on HTTPS gateways or MGM)
8443: https KRB5 port (only on HTTPS gateways or MGM)
8000: http port (only on MGMs)
8001: http port (only on FSTs)

Backup

You should do from time to time a backup of **/var/eos/md/files.*.mdlog** and **/var/eos/md/directories.*.mdlog** & **/var/eos/config/default.eoscf** containing the namespace and the active configuration file

Software Update

Automatic via YUM or manual ‘yum update’ on all nodes when desired.

ApMon

ApMon is already installed and configured, but needs to be started once on the FSTs:

```
service eosapmond start
```

Dual MGM Setup

Master MGM/MQ

Check that this host is the master

```
service eos status

eos -b ns
```

Slave MGM/MQ

```
service eos status

eos -b ns
```

The configuration and management of Master/Slave is described here:

https://eos.readthedocs.org/en/latest/configuration/master.html

MGM Alias

If possible on your site use a load-balanced alias for the two MGM nodes. The Master node can do read/write calls, the slave MGM can do read calls and redirects writes to the master. If not possible point to the master MGM (RW) and in case of failover change the DNS entry to point to the ex-SLAVE MGM after failover.

Information & Documentation

EOS Manual

`https://eos.readthedocs.org/en/latest/configuration.html`

EOS Webpage

`http://eos.cern.ch`

EOS External User List & Bugs

Get informed about releases, ask question to the community

Subscribe via e-groups.cern.ch : **`eos-operation-external@cern.ch`**

Bugs: **`https://savannah.cern.ch/projects/eos/`**

EOS CERN Admins

If you have a very detailed question and need expert advice write to **eos-admins@cern.ch**

EOS Puppetization / BDII

Get input/configs from **Jan Iven** (CERN) or **Jean-Michel Barbet** (Subatech)