

# **NDGF All Hands 2**

## **Report of Contributions**

Contribution ID: 1

Type: **not specified**

## Clom & Zanak replacement

*Tuesday, 6 November 2018 11:15 (20 minutes)*

Service for C&Z will run out on 30 September 2019. It is time to start deciding on and procuring the replacement. Discussion on sizing our production Ganeti and Postgresql cluster.

**Session Classification:** NT1 staff internal meeting

Contribution ID: 2

Type: **not specified**

## **dCache preprod monitoring**

*Tuesday, 6 November 2018 09:00 (20 minutes)*

Do we need more probes for the pre-production dCache setup on mordor?

**Session Classification:** NT1 staff internal meeting

Contribution ID: 3

Type: **not specified**

## **Tape Carousel Test - Lessons learned**

*Tuesday, 6 November 2018 14:30 (30 minutes)*

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: 4

Type: **not specified**

## Ceph for compute and storage

*Tuesday, 6 November 2018 15:50 (30 minutes)*

Discussion of experiences with using Ceph and ARC together and for using ceph with dcache.

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: 5

Type: **not specified**

## **New Ubuntu release readiness and upgrade plan**

*Tuesday, 6 November 2018 09:20 (30 minutes)*

**Session Classification:** NT1 staff internal meeting

Contribution ID: 6

Type: **not specified**

## Kafka for everything?

*Tuesday, 6 November 2018 10:45 (30 minutes)*

1. What can we put into kafka in addition to dcache billing.
2. What is our strategy towards logstash if we put everything into kafka: a. Ditch it and install es-consumer that feeds data directly to the ES. b. A separate instance on dedicated machine that queries kafka on different topics and submits them to the ES. c. Other.

**Session Classification:** NT1 staff internal meeting

Contribution ID: 7

Type: **not specified**

## ood/ooc scheduling

*Tuesday, 6 November 2018 09:50 (15 minutes)*

Please prep the common calendar the week before. Chrulle will then come up with a suggested schedule ahead of time. If you are not OoD or OoC you have an early break :)

**Session Classification:** NT1 staff internal meeting



Contribution ID: 8

Type: **not specified**

## Postgresql topics

*Tuesday, 6 November 2018 11:35 (20 minutes)*

OS / FS tuning: overcommit, dirty pages etc.

How are we handling DB backup today. Should we look into alternatives?

**Session Classification:** NT1 staff internal meeting

Contribution ID: 9

Type: **not specified**

## Site news and plans

*Tuesday, 6 November 2018 13:00 (1 hour)*

Sites will prepare a short presentation on news and plans.

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: **10**

Type: **not specified**

## Future tape performance

*Tuesday, 6 November 2018 15:30 (20 minutes)*

Future requirements for tape performance, with implications and discussions on tape pool hardware

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: **11**

Type: **not specified**

## Publishing our scripts

*Tuesday, 6 November 2018 10:05 (20 minutes)*

Discussion on publishing our ansible scripts etc on github.

**Session Classification:** NT1 staff internal meeting

Contribution ID: 12

Type: **not specified**

## Accounting/reporting of local jobs

*Tuesday, 6 November 2018 16:20 (20 minutes)*

Discussion on including local jobs in reporting.

**Presenter:** Mr SALMELA, Ville (csc.fi)

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: 13

Type: **not specified**

## LHC news and future

*Tuesday, 6 November 2018 14:00 (30 minutes)*

Oxana will give a presentation on LHC news and lead a discussion on future requirements.

**Presenter:** SMIRNOVA, Oxana (NeIC / Lund University)

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: **14**

Type: **not specified**

## Disk pool requirements

*Tuesday, 6 November 2018 16:40 (20 minutes)*

The usage of ATLAS and ALICE pools today look pretty different. Is the usage going to diverge further? Or is it going to converge? Are we going to need smaller pools more suited for random IO?

**Session Classification:** NDGF local sysadmins meeting

Contribution ID: 15

Type: **not specified**

## **Lunch venue decision**

*Tuesday, 6 November 2018 11:55 (5 minutes)*

**Session Classification:** NT1 staff internal meeting



Contribution ID: **16**

Type: **not specified**

## Next NDGF All Hands

*Tuesday, 6 November 2018 17:00 (10 minutes)*

**Session Classification:** NDGF local sysadmins meeting