Service report on accidental user removal of other users files in EOSCMS T3 (CAF User) at CERN

Disclaimer

This report is an extract of the original version, which contains names of all the involved person. This document is not intended to be an incident report since this is not considered to be a service incident.

Description

On November 5th in 2013, a CMS user script running on LXBATCH cluster (with at least 5 different jobs) caused file deletion on EOSCMS (mounted via FUSE), in directories that were group-writable (and files inside were group deletable). CERN EOS Operations team spot an anomalous amount of removal commands and "Operation not permitted" which led to the ban of that particular user.

Impact

28 EOSCMS CAF user and some cmgtool directories were involved with an overall of 78k files, or ~15TB of data removed (user issued 0.5M unique `rm` commands - 2.5M in totals). This erroneous rm command visited around 1.1k directories out of 11k possible deletable directories, open for group write and delete (EOSCMS itself has around 684k directories). The early detection has allowed to reduce the deletion effect of a factor of 10 and the previously defined protection policies of a factor of 7.

Time line of the incident

When	What
05.11.2013 ~16:00	A CMS user launched a recursive deletion on a EOSCMS FUSE mount point (by accident)
05.11.2013 17:00	Strange rm activity identified by the EOS Operations team
05.11.2013 17:00	User identified and banned (from EOS)
05.11.2013 17:26	CERN CERT team involved to exclude the possibility of a compromised account

Analysis

Deletion seems to have been part of some user analysis script (cleanup phase) - assume no malicious intent or compromised account. Identified sequence:

```
mount EOS; run; umount EOS; rm -r mountpoint
```

We assume that `umount` failed (i.e open files?) or was not done at all. Moreover if multiple jobs land on the very same batch node files deletion could be issued anyway since the mount point is shared and probably not protected via locking by the user.

Some EOS directories were changed by their owners to be group-writable, this action is probably a remnant from the CASTOR grid mapping which was mapping user certificate in pool accounts. EOSCMS maps grid user credentials to the correct local account, i.e. no need to open up personal space. Moreover EOS also supports a rich set of access controls beyond the UNIX-style "user:group:others" (including lists of users and E-Groups) .

Follow up

EOS Beryl has some protection against accidental deletions if the recycle bin is activated, EOSCMS update to EOS Beryl was postponed from 5th Nov 2013 to 12th Nov 2013 due to clash with CMS global run. After the update IT-DSS jointly with IT-SDC and CMS Computing decided to activate experimentally the recycle bin under the following trees in EOSCMS:

```
/eos/cms/store/user/
/eos/cms/store/group/
/eos/cms/store/caf/user/
/eos/cms/store/cmst3/user/
/eos/cms/store/cmst3/group/
```

IT-DSS (in agreement with IT-SDC and CMS Computing) offers to double-check the script that caused this issue and review acls and permission in the EOSCMS instance, reverting periodically directories with permission 777 back to 775 and adding deletion protection against other CMS users (add sys.acl="g:zh:!d") if a directory is open for group write. Furthermore all new users will inherit by default the flag g:zh:!d in their sys.acl which will prevent any unwanted deletion in case in the future they will change permission in one of their directories.