NIS -> LDAP migration

UNIX GID consolidation

Christoph Beyer, Yves Kemp

Linux User Meeting, 31.10.2013

Why move away from NIS?

- NIS has intrinsic group definition limit

```
> ypcat group
beatles::64:john,paul,george,ringo
```

<1024 characters!

- This limit has been reached at DESY for one group. Consequences:
  - ANY change in ANY group cannot be pushed to NIS master – and hence to clients
  - NIS basically unusable for everyone
  - Fortunately, some users in this group do not work under Unix. In agreement with group admins manually remove these users when feeding the NIS master from the registry. Dirty and ugly – but works.
- Splitting the group does not work: Master and clients do not work reliably
- **We have ~2-3 month at current growth of the affected group to find a permanent solution – for passwd and group in a first step**

Alternatives

- Generate /etc/passwd and /etc/group files and distribute e.g. via cron
  - E.g. done in Zeuthen
  - Only works for centrally managed and controlled systems
  - Will not work for commercial systems, e.g. Sonas
- LDAP
  - Established since years (actually first RFC 20y old)
  - Many sites using it – also some groups at DESY for particular purposes
  - Different server products, including commercial ones
- Microsoft AD
  - Containing an LDAP service
  - Used e.g. at CERN for Linux clients, including huge batch farms
  - Current Microsoft license situation does not encourage extension of AD usage

We opt for LDAP

- Checked that all currently supported Linux distributions support it
  - SL5 and above, Ubuntu 10.04 and above
  - Even some older Linux (unsupported☺) and even Solaris tested
  - Works against a test-LDAP server for passwd and group
- Currently investigating on best fit for LDAP server
  - Shortlisted two products from an extensive list
  - “389” (RedHat owned) and “OpenDJ” (Fork from codebase of SUN OpenDS product)
  - Test installation currently being setup
  - Gain insight in how to setup production infrastructure

Migration plan - tentative

- Passwd and group most pressing for some user groups
- Will change IT owned resources (PAL, BIRD,…) and affected group machines first to LDAP
  - NIS will stay at first, but once we run in "1024"-problem, simply drop the affected group from NIS as affected resources are migrated to LDAP by then
- NIS currently works well for netgroups.
  - In second step, move netgroups to LDAP
  - Problem: Registry is not only source for netgroups (in contrast to passwd and group)
  - Netgroup composed by registry, WBOOM, group admin tools
  - This needs to be implemented for LDAP
  - ... Once we master LDAP for passwd and group
- Once LDAP works for passwd, group and netgroup, will announce shutdown plan for NIS service … which should be in 2014

While we are at it: Get rid of another legacy

- https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/4/html/Introduction_To_System_Administration/s1-acctsgrps-rhlspec.html
- … GIDs below 500 are reserved for system use. … GIDs are never to be assigned to a user, as it is likely that some system component either currently uses or will use these UIDs/GIDs at some point in the future…
- Currently 85 groups with GID < 500 at DESY
  - And conflicts with existing groups exists!
  - This affects vital functions of Linux systems, e.g. sound or removable media
- Plan: Change groups <1000 to >1000 when migrating to LDAP
  - 1000 instead of 500 to be on safe side

<1000 GID migration. Steps:

- Current: beatles::**64**:john,paul,george,ringo
- Step zero: Clean up unused groups and check who can be migrated directly
- Step one: Two groups:
  - beatles::**1960**:john,paul,george,ringo - Managed via registry
  - beatles**_legacy**::**64**:john,paul,george,ringo - Created from above while pushing groups into NIS
- Step two: Change file and directory ownerships
  - Something around `chgrp –R 1960 /path/to/data`
  - No need for big-bang migration, each machine and path can be done when appropriate
  - No data access problems during migration phase
- Step three: Remove the _legacy group
- **… LDAP will only contain the >1000 groups – no legacy!!!**

Primary and secondary groups: The correct way

- For simplicity, we dropped the information about primary and secondary groups

- Old:
  - Passwd:

```
freddie:…:23:…
brian:…:23:…
john:…:23:…
roger:…:23:…
paul:…:42:…
```

  - Group:

```
queen:…:23:paul
badcompany:…:42:
```

`

- Intermediate & New:
  - Passwd:

```
freddie:…:2323:…
brian:…:2323:…
john:…:2323:…
roger:…:2323:…
paul:…:4242:…
```

  - Group:

```
queen:…:2323:paul
badcompany:…:4242:
queen_legacy:…:23:freddie,brian,john,roger,paul
badcompany_legacy:…:42:paul
```

- Groups are not used for access (should not)
  - Netgroups are used, and there is no GID clash with them
- Groups are (to our knowledge) only used for data permissions
  - … when different people are involved
  - E.g. local filesystem, NFS v3 (AFS uses different mechanisms)
- The proposed scheme will not work for very large groups:
  - … remember the 1024 character limit
  - Currently 7 groups are affected. Will contact these groups separately