TWiki > DICE Web > SystemInfoPostGDPR (13 Apr 2021, GrahamDutton)

Project 464: Restricting access to last/ps/w/etc (aka "System information post-GDPR").

This page is the 'project homepage' for Project 464. See also SystemInfoPostGDPRFinalReport (when complete)

Description

This project is to track the work (some of which is already complete) relating to:

- disabling access to system introspective functions which reveal PII (specifically targeting, `last`, `ps`, `w`, and related functions)
- identify what types of information are shared by these tools, and what must be restricted
- identify where blanket limitations will cause problems
- identify if our systems reveal this information in other ways which is deemed unacceptable
- identify if there are areas where it's not feasible to restrict this information

As part of this project we should also communicate some/any of this, and the rationale behind it, to users.

Deliverables

1. determine exclusions, eg single user machines
2. determine whether ps/w can be treated differently to last
3. disable as necessary

2021 Recap Report

Since last update we've moved to a new platform, and largely moved to working from home. The detail and questions from the previous reports have been digested and largely implemented. The complexity of categories of machine, etc. was largely deemed excessive, and so we now treat machines as either "multi user" or "single user" (though as ever there are quirks).

Single user

This consists of desktops and some special-purpose servers (usually VMs). There's no need to restrict access to anything, but could do so by default since changing things is simple. Lab desktops are specifically not multi-user.

Multi-user

These are all restricted as much as is possible, though there are quirks.

Login servers reveal extra PII such as remote IP and active / idle times which are more sensitive. This information is normally restricted, but see `loginctl` section below. This could be partially mitigated with user-facing documentation suggesting use of (either University or Informatics) VPN.

XRDP (in widespread use) was a new development since the introduction of remote / hybrid teaching. Surprisingly this presented no major categories of information, though it was noted that a set of logs produced by one of the XRDP components leaked login times / presence information. These logs were of no technical value - and in fact were causing service stability issues - and are now purged automatically.

Cluster hosts (head nodes and queue management tools specifically) are treated differently on the basis that they necessarily share PII, but of a very limited nature (i.e. although a 'queue record' contains a UUN, this information cannot be used to infer very much). It's been deemed prohibitively complex / expensive to restrict this further, since it would require complex patching, and unless done carefully would also potentially reduce the effectiveness of the clusters if users could not see the status of the system as a whole. Considering what could be done if somebody were to object to the status quo is a matter for a DPIA.

Technical Measures

Thanks to MPU we now have a `/privacy` subdirectory in LCFG which contains various blocking measures. These are already applied by default and cover `last`, `ps`, `dmesg` and `w*` tools. Notably since the move to Ubuntu we have (temporarily) lost control of the `proc` filesystem; this will be fixed as part of the DICE server porting project to Ubuntu, and is part of mount management so can be considered in progress / effectively guaranteed.

One major omission which was noted early on in SL7 was that `loginctl` leaks information equivalent to (and in fact more detailed) than all of the above tools. Multi-user privacy does not appear to have been one of the areas of interest of the developers and it's notable that there's no officially documented mechanism to restrict this despite investigations into `policykit` and similar tools. It would appear that the developers expect that on systems where privacy is critical, Linux namespaces (as a component of LXC) are the expected mechanism to achieve this. The alternative would be to patch the code but as this is a fairly core piece of the OS this comes at the usual cost and risks: chiefly that it can be time-consuming to do correctly, and can delay application of upstream security patches.

Conclusion

Taking a broader view, namespaces / limited containerisation is something which now warrants more investigation in any event: these technologies were in their infancy in SL7 and of little use, but Ubuntu Focal offers a good opportunity to revisit and research how we could use the technology. In summary, this project recommends some or all of the below:

- Consider the current measures the best we can reasonably do at the moment
- Research the use of namespaces on DICE Ubuntu Focal
- Draft a DPIA (or two, one for use of cluster(s)) for information that's not feasible to restrict
- Research the possibility of patching `loginctl` (low-priority)
- Restore `/proc` restriction (following Ubuntu Focal mount management changes)
- Document the implications of `loginctl` remaining user-accessible, and any mitigation (pending DPIA)

May 2019 update

Major changes since March:

- umask has been changed
- `/tmp` and `loginctl` are biggest remaining changes for shared machines
- classes of machine are largely correct.

Detail pending.

March 2019 update

Having noted that restrictions are in place (and incomplete) on the highest-priority systems, I decided to take a step back and review the whole landscape before applying further restriction. What follows is a partial brain-dump of thoughts, progress so far, and a list of actions which incorporate the original deliverables.

There are several breakdowns required to understand the restriction as it will be applied:

Categories of information

- historical presence information via `last`
- live presence information via `who`
- personal data via LDAP, `finger`, etc.
- live *process* information which does not imply presence via `ps`
- *implied* information? file metadata. etc.?
    - e.g. afs vos data

Questions:

- Are we disabling the tool or protecting the information source (i.e. don't forget systemd)
- Is there a distinction between any / all of these in legal terms?

Categories of host

single user machines: desktops (and VMs) for individual use.

It seems clear these can remain unrestricted (explore the CO edge case, but I think this is OK)

research servers: typically accessed by a closed user group

These are somewhat fuzzier but this (a) doesn't mean that `last` information is any more appropriate and (b) doesn't mean there's not a legitimate case for wanting to cooperate and share system usage data.

teaching machines: desktops or servers used for teaching, accessible to all/many students.

It seems clear these must be restricted (but explore the "teaching system" edge case)

gateway machines: accessible from outside, large combined user groups

These are already restricted, but the mechanisms must be clarified and improved.

Restriction Mechanisms

Current mechanisms aren't adequate:

- e.g. `w` is restricted, but `loginctl` still works

Full list is still to be enumerated:

- Some are already collected in `lcfg/options/privacy/`
- Some require better knowledge of systemd.
- Clusters share large amounts of process / job information, how much is covered by any/all of the above?
- Further isolation would involve cgroups (and a newer kernel, probably)
- Lots of mechanisms yet to be determined.

Other issues / edge cases

- Students are "taught" to use the UNIX environment and encouraged to learn how it works. ITO doesn't trump GDPR, but should we be consulting teaching staff about the (minimal) implications for the teaching environment?

- https://www.herald.ie/news/the-bins-are-back-in-town-gpo-trashes-gdpr-concerns-38077717.html

Computing staff

C(S)Os are special as we must access this information in the course of our work. We're obviously "people" for the purposes of non-sysadmin tasks. But:

- Should privilege escalations be mutually visible? i.e. when are we people, and when are we administrators?
- What about CO desktops?

Notification

A one-off system-wide email clarifying the situation would make sense.

Research servers presumably require consent - so depending on how active the level of consent required, we should resort to MOTD and/or agreement on gaining access to specific machines. This seems important for e.g. clusters where the technical means of restricting access to e.g. job information are beyond us (in the short term at least).

Actions

Notwithstanding all of the above discussion the following must take place ASAP:

1. a. document restrictions; document opt in/out (and consent-seeking) procedure; expected defaults for machine classes.
2. b. Provide operational guidance and LCFG mechanisms Expect macros!
3. a. apply restriction, where feasible, to shared machines
4. b. announce changes and rationale (for each restriction, explaining what we're targeting); advertise opt-in/out.
5. a. [ongoing] identify which information must be concealed, identify and close the gaps. -- updating (1) in the process
6. b. [ongoing] identify which information need not be concealed, identify and open the gaps. -- updating (1) in the process

-- GrahamDutton

Topic revision: r10 - 13 Apr 2021 - 09:06:07 - GrahamDutton