

Re: Metrics and your privacy

Source: <http://linux.derkeiler.com/Mailing-Lists/Fedora/2006-11/msg05219.html>

- *From:* Andy Green <andy@xxxxxxxxxxx>
 - *Date:* Wed, 22 Nov 2006 16:09:05 +0000
-

Craig White wrote:

On Wed, 2006-11-22 at 09:15 -0500, Jakub Jelinek wrote:

On Wed, Nov 22, 2006 at 02:02:18PM +0000, Andy Green wrote:

Probably the GUID is a bit of a red herring, since in the case there are millions of boxes it will be a ton of work to maintain the database of them and compare every log line against it, for limited hard information.

Why? Just running `uuidgen -r` once per install and saving it somewhere is good enough.

On the clients, yes. I was talking about the back end of processing the logs from the mirrors. Ignoring yum caches, because the activity of individual clients is not repetitive in terms of what is downloaded, you can infer from n downloads of the same package from the same IP that there are n clients behind the IP. If the IP is dynamic, makes no odds you again see n download actions. So you get limited additional information about a client being an individual from going to all the trouble of tacking the GUID in a database.

Because boxes will typically download specific updated RPMs just the once, you can get an idea of the number of active boxes just by filtering on packages that have been updated for a while.

That's wrong assumption, many people use proxies to avoid downloading updates for each box again and again.

or cases like I have just recently set up where I created my own repository (yam soon to be known as depo) and all computers will

Re: Metrics and your privacy

install/update from there and will never touch a fedoraproject repository at all.

Yes, that why I said this six hours ago on this thread:

"Information is still lost or degraded for

...

3. Machines behind a local yum cache

Whatever tools are provided to run the yum cache should have the repo log processing stuff folded into them, and report stats up to Fedora HQ by default. But a user should be able to turn it off.

..."

–Andy

—

fedora–list mailing list

fedora–list@xxxxxxxxxx

To unsubscribe: <https://www.redhat.com/mailman/listinfo/fedora–list>