



Contribution ID : 255

Type : **not specified**

libresource - Getting system resource information with standard APIs

Tuesday, 13 November 2018 14:55 (15)

System resource information, like memory, network and device statistics, are crucial for system administrators to understand the inner workings of their systems, and are increasingly being used by applications to fine tune performance in different environments.

Getting system resource information on Linux is not a straightforward affair. The best way is to collect the information from `procfs` or `sysfs`, but getting such information from `procfs` or `sysfs` presents many challenges. Each time an application wants to get a system resource information, it has to open a file, read the content and then parse the content to get actual information. If application is running in a container then even reading from `procfs` directly may give information about resources on whole system instead of just the container. Libresource tries to fix few of these problems by providing a standard library with set of APIs through which we can get system resource information e.g. memory, CPU, stat, networking, security related information.

Libresource provides/may provide following benefits:

- **Virtualization:** In cases where application is running in a virtualized environment using `cgroup` or `namespaces`, reading from `/proc` and `/sys` file-systems might not give correct information as these are not `cgroup` aware. Library API will take care of this e.g. if a process is running in a `cgroup` then library should provide information which is local to that `cgroup`.
- **Ease of use:** Currently applications needs to read this info mostly from `/proc` and `/sys` file-systems. In most of the cases complex string parsing is involved which is needed to be done in application code. With the library APIs application can get the information directly and all the string parsing, if any, will be done by library.
- **Stability:** If the format in which the information is provided in `/proc` or `/sys` file-system is changed then the application code is changed to align with those changes. Also if a better way to get information comes in future, like through a `syscall` or a `sysconf` then again application code needs to be changed to get the benefit of it. Library will take care of such changes and the application will never have to change the code.

I agree to abide by the anti-harassment policy

Presenter(s) : YADAV, Rahul (Oracle)

Session Classification : Containers MC