

"Paul Menage" <menage@google.com> writes:

> On 3/13/07, Dave Hansen <hansendc@us.ibm.com> wrote:
>> How do we determine what is shared, and goes into the shared zones?
>> Once we've allocated a page, it's too late because we already picked.
>> Do we just assume all page cache is shared? Base it on filesystem,
>> mount, ...? Mount seems the most logical to me, that a sysadmin would
>> have to set up a container's fs, anyway, and will likely be doing
>> special things to shared data, anyway (r/o bind mounts :).
>
> I played with an approach where you can bind a dentry to a set of
> memory zones, and any children of that dentry would inherit the
> mempolicy; I was envisaging that most data wouldn't be shared between
> different containers/jobs, and that userspace would set up "shared"
> zones for big shared regions such as /lib, /usr, /bin, and for
> specially-known cases of sharing.

Here is a wacky one.

Suppose there is some NFS server that exports something that most machines want to mount like company home directories.

Suppose multiple containers mount that NFS server based on local policy.
(If we can allow non-root users to mount filesystems a slightly more trusted guest admin certainly will be able to).

The NFS code as current written (unless I am confused) will do everything in it's power to share the filesystem cache between the different mounts (including the dentry tree).

How do we handle bit shared areas like that.

Dynamic programming solutions where we discovery the areas of sharing at runtime seem a lot more general then a priori solutions where you have to predict what will come next.

If a priori planning and knowledge about sharing is the best we can do it is the best we can do and we will have to live with the limits that imposes. Given the inflexibility in use and setup I'm not yet ready to concede that this is the best we can do.

Eric

Containers mailing list

