



Contribution ID: 159

Type: **not specified**

Firmware-Assisted Dump, a kdump alternative to kernel dump capturing mechanism

Friday, 20 September 2024 12:45 (45 minutes)

On almost all architectures, kdump has been the default or the only mechanism, to capture vmcore - used for debugging kernel crashes, for close to couple of decades. Fadump (Firmware-Assisted Dump [1], pronounced F-A-Dump) is being used as the alternative dump capturing mechanism on ppc64, for over a decade.

This talk gives a brief introduction of kdump, why fadump was introduced and how it is different from kdump, lists the advantages and pain points in both kdump and fadump dump capturing mechanisms. Then briefly talks about what pain points of fadump have been resolved in the past [and how]:

- relatively high memory reservation requirement for fadump [2]
- restrictions meant for kdump applied to fadump capture kernel, as it also uses /proc/vmcore [3]
- same initrd used for booting production kernel and fadump capture kernel [4]

Then gets into the crux of the talk by explaining:

- How two major pain points for fadump have been resolved recently (v6.10)
 - 1) Service downtime is needed to update resource information on CPU/Memory hot add/remove operations. [5] ensures that this downtime is eliminated completed by moving the resource information update to capture kernel.
 - 2) Fadump doesn't support passing additional parameters to capture kernel. Having that ability will help in disabling components that have high memory footprint and/or complicate capture kernel boot process, but have no real significance in capturing a vmcore. Memory preserving feature of fadump is used to pass additional parameters to dump capture kernel [6].
- The approach being considered to address the last major pain point for fadump - coming up with the right reservation size for fadump capture kernel that works for any system configuration. Explore if fixed reservation can be used for fadump capture kernel irrespective of what the system configuration is, by claiming additional memory required for capture kernel, if any, during capture kernel boot itself.

Lastly, looks at how fadump fares against kdump and what is the architecture support needed to enable/adapt fadump on other architectures.

[1] <https://github.com/torvalds/linux/blob/master/Documentation/arch/powerpc/firmware-assisted-dump.rst>

[2] <https://lore.kernel.org/all/153475298147.22527.9680437074324546897.stgit@jupiter.in.ibm.com/>

[3] <https://lore.kernel.org/all/20230912082950.856977-1-hbathini@linux.ibm.com/>

[4] <https://lists.fedoraproject.org/archives/list/kexec@lists.fedoraproject.org/thread/RPPFTZJMA6HTG3LIBQC7UHX3O27IPO42/>

[5] <https://lore.kernel.org/all/20240422195932.1583833-1-sourabhjain@linux.ibm.com/>

[6] <https://lore.kernel.org/all/20240509115755.519982-1-hbathini@linux.ibm.com/>

Primary author: BATHINI, Hari (IBM)

Presenter: BATHINI, Hari (IBM)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track