

Linux Plumbers Conference 2024



Contribution ID: 228

Type: **not specified**

Unification of RAS feature control - Enhancing EDAC

Thursday, 19 September 2024 11:05 (25 minutes)

Beyond simple error reporting, the CXL specification defines many features related to RAS. Examples being Memory Patrol Scrub and ECS control + features such as PPR directed at the runtime repair of memory. Whilst part of our motivation for looking at this area was to support the CXL features, moves such as OCP RAS API suggest there will be future opportunity for reuse.

There is considerable overlap with existing features distributed across the kernel so when we came to implement Scrub Control we proposed a new subsystem to unify the control interfaces, starting with driver support for the CXL feature and equivalent ACPI RAS2 feature. That proposal was intentionally separate from existing infrastructure to avoid legacy challenges and reflecting the fact that RAS in general has become highly distributed across kernel subsystems with the unification point being tools such as RASDaemon (user space). Perhaps this suffered from the 'lets make a new standard' to unify all these existing standards problem.

<https://lore.kernel.org/all/20240419164720.1765-1-shiju.jose@huawei.com/>

Feedback on that proposal included the question of why EDAC was not suitable, with one valid concern being that a separate overlapping subsystem would divide the review community and reduce quality for everyone.

EDAC is a very mature subsystem carrying a lot of legacy support that makes little sense if we are supporting new features (and hence avoid concerns about ABI breakage). The latest proposal is to use it as a 'home' bringing the benefits of a unified location in `/sys/bus/edac` and ensuring those most familiar with RAS features are heavily involved in the design, but use a modern subsystem design with simpler user of the kernel device model.

This session will include:

- * Motivations for exposing the features at all.
- * A brief summary of the proposed design (it's simple so won't take long!)
- * An outline of future roadmap (if time)

Inputs sought on:

- Relevance of use cases - which ones matter as we need motivating example / user space code.
- Does unification actually make sense?

Primary authors: JOSE, Shiju; SALVE, Vandana; CAMERON, Jonathan (Huawei Technologies R&D (UK))

Presenters: JOSE, Shiju; SALVE, Vandana; CAMERON, Jonathan (Huawei Technologies R&D (UK))

Session Classification: Compute Express Link MC

Track Classification: Compute Express Link MC