

---

Subject: Re: Sun UltraSPARC T1 CPU architecture compatibility

Posted by [dev](#) on Thu, 25 May 2006 06:29:56 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

I've just took a look at the recent patch-022stab077-core for arch/ia64 changes (i386 and x86\_64 have quite lots of OpenVZ unrelated mainstream fixes and 4GB split - noise) and can summary up the changes required for arch specific code:

- UBC: need to account any platform specific VMAs created by hand in arch specific code. i.e. if there are calls of `insert_vma_struct()` this should be accounted with `ub_memory_charge()`. Didn't find such this on sparc64.

- if there are user triggerable `printk()`'s (related to the user, not the system as a whole) better replace them with `ve_printk()`. Otherwise user can flood (DoS). minor actually.

- call to functions `find_task_by_pid()`, `for_each_process()` and `do_each_thread()/while_each_thread()` should be replaced with it's counterparts - `find_task_by_pid_XXX()`, `for_each_process_XXX()` and `do_each_thread_XXX()/while_each_thread_XXX()`, where XXX is 'all' or 've'. 'all' means that all system processes in the system will be scanned, while 've' means that only VE (VPS) accessible from this task (current context - `get_exec_env()`) will be visible. So you need to decide, whether the code in question is about system or user context.

- `task->pid` should be changed with `virt_pid(task)` in some places. The rule is simple: user should see only virtual pids, while kernel operate on global pids. e.g. in signals, virtual pid should be delivered to app.

- in interrupt handlers one need to set global host (VE0) context. i.e. `set_exec_env()`, `set_exec_ub()`. i.e. interrupt handlers are running in VE0 context.

- in `kernel_thread()` one needs to prohibit kernel threads in VE. mostly security related...

- `show_registers()` better to extend to show current VE.

- `utsname` should be virtualized. this mostly means that 'system\_utsnames' should be replaced with 've\_utsname'. See any arch code for this.

- some exports will be required. e.g. `show_mem()` and probably `cpu_khz`. easy.

- everything else are bugfixes.

all these are straightforward and really simple, so it should take a few hours to do so.

---