

ORACLE®



ORACLE[®]

One balloon for all - towards unified balloon driver

Daniel Kiper - daniel.kiper@oracle.com

Current situation

- Three separate balloon drivers for KVM, Xen, VMware
- Despite implementation differences their cores are almost identical
- There are differences in cooperation with other subsystems like memory hotplug or tmem
- Xen balloon driver has support for selfballooning
- Lots of duplicate code which is not that great

Do we have an example of a good solution?

- Generic memory hotplug subsystem
- Core functionality is hardware independent
- Architecture specific parts are minimal and provide direct communication with underlying hardware

Solution proposal

- Core functionality should be hardware independent and placed in MM subsystem
- It should cooperate on hardware independent level (if it is possible) with currently available features (e.g. memory hotplug, tmem) or with features currently under development (movable pages to minimize memory fragmentation which reduces large contiguous memory blocks)
- Selfballooning ready
- Support for hugepages
- Standard API and ABI (nice but not must) if it is possible
- Architecture specific parts should communicate with underlying hypervisor and hardware (if needed) only

Crazy idea

- Replace ballooning with memory hotplug offline/hotremove support
- However, ballooning operates on single pages when memory hotplug works on groups of pages known as sections (sometimes its size is measured in hundreds of MiB; it depends on architecture)

ORACLE®