



Kdump: Introduction and Overview

Cong Wang
Kernel Generalist
Xiyou.wangcong@gmail.com


What are kexec and kdump?

- **Kexec != kdump**
- **Kexec is a way to load another kernel from the current one.**
- **Kdump is a kexec based kernel crash dumping mechanism.**
- **Kdump boots the second kernel from reserve memory when crashing.**

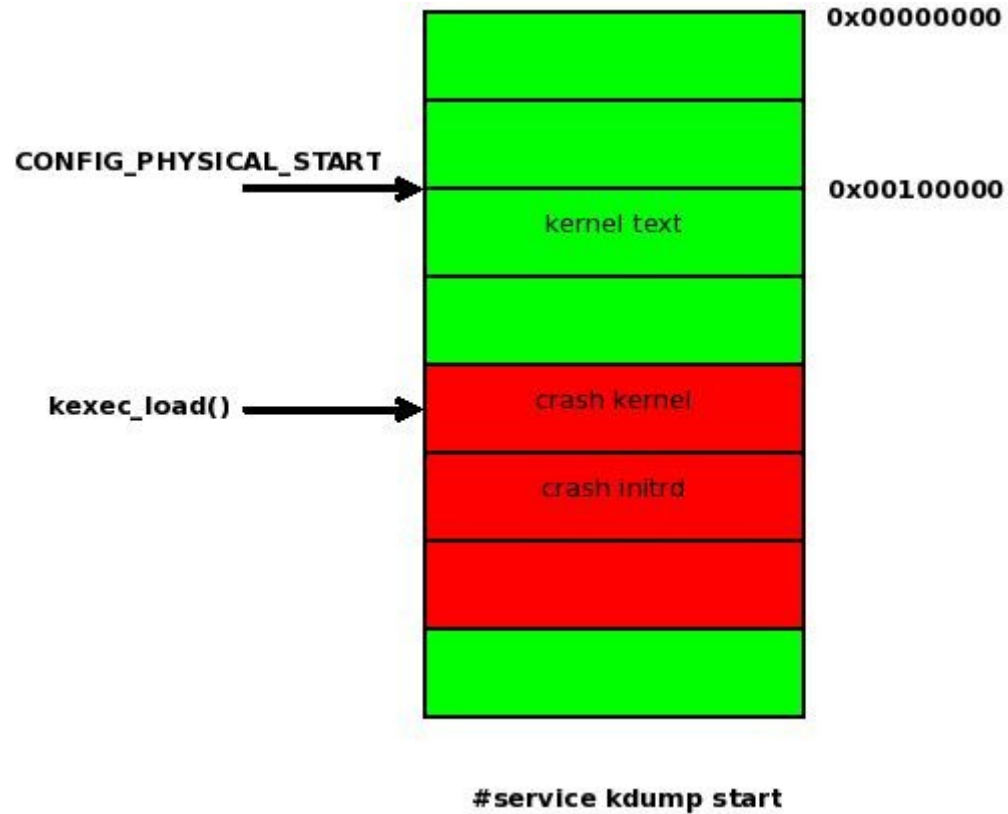
Using kdump

- Reserve memory for crash kernel: `crashkernel=128M@0`
- Load the crash kernel and initrd: `#service kdump start`
- Crash kernel will be automatically booted when crashing
- Core will be dumped as `/proc/vmcore` in the second kernel
- Core will be copied to desired place
- Use 'crash' to analyze the saved core.

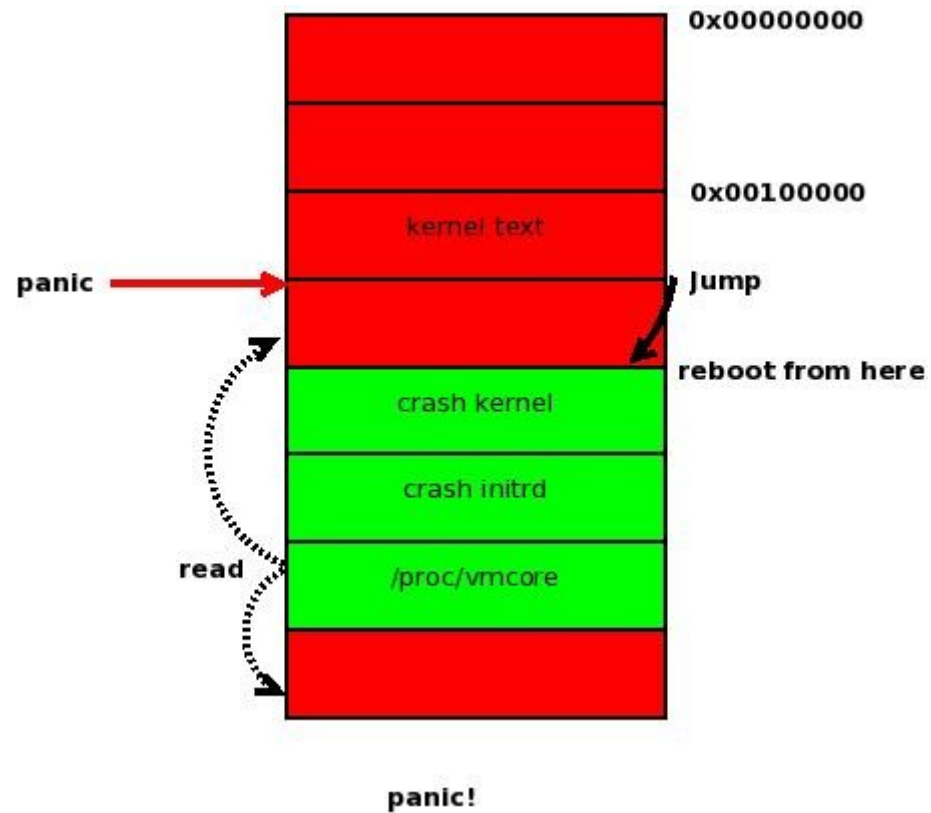
How it works?



How it works?



How it works?



Kernel configurations

- **CONFIG_RELOCATABLE=y**
- **CONFIG_CRASH_DUMP=y**
- **CONFIG_PROC_VMCORE=y**
- **CONFIG_KEXEC=y**
- **CONFIG_PHYSICAL_START=0x1000000**

`/etc/kdump.conf`

- **Basically, the second kernel will just copy the core and reboot.**
- **Various ways to copy the core.**
- **This can be tuned via `/etc/kdump.conf`**
- **RTFM**

The challenges

- **Driver init code needs to be reentrant!**
- **Kernel needs to be relocatable if using the same kernel.**
- **We need to construct a new initrd for the second kernel.**
- **We need to dump the core as the user specified.**

The limitations

- **Not all arch support kexec.**
- **64-bit kernel needs to go 32-bit first, 4G limited.**
- **Initrd can only be loaded to a limited address.**
- **Very limited memory in the second kernel.**
- **Virtualization, especially Xen.**
- **Some kernel parameters are harmful to kdump.**

The problems

- **crashkernel=X@Y does NOT always success.**
- **Check /proc/cmdline**
- **Check /proc/iomem**
- **Check dmesg: `dmesg | grep -i crashkernel`**
- **Check /sys/kernel/kexec_crash_size (RHEL6)**
- **Try bootmem_debug**

Crash utility

- **Kernel-debuginfo is needed.**
- **`/dev/crash`, `/proc/kcore`, `/dev/mem`**
- **Useful commands: `bt`, `whatis`, `struct`**
- **Read the white paper from David Anderson**

References

- http://people.redhat.com/anderson/crash_whitepaper/
- <http://lse.sourceforge.net/kdump/documentation/ols2005-kdump-presen>
- [Documentation/kdump/kdump.txt](#)
- [kdump.conf\(5\)](#)

Questions?

