

# To crash or not to crash: If you do, at least recover fast!



Guilherme G. Piccoli (Igalia)

2023-10-19 / XDC

# Context

- Steam Deck game console -> SteamOS Linux distro
  - Arch Linux based
- No official tool for Arch Linux to collect kernel crash logs / data - that would be nice to have
- What data do we want to collect? What's useful enough?
- Focus on being lightweight and don't waste resources

# Kernel oops / panic

- Oops is when kernel detects a bad condition / error
  - Panic is the kernel reaction, if set!
  - Potential panic causes:
    - NULL pointers (concurrency / SMP)
    - Lockups (soft/hard, hung tasks)
    - Artificial triggered (sysrq), BUG() clause
  - What can we do?
    - Reboot (and pretend nothing happened!)
    - Or...collect data: vmcore, dmesg (kernel log)

# Collecting (all the) data: kdump

- Jump to a new kernel (kexec), using pre-reserved RAM area
  - Preserves memory content of the crashed kernel
  - Collects broken kernel's core dump (vmcore)
- Pros: (Almost) full copy of crashed kernel memory
  - Post-mortem analysis (easy to share)
- Cons: Pre-reserved memory required (>200M usually)
  - Risks during kernel booting / vmcore collecting
  - PCI devices reset: potential nightmare

# Collecting (most) data: pstore

- Saves the kernel log in a persistent storage/backend during panic time
  - Multiple backends: RAM, UEFI, block device
  - Lightweight: used on embedded (and chromebooks)
- Pros: very fast and (hopefully) transparent process
  - {Small amount / No} memory required
  - Doesn't require kexec support
- Cons: does not save a full vmcore
  - Some backends may have limited space (UEFI)

# Presenting: kdumpst

- Kdump/pstore alone are kernel techs
  - They require userspace counter-parts
- **kdumpst** is an Arch Linux pstore/kdump tool
  - Available on **AUR**
  - Supports GRUB and initcpio / dracut (systemd-boot plans)
  - Defaults to pstore mode (but supports kdump); currently only ramoops backend (UEFI plans)
  - Customizable: sysctls, `crashkernel`, log compression
- Present on SteamOS as Deck's crash collecting tool



Feel free to reach me on IRC (**gpiccoli** - OFTC/Libera)

