dm-userspace

Dan Smith – danms@us.ibm.com
LTC Open Hypervisor Team
IBM Linux Technology Center

Somtime in 2006

Overview

- What it does
- Components
- IBM's goals
- Status
- Community interest

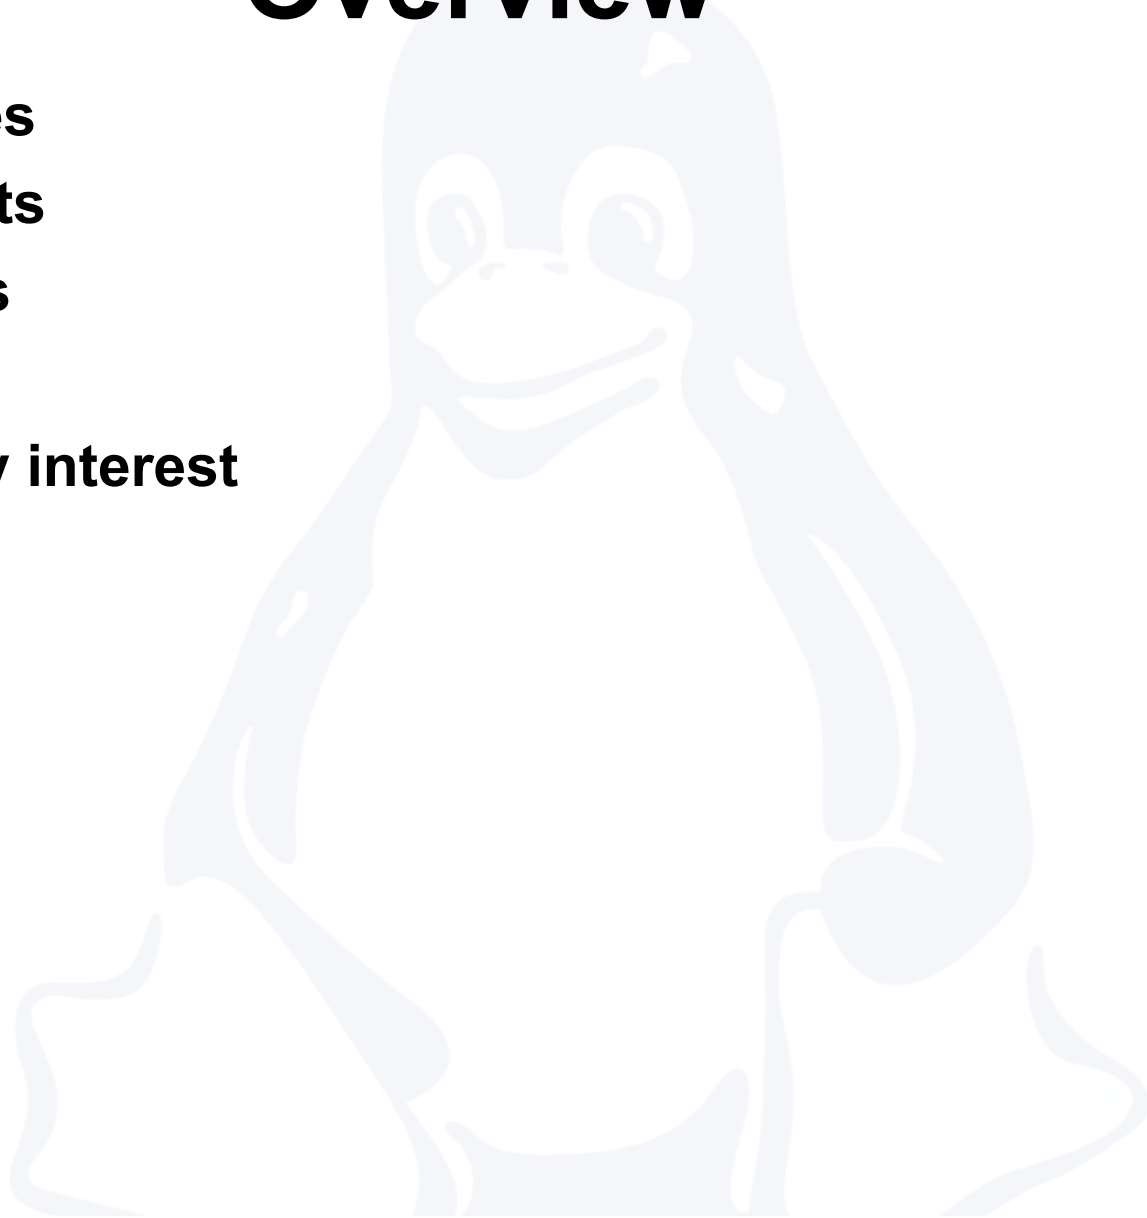

What it does

- Allows us to implement device-mapper functions in userspace
- Provides Xen-independent CoW functionality
- Allows mounting qcow,vmdk*,dscow,etc on a vanilla Linux system

*Event

Components

- Kernel module:
  - Allows userspace to implement a device-mapper extension
  - Control over where (device and sector) IO requests go
  - Faults reads and writes separately
  - Caches maps (or optionally, does not)
  - Optionally faults IO completion to userspace
- Userspace daemon: cowd
  - Provides plugin architecture for implementing formats
  - Currently provides a identity cow format

IBM's goals

- Provide CoW with block devices (iSCSI, nbd, LVM)
- Eventually more support for file-based cow (with loop-ng)
- Support normal persistent domains as well as transient domains
- Push upstream (kernel, and device-mapper)
- Other interesting tricks, like live LVM migration between machines

Status

- Kernel module is largely stable. Waiting final approval from DM maintainer
- Userspace daemon is still ugly, but works
- Format plugin (dscow) works
- Have run fsstress for correctness
- Have run dbench on domain0
- for performance:
  - dm-user and dscow, backed by LVM, cow on LVM, in fully-sync (safe) mode – 187 MB/s
  - dm-user and dscow, backed by LVM, cow on LVM, in transient mode – 248 MB/s
  - Native on LVM (no dm-user, no cow, straight to LVM) – 305 MB/s

Community Interest

- DM maintainer wants to push kernel changes upstream and integrate base userspace functionality into libdevmapper
- Several other groups (inside and outside IBM) are interested in (or are currently) using dm-userspace in their projects
- IBM wants dm-userspace/cowd in the Xen tree for exposure and testing. Expect it to mostly be removed when major parts hit upstream