

# Improving Xen security through domain-zero disaggregation

Derek Murray

University of Cambridge Computer Laboratory

Xen Summit

15<sup>th</sup> November 2007

Santa Clara, CA

# Outline

- What is disaggregation?
- What did we do?
- How does this affect you?

# Motivation

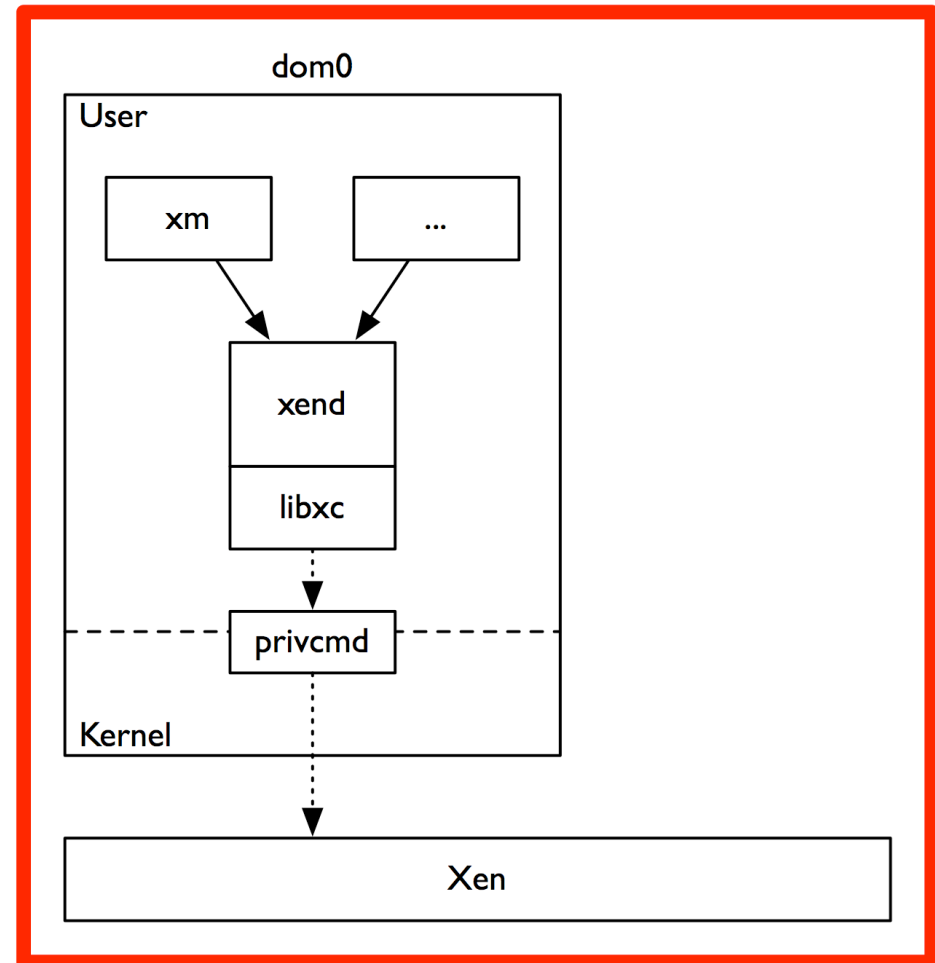
- Hypervisors are supposed to be secure...
- ...but Xen requires us to trust Dom0
- This results in a huge Trusted Computing Base
  - The hypervisor
  - The Dom0 kernel
  - *Everything running as root in Dom0 user-space*
- Why?
  - `xc_map_foreign_range`

# What is disaggregation?

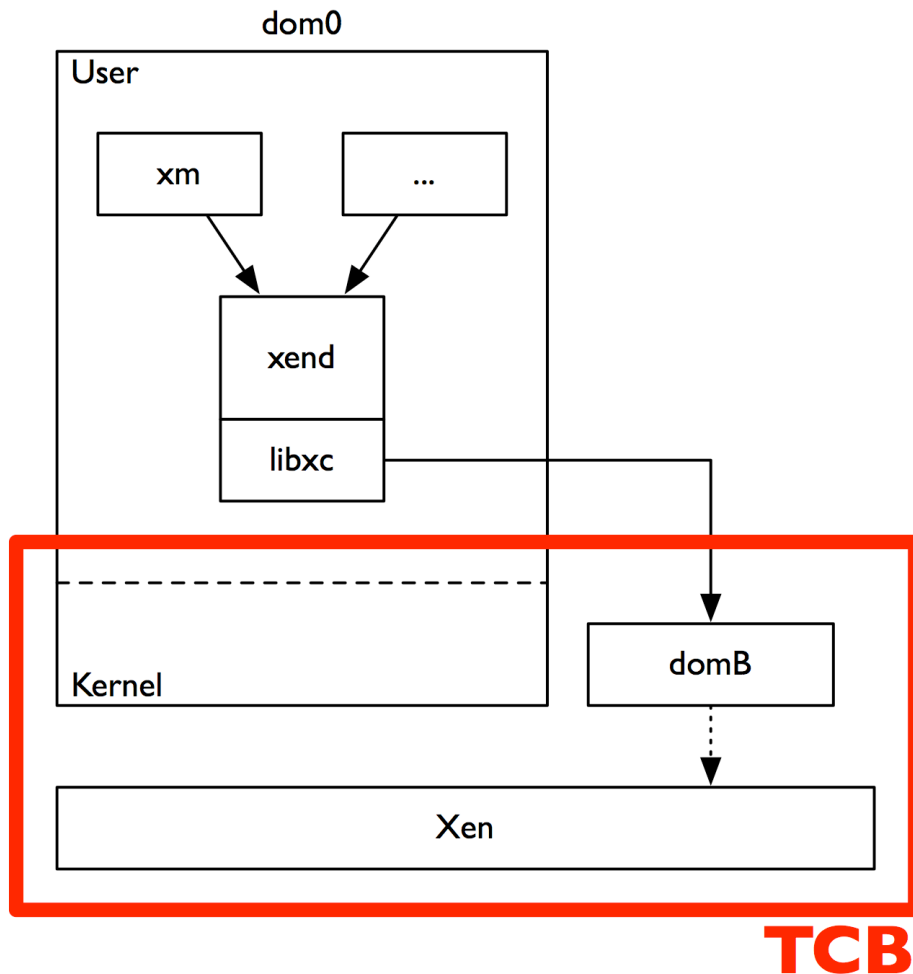
- Starts with a large, trusted piece of software
- Identify the components that require privileges
- Define a narrow interface between those components and the untrusted applications
- Move the trusted components to another domain
- Case study: the domain builder

# Disaggregation - Before

- Currently, the privcmd driver allows user-space tools to make privileged hypercalls
- Any code which can use privcmd must be trusted
  - Including anything that links against libxc



# Disaggregation - After



- Remove the privcmd driver from Dom0
- Only change is to libxc: user APIs unchanged
- Now we must trust only the hypervisor, DomB and the Dom0 kernel

# How is it implemented?

- Based on a port of libxc to MiniOS
- DomB booted alongside Dom0
- Inter-Domain Communication, based on DICE
- File system calls proxied to Dom0
- Minor changes to user-space tools

# Minor changes?

- Several tools use direct foreign mappings
  - Console, XenStore
- Want to retain functionality but control sharing

# Solution: gntdev

- Obvious solution is to use grant tables
- Only kernel can map grant references
  - blktap maps these pages to user-space
- gntdev generalises this approach

# xc\_gnttab API

- libxc bindings for the gntdev driver
  - xc\_gnttab\_open
  - xc\_gnttab\_close
  - xc\_gnttab\_map\_grant\_ref
  - xc\_gnttab\_map\_grant\_refs
  - xc\_gnttab\_munmap

# Porting to gntdev

- “Magic pages”
  - Use grant references instead of MFNs
- `xc_map_foreign_range`
  - Use `xc_gnttab_map_grant_ref(s)`

# A request

- When you write new tools that use shared memory
  - Use grants instead of MFNs
  - Use `xc_gnttab_*` instead of `xc_map_foreign_range`
  - Email me if you need help!
    - [Derek.Murray@cl.cam.ac.uk](mailto:Derek.Murray@cl.cam.ac.uk)

# Future work

- Refine and submit to xen-unstable
- HVM support
- Save/restore/migration support
- XSM integration
- gntdev
  - Granting *from* user-space
  - Solaris support

# Conclusions

- Existing architecture places all of Dom0 in TCB
- Disaggregation eliminates Dom0 user-space from TCB
- gntdev is the way forward for user-space tools