

Efficient Disk Space Management for Virtual Machines

Abhishek Gupta, Norman C. Hutchinson

{agupta, norm@cs.ubc.ca}

Department of Computer Science,
University of British Columbia, Vancouver,
Canada

Virtual Machines (VMs) and File Systems

- Issues of Interest
 - Common Disk Images and Software Configurations.
 - Fast Replication – Snapshots.
 - Efficiency in Data Access.
 - Space Reclamation.

Existing Solutions

- General Theme – Virtualize the block layer.
Use Copy-On-Write Logical Volumes.
 - Linux Logical Volume Manager
 - Kernel Space
 - Working copy-on-write snapshot feature
 - Available with latest Linux 2.6 release.
 - Parallax
 - User Space
 - Utilizes a radix tree data structure to enable fast copy-on-write snapshots.
 - Available in the tools directory under the Xen source tree.

Linux Logical Volume Manager (LVM2)

- Shortcomings
 - Implementation of snapshots – Aggressive copy-on-write scheme – Too intensive for a single machine hosting scores of VMs.
 - No support for recursive snapshots.

Parallax

- Shortcomings.
 - No benchmarks available – Need to establish the efficacy of the radix tree structure.
 - Block address translation mandates a traversal through the height of a radix tree.
 - Linear accessibility of blocks is not preserved.
 - No space reclamation.

Our Roadmap

- Quantify Existing Solutions
 - A Radix Tree target for LVM – Complete.
 - Performance Characterization – In progress.
- Fix problems with existing solutions or suggest a new data structure which preserves linear accessibility while supporting fast snapshots.