

# **Gazing Into Mirrors**

(Solaris)

**January 09, 2003**

Rodney P. Rutherford

Par3 Concepts

## What We'll Cover

- DRP – A brief overview
- Boot Drive Options
- Disk Duplication via native tools
- Solstice DiskSuite (SDS)
- Veritas Volume Manager (VXVM)

# DRP – Disaster Recovery Planning

- What is Disaster Recovery?
- Why Should I care?

# Boot Drive Options

Mirror, mirror, on my drive

What's the best way to keep me alive?

- Options we will discuss
  - Disk Duplication via scripts
  - Disk Mirroring via SDS
  - Disk Mirroring via VXVM

## Disk Duplication Overview

- Duplication done via scripts
- Can schedule via native cron utilities
- Uses native UNIX tools
- Consists of the following basic steps:
  - Prep destination drive
  - Copy data from source to destination drive
  - Make destination drive bootable

## Disk Duplication Pros/Cons

- Pros:
  - Changes are not automatically updated to mirror
  - Only run as often as you choose
- Cons:
  - Changes are not automatically updated to mirror
  - Requires manual maintenance

# Disk Duplication Tools

- Most common methods used:
  - dd – Disk to Disk Copy
  - ufsdump/ufsrestore – Inode level dump/restore
- Other options:
  - Pretty much any other UNIX tools
  - tar, cpio, etc.
  - Flash archives

# Disk Duplication Tool Comparison

- dd:
  - Block for block copy
  - Cannot copy larger drive to smaller drive
  - 2<sup>nd</sup> drive automatically has matching label to 1<sup>st</sup>
  - Any filesystem corruption is copied

# Disk Duplication Tool Comparison

(continued)

- **ufsdump/ufsrestore:**
  - Inode based copy
  - Must create desired label on destination drive
  - Can use any size drives, as long as destination filesystems are equal or larger than source filesystems
  - Filesystem must be valid prior to dump

# Disk Duplication Sample Script

## Solstice DiskSuite Overview

- Free RAID Software from Sun Microsystems
- Unbundled prior to Solaris 8
- Included with Solaris 8 and 9
- Requires 1 free slice on boot drive for meta db
- Packages
  - SUNWmdr Solstice DiskSuite Drivers (root)
  - SUNWmdu Solstice DiskSuite Commands
  - SUNWmdx Solstice DiskSuite Drivers (64-bit)

## Solstice DiskSuite Overview

- DiskSuite uses virtual disks to manage physical disks and their associated data. In DiskSuite, a virtual disk is called a *metadevice*. A metadevice is functionally identical to a physical disk in the view of an application. DiskSuite converts I/O requests directed at a metadevice into I/O requests to the underlying member disks.

## Solstice Disksuite Overview

- command line utilities in /usr/opt/SUNWmd/sbin
- driver modules in /kernel/drv and /kernel/misc
- daemons in /usr/opt/SUNWmd/sbin
- administrative files in /etc/opt/SUNWmd
- metadevices are named /dev/md/{dsk|rdsk}/dn, with n from 0 to 127 by default
- The /etc/opt/SUNWmd/md.tab file can be used to configure SDS automatically.

## SDS Commands

- `metastat -p > <filename>`  
This will output your configuration in md.tab format
- `metainit -a`  
This command reads the md.tab file and sets up the configuration accordingly

## SDS Commands

- **Metastat**  
shows the configuration and status of all metadevices
- **Metadb**  
tells the location and status of locally configured replicas

## SDS Pros/Cons

- Pros:
  - FREE!
  - Simple to install and configure
  - Doesn't mess with native Solaris slices
- Cons:
  - Not integrated with VXFS, VCS, etc.

# SDS Sample

## VXVM Overview

- Commercial product from Veritas
- Costs \$\$\$ and requires licensing – unless used with some Sun storage, such as SSA, A5xxx, etc.
- Sun resells as Sun Enterprise Volume Manager (SEVM)
- Tightly integrated with Veritas File System (VxFS) for data drives, VCS, etc.

## VXVM Terms

- **Volume - Virtual Disk Partition**
  - Virtual disk device (“/dev/vx/[r]dsk”)
  - Only object that the OS interacts with.
  - Made up of one or more **plexes**.
- **Plex - Mirrors.**
  - Made up of one or more subdisks.
  - Stripes, concatenations, or RAID-5

## VXVM Terms

- **Subdisks**
  - Allocated from public region of disks.
  - Make up **Plexes**.
- **Diskgroups**
  - Set of disks sharing a common configuration.
  - Owned by only one host at a time.
  - Default group is *rootdg*.
  - Diskgroups can be exported or imported among multi-connected hosts. (*except rootdg on each host*)

## VxVM - Required Packages

- Packages Required:
  - VRTSvxvm (watch for new versions)
- Typically included:
  - VRTSvmsa - GUI interface (was VRTSvxva)
  - VRTSvmman - Man pages
- Optional:
  - VRTSvmdev - Developer kit
  - VRTSvmdoc - Postscript/PDF documentation

## VXVM Sample

- Run ‘vxinstall’ to create rootdg
- Encapsulate root disk
  - Encapsulation requirements
- VXVM requires a license
- Manually input using ‘vxserial’ or ‘vxlicense’
- View with ‘vxserial -p’ or vxlicense -p
- Beware temp license expiration!

# VXVM Commands

# VXVM Example

# VXVM Best Practices

- Reference Sun Blueprint

# Mirroring Comparison Chart

# Summary

Mirror, mirror, on my drive

What's the best way to keep me alive?

- Answer:
  - Whatever you are most comfortable with and will actually use and maintain as an admin!
  - You want to set a standard and then implement that datacenter wide, and then DOCUMENT and TEST that standard.
  - Ideally, that should consist of either SDS or VXVM to mirror the boot drives, an SDS/VXVM controlled Hot Spare, plus a non-SDS/VXVM controlled clone drive that is duplicated via scripts whenever possible.

# Questions and Answers

***The End!***