Solaris Volume Manager : Mirror Hotsparing

SVM Mirror Hotsparing Example

- Metastat -p output

  mirror -m /dev/md/rdsk/stripe1 /dev/md/rdsk/stripe2 1
  stripe1 1 1 /dev/rdsk/c1t9d0s1 -h hspool1
  stripe2 1 1 /dev/rdsk/c1t10d0s1
  hspool1 c1t10d0s4

`

SVM Mirror Hotsparing Example

- Remove disk by cfgadm unconfigure

  bash-3.00# cfgadm -c unconfigure c1::dsk/c1t9d0

- Perform dd to error out the submirror and trigger hotspare replacement

  bash-3.00# dd if=/dev/zero of=/dev/md/rdsk/mirror bs=1024 count=1024

  Jan 17 13:33:45 dum md_mirror: WARNING: md: stripe1: /dev/dsk/c1t9d0s1 needs maintenance

  Jan 17 13:33:45 dum md_stripe: NOTICE: md: stripe1: hotspared device /dev/dsk/c1t9d0s1 with /dev/dsk/c1t10d0s4

  1024+0 records in

  1024+0 records out

  bash-3.00#

Detect error and initiate hs replacement

dd`_start+0x108

libc.so.1`_creat64+0x4

unix`syscall_trap32+0xcc

genunix`copen+0x20c

genunix`vn_openat+0x4c8

genunix`fop_open+0x78

specfs`spec_open+0x420

md`mdopen+0x250

md_mirror`mirror_internal_open+0xf0

md_mirror`mirror_open_all_devs

Detect error and initiate hs replacement

```
md_mirror`mirror_open_all_devs
    md`md_layered_open
        md_stripe`stripe_open
            stripe_open_all_devs
    set_sm_comp_state
    poke_hotspares
```

Detect error and initiate hs replacement

- Open each submirror
  - > Open the underlying components of each submirror
- If the submirror open fails, set the submirror's component state to CS_ERRED and submirror state to SMS_COMP_ERRED
- Call poke_hotspares to initiate hotspare replacement for the failed submirror

Let's look at each of these areas!

Open the submirrors

- mirror_open_all_devs
  - > md_layered_open
    - >md_stripe`stripe_open
      - – stripe_open_all_devs
        - md_resolve_bydevid
        - md_layered_open
    - >dev_lopen /* if it's a regular device */

- stripe_open_all_devs
  - > Open all component of the stripe
  - > Return ENXIO if it fails to open the device

Set submirror state

- mirror_open_all_devs
  - > md_layered_open
  - > set_sm_comp_state
    - >stripe_shared_by_indx
    - >stripe_get_dev
    - >md_getdevname
    - >mirror_set_sm_state
    - >mirror_commit

Set submirror state

- Setting the failed submirror to SMS_COMP_ERRED
  - Get the current state of the failed component
  - Get the devname for the failed component of the submirror
  - Generate a console message of the form
    - WARNING: md: stripe1: /dev/dsk/c1t9d0s1 needs maintenance
  - Set the state of the failed component to CS_ERRED
  - mirror_set_sm_state sets the submirror state to SMS_COMP_ERRED if any of the submirror's component is in CS_ERRED state. Submirror state is set to SMS_ALL_ERRED if all components of the submirror has CS_ERRED state.
  - commits changes for both mirror and submirror

Hotsparing the failed component

- poke_hotspares
  - > check_4_hotspares
    - > check_unit_4_hotspares
      - – check_comp_4_hotspares
        - stripe_hotspare_dev
          - md_hot_spare_ifc
            - hotspares_interface
              - reserve_a_hs
                - find_hot_spare_pool
                - usable_hs
                - set_hot_spare_state (HSS_RESERVED)
          - stripe_replace_dev
        - set_sm_comp_state
        - mirror_resync_unit

Hotsparing the failed component

- check_unit_4_hotspares: for submirrors not in SMS_INUSE state, call check_comp_4_hotspares for each comp of the submirror
- check_comp_4_hotspare
  - If the component is in CS_ERRED state
    - Find a hotspare (call stripe_hotspare_dev)
    - set_sm_comp_state to CS_RESYNC and commit the changes
    - Start a resync on the mirror
- stripe_hotspare_dev, md_hot_spare_ifc, and stripe_replace_dev
  - Get the size of the device to be replaced
  - Find a hotspare with the appropriate size & reserve it
  - Save devt of original component, set component's ms_comp_t values to that of the hotspare `

SVM Mirror Hotsparing