

Saturday, April 26. 2008

Less known Solaris Features: Point-in-time copy with AVS - Part 4: Dependent copy

The mechanism of dependent copies was introduced to get rid of this initial sync, as there are circumstances where this initial copy would pose too much load to the system.

The dependent copy uses the bitmap a little bit differently. The shadow disk doesn't contain all the data, it just contains the changed data. The bitmap is used to keep track of the blocks on the masters which have a copy on the shadow.

Deeper dive The dependent copy is one of the two mechanisms in AVS using the concept of the virtual shadow. Thus the model is a little more complex. Let's assume you create a dependent copy. The initialisation is simple. You move no data. You just initialize the data. When a user or application accesses the virtual shadow volume, it checks in the bitmap if the blocks have changed. If the bitmap signals no change, it just delivers the original block of the master volume.

When you change some data on the master volume, AVS starts to copy data. It copies the original content of the block onto the physical shadow volume at the same logical position in the volume. This is the reason, why master and shadow volumes have to have the same size when using dependent copies. Furthermore AVS logs in the bitmap that there is data on the shadow volumes, the block is dirty in the bitmap. When you access the virtual shadow volume now, the bitmap is checked again. But for blocks declared dirty in the bitmap, the data is delivered from the copy on the physical shadow volume, for all other "clean" blocks the data comes from the master volume.

Resyncing the shadow to the master is easy. Just reinitializing the bitmap. Now all data comes from the master volume again, until you change some data on it.

Advantages and Disadvantages The advantage of this method is its short time of initialisation and syncing from master to the shadow. It's virtually zero. But you lose the advantages of the independent copy. The dependent copy can only be used in conjunction with the original data. This has two main effects. You can't export the copied volume to a different server and you can't use it as a copy for disaster recovery when your master volume has failed.

Furthermore the master and the slave volume still have to be the same size. But the next method for making a point-in-time copy was derived from the dependent copy exactly to circumvent this problem.

Posted by Joerg Moellenkamp in English, Solaris at 21:34