

Saturday, April 26. 2008

Less known Solaris Features: Point-in-time copy with AVS - Part 5: Compact dependent copy

The compact dependent copy is similar to the normal dependent copy. But this dog knows an additional trick: The shadow and the master doesn't have to be at the same size.

Like the dependent copy this methods uses the concept of the virtual shadow volume. So the bitmap is really important. The bitmap of the compact dependent copy tracks the changes. This exactly as with the normal dependent snapshots. But the bitmap for compact dependent snapshots is enabled to store an important additional information. It's enabled to track, where the system has written the changed data blocks on the shadow volume. So you don't have to write it at the same logical position, you can write it at any position on your shadow volume and it can retrieve the old data with this information.

Deeper dive Okay, at start this picture looks pretty much the same like it's counterpart at the normal dependent snapshots. Please note, the additional information on the bitmap volume. To make the compact dependent copy we just initialize the bitmap again to it's "clean" state.

Let's assume that we change the fourth block on our disk. As with the normal copy, the block is declared as dirty. But now starts to work differently. The original data of the master volume is stored to the first free block on the physical shadow volume. In addition to that the position is stored at the bitmap.

The way to read from the shadow volume changes accordingly. When the bitmap signals, that a block is clean, it just passed the data from the master volume to the user or application. When the bitmap signals a dirty, thus changed block, it reads the position of the block on the physical shadow volume from the bitmap, reads the block from there and delivers it to the application or user.

When we change the next block, for example the third one, the same procedure starts. The original data is stored to the next free block, now the second one, on the physical shadow volume and this position is stored in the bitmap together with the dirty state of the block.

Okay, resyncing the shadow with the master is easy again. Just initializing the bitmap.

Advantages and Disadvantages The trick of storing the position with the dirty state has an big advantage. You don't need master and shadow volumes with the same size. When you know, that only a small percentage of blocks change between two point in time copies, you can size the shadow volume much smaller, thus saving space on your disk. In my opinion compact dependent copies are the only reasonable way to go when you want more than one copy of your master volume. The disadvantages are pretty much the same of the normal dependent copies.

Posted by Joerg Moellenkamp in English, Oracle at 22:19