



[□return to original page](#)

HP Surestore Virtual Arrays - Virtual Array Performance Using Large Capacity LUNs on HPUX

ISSUE:

A large capacity LUN configuration may need HPUX tuning for optimal performance using Virtual Array (VA) products.

CAUSE:

While especially in OLTP-environments where there are typically a large number of concurrent I/O-requests going on, having only a small number of very large LUNs may indeed lead to a performance problem. The root cause of this problem has never been on the VA-side. In other words: from the VA7400 or VA7410's perspective, the performance of 10 x 50 GB LUNs per redundancy groups is exactly the same as the performance of a single 500 GB LUN per redundancy group.

The reason for experiencing very severe throughput problems when using just the 2 x 500 GB LUNs is with HP-UX each disk device (more precisely each path to a LUN) has a default scsi queue depth of "8." This means that with just two giant LUNs there can only be a maximum of 16 outstanding I/O-requests to the VA at any time. While the VA may be able to process a lot more I/O requests in parallel, since it may be stuffed with 60 disks in the backend.

SOLUTION: The easiest solution to this problem does not necessarily consist in creating more and smaller LUNs. Use the "`scsictl -m queue_depth=<value> /dev/rdisk/...`" command. This will increase the queue depth for giant LUNs in order to allow sufficient concurrency (for example, the amount of concurrency that will allow the use of the full performance of all disks in the array).

With HP-UX 11.0, each new activation of an LVM volume group (after a reboot or an MC/SG-package start) will reset the scsi queue depth to its default value of "8" (or whatever the kernel setting for "scsi_max_qdepth"). This means that in order to ensure the appropriate scsi queue depth setting for giant LUNs, a script will be required which automatically adjusts the queue depth for these LUNs at boot time or at MC/SG-package start. As this may be perceived as cumbersome, using several smaller LUNs with the default scsi queue depth might be more convenient with HP-UX 11.0. With HP-UX 11i, on the other hand, the scsi queue depth set with "`scsictl -m ...`" will survive de-activation or re-activation of volume groups (including reboots) which is why with HP-UX 11i. It is recommended to use the alternative of using a minimal number of large LUNs (at least one per RG, of course) and then increasing the queue depth for those large LUNs.

content feedback

To help us improve our content, please provide your feedback below.

1. How does the information on this page help you?

very helpful somewhat helpful not helpful

2. Was it easy to find this document?

easy not easy

3. If you selected not easy for question 2, in which section did you expect to find it?

4. Comments:

[submit >>](#)

- Your feedback will be used to improve our content. Please note this form is for feedback only, so you will not receive a response.
- To obtain technical help, try [Contact hp for technical help.](#)

[privacy statement](#)

[using this site means you accept its terms](#)

[feedback to webmaster](#)

© 1994-2004 Hewlett-Packard Company