

Re: Are large RAID stripe sizes useful with FreeBSD?

## Re: Are large RAID stripe sizes useful with FreeBSD?

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2008-03/msg00638.html>

---

- *From:* "Ivan Voras" <[ivoras@xxxxxxxxxxx](mailto:ivoras@xxxxxxxxxxx)>
  - *Date:* Mon, 31 Mar 2008 23:34:04 +0200
- 

On 31/03/2008, Scott Long <[scottl@xxxxxxxxxxx](mailto:scottl@xxxxxxxxxxx)> wrote:

For writes, the performance penalty of smaller I/O's (assuming no RAID-5 effects) is minimal; most caching controllers and drives will batch the concurrent requests together, so the only loss is in the slight overhead of the extra transaction setup and completion. For reads, the penalty can be greater because the controller/disk will try to execute the first request immediately and not wait for the second part to be requested, leading to the potential for extra rotational and head movement delays. Many caching RAID controllers offer a read-ahead feature to counteract this. However, while my testing has shown little measurable benefit to this, YMMV.

Thank you, this is the kind of explanation I hoping for. One more thing: is TCQ (e.g. the SCSI variant) orthogonal to this?

---

[freebsd-current@xxxxxxxxxxx](mailto:freebsd-current@xxxxxxxxxxx) mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-current>

To unsubscribe, send any mail to "[freebsd-current-unsubscribe@xxxxxxxxxxx](mailto:freebsd-current-unsubscribe@xxxxxxxxxxx)"