
Subject: Re: point_lun is slow

Posted by [George McCabe](#) on Mon, 01 Nov 1999 08:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

hello and thanks for the reply Liam,

i reason that the record size has to be larger than the separation between consecutive extracted elements in order for there to be a gain in performance. the number of reads is reduced. also reading sequentially in a chunk is faster than seeking to each datum, an additional gain.

let's say the ratio of data read to data extracted is 'R'.

(the time it takes to seek and read the extracted data only) divided by (the time it takes to read sequentially and extract data) is 1 when R is what value? my guess is 100. what are your thoughts?

george

In article <38187A7E.8E60823A@ssec.wisc.edu>,

Liam Gumley <Liam.Gumley@ssec.wisc.edu> wrote:

> The following pseudo-algorithm reads records (chunks) of data from the
> disk in sequential order. Only records that cover the specified read
> locations are actually read from disk. Each record is only read once.
>

> Sort the array of read locations from lowest to highest
> Set the record size to 512 bytes (you can experiment with record sizes)
> Set the old record number to -1
> Start a loop over the read locations
> For this read location, compute the record number in the file
> If the record number is different than the old record number
> Read the current record
> Set the old record number to the current record number
> End If
> For this read location, compute the byte offset within the record
> Extract data from the record at the byte offset
> End Loop
>

> This method should be just as efficient for small or large numbers of
> read locations.

>
> Cheers,
> Liam.

>
> --
> Liam E. Gumley
> Space Science and Engineering Center, UW-Madison

> <http://cimss.ssec.wisc.edu/~gumley>

Sent via Deja.com <http://www.deja.com/>
Before you buy.
