Subject: Endian-ness
Posted by Jonathan Joseph on Fri, 08 Feb 2002 20:47:21 GMT
View Forum Message <> Reply to Message

Is there a system variable that gives the endian-ness of the current hardware? I am reading a file which tells me the endian-ness of the data, and I'd like to swap_endian if it is different from the current hardware. In lieu of finding a system variable to compare to, I have done this:

test_int = 1
byteorder,test_int,/ntohs
big_endian = test_int eq 1

"network" byte order is big-endian, so I convert a 1 to the host byte-order and see if it's still a 1.

This way works fine, but it seems as thought I'm missing something. Is there a better way? Using the byteorder routine to convert the data is not an option (unless it's been improved for 5.5).

Thanks.

-Jonathan

Subject: Re: Endian-ness
Posted by Liam E. Gumley on Mon, 11 Feb 2002 16:12:55 GMT
View Forum Message <> Reply to Message

Jonathan Joseph wrote:
[stuff deleted]
> Am I correct in assuming that Liam's functions don't need to be
> specifically cast to Long? In other words,
>
> byte(1,0) would yield the same result as byte(1L, 0L)

That's just my FORTRAN background showing: I prefer to use explicit casts.

Cheers,
Liam.
Practical IDL Programming
http://www.gumley.com/

Subject: Re: Endian-ness

Posted by Robert Stockwell on Mon, 11 Feb 2002 18:01:21 GMT
View Forum Message <> Reply to Message

Jonathan Joseph wrote:

> Thank you David and Liam,
>
> Liam's result seems more aesthetically pleasing (no offense David).
> Unfortunately, I don't think I can use the swap_if_big_endian and
> swap_if_little_endian keywords to OPEN, because whether I swap or
> not depends on BOTH the hardware and the file. I have no a-priori
> knowledge of the endian-ness of the file until I've already
> opened it. Although, I guess I could open the file, figure out
> what it is, close it and then re-open it.

This is exactly what I do. I have a rule that all data I write
shall be little endian, but there still can be files that are
big endian. I have a small range of possible values, so I read the
data in little endian, check the range, and if it exceeds the
range, reread the data as bigendian.
That seems to be the only way to tell if a file was written
as big/little endian. Of course, one could put a header
on the file indicating how it was written, but if someone else
writes the data, you are out of luck.

Cheers,
bob