Subject: Re: strings and memory usage
Posted by Conor on Wed, 22 Aug 2007 19:00:56 GMT
View Forum Message <> Reply to Message

On Aug 22, 2:54 pm, Conor <cmanc...@gmail.com> wrote:
> Does anyone know how IDL stores strings? I'm creating some very large
> string arrays and running out of memory when I shouldn't. So, for the
> following example I'm using the linux command 'top' to keep track of
> memory usage on a per-process basis. In the beginning, IDL is using
> 59 megabytes. Then, I create a string array with 5 million elements
> like this:
>
> test = strarr(5000000) + 'asdf'
>
> Now I have a string array with 5,000,000 elements, each with 4
> characters in it. According to top idl is now consuming 177
> megabytes! That means that each string takes up an average of 23
> bytes! To make matters worse, when I delete test (delvar,test) IDL
> drops back down to 120 megabytes!
>
> What in the world is going on? Naievly, I would expect a string array
> with strings 4 characters long to take up an absolute maximum of 8
> bytes per element (4 bytes for the characters, 2 bytes for the length,
> and maybe two bytes for pointers). Why is it taking up 23 bytes???
> Am I just confused about something? Also, why doesn't the memory
> usage drop back down to it's original value? I did notice one thing.
> When I then created more large variables, the memory usage didn't
> increase right away, so maybe IDL is clearing the memory but not
> releasing it to the operating system. Still, I find these problems
> very troubling. Is there something very wrong with the string arrays
> in IDL, or am I just being silly?

For comparison, when I execute:

test2 = fltarr(1000000)

memory usage for IDL goes up by 4 megabytes, according to top -
precisely what I would expect it to.