Subject: Re: Philosophical Question about NAN
Posted by David Fanning on Mon, 17 Nov 2008 16:02:05 GMT
View Forum Message <> Reply to Message

wlandsman writes:

> I agree with the sentiment but also note that always setting /NAN
> incurs a non-trivial performance penalty, e.g.
>
> IDL> a =3D randomn(seed,10000,2000)
> IDL> t =3D systime(1) & b =3D total(a) & print,systime(1)-t
> 0.25451803
> IDL> t =3D systime(1) & b =3D total(a,/nan) & print,systime(1)-t
> 0.35278893
>
> I've thought at times that arrays should carry a hidden bit saying
> whether or not they include NaN values, but this introduces other
> overhead problems.

I guess I would argue that in the overwhelming number of
cases in my experience, the performance penalty is trivial.
I'm calling these routines a couple of times at most. And
I am not arguing for the elimination of the keyword, only
that the default value could be changed. Thus, if I *was*
experiencing a performance penalty, and I was certain I
had good numbers, I could always set the NAN keyword to 0.

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Sepore ma de ni thui. ("Perhaps thou speakest truth.")