Subject: Re: getting structure elements by tag name string
Posted by Pavel A. Romashkin on Thu, 09 Jan 2003 17:20:33 GMT
View Forum Message <> Reply to Message

Hi JD,
Reding thriugh your nice explanation, I wondered: how are SQL engines implemented so that they look up pretty much random things but can be incredibly fast for even multi-million record, randomly organized arrays? They do create indices, true, but still they serch them somehow, and that is fast. I'd say the speed is size-dependent, but it doesn't seem to be linearly proportional.
Cheers,
Pavel

JD Smith wrote:
>
> Unfortunately, it's quite difficult to build good hash functions and
> matching hash table sizes which strike a decent balance between wasted
> memory, lookup speed, and flexibility for growing/shrinking the hash
> array. There's no reason a hash couldn't be implemented using regular
> IDL arrays and pointers, along with a suitable hash function, but I
> suspect it would be an underperformer, compared to more finely tuned
> offerings of other languages. Might be a good exercise to try though.