

---

Subject: Re: Large data ahead

Posted by [Tom1](#) on Wed, 30 Jan 2019 09:22:22 GMT

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

---

Hi,

How about making the containers internally int64 and then offering both int and int64 interfaces in parallel? I'm sure there would be some tight spots (with e.g. serialization to maintain binary compatibility with int indexed serialized content in a way that they maintain the old binary serialization format until the element count goes beyond reach of int.) In most cases we know for sure already that the item count will absolutely remain below 2G, so there is no point in using int64 indexing in those cases, and therefore int offsets can be retained throughout the rest of such code.

Anyway, I have no hurry for this at the moment, so let's just give it some time...

Best regards,

Tom