FastInfoset

FastInfoset

FastInfoset is an interestingly controversial topic.

The key idea is that the speed of marshalling and unmarshalling content can be improved by moving away from a textual markup language to a binary one.

This is done at the expense of human readability or accessibility. No longer can you look at messages through a logging proxy server. No longer can you write normative 'correct' requests and responses to XML files in your SCM repository, run ant's <schemavalidate> over them and check that your XSD matches expectations.

As such, it is firmly in the camp of *XML is a detail end users don't need to care about*, which is not far behind *WSDL and XSD are things the toolkit handles for you.* It is also not XML, not in its classic sense.

Between nodes that support it, FastInfoset may deliver tangible speedups through improvements in parse time. This may be at the expense of flexibility, but if the endpoint is written in a contract-last form, there is usually little flexibility in there anyway. Few Java methods are set up to deal with arbitrary amounts of incoming XML content, especially in unknown schemas, so will not be any less flexible by adopting FastInfoset.

The other issue with FastInfoset is security. Whoever implements the parser had better design it to resist malicious content.