













































Core Working Group Meeting Notes (25-Apr-2019)

Attendees:

Participants (11)

Q Type to filter...

	Trevor Conn (Host, me)			
	Lenny Goodell (Intel)			
	Akram Ahmad (Dell)			
	Anthony Bonafide			
	Brad Kemp			
	Eric Cotter			
	Jim Wang			
	Michael Estrin (Dell)			
	Rebekah			
	Steve Osselton			
	Toby Mosby (Intel)			

Old Business

- CBOR
 - Benchmarking
 - Checksum (Toby's data)
 - Unmarshaling / Marshaling (Anthony's data)
 - Revisit this with the group
 - There is a trade-off in the above approach – latency versus memory
 - What is the proper use approach for the most common use case we expect (looking to Intel for this)
 - Need a decision on this today.
 - *Preference is for the checksum solution*
 - *MD5 or better (xxHash is default)*
 - *Assuming no large 3rd party deps*
 - *Injection of hash algorithm via configuration*
 - *MD5 and xxHash*
 - *DEBUG level log message with metrics on payload size, hash calculation time, checksum value, etc.*
 - *Note confidence w/r/t collision probability*

New Business

- None scheduled