As part of the observatory project to upgrade the correlator, an assessment has been requested to different groups in the observatory that are affected or that would need to provide improvements to their products in order to be compatible with the proposed upgrades. In particular ACS will analyze the existing infrastructure and software to allow bulk data transfer to satisfy a higher throughput than what is currently needed and what efforts would be required to go beyond the current capabilities.

The existing infrastructure consists of 10 Gbps equipment including, Ethernet cards, switches and cables. In specific cases, the connection between some servers is handled by Cisco's fabric interconnect which has a throughput capacity of 10+ Gbps.

An empirical analysis was performed over the network using three approaches focused on checking the current performance of our existing implementation of the bulk data transfer system.

- We analyzed the network throughput using IPerf
- We analyzed the throughput and latencies over the same network using bulk data transfer system's underlying technology, RTI DDS
- We analyzed the throughput over the same network using the bulk data transfer system

All the details of this analysis can be found at: **ICT-19921** - Getting issue details... `STATUS`

IPerf Network Analysis (10 Gbps link)

Protocol	Sender	Receiver 1	Receiver 2
TCP	9.35 Gb/s	9.35 Gb/s	-
UDP Unicast	9.00 Gb/s	8.87 Gb/s	-
UDP Multicast	9.00 Gb/s	8.50 Gb/s	8.52 Gb/s

- TCP and UDP unicast used only one receiver
- The UDP protocol used is unreliable and is prone to datagram losses

RTI DDS Network Analysis (10 Gbps link)

Protocol	Sender	Receiver 1	Receiver 2
TCP*	-	-	-
UDP Unicast	969 s	8.19 Gb/s	-
UDP Multicast	3579 s	7.78 Gb/s	7.78 Gb/s

- Problems with TCP implementation of RTI DDS demo
- UDP unicast used only one receiver
- Sender information is the maximum latency identified during the transfers
- UDP uses a reliable protocol

BulkDataNT Network Analysis (10 Gbps link)

Protocol	Sender	Receiver 1	Receiver 2
TCP	1.59 Gb/s	-	-
UDP Unicast	3.06 Gb/s	-	-
UDP Multicast	2.66 Gb/s	-	-

- TCP and UDP unicast used only one receiver
- Sender identifies the slowest speed among its receivers
- UDP uses a reliable protocol

The limitations imposed by the existing infrastructure and technologies are as follows:

- Network (**x13**): Allows about thirteen times the current required bandwidth
- RTI DDS (**x12**): Allows about twelve times the current required bandwidth
- BulkDataNT (**x4**): Allows about four times the current required bandwidth

The BulkDataNT implementation is not effectively taking advantage of the underlying technology that is using, achieving around a 35% of what the underlying technology offers.

There are different alternatives to tackle this:

- #1: **0.00 FTE**: Change the underlying infrastructure to a faster link (i.e. 100 Gbps)
- #2: **0.25 FTE**: Investigate and redesign BulkDataNT to make better use of RTI DDS
- #3: **1.50 FTE**: Change the implementation of BulkDataNT to a different technology
    - **0.50 FTE** for investigation + **1.00 FTE** for implementation if an appropriate technology is found during the investigation

The expected bandwidth increases with each of the previous alternatives is as follows:

- #1: **x40**: We still expect inefficiencies in the BulkDataNT system, but should still achieve ~35% of the network capabilities
- #2: **x12**: This is what the underlying technology offers, so it's the limit we can aim towards
- #3: **x13**: It depends on the chosen technology, but to choose a change of technology, we should aim towards a higher throughput than the one offered by using RTI DDS efficiently
- #1+#2: **x120**: Although there are no formal analysis of RTI DDS over a 100 Gbps link, we expect it to scale in a similar fashion than it did on 10 Gbps
- #1+#3: **x130**: There are a lot of unknowns in this scenario, but again, it should only be followed if the chosen technology behaves better than an efficient RTI DDS implementation