Caltech

CHEP2023 — 26th International Conference on Computing in High Energy & Nuclear Physics

Scientific Community Transfer Protocols, Tools, and Their Performance Based on Network Capabilities

Raimondas Sirvinskas, Preeti Bhat, Justas Balcas, Harvey Newman

The efficiency of high energy physics workflows relies on the ability to rapidly transfer data among the sites where the data is processed and analyzed. The best data transfer tools should provide a simple and reliable solution for local, regional, national and in some cases intercontinental data transfers. This work outlines the results of data transfer tool tests using internal and external (simulated latency and packet loss) in 100 Gbps testbeds and compares the results among the existing solutions, while also treating the issue of tuning parameters and methods to help optimize the rates of transfers. Many tools have been developed to facilitate data transfers over wide area networks. However, few studies have shown the tools' requirements, use cases, and reliability through comparative measurements. Here, we were evaluating a variety of high-performance data transfer tools used today in the LHC and other scientific communities, such as FDT, WDT, and NDN in different environments. Furthermore, this test was made to reproduce real-world data transfer examples to analyse each tool's strengths and weaknesses, including the fault tolerance of the tools when we have packet loss. By comparing the tools in a controlled environment, we can shed light on the tool's relative reliability and usability for academia and industry. Also, this work highlights the best tuning parameters for WAN and LAN transfers for maximum performance, in several cases.

Procedure and Goals

The aim of this study was to compare the performance of various data transfer tools in high-speed 100 Gbps testbeds, under both normal and challenging network conditions. To achieve this, a series of tests were conducted using a structured methodology.

- ❒ First, the data transfer tools were tested under normal network conditions (Fig 2.) to establish a performance baseline. This involved measuring the transfer rates and evaluating any potential bottlenecks or limitations.
- ❒ Next the tools were tested under simulated network conditions including varying levels of latency (Fig. 3) and packet loss, to evaluate their ability to handle network challenges. This step was designed to replicate real-world conditions and assess each tool's resilience.
- ❒ Finally, the tests focused on optimizing the performance each data transfer tool by tuning its parameters and methods. This involved experimenting with various settings to identify the most efficient configurations.

The results of these tests were analyzed to compare the performance of the data transfer tools under various network conditions, and to identify any areas for improvement. By optimizing the tools' parameters and methods, the study aimed to help users achieve optimal data transfer rates for their specific use cases.

About the Tools

WDT, FDT, and NDN are data transfer tools that differ in their architecture, approach, and features.

WDT: Provides fast and reliable data transfer over various network paths, including long distances, with built-in support for encryption and authentication; but it is designed primarily for one-to-one transfers and may not be suitable for multicast or one-to-many transfers.

FDT: Offers parallel multi-streaming, error recovery, rate control, and buffer management, for efficient data transfer with support for various protocols; but it may require a learning curve for new users.

NDN (Named Data Networking): Uses content-based addressing instead of host-based addressing. It provides a scalable and efficient approach to data transfer, with built-in security and privacy features, and it could support multicast and one-to-many transfers. The NDN architecture is new and its adoption is still limited. Wide use of NDN could require significant infrastructure changes.

Experimental Environment

The environment used for the experiments consisted of two Dell PowerEdge R730xd servers equipped with two Intel(R) Xeon(R) E5-2667v4 3.20GHz processors, 128GB of DDR4 2133MHz RAM, a 6x NVME software RAID0 disk array, and a ConnectX-6 network card. The servers were connected at 100Gbps with the MTU set to 9000 bytes, and the interface TX and RX set to maximum performance. Additionally, the CPU governor was set to performance, and the ***systemctl*** parameters were optimized for the best performance.

To ensure the environment was in working order, we did a dry run (Fig. 1) with an FDT memory to memory with 16 threads.

Fig 1. Multi Thread WAN Transfer (75ms latency, 16 threads memory to memory)

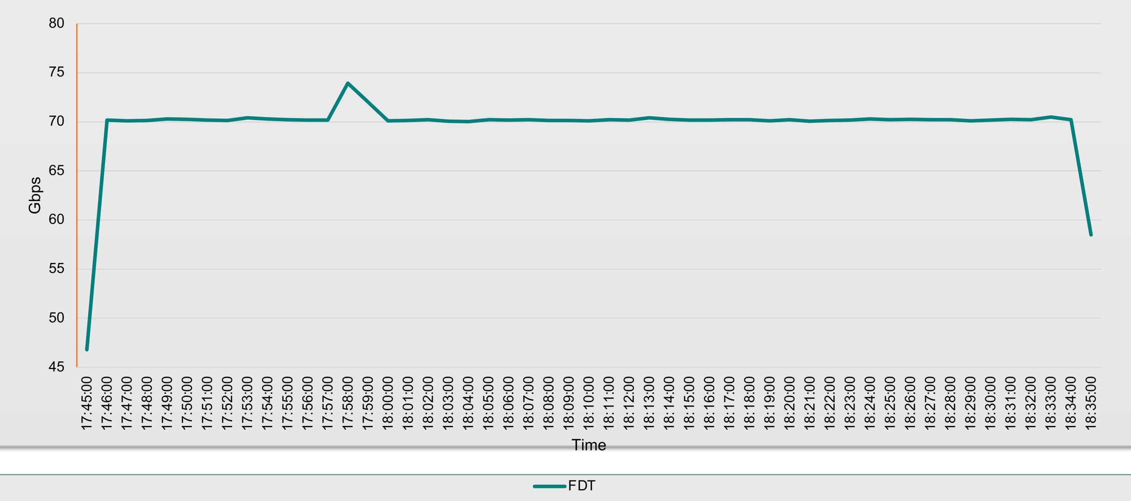

Fig 2. Single Threaded Local Transfer (<1ms latency)

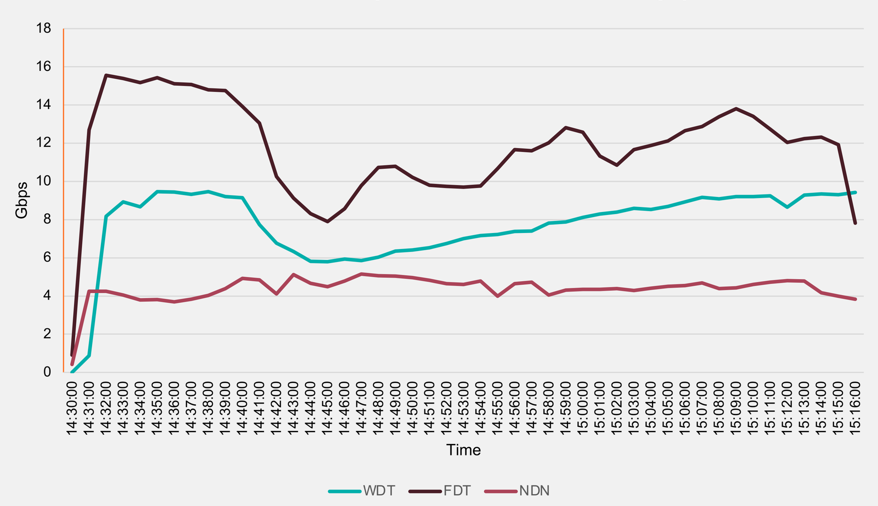

Fig 3. Single Threaded WAN Transfer (75ms latency)

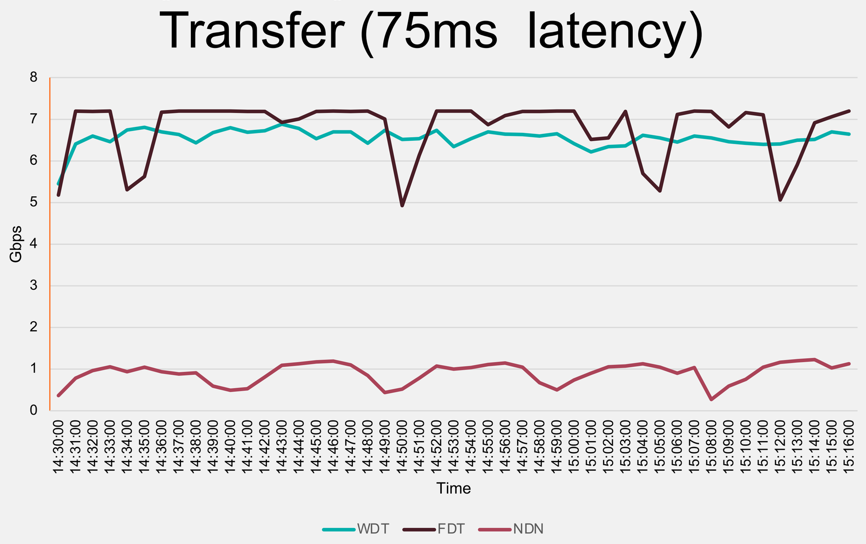

Fig 4. Multi Threaded WAN Transfer (75ms latency, 16 threads)

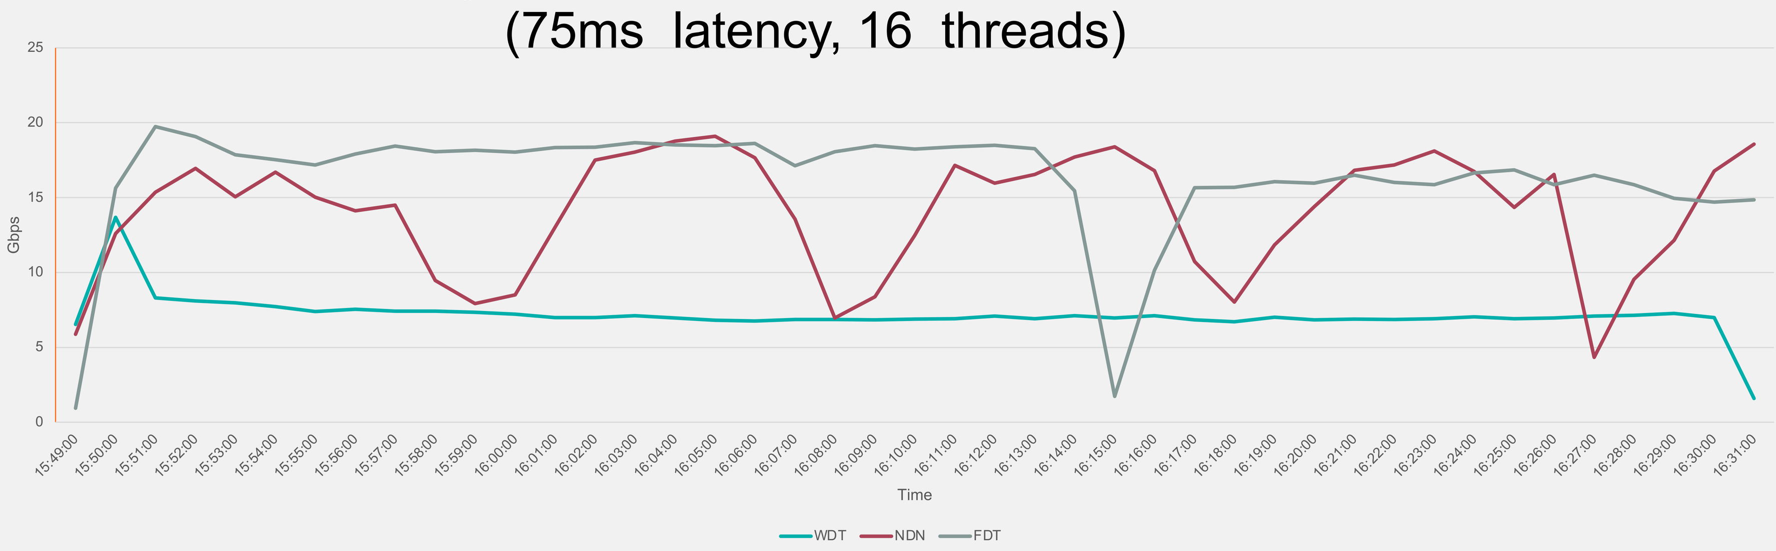

Conclusions

Based on our results, we can conclude that NDN and FDT are more suitable for high-speed data transfers as they achieved higher throughput compared to WDT. NDN was found to have the lowest latency, which is desirable in time-sensitive applications such as video streaming. However, it is important to note that NDN is more complex to set up and requires specialized software and hardware, which may not be feasible for all users. FDT offers a good balance between throughput and latency, making it a suitable option for general-purpose data transfers.

Summary

The results of our study showed (Fig 4.) that NDN achieved the highest throughput with an average of 15.8 Gbps on 75ms latency, followed by FDT with an average of 15.2 Gbps, while WDT achieved the lowest throughput with an average of 7.2 Gbps. In terms of latency, NDN had the lowest average latency with 11.3 ms, followed by WDT with an average of 19.2 ms, while FDT had the highest average latency with 58.8 ms. It is worth noting that during our tests with WDT, we observed an average of 15 packet drops every minute, which could potentially affect its performance.

Future Work

Our study provides a baseline comparison of the performance of these data transfer tools under specific conditions. Further, research can be conducted to evaluate their performance under different network conditions, such as varying the latency, bandwidth, and packet loss. Additionally, it would be interesting to compare the performance of these tools with other data transfer protocols, such as Aspera, UDT, and XRootD. Finally, it would be beneficial to investigate the impact of different file sizes and transfer distances on these tools' performance. Specifically:

1. Evaluate the impact is using the latest NVMe disks with Gen4 and Gen5 PCI express connections.
2. Add more tools for comparison.
3. Perform additional tests on the NDN and FDT protocols to assess their performance under different network conditions.
4. Explore the impact of varying network latency and bandwidth on the performance of the protocols.
5. Investigate the scalability of these protocols, especially with regards to the number of concurrent connections and data transfer rates.