


WHO SAYS TCP/IP IS SLOW?



For mainframe shops that need to move data on or off the mainframe, whether to the cloud or to an alternative on-premises destination, FICON® (Fibre Connection), the IBM mainstay for decades, is generally seen as the standard, and with good reason. When it was first introduced in 1998 it was a big step up from its predecessor, ESCON (enterprise system connector), that had been around since the early 1990s. Comparing the two was like comparing a firehose to a kitchen faucet.

FICON® is fast, in part, because it runs over Fibre Channel in an IBM proprietary form defined by ANSI FC-SB-3 Single-Byte Command Code Sets-3 Mapping Protocol for Fibre Channel (FC). In that schema it is an FC layer 4 protocol. As a mainframe protocol it is used on IBM zSystems™ to handle both direct-access storage device (DASD) and tape I/O. It is also supported by other vendors of disk and tape storage and switches designed for the IBM environment.

Over time, IBM has increased speeds and added features such as High Performance FICON® without significantly enhancing the disk and tape protocols that traverse over it, meaning these limitations on data movement remain. For this reason, the popularity and long history of FICON® do not make it the answer for every data movement challenge.

Stuck in the past

One challenge of particular concern today is that mainframe secondary storage is still being written

to tape via tape protocols, whether it is real physical tape or virtual tape emulating actual tape. With tape as a central technology, it implies dealing with tape mount protocols and tape management software to maintain where datasets reside on those miles of Mylar. The serial nature of tape and limitations of the original hardware required large datasets to often span multiple tape images.

Though virtual tapes written to DASD improved the speed of writes and recalls, the underlying protocol is still constrained by tape's serialized protocols. This implies waiting for tape mounts and waiting for I/O cycles to complete before the next data can be written. When reading back, the system must traverse through the tape image to find the specific dataset requested. In short, while traditional tape may have its virtues, speed—the 21st century speed of modern storage—is not among them. Even though tape and virtual tape are attached via FICON[®], the process of writing and recalling data relies on the underlying tape protocol for moving data, thus making FICON-attached less than ideal for many modern use cases.

Faster and better

There is an alternative that doesn't rely on tape or emulate tape because it does not have to. Instead, software generates multiple streams of data from a source and pushes data over IBM Open Systems Adapter (OSA) cards, using TCP/IP in an efficient and secure manner, to an object storage device, either on premises or in the cloud. The Open Systems Adapter functions as a network controller that supports many networking transport protocols, making it a powerful helper for this efficient approach to data movement. Importantly, as an open standard, OSA is developing faster than FICON[®]. While there is a belief common among many mainframe professionals that OSA cards are "not as good as FICON[®]," that is simply not true when the necessary steps are taken to optimize OSA throughput.

To achieve better overall performance, the data is captured well before tape handling, thus avoiding the overhead of tape management, tape mounts, etc. Rather than relying on serialized data movement, this approach breaks apart large datasets and sends them across the wire in simultaneous chunks, while also pushing multiple datasets at a time. Data can be compressed prior to leaving the mainframe and beginning its journey, reducing the amount of data that would otherwise be written. Dataset recalls and restores are also compressed and use multiple streams to ensure quick recovery of data from the cloud.

Having the ability to write multiple streams further increases throughput and reduces latency issues. In addition, compression on the mainframe side dramatically reduces the amount of data sent over the wire. If software is also designed to run on zIIP engines within the mainframe, data discovery and movement, as well as backup and recovery workloads, will consume less billable MIPS and TCP/IP cycles also benefit.

This approach delivers mainframe data to cloud storage, including all dataset types and historical data, in a quick and efficient manner. And this approach can also transform mainframe data into standard open formats that can be ingested by business intelligence (BI) and analytics off of the mainframe itself, with a key difference. When data transformation occurs on the cloud side, no mainframe MIPS are used to transform the data. This allows for the quick and easy movement of complete datasets, tables, image copies, etc. to the cloud, then makes all data available to open applications by transforming the data on the object store.

A modern, software-based approach to data movement means there is no longer a need to go to your mainframe team to update the complex extract, transform, and load (ETL) process on the mainframe side.

To address the problem of hard-to-move mainframe data, this software-based approach provides the ability to readily move mainframe data and, if desired, readily transform it to common open formats. This data transformation is accomplished on the cloud side, after data movement is complete, which means no mainframe resources are required to transform the data.

Other advantages include:

- Dedicated software quickly discovers (or rediscovers) all data on the mainframe. For example, even with no prior documentation or insights, BMC AMI Cloud solutions can rapidly assemble and map the data to be moved, expediting both modernization planning and data movement.
- Policies are defined to move either selected data sets or all data sets automatically, reducing oversight and management requirements dramatically as compared to other data movement methods.
- For the sake of simplicity, a software approach can be designed to invoke actions via a RESTful API, or a management UI, as well as from the Mainframe side via a traditional batch or command line.
- A software approach can also work with targets both on premises and in the cloud.

In summary, a wide range of useful features can make data movement with a software-based approach intuitive and easy. By avoiding older FICON® and tape protocols, a software-based approach can push mainframe data over TCP/IP to object storage in a secure and efficient manner, making it the answer to modern mainframe data movement challenges!