Data centre fabrics explained

A complex system of interconnections between switches and servers ultimately provides an efficient communications path.

By Paul Kish

There is a new standard under development in TIA TR 42.1 Subcommittee on "Cabling Guidelines for Data Centre Fabrics." It is currently going through a default balloting stage and should be published early next year as Addendum 1 to the TIA-942-A Telecommunications Infrastructure Standard for Data Centres.

A data centre fabric is a system of interconnections between switches and servers that are weaved together like a fabric. The purpose of a fabric to provide low-latency and high-bandwidth communication between any two points in the switch fabric.

In a recent article by Dhritiman Dasgupta of Juniper Networks Inc. on *Why Data Centres Need a Performance Review*, he explained that with the adoption of technologies such as server and storage virtualization, data centre traffic has changed direction over the past few years from predominantly north-south (into and out of the data centre) to predominantly east-west (server-to-server within the data centre).

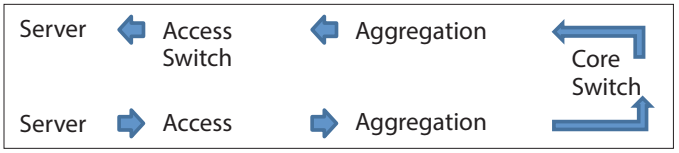

Figure 1 – Traditional Switching Architecture

A traditional three-tier switching architecture, as shown in Figure 1, is well-suited for traffic from servers to external destinations.

However, it is not suitable for large virtualized data centres where compute and storage servers may be located anywhere within the facility. For one server to communicate with another server it may need to traverse a hierarchical path through two aggregation switches and one core switch as shown in Figure 1, which adds to latency and can create traffic bottlenecks.

The new standard describes different implementations of data centre fabrics as follows:

- Data centre fabric fat-tree
- Data centre fabric full-mesh
- Data centre fabric interconnected meshes
- Data centre fabric centralized switch
- Data center fabric virtual switch

They all serve a similar function but are implemented differently. Due to the limited space available here, I will focus on describing the fat-tree architecture.

Data centre fabric fat-tree:

The connection path for the fat-tree architecture is illustrated in Figure 2a and 2b. In the fat-tree architecture, all access switches are connected to every interconnection switch.

An access switch may use a port extender to connect to a server as illustrated in Figure 2b. Port extenders are physical extensions of access switches to which they are attached.

Typically, they map many lower-speed ports for server communication to fewer higher-speed ports on the Access Switch.

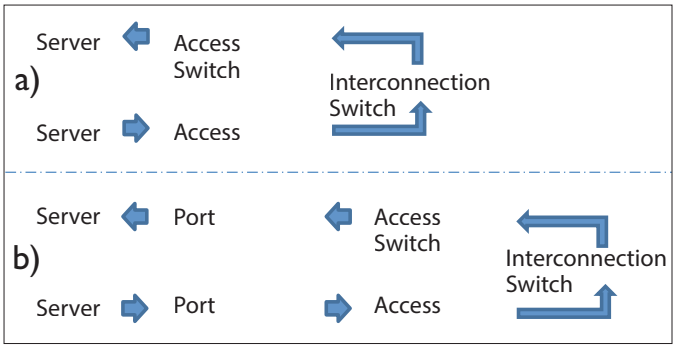

Figure 2 – Data Center Fabric Fat-Tree Architecture

Any server can communicate with any other server with no more than one interconnection switch path between any two access switches. This architecture can be non-blocking by providing sufficient bandwidth from each access switch to interconnection switches.

Location of Switches and Port Extenders:

- Interconnection Switches are typically located in the Main Distribution Area (MDA)
- Access Switches are typically located in the Horizontal Distribution Area (HDA) for "End-of-Row" or "Middle-of- Row" topologies or in the Equipment Distribution Area (EDA) for "Top-of-Rack" topologies.
- Port Extenders are typically located in EDAs at the top-of-rack.

Larger switch fabrics may be created by interconnecting pods where each pod is a fat-tree with a First Level of interconnection switches. A Second Level of interconnection switches is connected to each First Level interconnection switch in all Pods.

On the surface, data centre fabrics may appear to be a complex system of interconnections between switches and servers. When you look under the surface at the composition of data centre fabrics, they provide an efficient communications path between all the servers that are connected to the fabric. From a cabling perspective, it requires a lot more switch-to-switch connections than a traditional switching architecture in a hierarchical star topology.

In closing, the Data Centre Fabrics Addendum is a valuable addition to the Data Centre Standard to support current data centre applications. **CNS**

Paul Kish is Director, Systems and Standards at Belden. The information presented is the author's view and is not official TIA correspondence.