White paper

10/40/100 migration in data centers

The “three P’s”: Path, polarity and pins

Forward-looking data center designers and managers realize that, now that the standards are largely defined, migration of Ethernet links from 10Gb/s to 40 and 100Gb/s as well as Fibre Channel links from 4 to 8 to 16Gb/s is rapidly accelerating.

With further interesting wrinkles such as high density 10 Gigabit links (four 10Gb/s channels combined in a single 40Gb/s QSFP port and ten 10Gb/s channels combined in a single 100G port), physical aggregation layers, and direct-attach circuits, having a clear and defined plan is not just prudent, it is imperative. Fortunately a few simple rules can reduce much of the uncertainty in this process.

Background

Today, most fiber optic Ethernet channels operate at 10GB/s or less, and operate over a pair of multimode optical fibers, one for transmit and one for receive. The connector of choice for these channels is a duplex LC-type connector on active electronic ports as well as those in the physical cable plant.

In the LC connector, one fiber acts as a transmit fiber and one as a receive, schematically:

These pairs are often combined over trunk cables of 12, 24 or even much higher fiber counts using cassettes that combine and split out the pairs or other passive branching devices. (Some power users prefer 192 or even 432 fiber trunks to save space in their path.)

The next step in data rate, 40G, requires a different physical architecture. Because 40G channels are composed of four parallel duplex channels (a total of eight fibers) the preferred connector interface is now the MPO-type multiple fiber connector.

This connector contains one or multiple rows of parallel fibers, usually in multiples of twelve fibers. A single-row MPO connector would transmit over eight of the twelve fibers at 40G while leaving the remaining four dark.

Multiple channels can be combined on higher fiber count MPOs. For instance, three 8-fiber 40Gbs channels can be combined over a single two-row, 24 fiber MPO connector without abandoning any of the fibers in the connector or associated trunk cable.

The next step in Ethernet communication, 100G takes the parallel 10G methodology a step further, combining 10 parallel duplex channels over 20 fibers to form 100G connections. The connector of choice in this case is a two-row 24 fiber MPO connector, in which 20 of the 24 fibers are utilized and the four outermost fibers remain dark. (There is activity in some standards organization to implement 100G Ethernet over four 25Gbs parallel duplex links, or a total of eight fibers – in which case the 100G implementation looks very much like the 40G system noted above. At this time, however, 25G transmission has yet to be standardized. To minimize long term risk a prudent migration plan must take into account that one or even both solutions will be implemented somewhere in the future.)

Another leap, to 400G, is in the initial planning stages. Possible implementations include 32 multimode fibers in parallel or one or more single mode pairs to support transmission at this data rate. There are many technical and economic issues to be resolved, and these data rates will probably not be implemented in any significant volume until later in this decade.

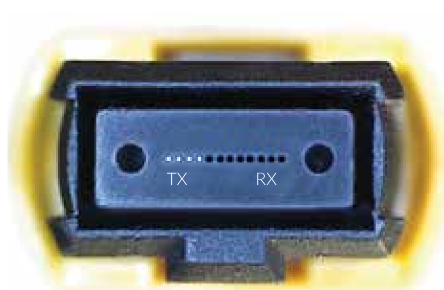

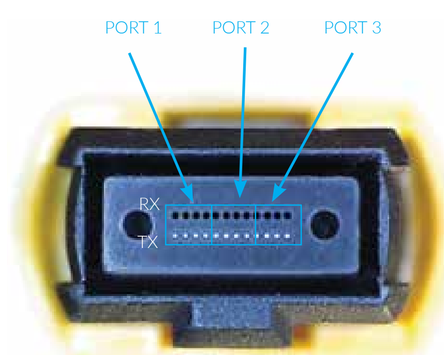

Physical layer planning

Often neglected, physical layer planning can mean the difference between a smooth, logical migration and disorder just this side of chaos. With multiple options for data rate, connectivity and trunking options, trying to configure a data center's physical layer mapping "on the fly" can rapidly turn into a self-defeating exercise.

Fortunately, by following a few simple rules, much of the guesswork can be eliminated, while maintaining flexibility and the ability to migrate a small or large portions of your data center as your demand and your budget dictate. No matter what you have to have a plan.

Rule 1: Have a plan, develop it carefully, and stick to it!

One of the simplest rubrics available is one called "The Three P's" – path, polarity and pins. Having a proper plan for the Three P's will get you a long way toward a smooth migration plan.

Path

Put simply, "Path" defines which groups of fibers at a given end of a link connect to which fiber groups at the other end. This includes all components in between, including patch cords, branching assemblies or devices, trunks and hydra assemblies, and all of the associated "hops" or transpositions.

An example path, in this case corresponding to that outlined in the TIA-568.3-D Draft 1.0 is shown below.

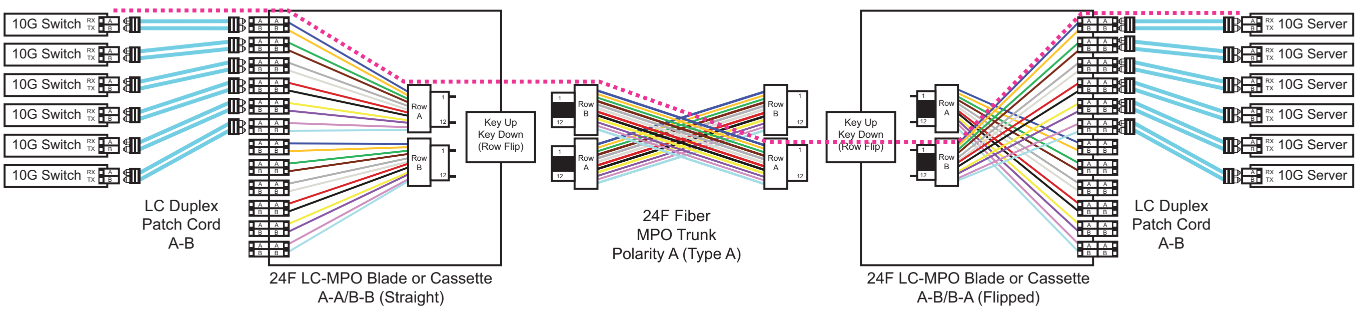

This path delineates an interconnection between transceivers on a piece of active electronics (for example a core switch) and a device at the other end (e.g. a server.) Components include patch cords, blades or cassettes (which could be replaced by hydra-type assemblies) and trunk cables.

In a similar manner, a cross-connect can be represented the same way:

In this case there are two trunk cables with a patching array between them (represented at the left). Draft and released standards exist for these implementations, from the TIA and other organizations.

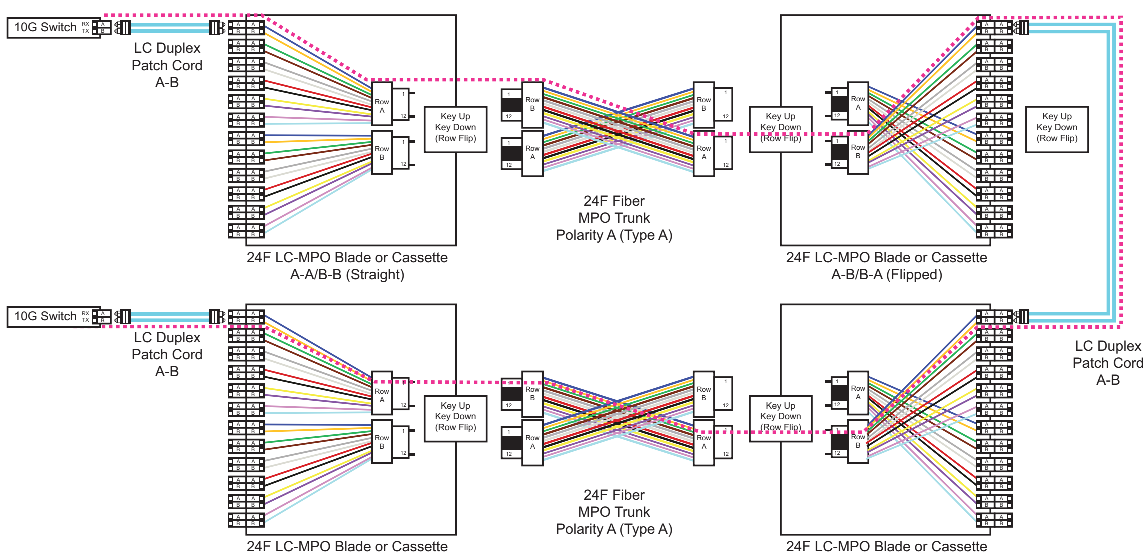

Rule 2: Develop your plan using industry standards.

Don't try to develop practices and plans on the fly. This is rarely an effective approach. Of course, standards often have different options, usually developed to support different design approaches or networking philosophies.

Rule 3: If the standard has options, pick the option that works best for you and stay with it.

Certainly different sections of a data center may require different practices. Interconnection at an aggregation point may be different than SAN connections, for instance. But for any given function within a data center, the designs and practices should remain constant. ANSI/TIA-942-A, as standard for data center design, provides excellent guidance in this area.

Consistency in choice of path removes a number of practical issues in data center operations. For example, if some cable plant uses single row 12 fiber MPO connectivity, and some use dual row 24 fiber connectivity, a problem can arise in that single and dual row MPO's share a standard outer mating mechanism, but the rows of fibers within the connector to not align between the two.

So a 12 fiber MPO can plug into a 24 fiber MPO, but no light will pass through as the fibers are not aligned. In some cases damage can result as well. While color coding of strain reliefs or other visual indicators may be used, they are not always effective, especially in heavily packed connection fields or poorly lit cabinets. Mixing of pathway types in the same field invites operational difficulties.

Polarity

Polarity differs from path in the sense that polarity defines the relationships of the fibers within the path. A specific transmit fiber at one end must connect to a specific receive fiber at the other end. In single pair applications, the transmit fiber of the pair on the active electronics is, by standard, always on the left looking into the transceiver (TX/RX).

So, the net effect of the polarity of the cabling infrastructure is that the pair must be "flipped"– in that the transmit at one end is connected to the receive at the other. A 10G channel, implemented with duplex LC connectors would be represented schematically this way, looking into the end of the connectors that plug into the active electronics.

There are a number of options to achieve this: One could make certain that the infrastructure (trunks plus cassettes, blades or other branching devices) always flips the pairs. A second would be to use a straight infrastructure with a patchcord that is straight through at one end and a second patchcord that is flipped at the other end. In a cross connect scenario, this can get even more complicated as there are more "flips" enforced by the patch fields. If the infrastructure is designed to migrate to higher data rates, where the fiber groupings are now multiples of eight or even twenty, yet more complications ensue.

Clearly, minimizing the variations in the infrastructure are important to design and installation, and especially during upgrades, migration and troubleshooting. Reducing variations is critical.

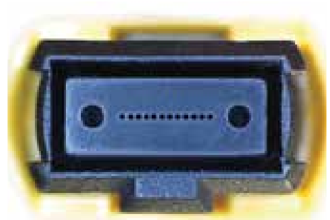

12 FIBER MPO

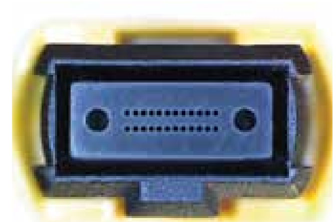

24 FIBER MPO

END "A"

END "B"

Rule 4: Minimize the cabling polarity variants within your infrastructure, and make certain polarity choice does not limit options of the devices you can employ within the infrastructure.

Which leads immediately to:

Rule 5: All trunks or components that act as trunks should have uniform polarity.

There are several reasons for this. If trunks are uniform, changes to pathways due to changes in topology can be implemented through the devices that plug into the trunks such as cassettes, blades or hydra assemblies. It is much easier to change out the device at the end of the trunk than to have to re-enter the pathways to replace a trunk. Upgrades to higher data rates can be accomplished through a reduced set of components at the end of the trunks as well. Finally, in cascaded topologies such as cross connects, where more than one trunk (or component that acts as a trunk) may be present, choices of end devices can still be standardized.

In 40 and 100Gbs applications any devices "self-discover" meaning that the electronics at the ends of the paths recognize these fiber interconnections. So why would a designer worry about polarity? There are a number of reasons:

- If the infrastructure must accommodate mixed data rates with different connectivity (LC vs. MPO) then duplex polarity is still important.
- Emerging applications such as high density 10G channels, where four 10G circuits are shared over a single, eight- ber 40G MPO port at the switch, require consistent polarity at the end devices where the 10G channels are broken out.
- In applications where multiple channels are shared over a higher fiber-count backbone, such as three 40G channels over a 24 fiber trunk, six over a 48 fiber etc. there must be consistency when these channels are assembled at one end and broken out at the other.
- In some cases, data rates may be lowered for a given backbone due to changes in topology or data center layout. In these cases there may be a transition from a higher speed and fiber count channels to multiple lower speed fiber-count channels.

Flexibility in choice of end devices, paradoxically, is enhanced by having greater uniformity in relatively fixed infrastructure elements such as trunks.

With this in mind:

Rule 6: Once you have chosen your trunk polarity, choose uniform designs and components for 10G, high-density 10G, 40G and 100G channels. Be certain that migration to higher and lower data rates is possible.

Which implies

Rule 7: Channel designs should account for future uncertainties (e.g. eight or twenty fiber solutions for 100G channels).

Polarity mapping can become quite complicated without uniform design practices. Trying to figure out polarities each time a network change occurs is time-consuming, frustrating and prone to error. Solid design practices eliminate these unnecessary distractions.

Pins

There is an additional wrinkle to the discussion. Because of the large number of fibers in MPO connectors, arrayed over a very small space, these connectors required enhanced alignment mechanisms. The MPO interface has a pair of high precision alignment pins on one half of the connection and mating high precision holes on the other half.

The connectors are the dimensionally identical otherwise, and a pinned connector can reside on either side of the MPO interface.

While the MPO is a time-proven and highly effective interface, this leads to a practical issues:

- Pinned and non-pinned connectors are not differentiated other than by the presence of pins.
- Transceiver ports on active equipment such as switches and servers virtually always have pinsinside the active port.
- Plugging pinned patchcord plugs into these pinned transceiver ports can, in some cases, cause pin damage.
- Plugging pinned plugs into pinned plugs can, in some cases, cause pin damage.
- There is often no way to know if the connector on the back of a passive panel interface has pins or not.
- Non-pinned connectors can plug into non-pinned connectors. While damage is unlikely, without the alignment pins, optical performance will be sharply degraded, usually to the point of loss of optical connection.

Shows schematically:

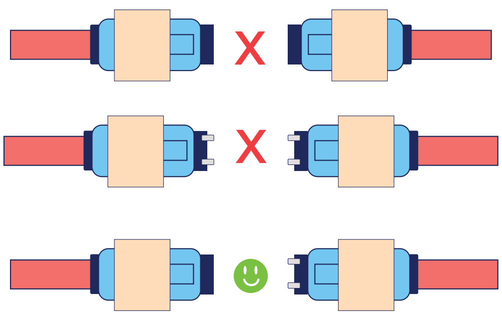

- Will plug
- High loss

- Won't plug
- Damage?

- Will plug
- Good loss

So keeping track of pins in one's infrastructure is important. The infrastructure needs to be uniform and predictable.

Rule 8: Have a uniform practice for MPO pins by data rate for both trunks and patchcords.

And installation should be simplified.

Rule 9: Have a uniform practice for MPO pins on trunks. Trunks should have the same connectors on both ends.

Suggestion: Non-pinned trunks reduce the risk of pin damage during installation. Having the pins on the mating devices, especially blades or cassettes, allow for greater pin protection.

Having the same connector on both ends of the trunk simplifies design and installation. It is not necessary to keep track of which end is which.

These uniform practices also reduce the number of patchcord variants required, which simplifies MACs and reduced inventory.

Keeping track of Path, Polarity and Pins greatly simplifies data center design and provides a framework for simplified, logical migration of data rates and changes in topologies. There is one additional rule that experience has taught is extremely helpful in structured design.

Rule 10: Once you have defined your practices and migration plan, diagram each variation and step in your migration and assemble and check a physical model of each to confirm your designs.

This is not complicated. While active switch ports would be the ideal, this could be costly and very time-consuming. Physical circuits can be verified with a simple light source, and proper plugging can be inspected visually. This one-time verification can save countless hours of design, MAC and troubleshooting time in the long run.

Summary

Agile design for a flexible data center takes planning, especially to support future migrations in topology and data rate. A few simple rules can help.

Rule 1: Have a plan, develop it carefully, and stick to it!

Rule 2: Develop your plan using industry standards.

Rule 3: If the standard has options, pick the option that works best for you and stay with it.

Rule 4: Minimize the cabling polarity variants within your infrastructure, and make certain polarity choice does not limit options of the devices you can employ within the infrastructure.

Rule 5: All trunks or components that act as trunks should have uniform polarity.

Rule 6: Once you have chosen your trunk polarity, choose uniform designs and components for 10G, high-density 10G, 40G and 100G channels. Be certain that migration to higher and lower data rates is possible.

Rule 7: Channel designs should account for future uncertainties (e.g. eight or twenty fiber solutions for 100G channels).

Rule 8: Have a uniform practice for MPO pins by data rate for patchcords.

Rule 9: Have a uniform practice for MPO pins on trunks. Trunks should have the same connectors on both ends.

Rule 10: Once you have de ned your practices and migration plan, diagram each variation and step in your migration and assemble and check a physical model of each to confirm your designs.

Disciplined planning is critical for robust designs. Simple rules for the "Three P's" help support this discipline both for initial designs as well as the long-term evolution and enhanced utility of the data center infrastructure.

Everyone communicates. It's the essence of the human experience. *How* we communicate is evolving. Technology is reshaping the way we live, learn and thrive. The epicenter of this transformation is the network—our passion. Our experts are rethinking the purpose, role and usage of networks to help our customers increase bandwidth, expand capacity, enhance efficiency, speed deployment and simplify migration. From remote cell sites to massive sports arenas, from busy airports to state-of-the-art data centers—we provide the essential expertise and vital infrastructure your business needs to succeed. The world's most advanced networks rely on CommScope connectivity.

COMMSCOPE®

commscope.com

Visit our website or contact your local CommScope representative for more information.

WP-111197-EN (02/17)