Power over Ethernet

The Story So Far

by Jason James, Technical Director, HellermannTyton

HellermanTyton chronicles the development of Power over Ethernet and makes best practice recommendations.

Introduction

Power over Ethernet (PoE) does exactly what its name suggests; it delivers usable power over Ethernet (twisted-pair) cabling. Providing both the power and data required by an attached device provides significant benefits such as ease of installation and relocation of the powered device without the need to consider the location of the power source. Remote control and monitoring are also easier to achieve with a remote power solution. Even though there are some obvious benefits to PoE, there are some things to know and others to be considered.

The standards

PoE has been in use in various forms for many years and the IEEE has standardised four versions. The first PoE standard IEEE 802.3af-2003 allowed for the powering of devices up to 15.4W. The second standard IEEE 802.3at- 2009, so called PoE plus or PoE+, was specified a maximum powering provision of 30W. The new standard IEEE 802.3bt-2018 specifies PoE type 3 and type 4 (sometimes referred to as PoE++) to deliver a maximum 60W and 100W respectively.

While these values are the reported maximum powers that can be provided by the power source equipment (PSE), it is not the actual power that is ultimately delivered to the powered device (PD). The resistance of the copper cabling converts some of the sourced power into heat and, as such, only 12.95W to 15.4W of power is actually delivered to the PD for type 1, 25.5W to 30W for type 2, 51W to 60W for type 3 and 71W to 100W for type 4.

Power delivery method

In addition to the four present power levels that are deliverable using PoE, there are numerous power delivery methods. However, only two power delivery methods have been standardised by the IEEE; they are referred to as Alternative A and Alternative B and describe the methods employed to deliver the power from the PSE to the PD.

The references to Alternative A and B only apply to the PSE since the PDs are normally designed to support both power delivery techniques, as shown in Figures 1-6 opposite.

It is generally understood that the two alternatives simply define whether the power is transmitted via used or unused pairs. This is the case for 10BASE-T and 100BASE-T Ethernet where only the orange and green pairs are used. However, this is not so for 1000BASE-T and 10GBASE-T, where all four pairs are used for data transmission. Figure 1 shows the endpoint PSE location of Alternative A for 10BASE-T/100BASE-TX using the

PoE has been in use in various forms for many years and the IEEE has standardised four versions.

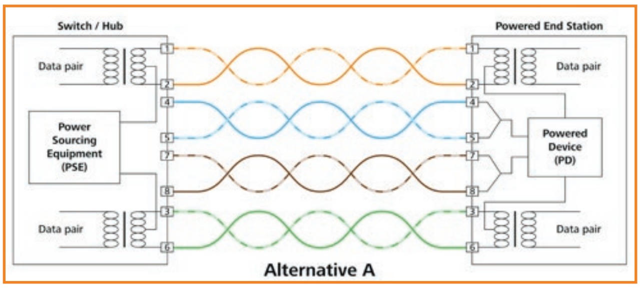

Figure 1

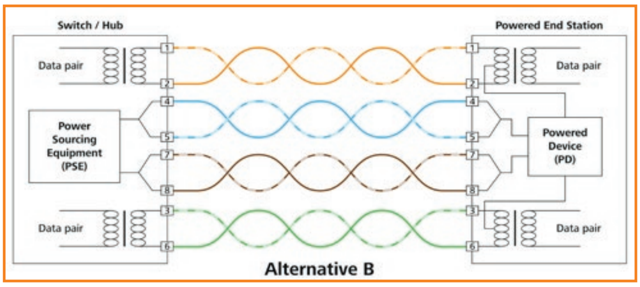

Figure 2

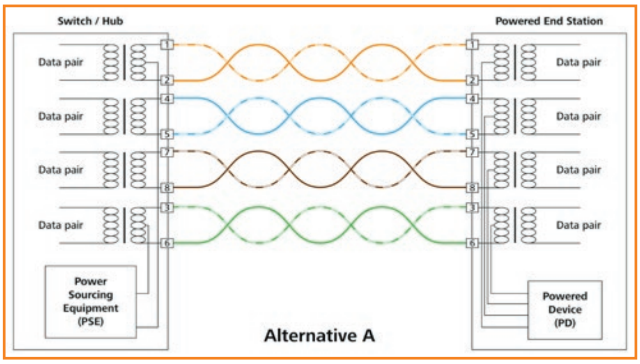

Figure 3

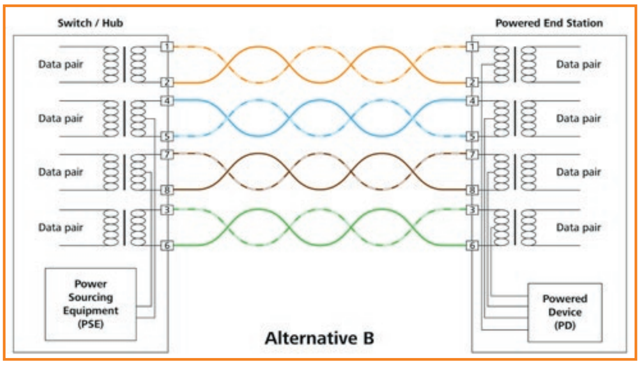

Figure 4

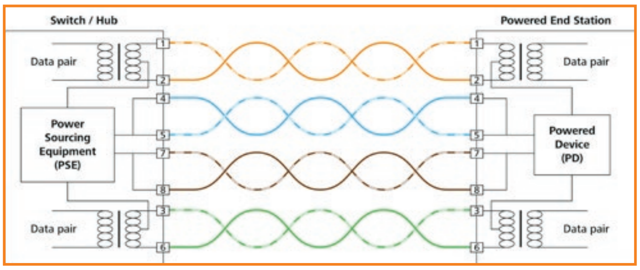

Figure 5

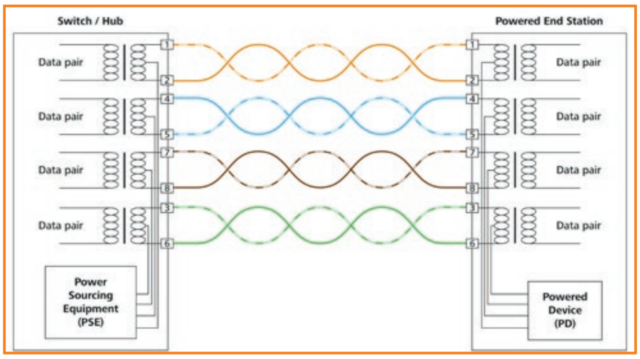

Figure 6

data pairs and Figure 2 shows the endpoint PSE location of Alternative B for 10BASE-T/100BASE-TX, using the unused (blue and brown) pairs.

For the four-pair signaling of 1000/2.5G/5G/10GBASE-T, the two alternatives simply determine which two pairs deliver the power. Figures 3 and 4 (on this page) show the endpoint PSE location of Alternative A and Alternative B respectively, with Figure 3 using the orange and green pairs, and Figure 4 using blue and brown.

PoE type 3 and type 4 allow all four pairs to carry a certain amount of current. Figure 5 and Figure 6 show the endpoint PSE locations for 4-pair PoE for 10BASE-T/ 100BASE-TX and 1G/2.5G/5G/10GBASE-T respectively.

Figure 1 through to Figure 6 depict the use of power source equipment within the endpoint (switch or hub). The same configurations are also available in a mid-span solution.

In addition to the four present power levels that are deliverable using PoE, there are numerous power delivery methods.

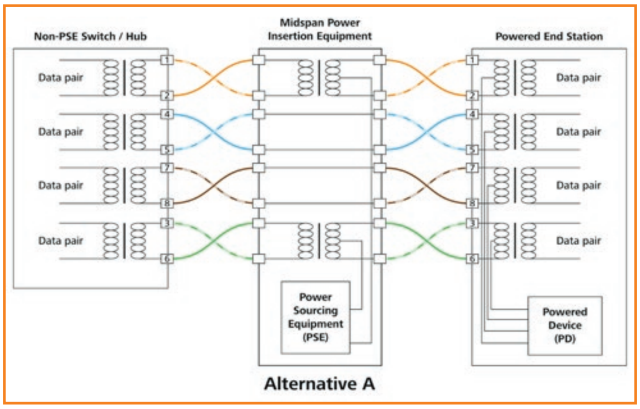

Figure 7

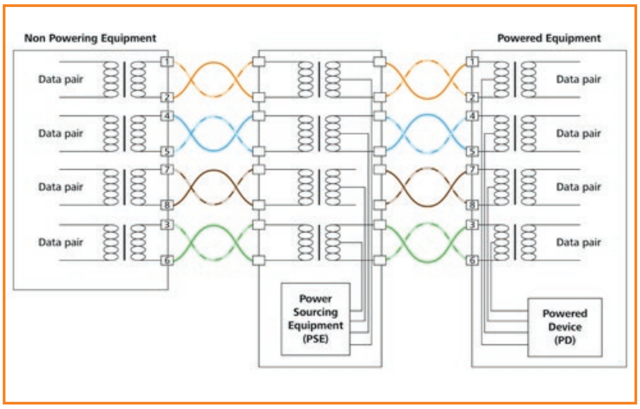

Figure 8

Endpoint PSE or mid-span PSE

Figure 7 above shows the 2-pair PoE mid-span configuration for Alternative A for a four-pair Ethernet system. Figure 8 above shows the 4-pair PoE mid-span configuration for a four-pair Ethernet system.

In this type of powering solution, the switch or hub can be of a standard design without integral power source equipment. While the switch cost will of course be lower, the combined cost of a separate switch and mid-span PSE will probably be greater and occupy more rack space. The benefits of keeping them separate enable them to be upgraded independently and to only provide power on those ports where it is needed.

Figure 9 shows a typical combined switch and PSE with 24 powered ports. Figure 10 shows a mid-span PSE device that supports 16 incoming Ethernet ports and 16 outgoing Ethernet + PoE ports.

Although referred to as mid-span, these devices are usually situated adjacent to the Ethernet switches.

Types of powered devices

Thousands of devices now support remote powering from the PSE. These PDs have the benefit of not requiring the installation of both dedicated power and network cables to the same location. Many domestic and commercial IT devices operate on low voltages and include power transformers to reduce mains voltages down to something more usable. These generally take the form of small power supply units.

Delivering power over the cabling reduces the need for this type of power provision and the significant inefficiencies they have in converting from mains voltages to much lower DC voltages. There are numerous applications for PoE, the most

Figure 9

Figure 10

Figure 11

Figure 12

Figure 13

Conductor	Alternative A (MDI-X)	Alternative A (MDI)	Alternative B (All)
1	Negative V_{PSE}	Positive V_{PSE}	
2	Negative V_{PSE}	Positive V_{PSE}	
3	Positive V_{PSE}	Negative V_{PSE}	
4			Positive V_{PSE}
5			Positive V_{PSE}
6	Positive V_{PSE}	Negative V_{PSE}	
7			Negative V_{PSE}
8			Negative V_{PSE}

Table 1A: PSE pinout alternatives

popular of which are wireless access points, IP cameras and IP telephones. Typical examples of these are shown in Figures 11-13 (on previous page)..

Devices such as IP telephones used in voice-over-IP (VoIP) networks often incorporate two RJ45 sockets, one used for the incoming powered Ethernet connection, back to the switch, and the other as a regular unpowered socket to allow a desktop or laptop computer to be connected, eliminating the need for a second RJ45 connection in the wall.

Connector impact

So, how does PoE work? Is the power always on? How does the PSE know how much power to supply? How do you turn the PSE off? This section will answer all these questions and more.

The low voltage PoE is not always on. If no device is detected, then the power is off. Only when a PD is attached to a PoE-enabled source is the power made available to the PD. This addresses any concerns over people receiving electric shocks from Ethernet sockets or patch cords if they are left accessible. There is a handshake that occurs between the PSE and the PD prior to the power being enabled. This only takes a fraction of a second and, during the handshake, the PD communicates its powering requirements back to the PSE so that the correct mode can be set in the PSE.

A PSE device can provide power using only one (at a time) four-wire connection, regardless of whether the Ethernet protocol is two- or four-pair. In either type, the two conductors of each utilised pair carry the same current. Table 33-2 is extracted from IEEE Std 802.3-2015.

Although the power is off when there is no PD connected to the link and the handshake between the two activates the power, the removal of the PD doesn't benefit from the same degree of control. In other words, the PSE cannot anticipate the removal of the PD through the unplugging of a patch cord and, as such, the voltage and current are still present when this action occurs.

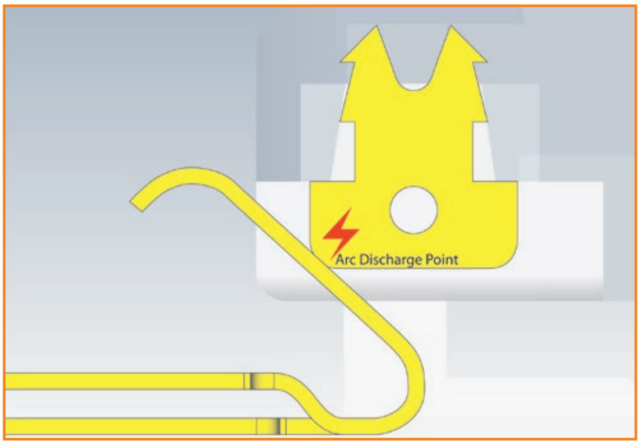

Figure 14

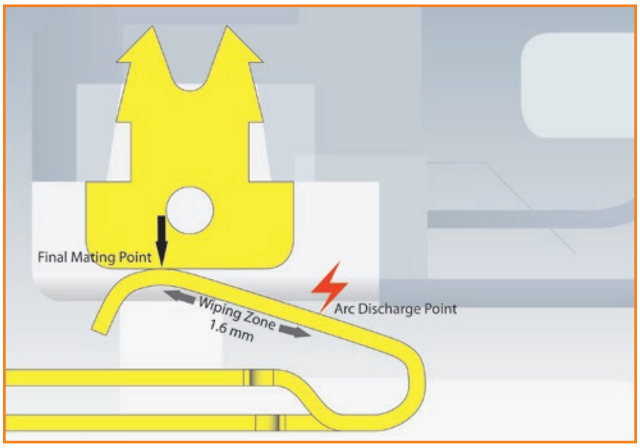

Figure 15

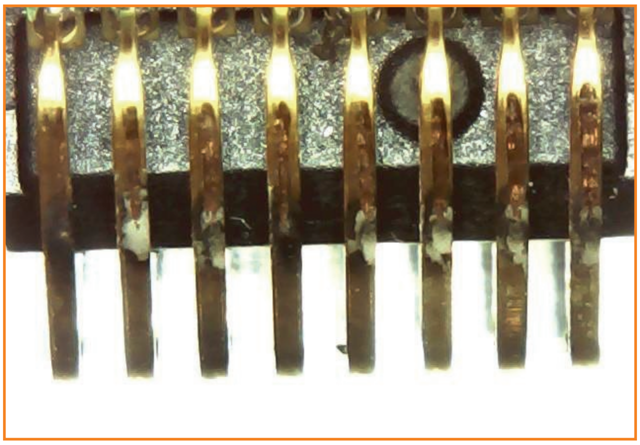

Figure 16

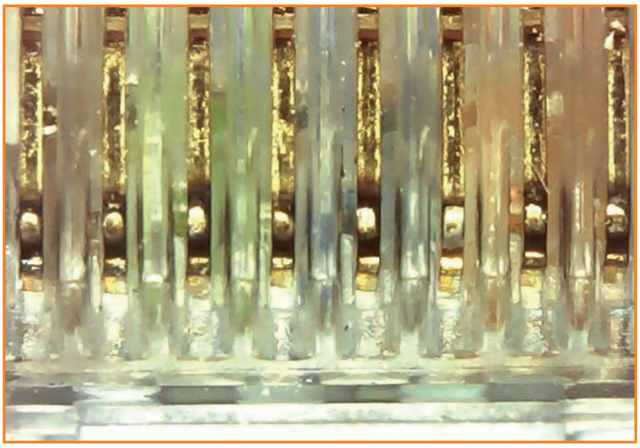

Figure 17

The net result of unplugging the cord is a small arc discharge between the removed plug and socket on one or more of the contacts. The discharge may cause a crater damage on the surface of the contacts which, in turn, increases the corrosion and over time will cause a connector failure.

Literature shows that when a plug is un-mating from a jack, if contacts are not opened simultaneously, the total power would be instantaneously carried by the last connected pair. Let's assume that a connector is delivering PoE type 4 applications at maximum 100W remote power. During un-mating operation, the last connected pair may carry 100W power at 55V open circuit voltage, which is equivalent to 2A current on each conductor of that pair. For this reason, IEC 60512-99-002 specifies a test schedule to validate the ability of a connector to withstand the un-intended un-mating operations when remote powering PoE type 4 applications.

The critical aspect of connector design to withstand this situation successfully is to ensure that the two electrical contacts from plug and jack have enough length of wiping zone so that the mating point is not the discharge point. Figure 14 and Figure 15 (previous page) demonstrate this mating operation using the CAD model of HellermannTyton's GST keystone jack.

Figure 16 and Figure 17 above show the arc discharge experiment results on a HellermannTyton GST Category 6A jack and a Category 6A plug after 50 mating cycles under 2A per conduct condition. It can be seen that the discharge caused crater damage on the copper plating. Also, the wiping trails indicate that the mating points are not the discharge points for this connector design.

The arcing only occurs for the initial disconnection within a single link and not any subsequent disconnection on the same link. While it is inconvenient if this pitting and deterioration occurs at the telecommunications outlet (TO), it is easier to replace it than an RJ45 soldered to the motherboard of a US$ 1300+ laptop. The decision about which connection to break first requires careful consideration.

This deterioration can be avoided, but not easily. While the PSE will only provide power to an appropriate PD, the power can be remotely switched off using, for example, a web-style interface, usually provided in the PSE. Using this approach, however, means that a call to the IT department would be required any time a user wanted to unplug a PD. This is clearly impractical especially in an open environment, such as a university.

The second solution to this problem involves the inclusion of a sacrificial connection between the TO and the PD. This through-coupler should be where the initial disconnection takes place, since it is the easiest and most cost-effective component to replace. The only requirement is that the user remembers to disconnect at this point before any other, as once the link is broken the power is removed and the risk of sparks and pin damage is eliminated altogether.

Moreover, when a connector continuously carries DC power, the connector starts to generate heat. This temperature rise inside the connector plus the ambient temperature negatively impacts the current-carrying capacity. The temperature rise inside the connector is characterised by the thermal properties of the materials that are used for the contacts, terminals and the insulating housing materials. IEC 60603-7 specifies a current-carrying derating curve to determine a connector's ability to carry DC current at different ambient temperatures. For example, a connector shall be able to carry 1A current at 60°C ambient temperature, which is the worst-case scenario for PoE type 4.

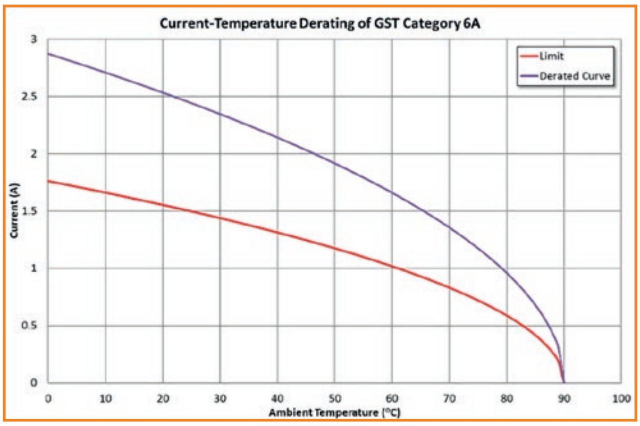

Figure 18

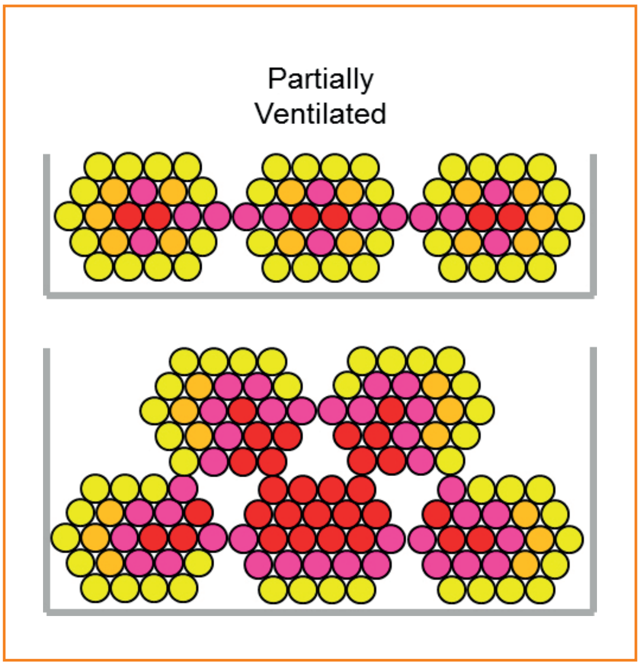

Figure 19

Figure 18 (left) shows a measured current-carrying derating curve of a HellermannTyton GST Category 6A connector. It is clear that the GST connector's capacity of current-carrying exceeds the standard specification.

Temperature rise of cabling

The effect of passing a current along a copper conductor causes it to heat up; the greater the current, the greater the rise in temperature. This is a well-known physical characteristic of copper cables. Maximum temperature rise typically occurs in or around the cable located closest to the general geometric centre of the grouping or bundle of cables.

Cables surrounding this "worst case" cable typically experience lower temperature rises with temperature rise decreasing towards the surface of the bundle. If the cable is insulated, then the heat will not dissipate freely, which would worsen the temperature rise compared to the cable in a ventilated environment. Figures 19 and 20 show the heat distribution of different cable bundles in different installation environments.

The increased amount of heat within the network adversely affects the resistance of the cabling which, in turn, causes more power to be lost as heat. While this chain reaction does reach equilibrium and stabilises, it has been proved that the temperature inside a bundle of cables, all delivering PoE, can rise by 40°C in some circumstances.

Therefore, standardisation organisations recommend maximum 15°C temperature rise in cable bundles when twisted-pair cables are energised by DC power. This recommendation

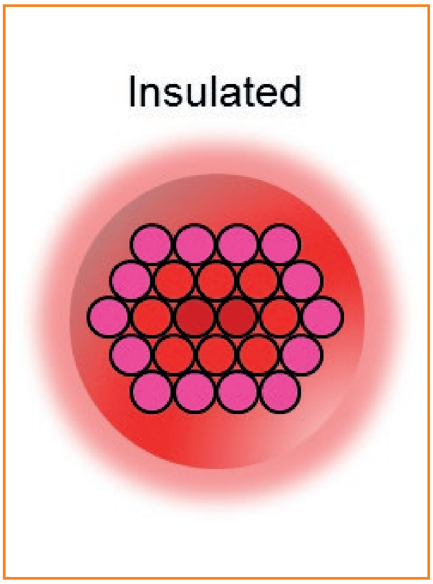

Figure 20

Figure 21

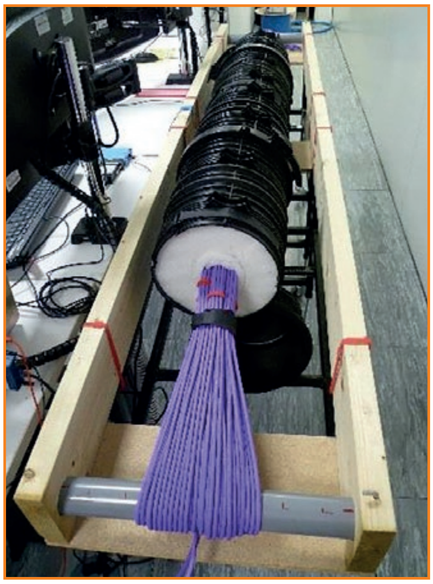

Figure 22

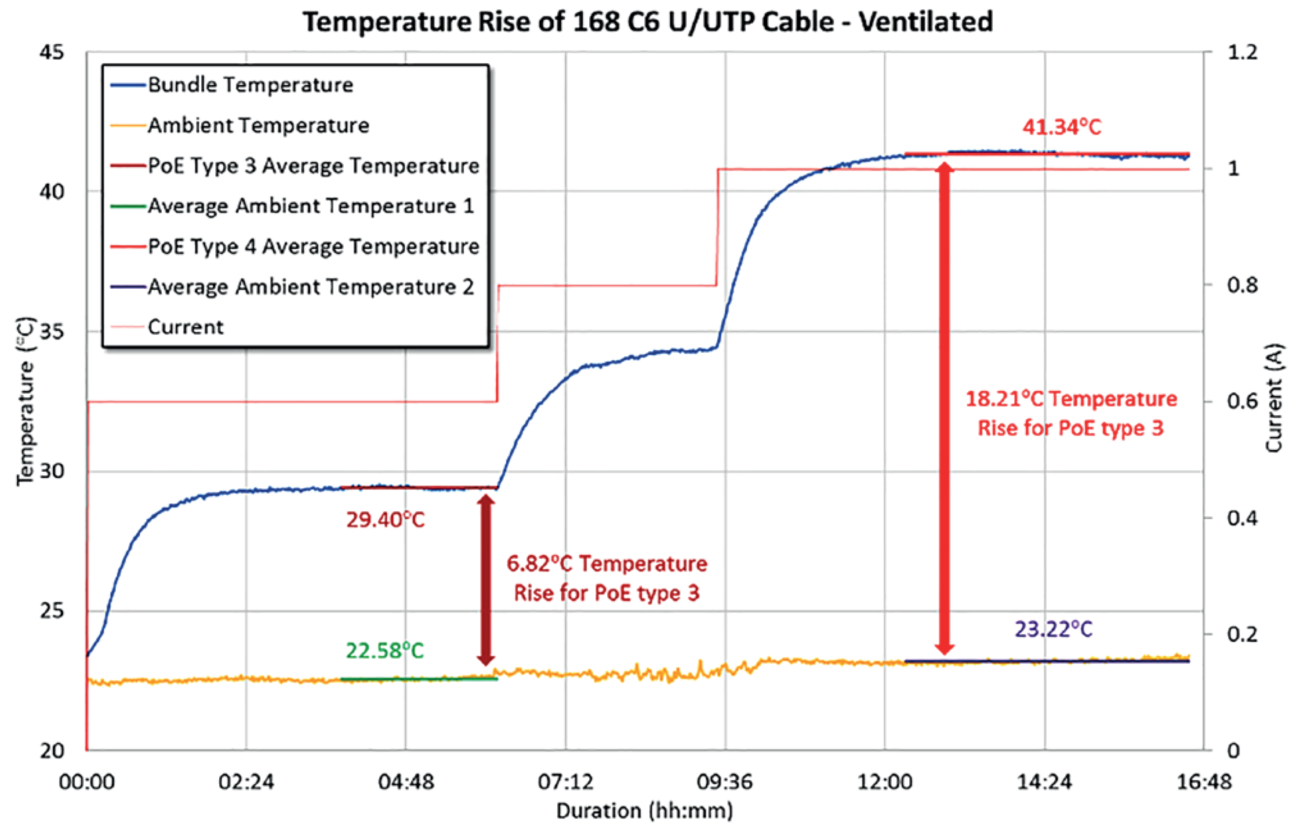

Figure 23

guarantees that cable bundles can operate at maximum 45°C ambient temperature while the highest local temperature inside bundles shall not exceed 60°C, which is the typical rated highest operation temperature of horizontal cable.

To give installation recommendations using HellermannTyton branded cabling products for PoE applications, several horizontal cable bundles were built in the HellermannTyton laboratory and the temperature rise was monitored under different conditions. Figure 21 (previous page) shows a cable bundle with 168 Category 6 U/UTP cable elements for temperature rise test in ventilated condition. Figure 22 (previous page) shows a cable bundle that was tested in conduit condition.

Figure 23 (above) depicts the maximum temperature rise of the cable shown in Figure 21, energised by 600mA (PoE type 3), 800mA and 1A (PoE type 4) current per pair.

After a good deal of experimentation and data analysis, theoretically maximum bundle sizes were determined for the mainstream HellermannTyton cabling products. Table 1B below lists the maximum cable bundle size for 15°C temperature rise in different installation environments and PoE applications at 20°C ambient temperature. In this assumption, the cable element at the geometrical centre shall operate at 35°C local temperature. Therefore, the maximum channel length is also recommended.

	Ventilated		**Conduit**		**Maximum Channel Length (m)**
	Type 4	Type 3	Type 4	Type 3	
C6A U/UTP	168	574	98	350	96
C6A U/FTP	161	525	105	378	97
C6 U/UTP	133	469	77	301	94

Table 1B

	Ventilated		**Conduit**		**Maximum Channel Length (m)**
	Type 4	Type 3	Type 4	Type 3	
C6A U/UTP	147	511	84	315	89
C6A U/FTP	140	476	91	343	93
C6 U/UTP	117	420	70	266	83

Table 2

Table 2 lists the similar information for 45°C ambient temperature. In this assumption, the cable element at the geometrical centre shall operate at 60°C local temperature. It can be noticed that the recommended maximum channel lengths are shorter than 20°C ambient temperature.

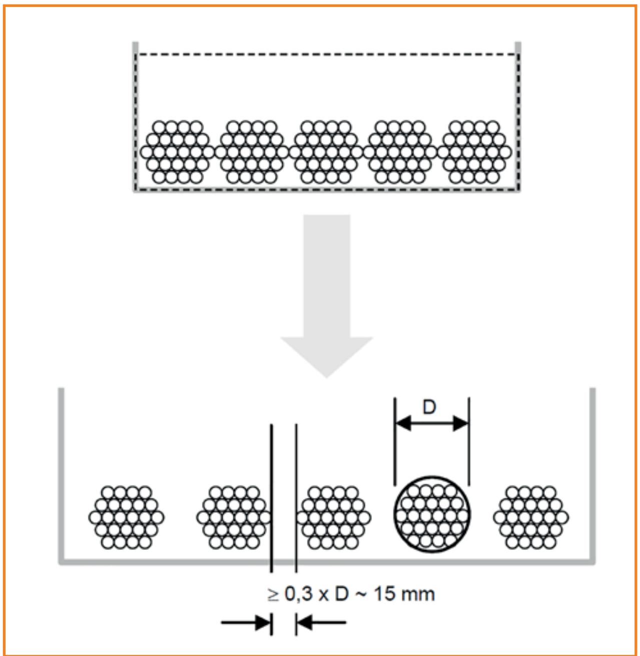

Figure 24

In field installation, it is rare that the cabling system is installed in a single giant bundle with several hundreds of cable elements. In ISO/IEC 14763-2, the size of cable bundles is restricted to a maximum of 24 4-pair balanced cables to support remote powering. Figure 24 illustrates this recommendation. If the maximum temperature rise still exceeds 15°C, the cable bundles should be separated by vertical "chimney" widths of 0.3 of bundle diameter, allowing each bundle to cool by convection.

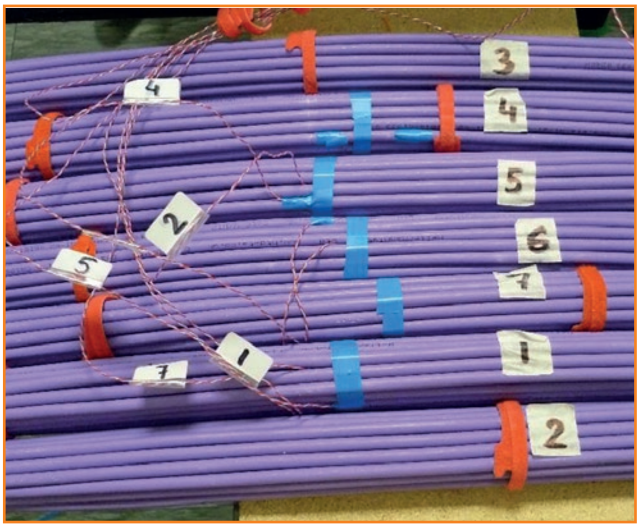

Figure 25

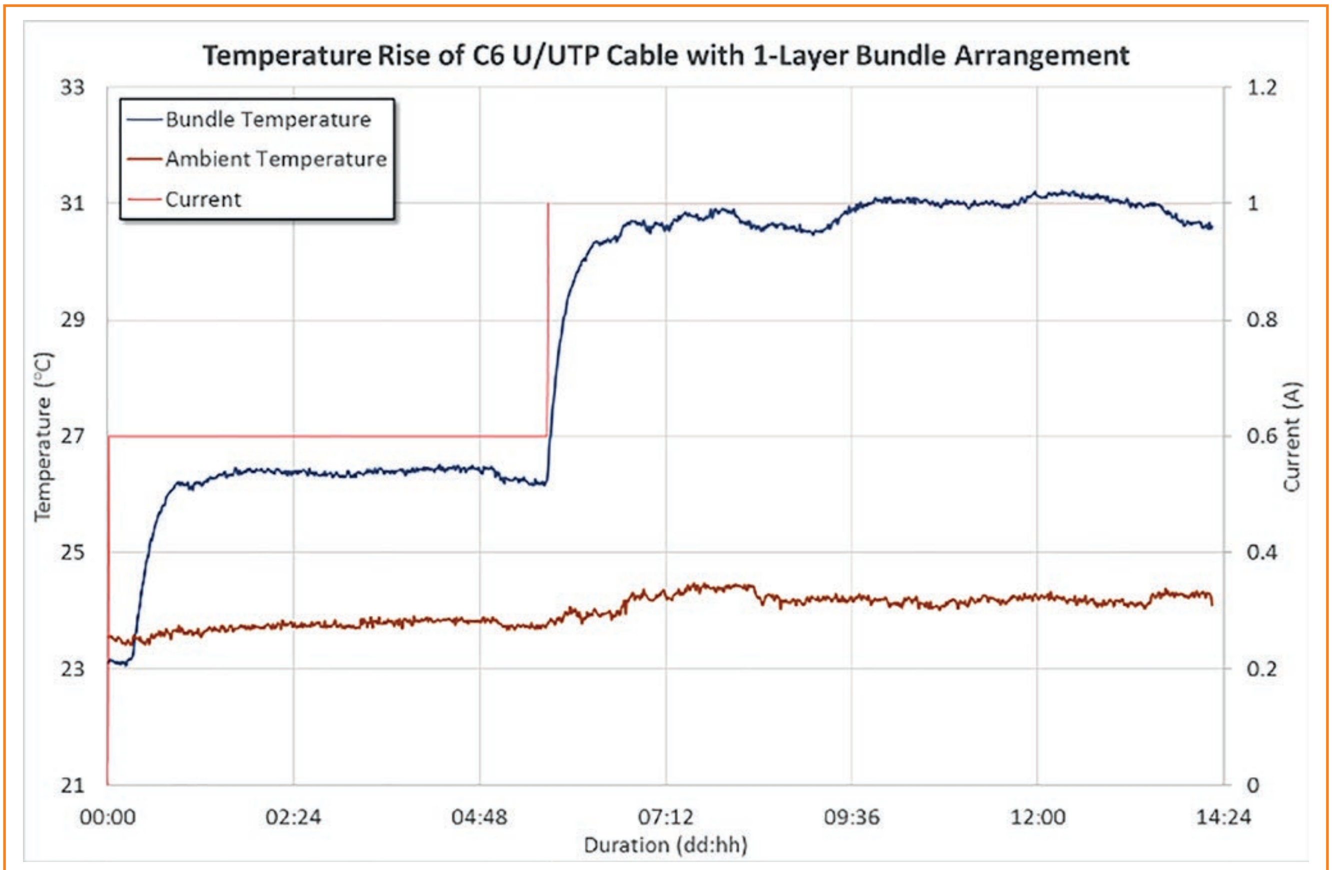

Figure 26

> **Managing the cabling used to deliver remote power on a day-today basis will be the key to successfully utilising higher-powered Ethernet networks.**

Figure 25 shows Category 6 U/UTP cable bundles were tested with this recommendation without separation. Figure 26 plots the monitored temperature rise at different positions. The maximum temperature rise is about 7°C.

For cable installation in a ventilated condition, HellermannTyton recommends maximum two-layer stack up of 24-element cable bundle to deliver PoE type 3 and maximum one-layer without separation for PoE type 4[1]. Table 3 lists this recommendation.

	Category 6	Category 6A
PoE Type 3	Two layers	Two layers
PoE Type 4	One layer	One layer

Table 3

The future of PoE

IEEE 802.3bt-2018 was published by IEEE which is to source between 71W and 100W, delivering a minimum power of 71W with up to 0.96A per conductor. This suggests, however, that up to 29W of power could be lost as heat in every PoE circuit.

Today's PDs don't require this degree of power to operate and at present, with the exception of VoIP telephones, wireless access points and the like, they are not deployed in large numbers. However, the significant increase in available power from the next generation of PSE will mean that more applications can take advantage of remote powering. Applications from 10GBASE-T wireless access points, supporting larger numbers of connected users, to building management and control systems and LED lighting are all valid reasons for increasing the amount of power that can be delivered using data cabling.

The use of PoE for commercial LED lighting solutions is still in its infancy. Figure 27 above shows a typical configuration for such a system where the power and control for the lighting is supplied over standard data cabling. Environmental sensors taking readings of ambient light levels or human activity are also connected to the network so that lights can be intelligently controlled and power consumption minimised.

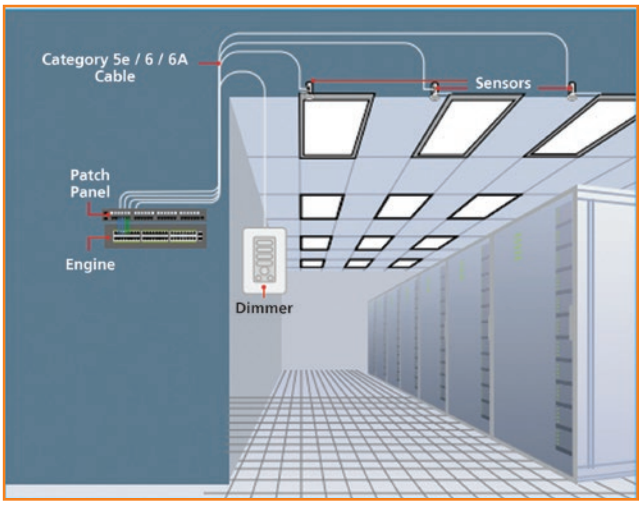

Figure 27

This increased usage of data infrastructure is not without its implications, however. With up to 29W of power being lost as heat for every active link, and the utilisation of a large number of links for non-traditional services such as LED lighting, it is now more important than ever to ensure that the type and performance of the chosen infrastructure products are given careful consideration.

As previously explained, these elevated powers will give rise to elevated temperatures which, in turn, will reduce the end-to-end performance of the data infrastructure and ultimately reduce either its length or data throughput. Knowing where these circuits are within your building infrastructure, and the circuits that they will probably affect as their temperature increases, will become even more important, especially as the lines blur between facilities and IT with respect to physical network design, installation and management.

Managing the cabling used to deliver remote power on a day-to-day basis will be the key to successfully utilising higher-powered Ethernet networks. Careful consideration should be given to those devices that utilise Ethernet cabling to deploy power but do not meet the requirements of the 802.3 range of standards. These non-standardised devices may source even higher power levels and be accompanied by greater temperature rises

1 Temperature rise performance in conduit/insulation condition is currently under study.

Keeping up to date and accurate records of power consumption and understanding how the cabling system is ventilated will ensure a healthy network infrastructure for all users.

and reductions in effective channel length which, in turn, will accelerate the performance decline of the infrastructure.

Further evaluation of the impact of 802.3bt and beyond is needed to determine the long-term effects on twisted-pair cabling systems.

Conclusion

The use of twisted-pair cables to deliver both power and data to remote devices will probably see an upturn in the future as the amount of available power increases to a useful level and more sophisticated and power-hungry devices utilise it. This increase comes with a greater responsibility to know how much power is being delivered over the cabling system and which routes it takes.

Keeping up to date and accurate records of power consumption and understanding how the cabling system is ventilated will ensure a healthy network infrastructure for all users. Ensuring that those departments responsible for the distribution of both power and IT services, over what has largely been considered an IT resource, should be clearly defined to avoid one powered service from disrupting another.

Adhering to industry standards and installation guidelines is an essential first step in ensuring a successful and high-performing remotely-powered building infrastructure.