

White Paper



10 Gigabit-Ethernet over Copper



Convincing cabling solutions

Table of Contents

1.	Introduction	3
2.	History of Ethernet	3
2.1.	ALOHA	3
2.2.	Start of Ethernet	3
2.3.	10 BaseT	4
2.4.	Fast Ethernet 100Base T	4
2.4.1.	ISO/IEC 11801	5
2.5.	Gigabit-Ethernet	5
2.5.1.	Alien NEXT and Echo	5
2.5.2.	PS- NEXT	6
3.	Standards activity at IEEE	6
3.1.	10 Gigabit-Ethernet over Copper	7
4.	New Challenges for 10 Gigabit Ethernet over Copper	8
4.1.	Line Coding	8
4.2.	Solution inside the cables	9
4.2.1.	Channel assumptions	9
4.3.	Parameters outside the cable	11
4.3.1.	External noise	11
4.3.2.	Alien crosstalk	12
4.3.3	Shielded cables	14
5.	Summary	14
6.	Further information	14

© Copyright 2003 Reichle & De-Massari AG (R&M). All rights reserved.

This publication or parts out of it may not be copied or distributed in any form or for any purpose without the express written authorisation of Reichle & De-Massari AG. Any information contained in this publication can be changed without prior notice.

Utmost care went into creating this document. The information it contains reflects product specifications at the time of publication. Subject to changes in technical specifications.

10 Gigabit Ethernet over Copper

1. Introduction

10 Gigabit Ethernet over copper! At first you may think we've made a mistake - shouldn't it say fiber? No, we mean twisted pair. IEEE started a study group last November which is looking at the feasibility of this technology. The following article will show the history of Ethernet, and discuss issues being evaluated and new approaches being considered in the 10 Gigabit sector.

Application:	Enterprise Cabling
Technology:	10 Gigabit Ethernet over Copper
Format:	White Paper
Topic:	Evaluation and new approaches in the 10 Gigabit sector
Objective:	Information on history of ethernet, standards activity at IEEE, new challenges for 10 Gigabit Ethernet over copper
Target group::	Distributor, Planner, Installer, End-user
Author:	René Trösch, R&M
Published:	July 2003

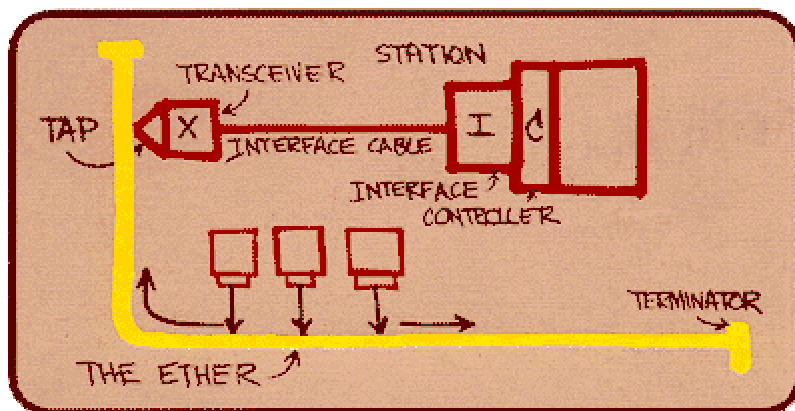
2. History of Ethernet

2.1. ALOHA

In 1970, the University of Hawaii started to link external computers together using radio transmissions. This system was called ALOHA. The first frequency channel was used for the uplink, the other channel for the downlink transmission. A problem occurred when more than one station was sending. The error could be detected but the whole transmission system slowed down. With the Carrier Sense Multiple Access (CSMA) technique the station listened to see if the radio channel was free. Only in this case would the station start to send. Because of the radio transmission and the delay of the signals it could occur that two stations were sending at the same time, causing errors. Nevertheless, when an error was detected the stations stopped transmitting and started again after a certain period. With this method, called Collision Detection (CD), the throughput could be optimized. CSMA/CD is still used in several protocols.

2.2. Start of Ethernet

At the "National Computer Conference" in 1976, employees from Xerox presented an Ethernet which could be 1km long and have 100 attached stations. The transmission speed was 2.94Mbit/s



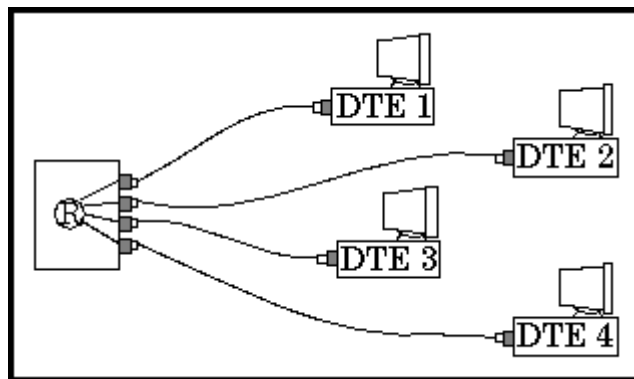
The terminology used in this first picture is still current today . The system used a coaxial cable (yellow cable) on which several stations could be attached.

The first IEEE 802.3 standard was published in 1985 and allowed transmission speeds up to 10Mbit/s over several coaxial cables. That Ethernet version was called 10Base5, where the 5 stands for 500m. The following 10Base2 (Thin Ethernet) allowed for 200m over RG 58 coaxial cable and was followed by the 10Broad 36 for 3600m coaxial length.

2.3. 10 BaseT

Twisted pair technology started in 1990 with the standardisation of 10BaseT. At that time a link length of 100m and a bit error rate of 10^{-8} were specified. It was decided to use common telephony cable. The electrical parameters of that twisted pair cable were designated as Cat. 3.

At that time the pair combination of Ethernet over twisted pair was defined and since then the pair combination 1,2-3,6 stands for Ethernet. The pair 1,2 was used for transmitting signals and the pair 3,6 for receiving. The data were transmitted in the half duplex mode, which means that only one station could send and the other could receive. The well-known telephony architecture was used. Later, the star architecture was used, where all stations were coupled over a repeater or a hub.



The advantage of twisted cable was that it could be used for telephony or data transmission and that the immunity against external sources could be improved. The disadvantage was that the length of the segment was reduced to 100m.

2.4. Fast Ethernet 100Base T

When the IEEE standardisation committee met to begin work on a faster Ethernet system, two approaches were presented. One approach was to speed up the original Ethernet system to 100-Mbps, keeping the original CSMA/CD medium access control mechanism. This approach is called 100BASE-T Fast Ethernet. Another approach presented to the committee was to create an entirely new medium access control mechanism, one based on hubs that controlled access to the medium using a "demand priority" mechanism. This new access control system transports standard Ethernet frames, but does it with a new medium access control mechanism. This system was further extended to allow it to transport token ring frames as well. As a result, this approach is now called 100VG-AnyLAN.

The IEEE decided to create standards for both approaches. The 100BASE-T Fast Ethernet standard described here is part of the original 802.3 standard. The 100VG-AnyLAN system is standardised under a new number: IEEE 802.12.

By 1997 three different fast Ethernet standards were published: The 100Base T2, a 2-pair transmission system over Cat. 3 cables, the 100Base T4, a 4-pair transmission system over Cat. 3 cables and the 100Base TX, a 2-pair transmission system over a Class D channel. These protocols and systems were backward compatible to 10 BaseT.

The 100BASE-TX and 100BASE-FX media standards used in Fast Ethernet were both adopted from physical media standards first developed by ANSI, the American National Standards Institute. The ANSI physical media standards were originally developed for the Fiber Distributed Data Interface (FDDI) LAN standard (ANSI standard X3T9.5), and are widely used in FDDI LANs. Rather than "re-inventing the wheel" when it came to signalling at 100-Mbps, the Fast Ethernet standard adapted these two ANSI media standards for use in the new Fast Ethernet medium specifications. This allowed old PC cards to still be used in parallel with newer ones.

2.4.1. ISO/IEC 11801

Remember, in 1995 the first edition of the international standard ISO/IEC 11801 was published. The length of the channels and the transmission rates matched those of Ethernet and token ring. With this standard it was now possible to create a passive telecommunication infrastructure on which different protocols could run over the same media.

Another important step was taken at this time. The transmission could be done in a full duplex mode, which meant that both stations could send and receive signals at the same time.

Today this protocol is the most used to connect work stations together and can be purchased for a few Euro.



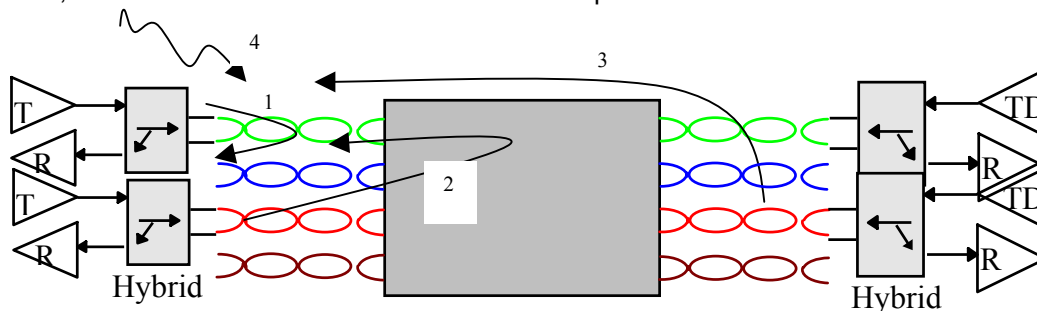
2.5. Gigabit-Ethernet

At the same speed that the computer industry put faster computers on the market, the need for more bandwidth and higher transmission speeds also increased. From experience gained from Fast Ethernet technology the following things are known:

- 100 Base-TX demonstrates that sending a 3-level symbol stream over Category 5 cables at 125Mbaud is possible.
- 100Base T4 demonstrates techniques for sending multilevel coded symbols over four pairs.
- 100Base T2 demonstrates the use of digital signal processing, five level coding and simultaneous two-way data streams while dealing with alien signals in adjacent pairs.

2.5.1. Alien NEXT and Echo

Due to the fact that the pair was now used for sending and receiving at the same time, the parameters NEXT and FEXT, as well as Alien NEXT and Echo became important.



The pair's own reflected signal (1) could be easily deleted in the processor. The signal from the same side (2), called NEXT, or near end crosstalk, was also known because the DSP was generating the signal and dividing the information on the two pairs. More difficult to detect were the wrong signals from the far end side (3), called FEXT or far end crosstalk, and the noise from outside the cable (4). Part of the noise (3, 4) could be cancelled with filters, allowing only the signal frequencies to pass. In the 100Base-T2 there is a fifth order filter which starts at 20MHz. This means that signals above 30 MHz are attenuated very strongly and can't disturb the receiver.

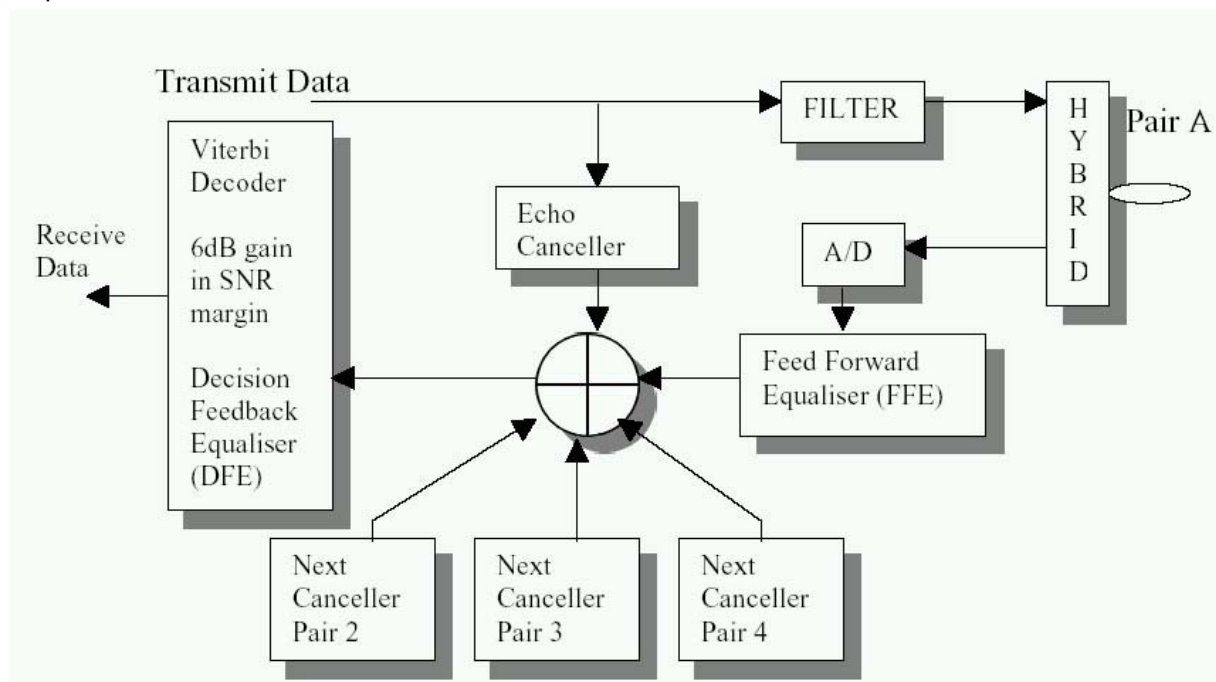
With Gigabit-Ethernet there are transceivers and receivers on all 4 pairs, so NEXT and Powersum NEXT are also important. The Powersum refers to the sum of the influence of the other 3 pairs in the cable on the fourth pair.

2.5.2. PS- NEXT

With Gigabit-Ethernet, transceivers and receivers are on all 4 pairs, resulting in increased noise. To meet the requirements new parameters had to be introduced, like Powersum NEXT and PS-ELFEXT. These parameters measure the disturbance from 3 pairs on one receiver at the same time. Because IEEE has taken a lot from Fast-Ethernet, the actual Powersum values are the same as the NEXT values from the first standard.

The DSP (Digital Signal Processor) also has a higher level of complexity. It uses 4 echo cancellers and 12 NEXT cancellers. They also had to put in a new signalling scheme (5 levels instead of 2), called PAM 5*5, and, to regain the loss of the SNR, a 4-D trellis code and a forward correction error were added.

The scheme of the Gigabit-Ethernet processor for only one pair is shown in the following picture. The whole chip will continue to contain this amount of cancellers and filters for some time.



This NEXT canceller in the Gigabit-Ethernet brings 10dB NEXT gain. This could be simplified if an STP system was used. The NEXT of such a system is much higher than a normal UTP system.

Before discussing 10 Gigabit Ethernet over copper, the IEEE and how it works should be explained.

3. Standards activity at IEEE

The IEEE's involvement in electrical standards dates back to 1890, when the AIEE (American Institute of Electrical Engineers) proposed a recommendation for the practical unit of inductance (henry). As a pioneer in voluntary electrical and information technology standards activity, IEEE became a founding member of ANSI (American National Standards Institute) in 1918.

In 1963, when the AIEE merged with the IRE (Institute of Radio Engineers) to form the IEEE, a formal standards body was established to support standards development. Envisioning the expanded role that standards were to play in the future and their impact on industry, IEEE formed its first Standards Board in 1973. As a standards body, the IEEE-SA (Standards Association) has responded to changes in the marketplace and as a result, the IEEE-SA of today is quite different and innovative, but still committed to providing the most current and reliable standards knowledge.

IEEE deals with the whole range of electronic and electrical engineering. In our case the LAN/MAN part, where the number 802 is assigned, is the most important. All the technologies which deal with the two lowest levels of the OSI model are under this project. The most important is 802.3 for Ethernet, and nowadays the 802.11 for WLAN's. The 802.5, token ring, is not active any more.

Projects for amendments have the numeric designation of the standard they amend plus a lowercase letter. Today they are working on the 802.3ak, 10GBase CX4, which means that so far 37 projects have been done to improve Ethernet performance (26 for a – z + 11 for a – k).

3.1. 10 Gigabit-Ethernet over Copper

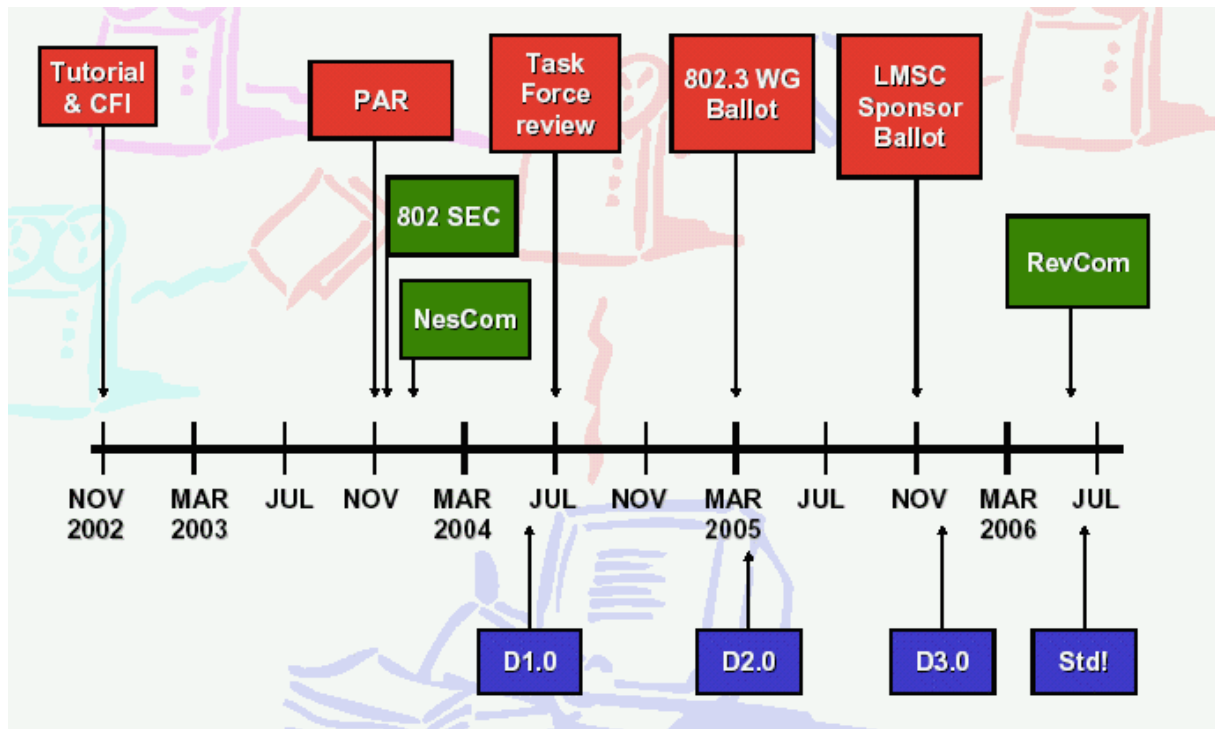
Last November, IEEE started a project team for 10Gigabit-Ethernet which has one year time to submit a project authorisation request (PAR) to the 802 plenary. This PAR has to be approved by the standards activity board and NESCOM.

The most important issue is to show evidence of the 5 PAR criteria:

- Broad market potential (broad sets of applicability, multiple vendors and users)
- Compatibility (conforms to the architecture of 802.1)
- Distinct identity (different from other standards, one solution per problem, easy for the reader)
- Technical feasibility (demonstrated system feasibility, proven technology, reliability)
- Economic feasibility (known cost factors, reasonable cost performance, installation costs)

If the PAR is submitted and approved, the IEEE Group has a maximum of 4 years to develop the whole standard. The current study group plans to submit the PAR and the criteria by this November. If the PAR is approved, they will have four years to develop the standard. This means that a standard will be available at the earliest in 2006.

The current plan is as follows:



The biggest effort in terms of the above mentioned five criteria will be in the area of technical feasibility. Is it actually possible to transmit 10Gbit over copper wires and a distance of 100m? Below we will review the topics which are being discussed and look at the possibilities for achieving 10 Gigabit Ethernet over copper.

4. New Challenges for 10 Gigabit Ethernet over Copper

The target is to transmit 10 Gigabit per second over a channel of 90m with a bit error rate of at least 10^{-10} . If we were to use the normal approach, where 10 Gigabit is equal to 10GHz of bandwidth, we would immediately see the impossibility of that approach.

We therefore need a method to reduce the amount of bandwidth. One way is to split the signal on different pairs and the second is coding. As on 1000Base-T, the signal for 10 Gigabit-Ethernet is split on four pairs. We therefore need to transfer 4 times 2.5 Gigabit instead of 1 time 10 Gigabit.

4.1. Line Coding

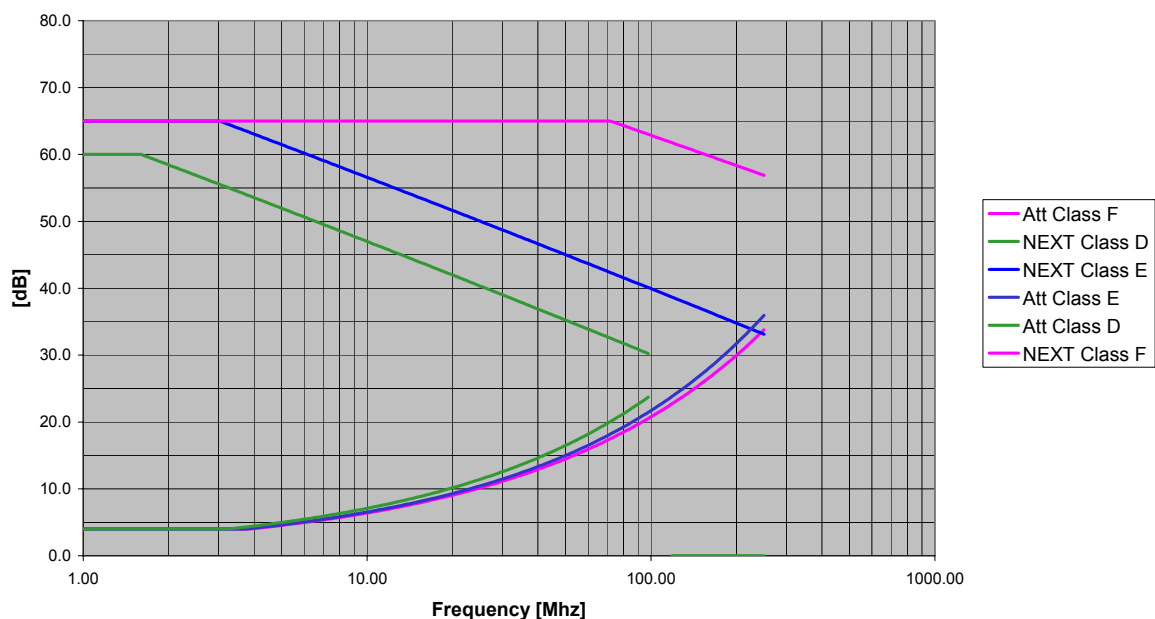
The second approach is coding, where more than $\frac{1}{2}$ bit is transmitted per pulse rate. The following table shows the currently discussed line codes for 10 Gigabit Ethernet and the needed bandwidth.

Line Code	Bits/ baud	Signal Bandwidth [MHz]	Baud Rate [MS/s]	ACR [dB]
MLT 3	1	1250	2500	21
PAM-5	2	625	1250	25
PAM-9	3	416	833	31
PAM-17	4	312.5	625	36
PAM-33	5	156.25	312.5	42

PAM-5 is used in 1000Base-T. The signal bandwidth in that region is 62.5 MHz which is enough to run Gigabit-Ethernet over class D channels.

We now see the first mismatch between the possible linecoding, the needed bandwidth and the current channel implementations. Class D stops at 100MHz, class E at 250MHz and class F at 600Mhz, but the ACR is not big enough. One of the other issues is that the capacity of the channel must be larger than the desired 10Gigabit. Imperfect equalization in the incoming processor, jitter and analog impairments will require more ACR or more possible capacity. The actual calculation from IEEE shows that the channel capacity must be approximately 16.6Gbit/s to allow a real 10Gigabit data transfer. The capacity for Cat. 5 is in the region of 2.5Gbit/s. Both headroom and capacity would then be in the same range, which seems reasonable. Due to the fact that the American market is mostly unshielded, their focus will be on the Cat. 6 Class E range. Cat. 7 currently has the problem that the balance of the cable itself is less stringent than Cat. 6 cables, making the coding and decoding more problematic.

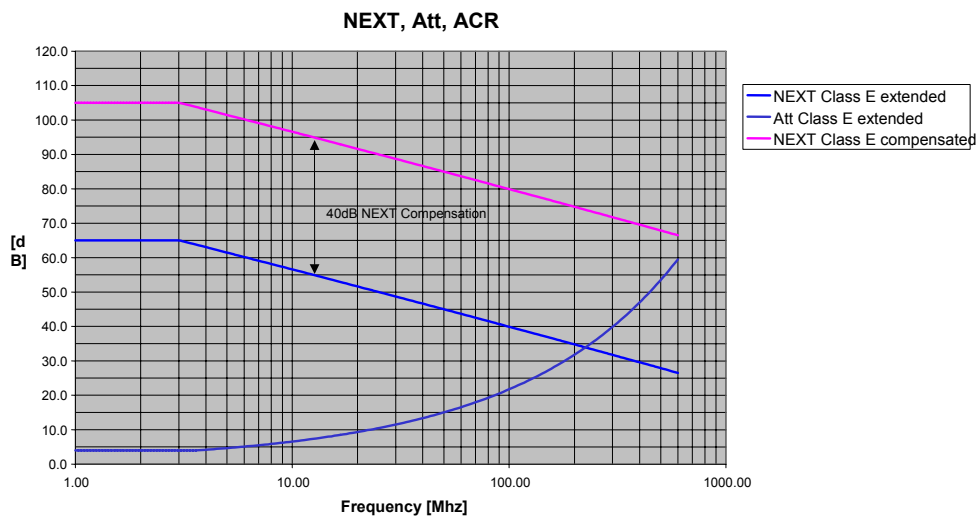
NEXT, Att, ACR



4.2 Solution inside the cables

We have seen that the ACR is not big enough on Class E channels. One solution is to look at the channels far beyond 250MHz to see what happens. This will give an extended class E channel up to 600MHz. The blue lines show that the ACR will be negative after 200MHz. The solution implemented for 100Base-T2 and 1000Base-T was to cancel the NEXT. In 100Base-T the cancellation was approximately 10dB, for 10 Gig Ethernet we will need at least 40 dB.

The same must be done for RL and FEXT. In those areas we are talking about 50-70dB echo cancellation and 15-25dB FEXT cancellation. Cancellation is dependent on the capacity of calculation power and clock speed. So it will not be a big issue. The issue is to know if it is possible to extend the graph values linearly to 600MHz.



4.2.1 Channel assumptions

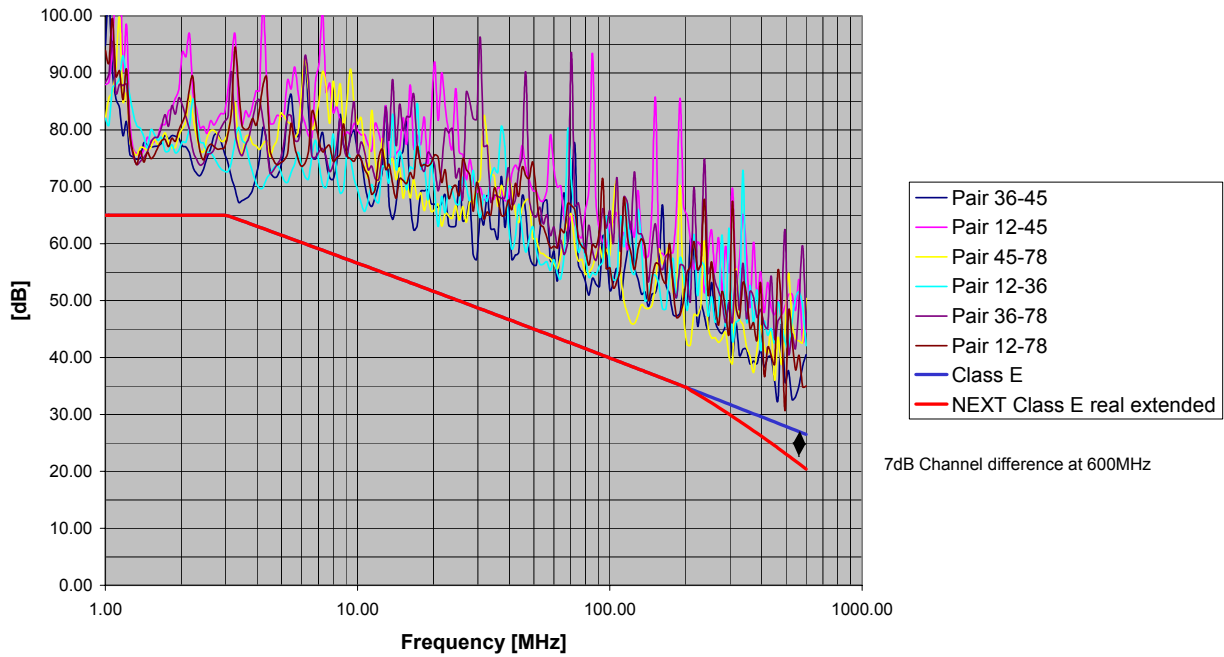
The channel calculation in the generic cabling standards are based upon a component up approach. This means that the components are categorized with formulas and then the channel is built up (summation of components).

Attenuation is a linear function which can be measured up to x MHz. One of the problems is the insertion loss deviation which starts at approx. 250MHz and gets bigger in the higher frequency range. The insertion loss deviation is a known factor which, depending on the balance of the cable, can be several dB's.

The NEXT of the cable also looks very linear up to a certain frequency where resonances (impedance mismatch) destroy the good values. The cut off frequency can be designed for the cables and therefore should not cause too many problems.

The NEXT for the connecting hardware is actually calculated with 20dB NEXT loss per decade of frequency. The reality shows that this value is true up to 200MHz and afterwards the NEXT loss changes from 20dB/decade to 40dB /decade. At 600MHz it makes a difference of approximately 7 dB. The measurement on a R&Mfreenet UTP Cat. 6 channel shows that from the cable and the connecting hardware enough headroom is available to meet the more stringent NEXT.

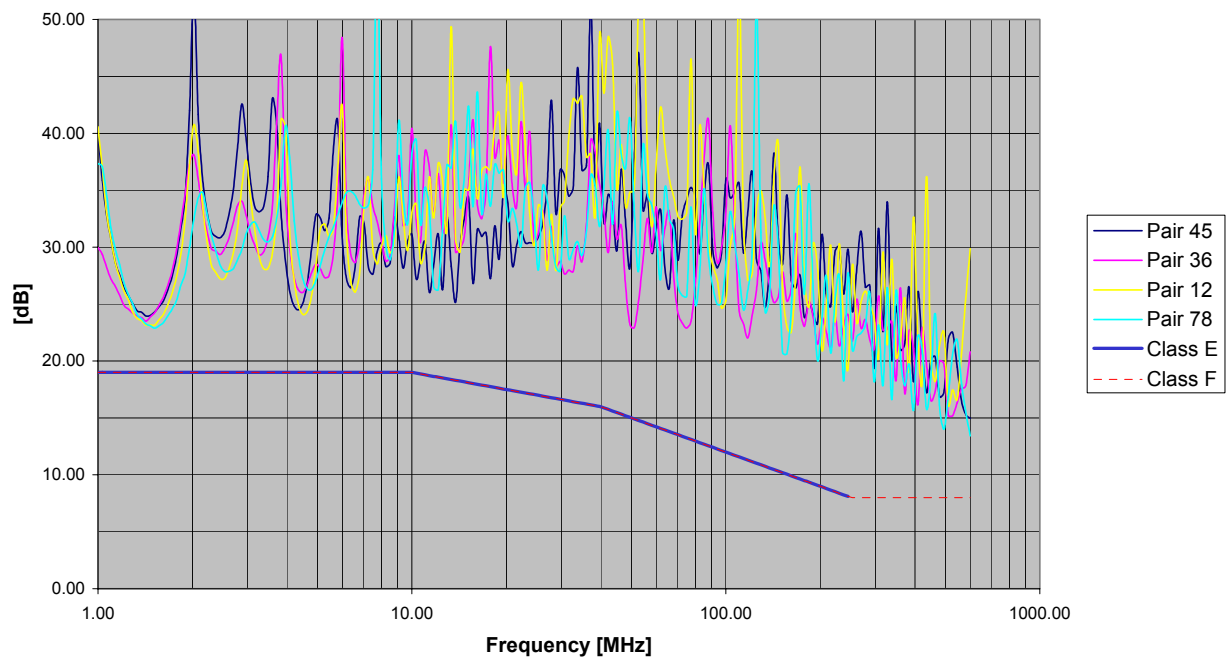
NEXT



The FEXT is similar to the NEXT and should also be met. Also RL is not a critical issue because the Class F standard follows the class E line and has a flat line of 10dB.

Measurements show that this can be easily achieved.

Return Loss



More critical are the delay parameters. The signal of 2.5Gbit is transmitted in parallel over all four pairs, so it is crucial to add the bits back together correctly. This means that the delay skew between the 4 pairs must be very low. The current standards specify a maximum of 50ns. Actual values on the channels are below 20ns. Otherwise big input buffers are needed to store the bits and put them together.

If only the internal parameters of the cables are taken into account it looks like it is possible to transmit 10 Gigabit-Ethernet over a Class E channel which is extended to 600MHz.

4.3. Parameters outside the cable

We have seen that it should be possible to transmit 10 Gig-Ethernet through a twisted pair cable. The thing which is always forgotten is that in a building not only data systems are running, but also light, power, machines, etc. In the standards a minimum distance between power and communication cables must be fulfilled, but in a lot of cases nobody cares. This noise is most often called external noise. A special noise is the noise which comes from other data communication cables which are bundled together. This noise is called alien crosstalk.

4.3.1. External noise

If we consider that noise can come from different sources, we must look at the frequency range of those services. For example, UKW radio stations transmit in a frequency range of 87.5-108MHz. TV signals transmit up to 862 MHz. A lot of other applications like navigation systems for cars, airplanes, airports and internal services in offices transmit signals in the same frequency range as we want to transmit.

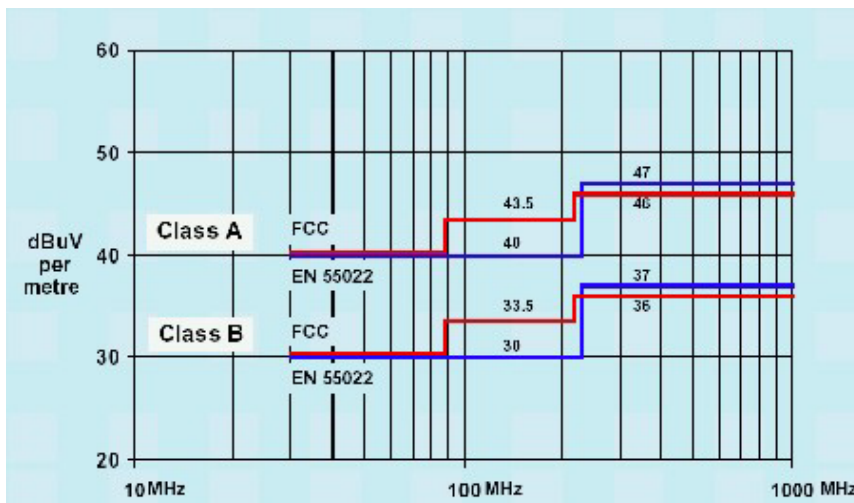
Another noise is the conducting noise on the power cables, which can be very high and also in the same frequency range. The good thing for both is that with distance to the source, the level of noise decreases. This fact is also written in the EN 50174-2:2000: Installation planning and practices inside buildings.

Type of installation	Distance between cables without metallic divider
unscreened power cable and unscreened IT cable	200mm
unscreened power cable and screened IT cable	50mm
screened power cable and unscreened IT cable	30mm
screened power cable and screened IT cable	0mm

If the noise is not known and the installation practice can't be assured, it would be better to use shielded systems and to build a correct earthing system.

It is important that for active components like PC hubs, as well as coffee machines, household appliances, welding machines, etc. emission and immunity standards exist, which must be met.

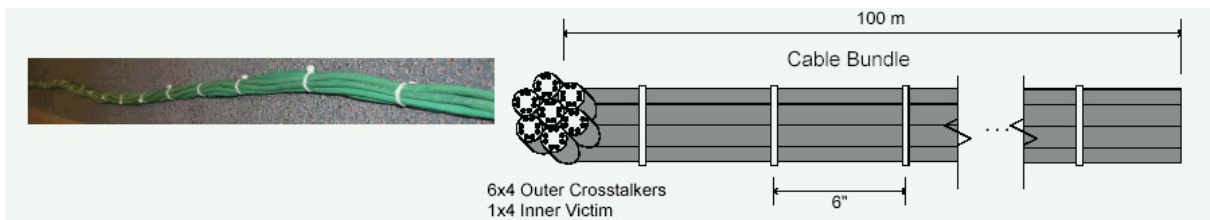
For IT equipment the EN 55022 for emission and the EN 55024 for immunity standards apply. In EU countries, the radiated emission limits for IT equipment must be met for the CE mark to be awarded.



The immunity criteria must also be fulfilled, which means that a device should continue to run without errors or crashes when a field of 3V/m is applied.

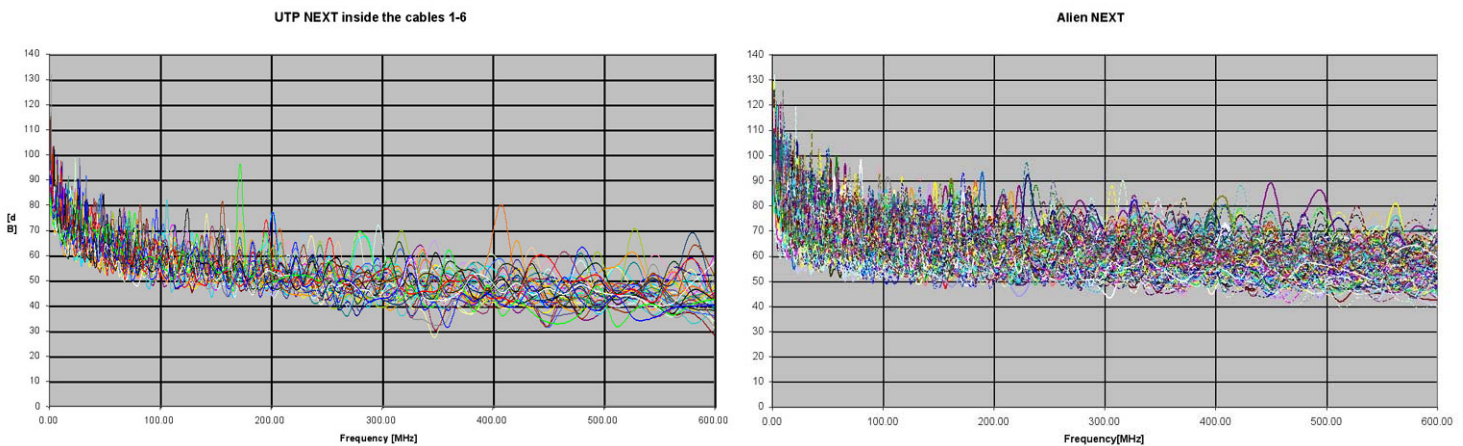
For transmitting 10Gig-Ethernet this external noise has to be taken into account. The IEEE currently calculates using a noise floor of -140dBm/Hz, which is the current limit. Better would be -150dBm/Hz if measurements show that this level can be reached. They will use this value for designing the processors or they will have to put in some noise cancellers, which will increase the amount of transistors on the chip. In all cases, an earthing system and the choice of a screened system is a safer way to obtain error-free transmissions.

4.3.2 Alien crosstalk



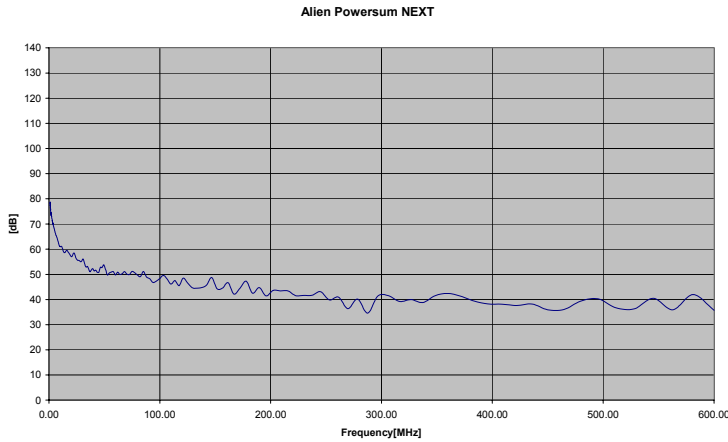
On good installations the cables are bundled together on ladders over several meters. This tight construction will cause some noise to travel from one cable to the other.

The measurements done in the R&M development and testlab look similar to those from other companies. For 6 unshielded Cat. 5e cables bundled together as shown above the following results were measured:



The alien NEXT is worse than the inside NEXT on low frequencies up to 50MHz. Above this frequency the alien NEXT is approximately 10dB better than the inside NEXT. The method of NEXT cancellation, which is very easy inside one cable, becomes very difficult with alien NEXT. The alien NEXT is seen with short ranges of bundled cables. The measurements show that direct bundling of only 3m is enough to get those results. Three meters is the common range where the cable is layed out and bundled in the rack.

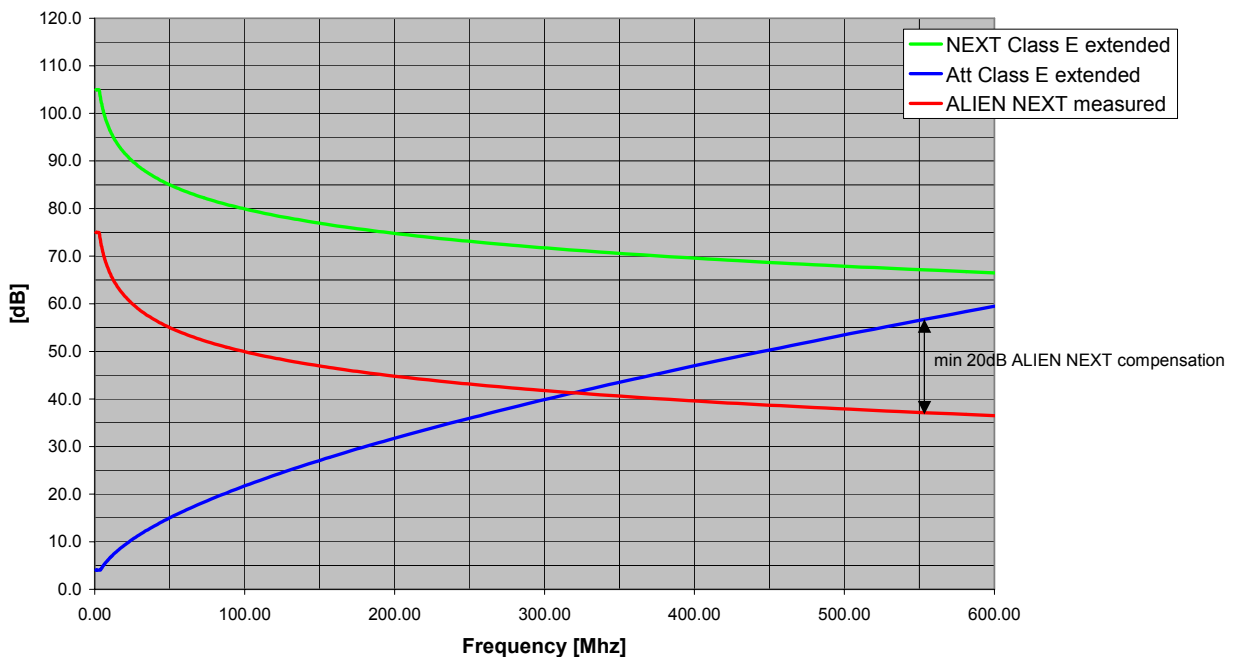
If the NEXT is cancelled with 40 dB the receiver will still see the alien NEXT. It becomes even worse if we calculate the Powersum Alien NEXT of 5 cables with 4 pairs disturbing one pair of the affected cable.



The Powersum Alien NEXT values are below the NEXT values inside the cable, which will be cancelled. For the alien ACR this value is clearly not enough, because the attenuation will be bigger than 30dB at 250MHz, making alien crosstalk the crucial point. The difference of the alien NEXT from Cat. 5e to 6 is very small, so we can use the values of Powersum Alien NEXT for both.

To summarize all NEXT into one plot:

Powersum



The Powersum Alien NEXT has to be compensated with at least 15-20 dB. Depending on the layout of the cables, the alien NEXT and Powersum Alien NEXT can vary over a range of 10dB. This makes it more difficult to find a good and valuable solution. IEEE is reviewing an alien NEXT compensation of 0-15 dB. Without this alien NEXT reduction, the transmission of 10Gig-Ethernet is not really possible with bundled cables.

Considering that the noise from other cables is not known and that other protocols can run over those cables, the following possibilities can optimize the alien NEXT:

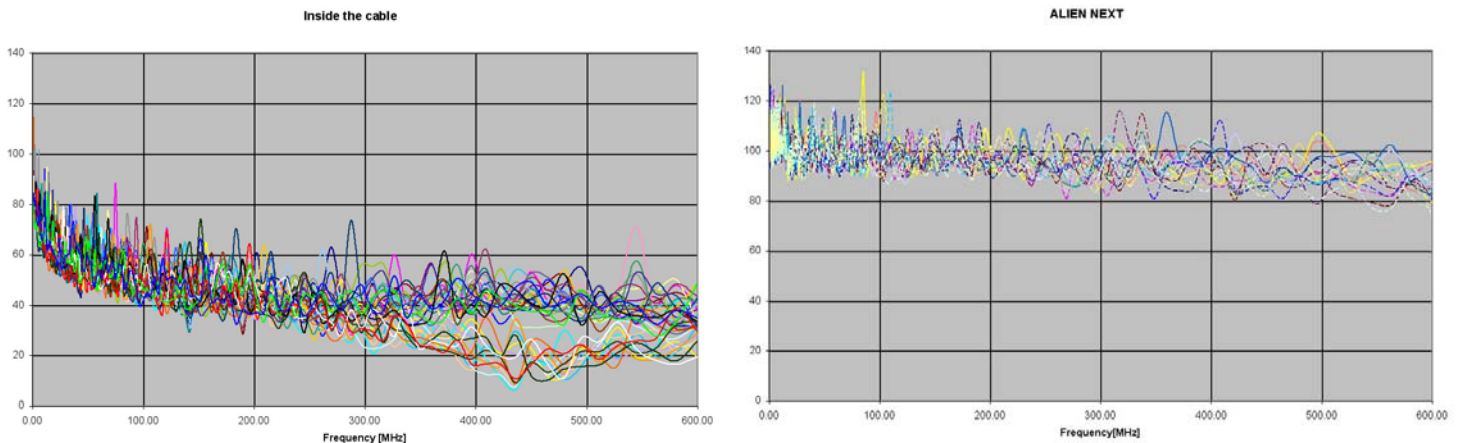
- Thicker outer jacket of the UTP cable. The cable will become maybe 8-10mm thick, which will influence the filling factor of a conduit. Also connecting a 10mm thick cable on a module will not be easy.
- Spaghetti cabling instead of parallel layouts could solve the problem, but the installation will be more time consuming and the result is questionable.
- Alien NEXT reduced cables. These are specially constructed cables which should reduce the effect of alien NEXT. I can't say what modifications they have made to reduce the alien NEXT, but a reduction of 20dB does not seem to be possible.

The last solution is to go away from unshielded cables and use shielded cables:

4.3.3 Shielded cables

The fact that central Europe is using shielded cables could be an advantage for 10 Gig-Ethernet. The IBM Type 1 cable was and is shielded and also the 10Gig CX4 standard uses a shielded cable to transmit the signal over 25m.

The NEXT and alien NEXT of an SF-UTP Cat. 5e cable (braid and foil shielded cable) is shown below:



In this case we can see two things: First the NEXT is very good for a normal Cat. 5e cable, but the cable is not optimized for frequencies higher than 250MHz. In this case even a NEXT cancellation of 40dB will not bring the necessary ACR value.

On the other hand we see that the alien NEXT is on a level of more than 80dB. Comparing this value with the 35dB of a UTP cable we see that the shielding of the cable brings a gain of 45 dB. Therefore alien NEXT compensation and special cabling are not an issue with good shielded cables. However, shielded cables need a good earthing, which must be planned, taught and implemented, to ensure that no current loops occur on the data cable and destroy the good shielding of the cables.

5. Summary

IEEE and their members always try to use state of the art technologies and evolve them. Currently they are using Fast and Gigabit-Ethernet technologies to go to 10 Gigabit-Ethernet. The current measurements and discussions in IEEE and ISO go in the direction of Cat 6 cabling which will be extended up to 600MHz. Those which are installing shielded Cat. 6 cabling systems will have a big advantage to really run 10 Gigabit-Ethernet over the cabling systems rather than a fraction of that speed because the rest of the bandwidth is used for retransmitting files due to errors. Especially for 10 Gig-Ethernet UTP cabling systems are not reliable enough any more. Installing a shielded R&Mfreenet Cat. 6 cabling system will assure future functionality up to at least 10 Gigabit-Ethernet

6. Further information

For further R&M product and solution information please visit our website at www.rdm.com.