



NBASE-T Downshift: Optimization of 2.5Gb/s and 5Gb/s Ethernet Data Rates Over Cat5e/Cat6 Cabling

This white paper describes NBASE-T Downshift and how it optimizes transmission rates to get the highest achievable data rate out of installed Category 5e and Category 6 structured cabling. NBASE-T Downshift augments the IEEE 802.3Bz™ standard by providing a mechanism to select the best speed available within the constraints of the cabling.

VERSION 1.0
DECEMBER 2017

*Ken Naumann, Technical Leader (Cisco Systems, Inc.)
Amrik Bains, Principal Engineer (Cisco Systems, Inc.)
NBASE-T Alliance Technical Working Group Chair*

Overview

The NBASE-T Physical Layer Specification version 2.3, and IEEE 802.3bz™ standard specify operation of 2.5Gb/s and 5Gb/s Ethernet over Category 5e or Category 6 (or better) structured cabling¹, applying sophisticated signaling techniques to exceed the 1Gb/s rate for which this cabling was originally designed. In the vast majority of cases, the technology is designed to deliver multigigabit speeds without any special considerations from the user. In some cases, however, the quality of the cable, how cables are bundled together, or “alien noise” (interference from sources external to the cable) affects the ability of a network link to reach these higher speeds. When a network node (e.g., a PC or a Wi-Fi access point) turns on, it automatically negotiates with the node at the opposite end of the cable, (e.g., an Ethernet switch) the fastest rate at which both nodes can operate. If the cabling environment is impaired such that operating at the highest negotiated speed may generate packet errors beyond the limit specified in the standard, the nodes will suspend communication (Link Down) and then repeat the Auto-Negotiation process.

If the nodes once more negotiate operation at the same highest advertised speed, errors may once again force the link to be down. Historically, only manual configuration of the network to force a lower speed, less susceptible to errors, could get the link out of this persistent link flap state.

With NBASE-T Downshift, the best available data rate is selected for the connected link. Downshift works for 10Gb/s and other rates in addition to 2.5Gb/s and 5Gb/s. Downshift enables a link that fails to automatically (no manual re-configuration required via a management interface) renegotiate to a lower link speed (using the IEEE 802.3bz standard Auto-Negotiation mechanism). NBASE-T Downshift enhances the user experience by smoothly reducing link rates, avoiding persistent Link Down.

Cabling Installed Base

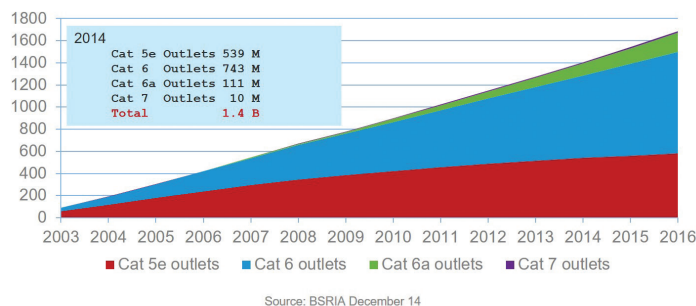
NBASE-T technology was designed to bring multigigabit speeds to the over 70 billion meters of installed structured cabling. The vast majority of this cabling is Category 5e and Category 6, and representing well over a billion outlets of infrastructure asset (Figure 1). NBASE-T development was initially driven by the backhaul requirements for 802.11ac Wave 2 wireless access points, to provide 2.5Gbps and 5Gbps service. The bandwidth

This white paper explains:

- Challenges of operating at various speeds over installed structured cabling due to frequency limitations and noise
- Types of noise conditions that can impact achievable link speed
- IEEE 802.3bz Auto-Negotiation
- How NBASE-T Downshift overcomes Auto-Negotiation’s limitations

requirements of NBASE-T were matched to the common capabilities of the infrastructure, and the coding and PHY technology was derived and improved from 10GBASE-T to provide noise immunity for the commonly found use cases. These are the same use cases and objectives that were later contributed to, and adopted by, the IEEE 802.3bz 2.5G/5GBASE-T standards project.

Figure 1: Installed Base of Cabling



The IEEE 802.3bz standard specifies operation at 2.5Gb/s and 5Gb/s over Cat5e, Cat6 (and better) cabling, provided the channel (i.e., the quality of signals over the cable) meets link-segment specifications. The 2.5Gb/s rate utilizes the maximum defined Cat5e bandwidth of 100MHz. The 5Gb/s rate exceeds the maximum bandwidth rating, utilizing up to 200MHz. Nonetheless, 5Gb/s data rates over Cat5e are achievable in most real life applications, including 100m cases in which alien crosstalk is not excessive. By comparison, the 10GBASE-T standard (IEEE 802.3 clause 55) specifies 10Gb/s rates over Cat6 (55m), Cat6A (100m) and requires 400MHz bandwidth.

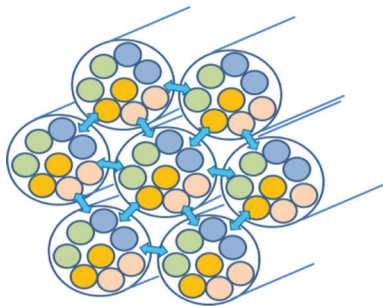
1. Class D and Class E channels specified in ISO/IEC 11801-1 are composed of Category 5e and Category 6 components, respectively. The ANSI/TIA-568 series of standards refers to both channels and components as “Category”.

Alien Noise

Reliable operation of a link at rates of 2.5Gb/s, 5Gb/s (and possibly 10Gb/s) over Cat5e channels is normally achievable, but some installations may have challenges. Links are potentially susceptible to noise from other cables in close contact. Unlike between the four twisted pairs of an Ethernet cable, the source of this noise is outside the control of the transceiver on the link, and is therefore known as “alien noise.”

Figure 2 shows an example of a challenging environment, depicting a worst case alien crosstalk noise configuration. The noise level is affected by tightness of the binding and the cable construction, both of which affect the separation of alien pairs. The worst-case happens when you have a dense bundle (pairs very close) running over a long distance. Cable-to-cable crosstalk may cause excessive bit errors, resulting in link failure even when the end devices support the rate.

Figure 2: Alien Cross-Talk in a 6-around-1 Bundle



The noise environment can change over time. An example is the daily cycle of office buildings or upgrade of the infrastructure over time. As the activity of the building varies, so does the alien crosstalk. A link that was previously running error free may see an increased number of bit errors, and possibly link failures. NBASE-T Downshift provides an elegant mitigation mechanism.

Another possible source of noise component is Radio Frequency Interference (RFI) (e.g. cell phone). RFI is known to affect 10GBASE-T, and a “fast retrain” mechanism was defined to address this (IEEE 802.3 Clause 55.4.2.2.2). Fast retrain avoids the need for normal Auto-Negotiation process. The link is briefly interrupted, but does not fail. If multiple fast retrain attempts fail in succession, the nodes use normal Auto-Negotiation and training process.

While bit errors can also be caused by other types of transient noise events (e.g., lightning, power surges, etc), they are short term and do not trigger Downshift.

With NBASE-T Downshift, the best available data rate is selected for the link. A link that fails due to a noisy channel re-negotiates to a lower speed without manual intervention.

IEEE 802.3bz Auto-Negotiation Detail

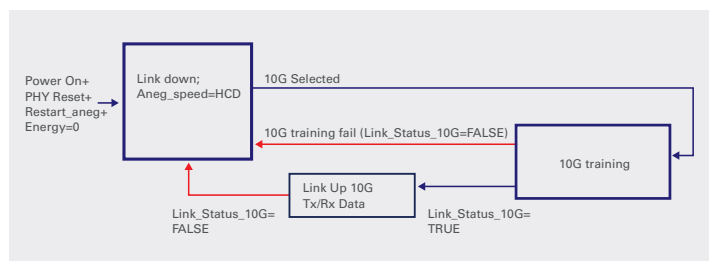
The 2.5GBASE-T and 5GBASE-T standards use IEEE 802.3bz Auto-Negotiation, specified in Clause 28, Clause 126.6, Annex 28B, Annex 28C, and Annex 28D.

As part of the link-training sequence, Auto-Negotiation supports interoperability with among nodes supporting networking standards ranging from 10Mb/s to 10Gb/s. Auto-Negotiation is performed as part of the initial set-up of the link and allows the nodes at each end to advertise their capabilities (e.g., speed, half or full duplex) and automatically select the operating mode for communication.

Auto-Negotiation selects the parameters that result in the fastest common rate and duplex type. At the end of link-training, the link is declared to be active (Link Up) if both end nodes observe a low rate of bit errors.

Figure 3 shows an example of 10GBASE-T Auto-Negotiation, started by power-up, operator commands or connection of a peer. The highest common speed advertised by both end devices is 10Gb/s. If the cabling infrastructure and noise environment can support the 10Gb/s speed (e.g. low bit error rate at that speed), Link Up is declared and data can be passed. If it cannot operate at 10Gb/s, then it returns to the Link Down state to try again. Without NBASE-T Downshift, the only way to resolve this scenario is to manually restrict the maximum speed.

Figure 3: Basic 10GBASE-T Auto-Negotiation



In our previous example where the 10Gb/s training failed, if 5Gb/s is advertised as the highest speed, the link may now be able to train and establish a clean link. With this link operating at 5Gb/s, consider the impact of additional links (know as aggressors) in the cable bundle starting operation. With more aggressors, noise on the victim link gets worse than it was during training, and could cause the victim link to fail.

IEEE 802.3bz can attempt fast retrain. If the noise change is relatively small, fast retrain can recover, or else Link Down is declared. Auto-Negotiation restarts and advertises 5Gb/s, as in the previous Auto-Negotiation cycle. Depending on the noise, renegotiation at 5Gb/s may or may not work. If not, Link Down is declared until the environment changes or re-configuration happens.

Downshift automates this process, addressing static and dynamic noise issues, and avoids the continuous Link Down described above.

NBASE-T Downshift

NBASE-T Downshift extends IEEE Auto-Negotiation and training capabilities to select the best speed supportable under both static and dynamic noise conditions. The Downshift is automatic, and deployment is plug-and-play. A network operator could also use Downshift events to deduce which links are subject to excessive noise.

Downshift uses of the following attributes:

Control:

- Downshift Enable/Disable
- Downshift Training Counter Threshold (dsh_thr): The number of training cycles before Downshift goes to the next speed

Status:

- Downshift Counter (dsh_cnt): The number of downshifts made to select the current speed
- Downshift History (Speed_dsh_status): Speeds tried while selecting the current speed

For simplicity, Figure 4 shows an example of 10Gb/s and 5Gb/s speed advertisement; however, the same steps are repeated for all advertised speeds from 100Mb/s to 10Gb/s.

If NBASE-T Downshift is enabled, the Auto-Negotiation and training procedure begins as discussed earlier, as specified by IEEE 802.3bz, but additionally resets Downshift Counter and sets

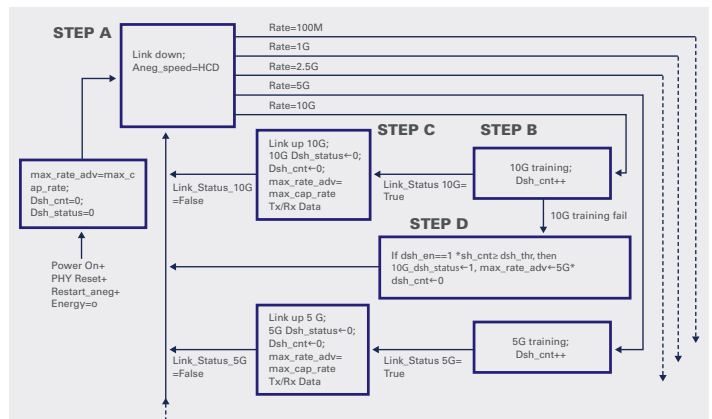
the Downshift status to false, as shown in Figure 4 (Step A). The main differences in using Downshift occur in the training phase.

In the training phase, the Downshift Counter is incremented with each training cycle, as shown in step B of Figure 4. When the training succeeds (at 10Gb/s in this case), Link Up is declared, and data transmission can begin (Step C).

However, if a training pass fails, it repeats (step B) and the Downshift Counter increments. If this happens more times than the Downshift Training Counter Threshold, training failure is declared (Step D) and Downshift takes the following automatic actions to establish a link:

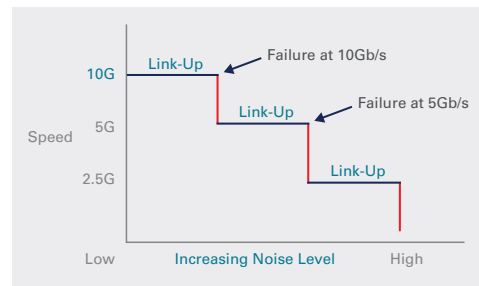
- Downshift_Status is set to True
- Downshift Counter is reset to Zero
- The current speed is removed from the advertisement list
- Auto-Negotiation is re-started (Step A)

Figure 4: Downshift State Machine and Procedure



NBASE-T Downshift reaction to increasing noise is shown in Figure 5. As noise increases and the link fails, Downshift selects the next lower speed. This process repeats, avoiding persistent Link Down.

Figure 5: Auto-Negotiation/Training Behavior with NBASE-T Downshift



SUMMARY

In summary, NBASE-T Downshift maximizes the sustainable data rate on a given link within the cable plant, and avoids persistent Link Down conditions by automatically:

1. Selecting the best speed possible for a given cable and environment
2. Changing speed when noise conditions change

NBASE-T Downshift is a key component enabling seamless deployment of millions of multigigabit Ethernet speeds, providing reliable connections over installed base structured cabling.



3855 SW 153RD DRIVE
BEAVERTON, OR 97003 USA
WWW.NBASET.ORG

The NBASE-T Alliance is an industry-wide cooperative effort of more than 45 companies focused on enabling the development and deployment of products that support 2.5G and 5GBASE-T Ethernet. To support early product development, the alliance developed and released the NBASE-T specification, which is compatible with the IEEE 802.3bz standard, and includes an added Downshift feature. The alliance focuses on publishing optimizations to the specification, facilitating interoperability and educating the market about the multiple applications of the NBASE-T technology.

© 2017 NBASE-T Alliance, Inc. All rights reserved. Unauthorized use strictly prohibited. NBASE-TTM, NBASE-T AllianceSM, and all NBASE-T logos shown in this document are registered or unregistered trademarks, service marks, and/or certification marks of NBASE-T Alliance, Inc. or its licensors in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited. Other names, marks, and logos shown in this document are trademarks and/or service marks of their respective owners. Specifications and content subject to change without notice.