

TONMEISTER ENGINEERING HANDBOOK

Ethernet Networking for High-Fidelity Audio

An Engineering Perspective on Switches, Cabling, Shielding and Electrical Noise

LOVE CABLE / TONMEISTER

Engineering first. Sound without interpretation.

CONTENTS

	Introduction — Digital Audio, Data Integrity and the Electrical Environment	4
1	Ethernet, Bit-Perfect Transmission and Why Music Does Not Play Directly From the Network	5
2	Electrical Noise, Galvanic Isolation and the Audio Environment	10
3	Managed and Unmanaged Ethernet Switches	17
4	Shielding, Grounding and Common-Mode Currents	25
5	System Design and Practical Network Topologies for High-Fidelity Audio	32
6	The Pursuit of Electrical Silence	37
7	Clocking, Jitter and Timing	41
8	Why Ethernet Became the Foundation of Modern Digital Communication	46
9	Power Supplies and the Electrical Environment	50
10	Measurement Before Assumption	56
11	Installation Practice and System Integration	61
12	Common Misconceptions	65
13	Engineering Recommendations	69

14	Engineering in Practice	73
-----------	-------------------------	----

REFERENCE

	References	77
--	------------	----

A	Key Engineering Standards Referenced in This Handbook	79
----------	---	----

B	Typical Engineering Measurements	80
----------	----------------------------------	----

INTRODUCTION

Digital Audio, Data Integrity and the Electrical Environment

The rapid adoption of network audio has transformed the way music is stored and reproduced. Today, music may originate from a local Network Attached Storage (NAS), a dedicated music server, or an online streaming service before traveling across an Ethernet network to a network streamer or integrated amplifier.

As network audio has become more widespread, discussions surrounding Ethernet switches, network cables and shielding have become increasingly polarized. Some argue that all Ethernet components sound identical because digital data is either correct or incorrect. Others maintain that every network component influences the final sound quality.

Both positions contain elements of truth, but neither tells the complete story.

To understand why, it is essential to distinguish between two fundamentally different aspects of Ethernet communication:

- The integrity of the digital data.
- The electrical environment in which that data is transported.

These are often confused, yet they are governed by entirely different physical principles.

The purpose of this document is to explain these principles from an engineering perspective, using established knowledge from digital communications, electromagnetic compatibility (EMC), transmission line theory and electronic circuit design. Throughout this document, the distinction between proven engineering principles and application-specific observations is carefully maintained.

The objective is not to support myths, nor to dismiss legitimate engineering concerns, but to explain how Ethernet functions within a modern high-fidelity audio system.

CHAPTER 1

Ethernet, Bit-Perfect Transmission and Why Music Does Not Play Directly From the Network

Digital Audio Is Data, Not a Continuous Signal

Unlike analogue audio, Ethernet does not transport a continuously varying electrical representation of music.

Instead, music is divided into digital information. This information is organized into packets that contain both the audio data and additional information required for reliable transmission across the network.

Each packet contains:

- Source and destination addresses.
- Timing and protocol information.
- User data, including the digital audio.
- Error detection information.

FIELD	PURPOSE
Source & Destination Address	Identifies the sender and receiver
Protocol Information	Identifies the type of data carried
Payload (Audio Data)	Contains the actual music data
Frame Check Sequence (CRC)	Detects transmission errors

The network itself has no knowledge that the data represents music. To the Ethernet infrastructure, a music file, a photograph, an email or a software update are simply collections of binary information.

The task of the network is therefore straightforward: deliver every packet accurately to its intended destination.

Ethernet Is Designed for Reliable Data Transport

Ethernet has evolved over several decades into one of the most reliable communication systems ever developed.

Modern Ethernet networks routinely transport enormous quantities of information in environments far more electrically demanding than a domestic listening room, including:

- Industrial automation.
- Medical imaging.
- Financial trading systems.
- Scientific laboratories.
- Data centres.
- Broadcasting facilities.

The technology achieves this reliability through multiple layers of protection.

At the physical layer, balanced differential signaling provides excellent immunity to external electromagnetic interference.

At the data link layer, every Ethernet frame includes a Frame Check Sequence (FCS) calculated using a Cyclic Redundancy Check (CRC) algorithm.

When a frame arrives, the receiving device performs the same mathematical calculation. If the calculated value differs from the transmitted value, the frame is immediately recognized as corrupted and discarded.

Corrupted information is therefore never accepted as valid data.

Depending on the higher-level communication protocol, the missing information is automatically retransmitted until a complete and correct copy has been received. This process is completely transparent to the user.

What Does “Bit Perfect” Really Mean?

The phrase bit perfect is frequently used within the audio industry, but it is often misunderstood.

A bit is the smallest unit of digital information and can have only one of two values: 0 or 1.

If every transmitted bit arrives exactly as intended, the transmission is considered bit perfect.

Unlike analogue signals, digital information does not gradually deteriorate with cable quality once transmission remains within specification.

Consider the hexadecimal number: 4F 2A B8 17

If a million Ethernet transmissions deliver exactly the same four bytes every time, the received information remains identical regardless of cable brand, cable price or switch manufacturer, provided the network operates correctly.

The Ethernet cable does not alter the numerical values. It either transports the data correctly or the communication protocol identifies an error.

There is no intermediate condition in which the musical information becomes subtly warmer, brighter or more detailed because individual bits have changed in a controlled manner. Such behavior is incompatible with the design of modern digital communication systems.

Why Music Does Not Play Directly From the Network

One common misconception is that the Ethernet connection directly determines the timing of digital-to-analogue conversion.

Signal path: Network Packets → Buffer (Memory) → Streamer's Playback Clock → DAC → Analog Output

In reality, the network and the audio playback process operate independently.

When a network streamer receives audio packets, it does not convert them immediately into music. Instead, the packets are temporarily stored in memory. This memory is commonly referred to as a buffer.

The buffer performs several important functions:

- It compensates for small variations in packet arrival time.
- It prevents interruptions caused by normal network activity.
- It separates network timing from audio playback timing.

Once sufficient data has accumulated inside the buffer, the streamer retrieves the information using its own internal clock. This clock, together with the digital-to-analogue converter (DAC), determines the timing of audio reproduction.

Consequently, minor variations in Ethernet packet arrival, often referred to as packet jitter, do not become sampling jitter at the DAC.

The network supplies data when required. The streamer determines when that data is converted into analogue music.

This distinction is fundamental to understanding network audio.

When Does Ethernet Actually Fail?

Under normal operating conditions, Ethernet either functions correctly or it does not. If network conditions become sufficiently poor, the result is not a gradual reduction in audio quality but a measurable communication failure.

Typical symptoms include:

- Playback interruptions.
- Buffer underruns.
- Audible dropouts.
- Loss of network connection.
- Complete interruption of the music stream.

These failures occur only when the streamer can no longer maintain sufficient buffered data for uninterrupted playback.

Before this point is reached, Ethernet continues to deliver perfectly accurate digital information.

For this reason, differences in perceived sound quality cannot reasonably be attributed to changes in the digital data itself when the network operates within specification.

The explanation, if one exists, must therefore be sought elsewhere.

Beyond the Bits

If Ethernet already delivers the music with complete digital accuracy, why do discussions concerning switches, cables and network hardware continue within the audio community?

The answer lies not in the digital information itself, but in the electrical behavior of the hardware that transports it. Every electronic device generates electrical noise. Network switches, computers, NAS devices, routers and streaming transports all contain high-speed digital circuitry, switching power supplies, oscillators and processors. These components generate high-frequency electrical energy that exists independently of the digital audio data.

Although Ethernet's galvanic isolation effectively prevents direct current from flowing between connected devices, certain forms of high-frequency common-mode energy can still couple between equipment through unavoidable parasitic capacitances and other electromagnetic mechanisms.

Unlike corrupted data, this electrical noise is not corrected by CRC checks or packet retransmission because it is not part of the digital information. It is simply unwanted electrical energy accompanying otherwise perfect data transmission.

Understanding how this electrical environment interacts with sensitive audio electronics forms the basis of the remainder of this paper.

CHAPTER 2

Electrical Noise, Galvanic Isolation and the Audio Environment

The previous chapter established that Ethernet is an exceptionally reliable method of transporting digital information. Modern Ethernet networks deliver data with extremely high integrity, and under normal operating conditions the audio data arriving at the network streamer is identical to the data originally transmitted.

This naturally raises an important question.

If the digital information is unchanged, why do engineers continue to investigate network switches, cable construction and electromagnetic compatibility within high-performance audio systems?

The answer lies not in the data itself, but in the electrical behavior of the equipment carrying that data.

Every Electronic Device Generates Electrical Noise

No electronic device is electrically silent.

Regardless of manufacturer or price, every digital product generates unwanted electrical energy as a by-product of normal operation.

Within a typical network switch this includes:

- High-speed Ethernet transceivers.
- Digital processors.
- Clock oscillators.
- Voltage regulators.
- Switching power supplies.
- Memory devices.
- Internal DC/DC converters.

Each operates by switching electrical currents at very high speeds. Every transition produces small electromagnetic fields that extend beyond the individual circuit components.

Most of this energy remains confined within the equipment and is harmless. However, some inevitably appears as electromagnetic interference (EMI) or radio-frequency interference (RFI). Good engineering is therefore not about eliminating electrical noise, which is impossible, but about controlling where it flows and preventing it from reaching circuits where it can become problematic.

Differential Signals and Common-Mode Noise

To understand why Ethernet is remarkably resistant to interference, it is necessary to distinguish between two fundamentally different types of electrical signals.

Differential Signals. Ethernet transmits information using balanced differential signaling. Each data signal travels over a twisted pair of conductors. One conductor carries the positive component of the signal while the other carries an equal but opposite component.

The receiving circuit measures only the voltage difference between the two conductors. Any electrical interference that affects both conductors equally is largely rejected.

This principle is known as common-mode rejection and is one of the primary reasons Ethernet performs reliably even in electrically noisy industrial environments.

Common-Mode Noise. Not all electrical energy appears as useful differential signals. Electrical noise may also exist equally on both conductors with respect to ground. This is known as common-mode noise.

Unlike the differential Ethernet signal, common-mode currents do not contribute to data transmission. Instead they represent unwanted electrical energy generated by the connected equipment itself or induced from the surrounding environment. Typical sources include:

- Switching power supplies.
- Computer processors.
- Wi-Fi transmitters.
- LED lighting.
- Solar inverters.
- Variable-speed motor drives.
- Mains-borne interference.
- Other nearby digital electronics.

Although Ethernet receivers reject most common-mode noise during data recovery, the noise itself still exists as electrical energy.

This distinction is extremely important. The Ethernet receiver may successfully ignore the noise while nearby analogue circuitry remains exposed to its electromagnetic effects.

ASPECT	DIFFERENTIAL SIGNAL (DATA)	COMMON-MODE NOISE
Behavior on the two wires	Equal and opposite	Identical on both wires
Carries audio data	Yes	No
Rejected by the receiver	No (wanted signal)	Largely yes
Typical source	Ethernet transmitter	Power supplies, clocks, EMI

Galvanic Isolation

One of Ethernet's most elegant design features is galvanic isolation.

Every standard Ethernet interface contains isolation transformers, commonly referred to as Ethernet magnetics. These transformers perform several essential functions:

- Electrical safety.
- Isolation between connected equipment.
- Impedance matching.
- Improved common-mode rejection.
- Reduction of conducted interference.

Most Ethernet interfaces provide isolation exceeding one thousand volts. Consequently, direct current cannot normally flow through an Ethernet cable from one device to another.

This isolation protects equipment against ground potential differences while allowing high-frequency differential data to pass efficiently.

For this reason Ethernet is fundamentally different from interfaces such as USB, where the ground conductors of both devices are directly connected.

Isolation Is Not Infinite

Although galvanic isolation is extremely effective, it is not perfect. No transformer can completely eliminate all forms of electrical coupling.

Every transformer contains unavoidable parasitic capacitance between its windings. At mains frequency this capacitance has almost no practical effect. At radio frequencies, however, even a very small capacitance can provide a path through which tiny high-frequency currents may pass.

The magnitude of these currents is generally very small, but modern digital electronics operate over extremely wide frequency ranges extending well into the hundreds of megahertz.

Consequently, high-frequency common-mode energy can still couple across the isolation barrier without compromising the integrity of the transmitted data.

This does not represent a failure of Ethernet. It is simply a consequence of fundamental electromagnetic physics.

Leakage Currents

Modern electronic equipment frequently contains switching power supplies. To comply with international electromagnetic compatibility regulations, these supplies often include suppression capacitors connected between the mains input and the equipment chassis or internal reference ground.

These components help reduce conducted electromagnetic emissions. They also create very small alternating leakage currents.

The currents involved are well within international safety limits and are entirely normal. Nevertheless, when multiple devices are interconnected, these tiny currents can interact with one another through available conductive or capacitive paths. Examples include:

- Protective earth conductors.
- Cable shields.
- Chassis connections.
- Parasitic capacitances.
- External grounding systems.

The resulting current paths are usually insignificant for ordinary computing applications. High-resolution analogue circuitry, however, often operates at signal levels measured in microvolts. At these levels, careful control of unwanted electrical currents becomes increasingly important.

Why Noise and Data Are Different

A useful analogy is to imagine a motorway carrying both vehicles and wind. The vehicles represent the Ethernet packets. The wind represents the surrounding electrical environment.

The motorway successfully delivers every vehicle to its destination. Nothing has altered the vehicles themselves. However, the strength and direction of the wind may vary considerably.

Similarly, Ethernet delivers the digital information correctly while the accompanying electrical environment may differ depending on the design of the connected equipment.

This distinction explains why engineers studying electromagnetic compatibility focus not only on signal integrity but also on current paths, shielding, grounding and radiated emissions.

The digital information remains unchanged. The surrounding electrical environment may not.

Why This Matters in Audio Equipment

A network streamer consists of far more than an Ethernet receiver. Internally it contains:

- Voltage regulators.
- Precision oscillators.
- Clock distribution circuitry.
- Digital signal processors.
- Digital-to-analogue converters.
- Low-noise analogue output stages.

Many of these circuits operate with extremely small voltage levels. The analogue output stage of a high-performance DAC may resolve signals approaching the thermal noise limits of modern electronic components.

Designers therefore devote considerable effort to:

- Power supply filtering.
- Ground plane design.
- Clock isolation.
- PCB layout.
- Electromagnetic shielding.

- Separation of digital and analogue circuitry.

The objective is to prevent internally generated digital noise from degrading analogue performance. External electrical noise entering through interconnected equipment represents another aspect of the same engineering challenge.

Engineering Rather Than Speculation

It is important to distinguish carefully between established engineering principles and unsupported conclusions.

The following statements are well established:

- Ethernet delivers data with extremely high integrity.
- CRC detects transmission errors.
- Network streamers buffer incoming data.
- Ethernet interfaces provide galvanic isolation.
- High-frequency common-mode coupling can occur across isolation transformers.
- All electronic equipment generates electromagnetic noise.
- Good grounding and EMC design improve electronic performance.

What cannot automatically be concluded is that every Ethernet cable, every switch or every network accessory will produce audible differences.

Whether unwanted electrical energy influences the final analogue output depends upon the complete system design, including the streamer, DAC, power supplies, grounding architecture and analogue circuitry. It is also worth stating plainly that this remains an area where controlled listening evidence is limited. While the physical mechanisms described in this chapter are real and measurable, well-controlled blind comparisons of Ethernet switches and cables have not consistently demonstrated audible differences once bit-perfect data delivery is confirmed.

This does not invalidate the underlying electrical engineering, but it does mean that claims of audibility should be treated as a separate, and so far less settled, question from the electrical mechanisms themselves. This distinction, between a measurable electrical mechanism and a demonstrated audible effect, is one this handbook returns to throughout, and it should be assumed to apply even where a later chapter does not restate it explicitly.

Consequently, Ethernet optimization should never be viewed as an attempt to improve the digital information. The digital information is already correct. The objective, where appropriate, is to minimize unnecessary electrical interference before it reaches the sensitive analogue electronics responsible for recreating the original musical performance.

CHAPTER 3

Managed and Unmanaged Ethernet Switches

The Ethernet switch is often regarded as the centre of a modern digital audio network. Whether music originates from a NAS, music server or online streaming service, virtually all data passes through one or more switches before reaching the network streamer.

As a result, Ethernet switches have become the subject of considerable debate within the audio community.

Some believe that every switch sounds different. Others argue that all compliant switches must produce identical results.

From an engineering perspective, neither position is entirely correct.

A properly functioning Ethernet switch will deliver the digital information with complete accuracy. However, switches are also electronic devices that generate electromagnetic energy, consume power, contain oscillators and processors, and interact electrically with the equipment connected to them.

To understand where differences may arise, it is first necessary to understand what a switch actually does.

What Is an Ethernet Switch?

An Ethernet switch is a specialized network device that forwards Ethernet frames between connected devices.

Unlike an older network hub, which transmitted every incoming frame to every connected port, a switch intelligently forwards each frame only to its intended destination. This significantly reduces unnecessary network traffic and improves both efficiency and reliability.

Modern switches perform millions of forwarding operations every second while introducing extremely little latency. For audio streaming, this latency is insignificant because the network streamer buffers incoming data before playback begins.

Switching Methods

Ethernet switches use one of several internal forwarding techniques. The most common in consumer and small-business Gigabit equipment is store-and-forward switching, though cut-through and fragment-free switching are also used in some managed and enterprise-oriented products, and the choice affects latency and error handling rather than data integrity for a correctly operating network.

In store-and-forward switching, rather than immediately forwarding an incoming frame, the switch first stores the complete frame in internal memory. Only after the entire frame has been received does the switch verify the Frame Check Sequence (FCS). If the frame contains an error, it is discarded. If the frame is valid, it is forwarded to the correct destination port.

This approach offers several important advantages:

- Corrupted frames are prevented from propagating through the network.
- Different Ethernet speeds can be interconnected efficiently.
- Temporary bursts of traffic can be buffered.
- Reliable operation is maintained under varying network conditions.

Cut-through switching, by contrast, begins forwarding a frame as soon as the destination address has been read, without waiting for the complete frame or verifying the FCS first. This reduces latency but means a corrupted frame may briefly propagate before being caught downstream. Fragment-free switching is a compromise between the two.

For audio playback, none of these differences matter in practice, since the streamer's playback buffer easily absorbs latency differences measured in microseconds, and any switch operating within specification still delivers frames with verified data integrity by the time they reach the application layer.

Buffer Memory

Every Ethernet switch contains memory used for temporary packet storage. This memory is often referred to as the packet buffer.

The buffer allows the switch to accommodate brief periods of heavy network activity without immediately dropping packets.

Unlike the playback buffer inside a network streamer, the switch buffer has no relationship to audio timing. Its purpose is simply to regulate data flow through the network.

Whether the switch temporarily stores a frame for a few microseconds has no influence on the timing of digital-to-analogue conversion. The streamer receives correctly ordered data regardless.

Packet arrival variation is an inherent characteristic of packet-switched networks and is anticipated by Ethernet protocol design. Modern playback devices employ buffering specifically to absorb this variation before audio samples are reconstructed. Packet timing should therefore not be confused with the sampling-clock stability that governs digital-to-analogue conversion inside the DAC, since these represent entirely different engineering domains, a distinction explored fully in Chapter 7.

MAC Address Learning

Each Ethernet device possesses a unique hardware identifier known as a Media Access Control (MAC) address.

As frames pass through the switch, it automatically learns which MAC addresses are connected to which ports. This information is stored in an internal forwarding table.

Once learned, subsequent frames are sent only to the appropriate destination. The switch therefore reduces unnecessary traffic while improving network efficiency.

This learning process occurs continuously and automatically. No user intervention is required.

The Internal Architecture of a Switch

Although Ethernet switches differ widely in cost and features, their basic internal architecture is remarkably similar. Most contain the following functional blocks:

- Ethernet PHY transceivers.
- Switching ASIC (Application Specific Integrated Circuit).
- Packet memory.
- Clock oscillators.
- Voltage regulation.
- Power supply.
- Management processor (where applicable).

The switching ASIC performs the vast majority of forwarding operations in dedicated hardware. This is one reason Ethernet switches can process extremely large volumes of traffic with minimal latency.

The audio data itself is not interpreted by the switch. It is simply forwarded according to standard Ethernet protocols.

Unmanaged Switches

An unmanaged switch performs only the essential Ethernet functions required for normal network communication. Characteristics include:

- Automatic configuration.
- Automatic speed negotiation.
- Automatic MAC address learning.
- Store-and-forward packet switching.
- No user configuration.

From an engineering perspective, unmanaged switches possess several advantages for domestic audio systems. Their simplicity generally results in:

- Lower processor activity.
- Reduced software complexity.
- Lower power consumption.
- Straightforward installation.
- Excellent reliability.

For the vast majority of home music systems, an unmanaged Gigabit Ethernet switch is entirely sufficient. Provided the switch complies with Ethernet standards and operates correctly, it will deliver the audio data without alteration.

Managed Switches

A managed switch performs all the functions of an unmanaged switch while adding configuration and monitoring capabilities. Typical features include:

- Virtual Local Area Networks (VLANs).
- Quality of Service (QoS).

- Port statistics.
- Traffic monitoring.
- Port mirroring.
- Remote management.
- Link aggregation.
- Spanning Tree Protocol.
- Access control.

These features are indispensable within enterprise environments where hundreds or thousands of network devices must operate efficiently.

In a typical domestic music system, however, many of these functions remain unused. Their presence neither improves nor degrades the accuracy of Ethernet data transmission. They simply provide greater administrative control over the network.

Does a Managed Switch Sound Better?

From the perspective of digital communications, the answer is no.

A managed switch does not transmit more accurate Ethernet frames than an unmanaged switch. Both either deliver valid Ethernet data or they do not. The digital information arriving at the network streamer remains identical.

Where differences may exist is in the electrical design of the switch itself. A managed switch generally contains:

- A more powerful processor.
- Additional memory.
- More complex firmware.
- Greater internal activity.
- Increased power consumption.

These factors may increase internally generated electromagnetic noise. However, this should not be interpreted as evidence that managed switches inherently perform worse.

A well-designed managed switch may exhibit lower conducted and radiated emissions than a poorly designed unmanaged switch. Overall engineering quality remains far more important than whether the switch is managed or unmanaged.

FEATURE	UNMANAGED SWITCH	MANAGED SWITCH
Configuration	None required	VLANs, QoS, monitoring, etc.
Complexity	Low	Higher
Power consumption	Lower	Generally higher
Ethernet data delivered	Identical	Identical
Best suited for	Home audio systems	Larger or business networks

Internal Clocking

One frequently misunderstood subject concerns switch clocks.

Every Ethernet switch contains one or more precision oscillators. These clocks regulate the internal operation of the switching electronics. They do not determine the sampling clock of the DAC.

Once Ethernet frames have been received and stored within the streamer's playback buffer, the DAC reconstructs the audio using its own local master clock. Consequently, improving the clock accuracy inside an Ethernet switch does not directly improve the sampling accuracy of the digital-to-analogue converter. The two timing domains are independent.

This is an important distinction because Ethernet is an asynchronous communication system. Unlike interfaces such as AES3 or S/PDIF, Ethernet does not transport the DAC sampling clock.

Switch Power Supplies

One aspect often overlooked is the switch power supply.

Nearly all modern Ethernet switches employ compact switch-mode power supplies. These offer excellent efficiency but inevitably generate high-frequency switching noise.

Well-designed switches minimize this through:

- Careful PCB layout.
- Effective grounding.
- Low-impedance power distribution.
- EMC filtering.

- Shielding.
- Compliance with international EMC regulations.

The quality of the power supply therefore influences the overall electrical behavior of the switch, although not the integrity of the Ethernet data itself.

Some enthusiasts choose to replace external switch-mode adapters with well-designed linear power supplies. Whether this provides measurable benefit depends entirely upon the electrical design of the switch and the connected audio equipment. No universal conclusion can be drawn without measurement.

Practical Considerations for Audio Systems

For most domestic installations, the network requirements of audio streaming are extremely modest. Even high-resolution PCM and DSD streams utilize only a small fraction of the available bandwidth of a Gigabit Ethernet network.

Consequently, the primary selection criteria for an audio switch should be:

- Reliable standards-compliant operation.
- Good EMC performance.
- Stable power supply design.
- High-quality construction.
- Appropriate port count.
- Quiet electrical behaviour.

Advanced management features should be selected only when the network itself requires them. Their presence does not improve the fidelity of the transmitted digital information.

ENGINEERING SUMMARY

A network switch has one primary responsibility: to forward Ethernet frames accurately and reliably. Both managed and unmanaged switches accomplish this task exceptionally well when operating within specification. The decision between them should therefore be based upon networking requirements rather than expectations of improved digital audio quality.

From an engineering perspective, greater attention should instead be directed toward overall electromagnetic compatibility, grounding strategy, power supply quality and system integration. These factors influence the electrical environment surrounding the audio system, while Ethernet itself continues to perform the task for which it was designed: transporting digital information with remarkable reliability and precision.

CHAPTER 4

Shielding, Grounding and Common-Mode Currents

Among all aspects of Ethernet networking, few subjects generate more misunderstanding than shielding and grounding.

Discussions often focus on whether a cable is shielded or unshielded, whether a shield should be connected at one end or both ends, or whether one cable category sounds better than another. These questions cannot be answered meaningfully without first understanding the electrical principles involved.

Compliance with a cable category specification establishes minimum electrical performance under defined test conditions, but it does not by itself describe the precision with which those requirements are achieved during manufacturing. Conductor concentricity, twist consistency, pair geometry, dielectric uniformity and connector termination all contribute to impedance stability and long-term electrical performance, and these are matters of manufacturing quality rather than category designation alone. This point becomes relevant below when comparing shielding approaches and cable grades.

Shielding is not intended to improve digital data. Grounding is not intended to improve bandwidth. Both exist to control the flow of unwanted electrical energy, and, as established in Chapter 2, whether that unwanted energy becomes audible in a given system is a separate and less settled question from whether it can be measured.

Only when this distinction is understood does it become possible to make informed decisions about cable construction and system design.

What Is Ground?

The term ground is widely used, yet it can refer to several different electrical concepts. In engineering, these must never be confused.

Protective Earth (PE). Protective Earth is the safety conductor connected to the mains installation. Its sole purpose is electrical safety.

Under normal operating conditions, no significant current should flow through the protective earth conductor except unavoidable leakage currents and fault currents.

Protective Earth is not an audio reference. It is not intended as a signal return path.

Signal Ground. Signal ground is the electrical reference used by electronic circuits. It provides the voltage reference from which analogue and digital signals are measured.

Signal ground exists entirely within the equipment and may or may not be connected directly to chassis or Protective Earth. Different manufacturers implement this relationship differently depending upon the equipment design.

Chassis Ground. The metal enclosure surrounding electronic equipment is commonly referred to as the chassis. The chassis acts as an electrical shield, helping to reduce electromagnetic emissions while protecting sensitive circuitry from external interference.

Depending upon the design, the chassis may be:

- Connected directly to Protective Earth.
- Connected through passive components.
- Connected to signal ground.
- Electrically isolated from both.

No single arrangement is universally correct. Each represents a different engineering compromise based upon safety, electromagnetic compatibility and circuit topology.

Why Two Grounds Are Never Exactly the Same

It is tempting to imagine that all grounded equipment shares exactly the same electrical potential. In reality, this is never completely true.

Every conductor possesses resistance, inductance and capacitance. Whenever current flows through these conductors, small voltage differences develop.

Even within the same equipment rack, two chassis may differ by a few millivolts at low frequencies and by considerably more at radio frequencies. These differences are entirely normal. They are a consequence of the finite impedance present in every conductor.

What Is a Current Loop?

Electric current always flows in a complete circuit. Whenever two pieces of equipment are connected by more than one conductive path, a loop is created.

For example: Protective Earth, equipment chassis, and cable shields together may form a closed conductive path through which unwanted current can circulate. Such current is commonly referred to as a ground loop current.

Ground loops are primarily a concern for low-frequency analogue systems, where they may introduce audible hum. At much higher frequencies, however, the same physical structure behaves differently. The conductors now possess significant inductance and capacitance. Current no longer follows only the path of lowest resistance. Instead, it follows the path of lowest impedance.

This distinction is fundamental to understanding electromagnetic compatibility.

Common-Mode Currents

Not all unwanted currents are caused by classical ground loops.

Modern electronic equipment generates substantial high-frequency energy through digital switching, switching power supplies and clock oscillators. Some of this energy appears as common-mode current.

Unlike the useful differential Ethernet signal, common-mode current flows in the same direction on both conductors with respect to the surrounding environment. It contributes nothing to data transmission. Instead, it represents unwanted electrical energy that seeks a return path through the interconnected system.

Potential return paths include:

- Cable shields.
- Chassis connections.
- Protective Earth.
- Parasitic capacitances.
- Nearby conductive structures.

The amount of current is typically very small. Nevertheless, modern analogue circuitry may operate with signal levels measured in microvolts. Consequently, even very small high-frequency currents deserve careful attention during equipment design.

Galvanic Isolation and Its Limits

As discussed previously, every standard Ethernet interface incorporates isolation transformers. These transformers prevent direct current from flowing between connected devices. This galvanic isolation is one of Ethernet's greatest strengths.

However, isolation is not the same as infinite electrical separation. Every transformer contains parasitic capacitance between its windings. At low frequencies, this capacitance presents an extremely high impedance. At radio frequencies, its impedance becomes much lower.

Consequently, very small high-frequency common-mode currents may couple across the transformer without affecting the integrity of the Ethernet data.

This phenomenon is entirely consistent with Maxwell's equations and with established electromagnetic theory. It does not indicate a failure of the Ethernet interface. Rather, it illustrates that perfect electrical isolation does not exist in physical systems.

What Does a Cable Shield Actually Do?

One of the most persistent misconceptions in audio is that a shield somehow improves the digital signal itself. It does not.

The Ethernet data is transmitted by the balanced twisted pairs. The shield performs a different function. It acts as a conductive barrier between the cable and its electromagnetic environment. Its primary purpose is to:

- Reduce electric-field coupling.
- Reduce radiated emissions.
- Improve electromagnetic compatibility.
- Provide a controlled path for induced high-frequency energy when correctly terminated.

A shield should therefore be regarded as part of the system's EMC strategy rather than as part of the data transmission path.

Foil Shields and Braided Shields

Not all cable shields are identical. Two constructions are commonly used.

Foil Shield. An aluminum foil provides almost complete coverage around the conductors. Advantages include excellent high-frequency screening, very high coverage, and lightweight construction.

Because the foil is fragile, it is normally accompanied by a drain wire that provides a reliable electrical connection during termination.

Braided Shield. A braided shield consists of many fine copper strands woven around the cable. Advantages include excellent mechanical durability, low resistance, good low-frequency performance, and greater flexibility.

Its coverage is typically lower than that of foil. For this reason, many high-performance Ethernet cables combine both methods.

PROPERTY	FOIL SHIELD	BRAIDED SHIELD
Coverage	Near-complete	Lower than foil
High-frequency screening	Excellent	Good
Low-frequency performance	Moderate	Good
Mechanical durability	Lower (requires drain wire)	High
Flexibility	Lower	Higher

Individual Pair Shielding

Many Cat6A and most Cat7 cables employ individual foil shields (and sometimes an additional braided shield) around each twisted pair.

This provides several engineering advantages. The individual shields reduce coupling between adjacent pairs while also improving rejection of external electric fields.

Combined with an overall braid, this creates a cable possessing excellent electromagnetic compatibility without altering the digital information being transmitted. The signal continues to travel entirely within the balanced twisted pairs.

One-End Versus Two-End Shield Termination

Perhaps no aspect of Ethernet cabling generates more discussion within the audio community than shield termination.

The internationally recognized Ethernet cabling standards specify that shielded structured cabling should normally be bonded at both ends. This provides a continuous low-impedance shield intended to maximize electromagnetic compatibility throughout the installation.

Within specialist fields such as instrumentation, measurement systems and certain high-performance audio installations, an alternative approach is sometimes employed. The shield is bonded at one end only.

The objective is not to improve the Ethernet signal. Instead, the intention is to interrupt low-frequency conductive shield currents while retaining much of the shield's ability to intercept external electric fields.

Because the shield no longer forms a continuous conductive path between the two pieces of equipment, the possibility of circulating shield currents is reduced. At the same time, it should be recognized that the shielding effectiveness at higher frequencies differs from that of a shield bonded at both ends, and small currents may still couple through unavoidable parasitic capacitances.

For this reason, single-end shield termination should be regarded as an application-specific EMC technique rather than a requirement of the Ethernet standard. Whether it offers an advantage depends upon the electrical design of the complete system, including the grounding architecture of the connected equipment.

Choosing the Appropriate Cable

For most domestic music systems, both high-quality UTP and properly constructed shielded Ethernet cables are capable of transporting digital audio with complete accuracy.

The choice should therefore be guided by the electromagnetic environment rather than by expectations of changing the digital data.

As a general engineering guideline:

- Typical residential installation with minimal electromagnetic interference: high-quality Cat5e or Cat6 UTP.
- Systems installed near multiple switch-mode power supplies, servers, wireless equipment or other significant sources of EMI: high-quality Cat6A or Cat7 S/FTP, provided the grounding strategy has been carefully considered.
- Professional or industrial environments requiring maximum electromagnetic compatibility: shielded structured cabling installed in accordance with recognized Ethernet cabling standards.

No single cable is universally superior. The objective is always the same: to preserve reliable Ethernet communication while controlling the unwanted electrical energy that inevitably accompanies every modern electronic system.

ENGINEERING SUMMARY

An Ethernet cable transports digital information through balanced differential signalling. Its shield, when present, does not carry the data. Instead, it forms part of the electromagnetic compatibility strategy of the installation.

Understanding the distinction between differential signals, common-mode currents and shield currents allows cable selection to be based upon engineering principles rather than assumption.

In high-fidelity audio systems, the role of the Ethernet cable is therefore twofold. It must transport data with complete integrity. It should also contribute to a stable and well-controlled electrical environment in which sensitive analogue circuitry can perform to its full potential.

CHAPTER 5

System Design and Practical Network Topologies for High-Fidelity Audio

Understanding how Ethernet functions is only the first step in designing a high-performance audio network. Equally important is understanding how the various components interact electrically.

A network consists of far more than cables and switches. Every connected device contributes to the overall electromagnetic environment in which the audio system operates.

The objective of good system design is therefore not simply to establish network communication, but to create an electrically quiet environment in which sensitive analogue circuitry can perform at its highest level.

A Typical Audio Network

A modern music playback system often consists of the following components, connected in sequence: Internet, router, Ethernet switch, then in parallel a NAS and a music server, both feeding a network streamer, which connects to a DAC, then a preamplifier, a power amplifier, and finally the loudspeakers.

Although music travels as Ethernet packets from the NAS or streaming service toward the streamer, unwanted electrical noise may be generated by every device connected to the network.

The engineering challenge is therefore to manage both the flow of data and the flow of electrical energy. These are related, but fundamentally different processes.

Signal Path: Internet → Router → Ethernet Switch → NAS / Music Server → Network Streamer → DAC → Preamplifier → Power Amplifier → Loudspeakers

Where Does Electrical Noise Originate?

Every digital device generates electromagnetic energy. Typical sources include:

- Central processing units (CPUs).
- Memory devices.

- Ethernet PHY transceivers.
- Voltage regulators.
- Switch-mode power supplies.
- Clock oscillators.
- DC/DC converters.
- Wireless transmitters.

These devices operate continuously, regardless of whether music is playing. The electrical noise they generate exists independently of the audio signal itself.

Consequently, reducing unnecessary electrical noise throughout the network can be regarded as good engineering practice, even though the digital music data remains entirely unchanged.

The Network Switch as an Electrical Junction

The Ethernet switch occupies a central position within most networks. Its primary task is packet forwarding.

Electrically, however, it also becomes the point where several independent devices converge. A switch may simultaneously connect a router, a NAS, a personal computer, a television, a home automation system, and one or more audio streamers.

Each connected device possesses its own power supply, internal clocks and electromagnetic emissions. Although Ethernet provides galvanic isolation at every port, the switch remains an important junction within the overall electromagnetic environment.

For this reason, locating the switch away from particularly noisy equipment may be advantageous from an EMC perspective.

Locating the NAS

The physical location of the Network Attached Storage device is generally of little importance from the standpoint of data integrity. A NAS located five meters away transports exactly the same digital information as one located twenty meters away.

However, NAS units often contain high-speed processors, large switching regulators, cooling fans, hard drives or SSD controllers, and multiple Ethernet interfaces. Collectively these components generate substantial high-frequency activity.

Where practical, many engineers prefer locating the NAS outside the immediate listening environment. This reduces acoustic noise from cooling systems while also physically separating one source of electromagnetic energy from the audio electronics.

Music Servers

Dedicated music servers occupy an intermediate position between a conventional computer and a simple NAS. In addition to storing music, they often perform library management, streaming service integration, metadata retrieval, digital signal processing, and user interface functions. Some also perform digital upsampling or room correction.

Because their internal processing activity is generally greater than that of a simple storage device, careful attention to power supply design, grounding and electromagnetic compatibility becomes increasingly important.

Again, the objective is not to improve the digital data. It is to minimize unnecessary electrical noise generated by the supporting electronics.

The Streamer

Among all network components, the streamer occupies a unique position. It forms the boundary between the digital network and the audio reproduction chain.

Within the streamer several critical processes occur: Ethernet reception, packet buffering, clock generation, audio data reconstruction, digital signal processing, and digital-to-analogue conversion, or digital output to an external DAC.

This is the first location within the signal chain where electrical behavior may directly influence the analogue performance of the system. Consequently, the streamer is often the component that benefits most from careful electromagnetic design.

High-quality streamers typically devote considerable attention to PCB layout, ground-plane architecture, power supply regulation, clock isolation, physical separation of digital and analogue circuitry, and mechanical shielding.

These design choices are largely invisible to the user, yet they often determine how effectively the streamer rejects internally and externally generated electrical noise.

Cable Routing

Ethernet is remarkably resistant to interference. Nevertheless, sensible cable routing remains good engineering practice. Where practical:

- Avoid prolonged parallel runs alongside mains cables.
- Maintain reasonable separation from large transformers.
- Avoid tightly bundling Ethernet cables with AC power wiring.
- Prevent unnecessary mechanical strain on connectors.
- Respect the manufacturer's minimum bend radius.

These recommendations are not unique to audio. They represent standard practice for reliable network installations.

The Importance of Power Supplies

Every network component requires electrical power. The quality of that power influences the behavior of the electronics contained within the device.

Most modern networking equipment employs highly efficient switch-mode power supplies. Properly designed switch-mode supplies are capable of excellent electrical performance while complying with stringent international EMC regulations.

Poorly designed supplies, however, may generate elevated conducted or radiated emissions. For this reason, attention should be directed toward the overall quality of the power supply rather than simply its operating principle.

Neither linear nor switch-mode designs are inherently superior. Both may be implemented exceptionally well or exceptionally poorly. Good engineering always takes precedence over generalisation.

Fibre Optics

Some audio enthusiasts employ optical fibre between sections of their network. Unlike copper Ethernet, fibre optic cable transports information using light rather than electrical conductors.

This offers one important advantage. No conductive electrical path exists between the connected devices. As a result, ground currents cannot flow through the fibre, conducted common-mode noise is interrupted, and complete galvanic isolation is achieved between the fibre interfaces.

The optical transceivers themselves remain electronic devices requiring power, but the fibre link eliminates any direct electrical connection between the two ends of the cable.

In electrically complex installations, fibre may therefore form an effective means of separating different parts of the network while maintaining complete data integrity.

Simplicity Is Often Good Engineering

As systems become more complicated, unintended electrical interactions also increase. Additional switches, media converters, power supplies and accessories should therefore be introduced only when they serve a clear engineering purpose.

Complexity alone is rarely beneficial. A carefully designed network consisting of a reliable router, one high-quality Gigabit Ethernet switch, well-constructed Ethernet cabling, and a thoughtfully engineered streamer will often perform every required function while maintaining an electrically well-controlled environment.

Good engineering values simplicity wherever simplicity fulfills the technical requirements.

ENGINEERING SUMMARY

An Ethernet network should be viewed as an interconnected electrical ecosystem rather than as a collection of isolated components. Every device contributes to the overall electromagnetic environment.

The task of the network remains unchanged: to deliver digital information accurately and reliably. The task of the system designer is different. It is to minimize unnecessary electrical noise while preserving the integrity of that communication.

CHAPTER 6

The Pursuit of Electrical Silence

One of the most persistent misconceptions in high-fidelity audio is that digital technology has eliminated the importance of electrical engineering.

Nothing could be further from reality.

Modern digital audio systems contain billions of switching events every second. Microprocessors, memory devices, Ethernet controllers, clock oscillators and power converters all operate simultaneously within an increasingly complex electromagnetic environment.

The challenge facing today's audio designer is therefore fundamentally different from that of the analogue era. The digital information itself is remarkably robust. The surrounding electrical environment is not.

Every Circuit Produces Both Signal and Noise

No practical electronic circuit generates only the desired signal.

Every resistor generates thermal noise. Every semiconductor exhibits junction capacitance. Every conductor possesses resistance, inductance and capacitance. Every clock oscillator produces harmonics extending far beyond its fundamental frequency. Every switching regulator generates transient currents.

These are not design flaws. They are unavoidable consequences of the laws of physics.

Engineering does not eliminate these phenomena. Engineering manages them.

Noise Cannot Be Destroyed

A useful engineering principle is that unwanted electrical energy cannot simply be made to disappear. It must always be reduced, absorbed, redirected, filtered, or prevented from coupling into sensitive circuitry.

This philosophy underlies virtually every discipline of electromagnetic compatibility engineering. The objective is not perfection. The objective is control.

The Importance of Current Paths

Engineers often speak about voltage. Electromagnetic compatibility engineers frequently think first about current. This distinction is significant.

Voltage differences become important only because current is capable of flowing. Whenever unwanted current finds a return path through an audio system, that path deserves careful examination.

Questions naturally follow. Where did the current originate? Why did it choose this path? Can the impedance of the path be changed? Can the current be diverted elsewhere? Can the loop area be reduced?

These questions form the basis of practical EMC design.

Ground Is a System, Not a Point

Many diagrams represent ground using a simple symbol. Reality is considerably more complex.

Ground should not be imagined as an infinitely large electrical reservoir having identical potential everywhere. Instead, every conductor possesses finite impedance. Consequently, every ground system exhibits small but measurable voltage differences whenever current flows.

This is true whether the conductor is a printed circuit board trace, a cable shield, a chassis connection, a protective earth conductor, or a copper busbar.

Good grounding therefore seeks to control where currents flow rather than assuming that all grounds remain electrically identical.

The Quietest System Is Not Always the Most Complex

As audio systems become increasingly elaborate, opportunities for unwanted electrical interaction also increase. Additional switches, power supplies, clocks, cables and accessories may each serve a legitimate purpose, but every component also contributes its own electrical behaviour.

The objective should therefore never be complexity for its own sake. Instead, each additional component should demonstrate a clear engineering benefit. Simplicity frequently represents the most elegant engineering solution.

Engineering Before Optimization

Within the audio community, considerable attention is often directed toward optimizing individual components. Yet optimization cannot compensate for poor system design.

A carefully engineered streamer connected through a poorly considered grounding arrangement may never achieve its full performance. Likewise, an expensive Ethernet cable cannot compensate for inadequate power supply design inside the playback equipment itself.

The greatest improvements usually result from addressing the system as a whole. Engineering proceeds from fundamentals toward refinement. Not the other way around.

The Role of Measurement

Good engineering begins with observation. Whenever possible, design decisions should be supported by measurement rather than assumption.

Modern instrumentation allows engineers to examine conducted emissions, radiated emissions, common-mode currents, ground potential differences, power supply noise, oscillator stability, network performance, and analogue output performance. A summary of the instruments most commonly used for this purpose appears in Appendix B.

Measurements cannot answer every question. They do, however, provide an objective foundation upon which further investigation can proceed.

Listening remains essential when evaluating a music reproduction system. Measurement remains essential when understanding why that system behaves as it does. The two approaches complement rather than oppose one another.

The Philosophy of Transparency

Within high-fidelity audio, transparency is often described as a subjective impression. From an engineering perspective, transparency may be understood differently.

A transparent system is one that contributes as little of its own behavior as reasonably possible. It neither alters the digital information nor unnecessarily contaminates the electrical environment in which that information is processed.

Transparency therefore represents the absence of avoidable error rather than the addition of desirable character. This distinction is subtle, yet profound.

The engineer does not seek to create a particular sound. The engineer seeks to remove mechanisms capable of altering the original recording.

Why This Matters

Every stage of the recording and reproduction chain represents an opportunity either to preserve information or to obscure it. Musicians create the performance. Recording engineers capture it. Mixing and mastering engineers shape the final production.

Playback equipment has only one responsibility: to reproduce that information as faithfully as possible.

The network, the cables, the switch, the streamer and the digital-to-analogue converter should therefore function as transparent engineering systems rather than as artistic interpreters. The closer each component approaches that objective, the closer the listener comes to experiencing the artistic intentions preserved within the recording.

ENGINEERING SUMMARY

The purpose of network optimization is not to alter music. It is not to create a warmer presentation, greater impact or a more impressive soundstage. Its purpose is considerably simpler: to reduce avoidable electrical interference while preserving the integrity of both the digital information and the analogue circuitry responsible for recreating it.

This philosophy lies at the heart of modern electromagnetic compatibility engineering. It also lies at the heart of truly transparent music reproduction.

CHAPTER 7

Clocking, Jitter and Timing

Few technical subjects within digital audio have generated more discussion than clocking and jitter.

The terms are frequently used to explain perceived differences between digital components, yet they are often applied without clearly distinguishing the various forms of timing that exist within a modern playback system.

This distinction is essential. Not every clock performs the same function. Not every form of jitter influences the reproduced music. Understanding where timing matters begins with understanding how digital audio systems are organised.

What Is a Clock?

Every digital electronic system requires a clock. A clock is simply a highly stable periodic electrical signal that coordinates the operation of digital circuits.

Without a clock, processors could not execute instructions, memory could not be read correctly and digital interfaces could not exchange information reliably.

A modern network audio system therefore contains many independent clocks. Typical examples include CPU clocks, memory clocks, Ethernet PHY clocks, USB interface clocks, network switch clocks, Digital Signal Processor (DSP) clocks, and Digital-to-Analogue Converter (DAC) master clocks.

Each performs a different task. The presence of multiple clocks should not be interpreted as multiple timing references for the music itself. Only one clock ultimately determines when each analogue sample leaves the DAC.

What Is Jitter?

Jitter is the deviation of a timing event from its ideal position in time.

An ideal clock produces transitions at perfectly uniform intervals. Real clocks do not. Every practical oscillator exhibits extremely small variations caused by thermal noise, power supply fluctuations, semiconductor noise and numerous other physical mechanisms. These variations are collectively referred to as jitter.

The magnitude of jitter is typically measured in picoseconds (10^{-12} seconds) or femtoseconds (10^{-15} seconds). Although these values are extraordinarily small, modern high-resolution digital systems are capable of measuring them with remarkable precision.

Not All Jitter Is Equal

The word jitter is often used as though it describes a single phenomenon. In reality, several different forms exist, among the most important being interface jitter, transport jitter, clock jitter, sampling jitter, and phase noise.

Each influences a different part of the system. Confusing them frequently leads to incorrect conclusions.

Ethernet Timing

Ethernet is an asynchronous communication system. This statement cannot be overemphasised.

The Ethernet network does not determine when the DAC converts digital samples into analogue voltages. Instead, Ethernet transfers packets whenever network bandwidth is available. Packets may arrive slightly earlier, slightly later, in groups, or with variable spacing.

Provided the streamer's playback buffer remains adequately filled, these variations have no influence upon the timing of audio conversion. The streamer simply stores the incoming data until it is required.

The Playback Buffer

The playback buffer forms one of the most important boundaries within the entire system. On one side lies the network. On the other lies the audio reproduction chain.

The network delivers packets whenever they arrive. The streamer retrieves data from the buffer according to its own internal timing. Consequently, network timing and audio timing become separate processes.

This is precisely why packet arrival jitter does not automatically become sampling jitter. The buffer decouples the two timing domains.

The DAC Master Clock

The master clock inside the DAC determines precisely when each digital sample is converted into an analogue voltage. This is the timing reference that ultimately influences waveform reconstruction.

For this reason, considerable engineering effort is devoted to low phase-noise oscillators, stable power supplies, careful PCB layout, short clock distribution paths, ground isolation, and thermal stability.

Unlike Ethernet transport timing, the quality of the DAC master clock directly affects analogue performance. It is therefore appropriate that high-performance DAC designers devote significant attention to this aspect of the system.

Phase Noise

No oscillator is perfectly stable. Rather than changing abruptly, its timing fluctuates continuously by extremely small amounts. In the frequency domain these fluctuations appear as sidebands surrounding the oscillator frequency. These sidebands are collectively referred to as phase noise.

Phase noise is particularly important because it represents the spectral purity of the clock rather than simply its average frequency accuracy. Two oscillators may both operate at exactly the same nominal frequency while exhibiting very different phase-noise characteristics.

For precision digital-to-analogue conversion, low phase noise is generally more important than absolute frequency accuracy alone.

Phase-Locked Loops

Many digital systems employ Phase-Locked Loops (PLLs). A PLL compares an incoming clock with a local oscillator and continuously adjusts the local oscillator until both remain synchronised.

PLLs are widely used within digital receivers, communication systems, clock recovery circuits, and frequency synthesisers.

The performance of a PLL depends upon numerous design parameters including loop bandwidth, noise characteristics, lock stability, and reference oscillator quality.

A well-designed PLL suppresses unwanted timing fluctuations while maintaining reliable synchronisation.

Why the Switch Clock Does Not Clock the DAC

One of the most common misconceptions within network audio is that improving the clock inside an Ethernet switch somehow improves the timing accuracy of the DAC.

This is not how Ethernet operates.

The switch requires accurate clocks to transmit Ethernet frames according to the IEEE specification. Once those frames have been received, checked for errors and stored within the playback buffer, their transport timing has fulfilled its purpose.

The DAC does not continue using the switch clock. Instead, it reconstructs the audio waveform according to its own local master clock. These are separate timing domains.

Consequently, improving the stability of the switch clock does not automatically improve the timing accuracy of the DAC.

If differences are observed between network components, they are more appropriately investigated from the perspective of electromagnetic compatibility, conducted emissions, common-mode currents and power supply behavior than from the assumption that the switch is directly clocking the audio converter.

When Timing Really Matters

The closer a clock lies to the point of digital-to-analogue conversion, the greater its potential influence upon the analogue waveform.

From an engineering perspective, timing importance generally increases in the following order: router CPU (negligible), network switch (negligible for audio timing), NAS processor (negligible for audio timing), Ethernet PHY (negligible for audio timing), streamer playback clock (important), and DAC master clock (critical).

This hierarchy reflects the functional architecture of modern asynchronous audio systems. It also explains why apparently similar clocks may differ greatly in their relevance to final audio performance.

CLOCK LOCATION	IMPORTANCE FOR AUDIO TIMING
Router CPU	Negligible
Network Switch	Negligible

CLOCK LOCATION	IMPORTANCE FOR AUDIO TIMING
NAS Processor	Negligible
Ethernet PHY	Negligible
Streamer Playback Clock	Important
DAC Master Clock	Critical

ENGINEERING SUMMARY

Every digital audio system contains numerous clocks, each serving a specific purpose. Confusing these independent timing domains has contributed significantly to misunderstanding within the audio community.

Ethernet timing governs packet transport. The playback buffer separates network timing from audio reconstruction. The DAC master clock determines when each sample becomes an analogue voltage.

Understanding this hierarchy allows engineering effort to be directed where it is most effective. Rather than assuming that every clock influences the reproduced music equally, good engineering identifies the clocks that genuinely determine analogue performance and focuses attention accordingly.

In doing so, the discussion shifts away from assumption and toward the physical mechanisms that can be measured, understood and optimised.

CHAPTER 8

Why Ethernet Became the Foundation of Modern Digital Communication

Today, Ethernet connects billions of devices across the world. It links personal computers, industrial control systems, hospitals, financial institutions, research laboratories, broadcasting facilities and data centers. It also forms the backbone of countless professional audio and video networks.

This widespread adoption was not accidental. Ethernet became the dominant communication technology because it combines reliability, scalability, interoperability and cost-effectiveness within a carefully defined engineering standard.

For digital audio, these same characteristics provide an exceptionally robust foundation for transporting music data. Understanding why Ethernet became the global standard helps explain why it is equally well suited to high-fidelity music reproduction.

The Importance of Standardization

One of Ethernet's greatest strengths is that it is defined by internationally recognized standards. Manufacturers throughout the world design their equipment according to the IEEE 802.3 family of specifications, alongside related structured cabling standards such as ISO/IEC 11801 and ANSI/TIA-568. A full list of the standards referenced throughout this handbook appears in Appendix A.

These standards define electrical characteristics, data encoding, timing requirements, connector interfaces, error detection, interoperability, cable performance, and electromagnetic compatibility.

As a result, equipment produced by different manufacturers is expected to communicate reliably without requiring proprietary protocols or specialized hardware. This interoperability has been one of the principal reasons for Ethernet's worldwide success.

Reliability by Design

Unlike many consumer technologies, Ethernet was developed with reliability as a primary objective.

Every transmitted frame includes error detection information. Every receiving device verifies the integrity of the frame before accepting it. Frames containing errors are rejected. Higher-level communication protocols determine whether retransmission is required.

This layered approach allows the network to maintain extremely high reliability, even in electrically demanding environments. Rather than assuming perfection, Ethernet is designed to detect and manage the occasional error that may occur in any practical communication system.

Differential Signaling

One of Ethernet's defining engineering features is the use of balanced differential signaling. Information is transmitted as the voltage difference between two conductors rather than as the voltage on a single conductor referenced to ground.

This provides several important advantages: excellent rejection of external interference, reduced electromagnetic emissions, improved noise immunity, greater transmission distances, and reliable operation in industrial environments.

Differential signaling has become one of the cornerstones of modern high-speed digital communication. It is also employed in numerous professional audio interfaces.

Galvanic Isolation

Every standard copper Ethernet interface incorporates isolation transformers. These transformers interrupt direct current flow between connected devices while allowing the differential data signal to pass.

This offers several significant benefits: improved electrical safety, reduction of ground-related problems, greater equipment compatibility, and enhanced system robustness.

Although, as discussed earlier, extremely small high-frequency currents may still couple through unavoidable parasitic capacitances, galvanic isolation remains one of Ethernet's most valuable engineering features.

Scalability

Ethernet has evolved continuously since its introduction. Network speeds have increased from megabits per second to hundreds of gigabits per second without abandoning the underlying communication principles.

This evolution has been possible because the standard was designed with scalability in mind.

For audio, this means that improvements in networking technology naturally become available without requiring the audio industry to develop entirely new transmission systems.

Economics Through Standardization

The widespread adoption of Ethernet has resulted in enormous global investment in switches, connectors, integrated circuits, measurement equipment, test procedures, and manufacturing techniques.

Every year, billions of Ethernet components are produced worldwide. This scale of production has driven remarkable advances in quality while simultaneously reducing manufacturing costs.

The audio industry benefits directly from these developments. Rather than creating proprietary communication systems, it can build upon decades of engineering refinement supported by one of the largest technology sectors in the world.

Ethernet in Professional Audio

Professional audio increasingly relies upon Ethernet-based technologies. Modern recording studios, broadcast facilities, concert venues and post-production environments routinely transport hundreds or even thousands of audio channels across standard network infrastructure.

This has been made possible through protocols specifically developed for professional audio transport, examples including Dante, Ravenna, AES67, and AVB/TSN.

Although these systems employ specialized protocols to meet demanding real-time requirements, they all depend upon the robustness and flexibility of Ethernet as the underlying transport technology.

This illustrates an important principle. Professional audio has not replaced Ethernet. It has extended Ethernet.

Lessons for High-Fidelity Audio

Domestic music systems operate under considerably less demanding conditions than professional production facilities. A single high-resolution stereo music stream represents only a tiny fraction of the available bandwidth of a modern Gigabit Ethernet network.

Consequently, network capacity is rarely the limiting factor. Instead, the engineering emphasis shifts toward reliable implementation, good electromagnetic compatibility, stable equipment design, appropriate grounding, and careful system integration.

These are precisely the themes explored throughout this handbook.

Engineering Is an Accumulation of Knowledge

One of the reasons Ethernet has remained successful for decades is that it has evolved through continuous engineering refinement. Every new generation has benefited from the lessons learned by the previous one.

Standards have improved. Semiconductors have advanced. Measurement techniques have become more sophisticated. Manufacturing tolerances have tightened.

The result is a communication system of exceptional maturity and reliability. High-fidelity audio benefits from this accumulated engineering knowledge.

Rather than viewing Ethernet as an external technology adapted for music reproduction, it is more accurate to regard it as a highly refined communication system whose proven capabilities provide an ideal foundation for transporting digital audio.

ENGINEERING SUMMARY

Ethernet did not become the world's dominant networking technology by chance. Its success is the result of decades of international collaboration, rigorous standardization and continual engineering refinement.

The audio industry has wisely adopted this mature technology rather than attempting to reinvent it.

For the listener, this provides an important reassurance. When correctly implemented, Ethernet offers an exceptionally reliable means of transporting digital music.

The engineering challenge is therefore not to improve Ethernet itself, but to integrate it into the playback system in a manner that preserves both communication integrity and a well-controlled electromagnetic environment.

CHAPTER 9

Power Supplies and the Electrical Environment

Every component within a digital audio system depends upon a power supply. Without a stable and well-regulated source of electrical energy, neither analogue nor digital circuits can perform as intended.

Despite this, power supplies are often discussed only in terms of voltage and current ratings. From an engineering perspective, their role extends considerably further.

A power supply also influences the electrical environment in which the circuitry operates. It contributes to conducted emissions, radiated emissions, common-mode currents, transient behaviour and overall electromagnetic compatibility.

Consequently, the quality of a power supply cannot be judged solely by its topology. It must be evaluated as part of the complete electrical system.

The Function of a Power Supply

The primary purpose of any power supply is straightforward. It converts electrical energy from one form into another while providing the voltage and current required by the connected equipment.

To accomplish this reliably, a well-designed power supply should provide stable output voltage, low output impedance across a wide frequency range, good load regulation, low residual noise, minimal conducted emissions, minimal radiated emissions, and reliable transient response.

Achieving all of these objectives simultaneously requires careful engineering. No single circuit topology guarantees success.

Linear Power Supplies

For many years, linear power supplies represented the standard solution for audio equipment. A conventional linear supply consists of a mains transformer, a rectifier, reservoir capacitors, voltage regulators, and additional filtering where required.

Because the mains frequency is relatively low, the circuit generates little high-frequency switching energy. This characteristic has contributed to the widespread perception that linear supplies are inherently quieter.

In practice, however, every design involves compromises. Linear supplies may exhibit transformer magnetic fields, rectifier switching transients, ripple voltage, thermal drift, lower conversion efficiency, and greater size and weight.

A well-designed linear supply can achieve excellent electrical performance, but this performance results from careful engineering rather than from the topology alone.

Switch-Mode Power Supplies

Switch-mode power supplies operate according to a different principle. Rather than regulating voltage through continuous dissipation, they convert electrical energy at high switching frequencies before filtering it to produce the required output.

Their principal advantages include high efficiency, compact size, low weight, excellent voltage regulation, wide input voltage range, and reduced heat generation.

These characteristics explain why switch-mode supplies are now used extensively in medical equipment, telecommunications, aerospace systems, industrial automation, computing, and professional broadcasting.

Modern switch-mode designs can achieve outstanding electrical performance while satisfying some of the world’s most demanding electromagnetic compatibility standards.

PROPERTY	LINEAR SUPPLY	SWITCH-MODE SUPPLY
High-frequency switching noise	Minimal	Present (requires filtering)
Efficiency	Lower	Higher
Size and weight	Larger and heavier	Compact and lightweight
Can be engineered to be very quiet	Yes	Yes
Overall performance depends on	Design and components	Design and components

Why Switch-Mode Supplies Generate High-Frequency Energy

The rapid switching action that makes these supplies highly efficient also produces fast voltage and current transitions.

According to Fourier analysis, fast transitions contain energy distributed across a broad frequency spectrum. Without appropriate filtering, some of this energy may appear as conducted emissions, radiated emissions, common-mode noise, or differential-mode noise.

For this reason, every well-designed switch-mode supply incorporates carefully engineered filtering and shielding to minimize these effects.

The presence of switching frequencies should therefore not be regarded as a flaw. It is simply one aspect of the design that must be managed correctly.

Conducted and Radiated Emissions

Electrical noise generated within a power supply may propagate in two principal ways.

Conducted emissions travel through electrical conductors such as power cables, ground connections, and chassis structures.

Radiated emissions propagate through space as electromagnetic fields.

Good electromagnetic compatibility requires attention to both mechanisms. Reducing one while ignoring the other rarely produces an optimal result.

Common-Mode Chokes

Many modern power supplies employ common-mode chokes as part of their input filtering.

A common-mode choke presents a high impedance to unwanted common-mode currents while allowing the desired mains current to pass with minimal attenuation.

This contributes significantly to reducing conducted electromagnetic interference entering or leaving the equipment. The effectiveness of the choke depends upon its design, operating frequency and integration within the overall filter network.

Y-Capacitors

One component often overlooked outside engineering circles is the safety-rated Y-capacitor.

These capacitors are connected between the mains conductors and protective earth, or between primary and secondary circuits where permitted by safety standards. Their purpose is to provide a controlled path for high-frequency interference, thereby improving electromagnetic compatibility.

However, because capacitors pass alternating current, they also permit extremely small leakage currents to flow. These leakage currents are entirely normal and are tightly limited by international safety standards.

In interconnected audio systems, they may contribute to small potential differences between equipment chassis. Understanding this behavior is important when considering grounding and common-mode currents.

Leakage Current

The term leakage current often causes unnecessary concern. Within properly designed equipment, leakage current is neither a fault nor an indication of poor quality.

Instead, it is an unavoidable consequence of safety filtering, parasitic capacitance, electromagnetic compatibility measures, and transformer construction.

The magnitude of these currents is carefully regulated to ensure both safety and compliance with applicable standards. From an engineering perspective, the objective is not to eliminate leakage current entirely, which is generally impossible, but to understand and manage its effects within the complete system.

Linear Versus Switch-Mode

The question is often asked: which topology is superior?

From an engineering standpoint, the answer is that neither topology is inherently superior. A poorly designed linear supply may exhibit greater residual noise than an expertly designed switch-mode supply. Conversely, a poorly filtered switch-mode design may generate unnecessary electromagnetic emissions despite its excellent efficiency.

The determining factor is therefore not the operating principle itself. It is the quality of the engineering.

Good design, careful component selection, thoughtful PCB layout and rigorous electromagnetic compatibility practices influence performance far more than the simple distinction between linear and switch-mode operation.

Power Distribution Within Equipment

The external power supply represents only the first stage of electrical regulation. Inside modern audio equipment, the incoming voltage is often converted several more times before reaching individual circuits.

Separate regulators may be provided for digital processors, memory, clock oscillators, Ethernet interfaces, analogue output stages, and display electronics.

This local regulation reduces interactions between functional blocks and allows each circuit to operate under conditions appropriate to its requirements. Consequently, internal power distribution is frequently just as important as the external power supply itself.

Engineering as a Complete System

No power supply operates in isolation. Its electrical behavior interacts continuously with the mains installation, connected equipment, grounding architecture, shielding, cable routing, and load characteristics.

For this reason, evaluating a power supply independently of the system in which it operates provides only part of the engineering picture.

The objective is not to identify a universally superior topology. The objective is to design a complete electrical system in which every component contributes to stable, predictable and electromagnetically compatible operation.

ENGINEERING SUMMARY

Power supplies provide considerably more than electrical energy. They influence the electromagnetic environment throughout the entire playback system.

Both linear and switch-mode topologies are capable of excellent performance when properly engineered. Likewise, both may perform poorly if implemented without sufficient attention to electromagnetic compatibility, filtering and regulation.

For the audio engineer, the relevant question is therefore not whether a power supply is linear or switch-mode. The relevant question is whether it provides a stable electrical foundation while minimizing unnecessary interactions with the surrounding system.

Ultimately, the objective remains unchanged: to support transparent music reproduction by creating an electrically stable environment in which every stage of the playback chain can operate according to its intended design.

CHAPTER 10

Measurement Before Assumption

Engineering begins with observation. Before proposing a solution, the engineer first seeks to understand the behavior of the system. This principle has guided scientific and technical progress for centuries: observe, measure, analyze, and only then modify.

Within high-fidelity audio, however, discussions sometimes proceed in the opposite direction. Components are replaced before the underlying behavior of the system has been investigated.

Engineering encourages a different approach. Rather than asking, “What should I buy?”, the engineer first asks, “What is actually happening?”

Why We Measure

Every electronic system generates electrical phenomena that cannot be observed directly by the human senses: voltages, currents, electric fields, magnetic fields, noise spectra, clock stability, and signal integrity.

Without suitable instrumentation, these remain invisible. Measurement allows these phenomena to be examined objectively. It transforms assumption into knowledge.

Measurement Is Not the Enemy of Listening

Within audio, measurement and listening are sometimes presented as opposing philosophies. They are not. They answer different questions.

Listening evaluates the artistic outcome of the reproduction chain. Measurement evaluates the physical behavior of the engineering system. Neither replaces the other.

A violinist tunes by listening. The engineer designs the microphone by measuring. Both contribute to the final musical experience.

The Oscilloscope

Perhaps the most familiar engineering instrument is the oscilloscope. An oscilloscope displays voltage as a function of time.

It allows engineers to examine waveform shape, rise and fall times, overshoot, ringing, noise, timing relationships, and power supply ripple.

Modern digital oscilloscopes are capable of capturing events lasting only a few picoseconds. They reveal behavior entirely invisible to the unaided eye.

The Spectrum Analyzer

Where the oscilloscope examines signals in the time domain, the spectrum analyzer examines them in the frequency domain. This distinction is fundamental.

Many forms of electrical interference become far easier to understand when viewed as frequency components. Spectrum analyzers allow engineers to investigate switching frequencies, harmonics, spurious emissions, clock sidebands, broadband noise, and electromagnetic interference.

Rather than asking whether noise exists, engineers can determine at which frequencies, at what level, and under what operating conditions.

Vector Network Analyzers

High-speed digital communication depends upon controlled transmission lines. Vector Network Analyzers (VNAs) measure how signals propagate through cables, connectors and electronic circuits.

Typical measurements include characteristic impedance, return loss, insertion loss, phase response, and reflection coefficients.

These measurements are essential when developing Ethernet hardware operating at gigabit data rates and beyond.

Time-Domain Reflectometry

A transmission line should present a uniform impedance throughout its length. Time-Domain Reflectometry (TDR) provides a means of examining whether this condition has been achieved.

A fast electrical pulse is launched into the cable. Any impedance discontinuity reflects part of the pulse back toward the source.

By analyzing these reflections, engineers can identify connector imperfections, manufacturing defects, impedance variations, cable damage, and termination quality.

TDR has become one of the most valuable tools in transmission-line engineering.

Current Probes

Voltage measurements alone rarely provide the complete picture. Electromagnetic compatibility frequently concerns current flow.

Specialized current probes allow engineers to measure common-mode currents, leakage currents, ground currents, and conducted emissions.

Because unwanted electromagnetic interactions often arise from current rather than voltage alone, these instruments play an important role in EMC investigations.

Near-Field Probes

Not every source of interference travels through cables. Some is radiated directly from electronic circuits.

Near-field probes allow engineers to examine local electric and magnetic fields around printed circuit boards, voltage regulators, clock oscillators, Ethernet interfaces, and power supplies.

These probes help identify the physical origin of radiated emissions before more comprehensive EMC testing is undertaken.

EMC Receivers

Professional electromagnetic compatibility laboratories employ specialized receivers designed specifically for compliance testing.

These instruments measure conducted and radiated emissions according to internationally recognized standards. They determine whether equipment satisfies the regulatory limits established for commercial products.

Although such measurements are primarily associated with legal compliance, they also provide valuable engineering insight during product development.

A summary table of these instruments and their typical applications appears in Appendix B for quick reference.

Interpreting Measurements

Perhaps the most important aspect of engineering measurement is interpretation. Collecting data is only the beginning.

The engineer must also understand what was measured, under what conditions, with which instrument, according to which standard, and what uncertainties exist.

No measurement should ever be interpreted outside the context in which it was obtained. Sound engineering depends as much upon understanding limitations as upon obtaining results.

Correlation

A recurring challenge within audio engineering is establishing meaningful correlation. If an electrical characteristic changes, does that change influence measurable analogue performance? If analogue performance changes, does that influence perception?

These questions deserve careful investigation. Correlation should never be assumed. Neither should it be dismissed without evidence.

Engineering advances by establishing reliable relationships between physical behavior and observable outcomes.

Measurement Supports Engineering

Measurement does not exist to prove preconceived conclusions. Its purpose is to improve understanding.

Occasionally, measurements confirm an existing hypothesis. Sometimes they reveal an unexpected mechanism. Occasionally, they demonstrate that a widely held assumption is incorrect.

This is not failure. It is progress. Engineering evolves by replacing assumption with evidence.

ENGINEERING SUMMARY

The purpose of measurement is not to replace listening. Nor is it to reduce music to numbers.

Its purpose is to understand the physical behavior of the systems responsible for reproducing recorded performances.

Good engineering therefore begins with observation, continues through measurement and analysis, and only then proceeds toward optimization. In this way, every design decision becomes an informed response to an understood problem rather than an attempt to solve an assumed one.

This approach has guided generations of engineers. It remains just as relevant to modern high-fidelity audio as it is to every other field of technical endeavour.

CHAPTER 11

Installation Practice and System Integration

Even the most carefully engineered components cannot perform as intended if they are integrated poorly. Installation is therefore not the final step of a system. It is an integral part of the engineering process.

A thoughtfully designed installation preserves the intended electrical behavior of every component. Conversely, poor installation practice may introduce unnecessary electromagnetic interactions, increase susceptibility to interference or complicate future maintenance.

Good installation practice is not about perfection. It is about consistency, predictability and attention to detail.

Designing the System Before Connecting It

Experienced engineers rarely begin by connecting cables. Instead, they first consider the complete system.

Questions typically include: which devices will communicate, where will they be located, what sources of electromagnetic interference are present, which cables will share the same physical routes, how will power be distributed, and is future expansion anticipated.

Answering these questions before installation often prevents problems that would otherwise prove difficult to diagnose later. Good planning is frequently the most effective engineering tool available.

Equipment Placement

The physical arrangement of equipment influences both practical operation and electromagnetic compatibility. Where possible, keep network switches away from large mains transformers, avoid placing switch-mode power supplies directly adjacent to sensitive analogue equipment, allow adequate ventilation around networking equipment, and minimize unnecessary cable lengths while maintaining convenient service access.

Mechanical organization contributes directly to electrical organization. A well-ordered installation is often easier to understand, maintain and optimise.

Cable Separation

Although modern Ethernet is highly resistant to interference, sensible cable management remains good engineering practice. Where practical, separate Ethernet cables from mains power cables, avoid long parallel runs alongside high-current AC wiring, cross power and signal cables at approximately ninety degrees when necessary, avoid tightly bundling data and power cables together over long distances, and respect the manufacturer's specified bend radius.

These recommendations reduce opportunities for unwanted electromagnetic coupling while preserving the mechanical integrity of the installation.

Connector Quality

Every connection represents a potential point of mechanical and electrical variation. Reliable connectors should provide secure mechanical retention, consistent contact pressure, accurate termination, and stable long-term performance.

The objective is not decorative construction. It is reliable electrical performance over many years of operation. A correctly terminated connector is generally more important than an expensive connector that has been installed poorly.

Shield Continuity

When shielded Ethernet cabling is employed, the shielding system should be considered as part of the overall electromagnetic compatibility strategy.

Shield continuity should be intentional rather than accidental. Before selecting shielded cabling, it is appropriate to understand how the connected equipment implements grounding, whether shielded connectors are required, how the installation manages common-mode currents, and whether the intended shield termination strategy is compatible with the overall system design.

Shielding should never be regarded as an isolated property of the cable itself. It forms part of the complete electrical architecture.

Power Distribution

Stable power distribution contributes significantly to reliable system operation. Where possible, avoid overloading distribution circuits, maintain secure earth connections, use appropriately rated power distribution equipment, and avoid unnecessary adaptors and temporary wiring.

Electrical safety should always take precedence over convenience. A properly installed mains distribution system provides the foundation upon which the remainder of the installation depends.

Network Simplicity

Modern Ethernet is remarkably capable. Consequently, additional network hardware should be introduced only when it serves a clearly defined purpose.

Every additional device contributes another power supply, additional electronic circuitry, more cable connections, and increased installation complexity.

Simplicity frequently improves reliability. Where a straightforward network satisfies the operational requirements, unnecessary complexity rarely provides engineering benefit.

Documentation

Professional engineering projects are carefully documented. Even relatively simple domestic installations benefit from basic records.

Useful documentation may include network topology, cable routes, equipment locations, IP addressing where applicable, power distribution, and equipment serial numbers.

Accurate documentation simplifies maintenance, troubleshooting and future system expansion.

Troubleshooting

Should unexpected behavior occur, systematic investigation generally proves more effective than random component replacement.

A structured troubleshooting process may include the following steps: confirm the observed behaviour, isolate the affected equipment, verify power distribution, examine cable integrity, check network operation, measure where appropriate, introduce changes individually, and confirm the result before proceeding further.

This methodical approach reduces uncertainty while avoiding unnecessary expense.

Engineering Is Often Conservative

One characteristic shared by experienced engineers is restraint. Rather than changing several variables simultaneously, they typically modify one parameter at a time.

This approach allows cause and effect to remain identifiable. Although this process may require additional patience, it usually produces a clearer understanding of the system.

Engineering values knowledge over speed.

Long-Term Reliability

A successful installation should remain reliable for many years. Mechanical support, proper ventilation, secure connectors, thoughtful cable routing, and appropriate environmental conditions often contribute more to long-term performance than small differences between individual components.

Reliability is seldom dramatic. Its value becomes evident through years of trouble-free operation.

ENGINEERING SUMMARY

Installation is not merely the process of assembling equipment. It is the practical application of engineering principles developed throughout the design process.

A well-installed system promotes reliable communication, good electromagnetic compatibility and straightforward maintenance.

Perhaps more importantly, careful installation reduces unnecessary variables. When the electrical environment has been considered thoughtfully from the outset, every component is given the opportunity to operate according to its intended design.

Good installation therefore represents not the final stage of engineering, but its successful implementation.

CHAPTER 12

Common Misconceptions

Engineering progresses by questioning assumptions. Many widely held beliefs contain an element of truth, yet become misleading when applied without considering the broader engineering context.

The purpose of this chapter is not to dismiss commonly expressed views, but to examine them in light of established engineering principles.

In many cases, the answer is not simply “true” or “false.” It is “under these conditions.” Understanding those conditions is the essence of engineering.

Misconception 1: “Bits are bits, therefore Ethernet cables cannot matter.”

This statement contains an important truth. Ethernet is specifically designed to deliver digital information with a very high degree of integrity. Error detection, controlled signaling, differential transmission and well-defined physical layer specifications make Ethernet remarkably robust. If data are transmitted successfully, the received information is identical to the transmitted information.

However, the cable forms more than a digital transport path. It also exists within the electromagnetic environment of the playback system. Its construction influences characteristics such as characteristic impedance, shield effectiveness, pair balance, crosstalk, common-mode behaviour, return loss, and mechanical durability.

These properties do not change the binary value of the transmitted data when the link operates correctly. They may, however, influence how effectively the cable rejects or contains unwanted electromagnetic energy.

Data integrity and electromagnetic compatibility are related engineering subjects, but they are not identical.

Misconception 2: “All certified Ethernet cables perform identically.”

Compliance with an Ethernet category specification is an essential starting point. It is not the complete engineering story.

International standards define minimum electrical performance requirements. Manufacturers remain free to exceed those requirements through improved manufacturing tolerances, better conductor consistency, superior dielectric materials, more accurate pair geometry, higher-quality connector termination, and improved mechanical construction.

A well-engineered cable is distinguished not by marketing claims, but by consistent manufacturing, verified performance and long-term reliability.

Quality of construction matters. Quality of assembly matters. Quality of materials matters. Not because they alter the digital information, but because they contribute to predictable electrical behavior, mechanical integrity and long-term stability.

Engineering excellence is achieved through attention to detail.

Misconception 3: "Shielding is always better."

Shielding is one of the most misunderstood subjects within networking.

A correctly implemented shield may reduce susceptibility to external electromagnetic interference while also reducing emissions from the cable itself. However, shielding also introduces an additional conductive path between interconnected equipment. Whether this is advantageous depends upon the grounding architecture of the complete installation.

For this reason, shielding should never be considered independently of equipment design, grounding strategy, installation practice, and the electromagnetic environment.

The appropriate solution depends upon the system. Not every installation benefits from the same approach.

Misconception 4: "Linear power supplies are always quieter."

Linear power supplies do not generate high-frequency switching energy in the same manner as switch-mode designs. This observation is correct.

It does not automatically follow that every linear supply exhibits lower overall electrical noise. Transformer design, rectification, voltage regulation, filtering, PCB layout, grounding and component quality all contribute significantly to final performance.

Likewise, modern switch-mode power supplies can achieve excellent electromagnetic compatibility when properly engineered.

The topology establishes the operating principle. The implementation determines the outcome.

Misconception 5: “The Ethernet switch clocks the music.”

This misunderstanding arises because every digital device contains clocks. An Ethernet switch requires accurate timing to transmit and receive network frames according to the IEEE specification. The DAC requires accurate timing to reconstruct the analogue waveform. These are different timing domains.

Between them lies the playback buffer, which separates packet transport from sample conversion. Consequently, the switch does not directly determine the timing of digital-to-analogue conversion.

Engineering attention should therefore focus upon the clocks that actually govern audio reconstruction while also recognizing that the overall electrical environment remains important for optimum system performance.

Misconception 6: “Expensive automatically means better.”

Price alone provides little engineering information. A product should be evaluated according to electrical performance, mechanical construction, manufacturing consistency, compliance with relevant standards, long-term reliability, and suitability for the intended application.

Some products justify higher cost through careful engineering and superior manufacturing. Others derive their price primarily from appearance, exclusivity or marketing.

Engineering evaluates measurable design quality rather than price alone.

Misconception 7: “More complexity produces better results.”

Additional components inevitably increase system complexity. Each introduces another power supply, additional grounding paths, more cable connections, greater installation complexity, and more potential sources of electromagnetic interaction.

Complexity should therefore be introduced only where it serves a clearly defined engineering purpose. Simplicity often contributes to greater reliability and easier diagnosis.

The most elegant engineering solution is frequently the simplest one that satisfies the technical requirements.

Misconception 8: “If a system sounds different, the explanation must already be known.”

Engineering history teaches a different lesson. Many phenomena were observed experimentally long before their underlying mechanisms were fully understood.

This does not justify accepting every explanation uncritically. Neither does it justify dismissing every observation simply because the mechanism has not yet been completely characterised.

Good engineering remains open to investigation while insisting upon careful experimentation, repeatability and objective analysis. Curiosity and skepticism are complementary qualities. Together, they drive scientific and engineering progress.

ENGINEERING SUMMARY

Most misconceptions arise because they simplify complex subjects into absolute statements. Engineering rarely operates in absolutes. Instead, it considers context: operating conditions, system architecture, implementation quality, measurement, and verification.

This approach replaces certainty with understanding. Rather than asking whether a statement is universally true or universally false, the engineer asks a different question: under which conditions is this statement valid?

That question has guided every chapter of this handbook. It remains the foundation of thoughtful engineering and, ultimately, of transparent music reproduction.

CHAPTER 13

Engineering Recommendations

Throughout this handbook we have examined Ethernet networking from the perspective of engineering rather than assumption. We have explored digital communication, electromagnetic compatibility, grounding, shielding, power supplies, clocking, measurement and installation practice.

Each chapter has highlighted an important principle. No single component exists in isolation. Every component forms part of a complete electrical system.

For that reason, engineering recommendations should never be reduced to universal rules. Instead, they should be based upon an understanding of the application. The following recommendations summarize the principles discussed throughout this handbook.

Begin with the System

The first responsibility of an engineer is to understand the complete installation. Before recommending changes, establish the system architecture, the signal path, the network topology, the grounding arrangement, the power distribution, and the electromagnetic environment.

Only after these factors have been considered can meaningful engineering decisions be made.

Choose the Appropriate Cable Category

Select Ethernet cables that comply with recognized international standards. For most modern audio installations, Category 6A provides an excellent balance between electrical performance, mechanical robustness and long-term compatibility.

Where additional shielding or specialized installation requirements exist, Category 7 or Category 7A may also be appropriate, provided the complete grounding strategy has been considered.

The category should always suit the intended application rather than exceed it simply for the sake of specification.

Prioritize Construction Quality

The electrical category defines minimum performance requirements. The quality of construction determines how consistently those requirements are achieved.

Consider manufacturing consistency, conductor quality, pair geometry, dielectric stability, shield construction where applicable, connector termination, and mechanical durability.

Well-controlled manufacturing contributes to long-term electrical stability and reliable operation.

Consider Shielding as Part of the System

Shielding should never be selected independently of the installation. Before choosing shielded cabling, evaluate equipment grounding, chassis bonding, existing common-mode current paths, sources of electromagnetic interference, and installation routing.

Where shielded cabling is appropriate, shield termination should follow a clearly defined engineering strategy rather than occur accidentally through mixed components or inconsistent installation practice.

Keep Network Architecture Simple

Modern Ethernet networks are highly reliable. Additional switches, media converters or accessories should therefore be introduced only where they provide a clearly identifiable engineering benefit.

Unnecessary complexity increases installation effort, power consumption, grounding paths, maintenance requirements, and opportunities for unwanted electrical interaction.

Simplicity frequently represents the most robust engineering solution.

Maintain Good Installation Practice

Attention to installation details contributes directly to predictable electrical performance. Where practical, separate data and mains cabling, avoid unnecessary cable stress, respect specified bend radii, maintain secure connector termination, provide adequate equipment ventilation, and support cables mechanically where required.

Good installation practice preserves the electrical characteristics established during product design.

Evaluate Power Supplies Carefully

Avoid judging a power supply solely by its operating topology. Instead, consider voltage stability, output regulation, conducted emissions, radiated emissions, leakage currents, electromagnetic compatibility, and overall implementation quality.

Good engineering is determined by performance rather than by classification.

Measure Where Possible

Whenever practical, support engineering decisions with measurement. Measurements assist in identifying ground-related problems, electromagnetic interference, conducted emissions, cable faults, power supply behavior, and installation issues.

Objective evidence provides a stronger foundation than assumption.

Change One Variable at a Time

When investigating system behavior, avoid modifying multiple parameters simultaneously. Introducing one change at a time allows the effect of each modification to be evaluated more reliably.

This disciplined approach simplifies troubleshooting while improving confidence in the conclusions reached.

Remember the Objective

Every engineering decision should support the same objective: not to change the music, not to create a particular sonic character, but to preserve the integrity of the recorded information while providing the playback equipment with the most electrically stable operating environment reasonably achievable.

This objective remains consistent regardless of equipment manufacturer, system complexity or budget.

ENGINEERING SUMMARY

Engineering recommendations are not recipes. They are the practical application of established principles to individual systems.

No single solution is universally correct. The most appropriate solution depends upon the complete electrical environment, the equipment involved and the intended application.

For this reason, engineering begins not by selecting products, but by understanding the problem. Only then can appropriate technical solutions be proposed.

This systematic approach has guided every chapter of this handbook. It remains the most reliable path towards transparent, predictable and repeatable music reproduction.

CHAPTER 14

Engineering in Practice

Every engineering project begins with a question. Not an answer.

The purpose of engineering is not to confirm assumptions, but to understand reality. Only after understanding the problem can an appropriate solution be proposed.

This principle has guided every chapter of this handbook. It also guides the way we approach every music system.

Every System Is Unique

No two audio systems are identical. They differ in equipment, installation, grounding architecture, power distribution, electromagnetic environment, listening space, and user requirements.

For this reason, identical recommendations cannot reasonably be expected to produce identical outcomes. Engineering begins by recognizing these differences rather than ignoring them.

Products Are Not Solutions

A cable is a component. A switch is a component. A power supply is a component. None of these, by themselves, constitutes an engineering solution.

A solution is the result of understanding how all components interact within the complete system.

The objective is never to recommend more components than necessary. The objective is to recommend the most appropriate solution for the application.

Sometimes that solution involves replacing a component. Sometimes it involves improving the installation. Sometimes it requires a different grounding strategy. Sometimes the most appropriate recommendation is that no change is required.

Good engineering accepts all of these possibilities equally.

Asking the Right Questions

Before considering products, experienced engineers ask questions. Examples include: what equipment is being used, how is the system interconnected, which Ethernet category is installed, is shielding appropriate for this installation, how are power and signal cables routed, what is the grounding arrangement, are there identifiable sources of electromagnetic interference, has the observed behavior been measured, and which aspect of system performance is the user seeking to improve.

The answers to these questions define the engineering problem. Only then can meaningful recommendations follow.

Simplicity Before Complexity

Throughout this handbook one principle has appeared repeatedly: simplicity.

Additional components should not be introduced merely because they are available. Each additional device contributes another power supply, additional electrical connections, additional grounding paths, greater installation complexity, and more variables.

Engineering seeks the simplest solution that satisfies the technical requirements. This approach generally improves reliability, simplifies maintenance and reduces unnecessary uncertainty.

Consultation Rather Than Assumption

Every recommendation represents a technical judgement based upon the available information. For this reason, consultation forms an essential part of responsible engineering.

Understanding the installation often reveals opportunities that are not immediately apparent: a change in cable routing, a grounding improvement, a more appropriate cable category, a simplified network layout, or occasionally, confirmation that the existing installation already represents good engineering practice.

Consultation therefore serves a purpose beyond selecting products. It helps identify the most appropriate engineering response.

Engineering Is a Continuous Process

Technology evolves continuously: new semiconductor processes, improved measurement techniques, revised standards, and better manufacturing methods. Our understanding also evolves.

Engineering is therefore not a fixed body of knowledge. It is an ongoing process of observation, measurement and refinement.

Good engineers remain willing to question their own assumptions whenever new evidence becomes available. This willingness is not uncertainty. It is professional discipline.

Respecting the Recording

Every recording represents the combined work of musicians, producers and recording engineers. Their decisions define the artistic content.

The playback system has a different responsibility. Its purpose is not to reinterpret those decisions. Its purpose is to reproduce them as faithfully as reasonably possible.

The engineer therefore seeks to minimize unnecessary electrical interactions that may influence the behavior of the playback equipment. This philosophy has shaped every subject discussed throughout this handbook.

The Engineer's Responsibility

Engineering carries a responsibility beyond technical competence. It requires honesty, precision, curiosity, and humility.

An engineer should be prepared to say, "I do not yet know." Equally, an engineer should be prepared to say, "No change is necessary."

Neither response represents weakness. Both demonstrate confidence in evidence rather than assumption. Trust is built not by always providing an answer, but by providing the correct answer.

A Philosophy of Transparency

Transparency is often described as a characteristic of audio equipment. Perhaps it is more accurately described as a characteristic of engineering.

Transparent engineering avoids introducing unnecessary variables. It avoids unnecessary complexity. It respects established science while remaining open to continued investigation.

Above all, it recognizes that the purpose of the playback system is not to become part of the performance. Its purpose is to disappear.

When engineering succeeds, attention moves away from the equipment and returns to the music. There can be no higher objective.

Final Thoughts

The subjects explored throughout this handbook share a common theme: Ethernet, grounding, shielding, power supplies, clocking, measurement, and installation.

Each contributes to the creation of a stable and predictable electrical environment. No single topic provides every answer. Together, however, they illustrate an important principle.

Faithful music reproduction is achieved not through isolated optimization, but through thoughtful system engineering.

The purpose of engineering is therefore not to change the recording. It is to preserve it.

The recording engineer safeguards the musical performance during its creation. The playback engineer safeguards its integrity during reproduction. Both serve the same music. Both deserve the same respect.

Engineering is never truly finished. As understanding improves, so too do the tools, methods and technologies available. The products discussed today will inevitably evolve: new networking hardware will appear, measurement techniques will improve, and today's open questions will give way to new ones. The underlying principles, however, remain remarkably constant.

Transmission lines, differential signaling, grounding, shielding, electromagnetic compatibility, careful measurement and thoughtful system integration rest upon physical laws refined through decades of scientific research and practical engineering experience.

This handbook is offered not because it describes only today's products, but because it explains the enduring principles upon which tomorrow's technologies will continue to be built, and it stands on the accumulated knowledge of the many scientists, engineers and standards bodies whose work made modern digital audio and network communication possible.

And both remind us that the highest achievement in audio engineering is not to hear the equipment. It is to hear the artistry preserved within the recording.

REFERENCE

References

The following sources provide the established engineering and standards basis for the principles discussed throughout this handbook. Readers wishing to explore any subject in greater depth are encouraged to consult these directly.

Ott, H. W. *Noise Reduction Techniques in Electronic Systems*. John Wiley and Sons. The standard reference on grounding, shielding, common-mode noise and electromagnetic compatibility in electronic system design.

Paul, C. R. *Introduction to Electromagnetic Compatibility*. John Wiley and Sons. A core academic text covering transmission line theory, differential and common-mode coupling, and radiated and conducted emissions.

Johnson, H. and Graham, M. *High-Speed Digital Design: A Handbook of Black Magic*. Prentice Hall. Widely used reference on signal integrity, transmission lines, PCB layout and clock distribution in high-speed digital systems.

IEEE 802.3 Standard for Ethernet. Institute of Electrical and Electronics Engineers. The governing specification for Ethernet physical layer characteristics, frame structure, error detection and interoperability.

ISO/IEC 11801 Information Technology, Generic Cabling for Customer Premises. International Organization for Standardization / International Electrotechnical Commission. The international standard governing structured cabling performance, including shielded and unshielded twisted pair categories and shield bonding practice.

TIA/EIA-568 Commercial Building Telecommunications Cabling Standard. Telecommunications Industry Association. The North American counterpart to ISO/IEC 11801, defining Ethernet cabling categories and installation practice.

AES67-2018, AES Standard for Audio Applications of Networks: High-Performance Streaming Audio-over-IP Interoperability. Audio Engineering Society. The interoperability standard underlying professional audio-over-Ethernet transport, referenced in the discussion of professional audio networking.

CISPR 32 and CISPR 35, International Special Committee on Radio Interference. International Electrotechnical Commission. Governing standards for electromagnetic emissions and immunity of multimedia equipment, referenced in the discussion of conducted and radiated emissions testing.

IEC 60950-1 / IEC 62368-1, Safety of Information Technology and Audio/Video Equipment. International Electrotechnical Commission. Relevant to the discussion of protective earth, leakage current limits and Y-capacitor safety filtering.

IEEE 1588, Precision Time Protocol. Institute of Electrical and Electronics Engineers. Referenced in the discussion of network timing standards used in professional media networking.

Manufacturer and industry documentation for Dante (Audinate), Ravenna, and AVB/TSN (IEEE 802.1), consulted for the discussion of professional audio-over-Ethernet protocols in Chapter 8.

Note on empirical listening evidence: where this handbook distinguishes between established electrical mechanisms and their audibility in a properly functioning system, this reflects the current absence of a consistent body of controlled, blind listening evidence demonstrating audible differences between compliant Ethernet switches or cables once bit-perfect data delivery is confirmed. Readers interested in this specific question are encouraged to consult peer-reviewed audio engineering literature on blind listening methodology, such as papers published by the Audio Engineering Society, rather than manufacturer claims alone.

APPENDIX A

Key Engineering Standards Referenced in This Handbook

STANDARD	SUBJECT
IEEE 802.3	Ethernet physical and data link layers
ISO/IEC 11801	Generic structured cabling
ANSI/TIA-568	Commercial building telecommunications cabling
IEC 61000 Series	Electromagnetic compatibility
CISPR 32	Multimedia equipment emissions
AES67	Audio-over-IP interoperability
IEEE 1588	Precision Time Protocol

APPENDIX B

Typical Engineering Measurements

MEASUREMENT	INSTRUMENT
Characteristic Impedance	Vector Network Analyser
Return Loss	Vector Network Analyser
Near-End Crosstalk (NEXT)	Cable Certifier
Alien Crosstalk	Cable Certifier
Conducted Emissions	EMI Receiver
Radiated Emissions	EMC Chamber
Common-Mode Current	RF Current Probe
Leakage Current	Leakage Current Meter
Clock Phase Noise	Phase Noise Analyser
Jitter	High-Speed Oscilloscope or Jitter Analyser