Lessons Learned From Real World IP Deployments

Introduction

At a recent meeting of Artel engineers, we decided to collect some "lessons learned" stories. We figured we'd share these stories — focused around IP-based media workflows — in case they can help you to prevent unwanted downtime or avoid a painful failure.

This eBook contains some of the top lessons learned by Artel engineers in building IP-based solutions for our customers, in supporting deployments in IP environments, and in responding to evolving customer requirements as they extend their use of IP networks for media transport.

The Importance of PTP

The first of these lessons learned? Go with a PTP-aware managed switch. Realistically, this choice can solve or prevent a lot of problems for you, so we'll start by talking about what PTP-aware means and why it's so important.

As you know already, Precision Time Protocol (PTP) plays a central role in standards-based IP media transport, typically AES67 audio and SMPTE ST2110 video. PTP data is critical because it provides the timing information needed to synchronize devices to a shared clock across IP (packet-based) networks, including Ethernet switches and IP routers. When switches are not PTP-aware, they don't participate in the delay corrections called by the protocol, in turn compromising the accuracy of the downstream clocks.

If you work with a small and simple network with few endpoints, you may be able to get by without a PTP-aware managed switch. But as your network gets more complicated and the number of end points increase — as it inevitably will — you will benefit tremendously from the added traffic control/management and PTP functionality that comes with your PTP-aware switch.

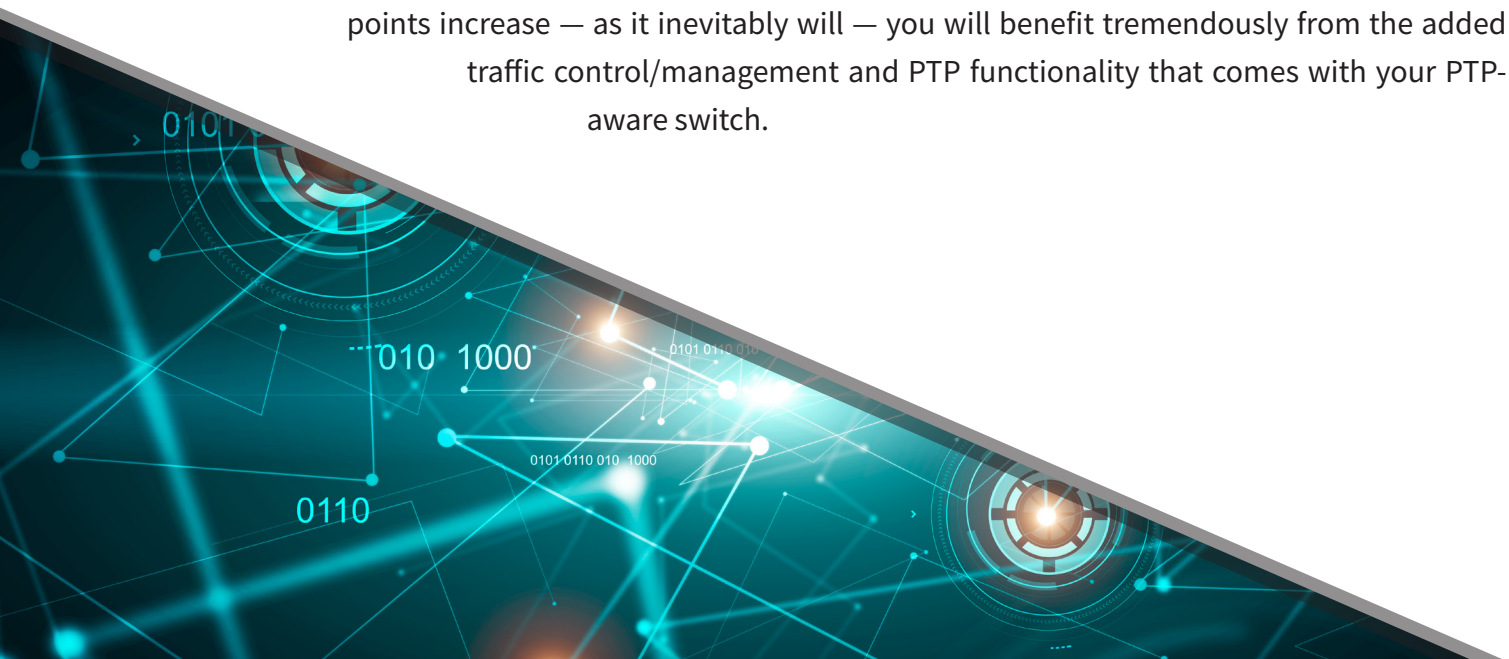

In handling PTP messages, a PTP-aware switch should give you the option of using transparent and boundary clocks. Both are used to distribute PTP messages properly to multiple devices, and you'll often find both at work across different layers of the PTP clock hierarchy when it's necessary to keep a large number of devices synchronized. Beyond correcting for the transit time of PTP messages across the switch, transparent clocks more or less act as though they are invisible. Boundary clocks, on the other hand, may be synchronized to a PTP timing server on one port and serve as a timing server on all the other ports supporting a large groups of client devices, effectively taking some of the load off of that timing server.

A managed switch will support useful functions such as VLANs that logically separate and prioritize various traffic flows, thereby protecting mission-critical media and data flows. You can even set up multiple boundaries by port, such as when you have two VLANs and two boundary clocks talking to the same timing server. Similarly, support of robust IGMP functionality allows the switch to receive multicast group memberships identifying the traffic that should be forwarded from a specific group to the client.

With quality of service (QoS) functionality, a switch can be configured to direct traffic (packets) with time-sensitive packets taking priority and, if specified, bandwidth constraints preventing unexpected traffic from interfering with the delivery of those packets.

One further thing to consider in choosing a switch is that you can use a common API (e.g., JSON) for control and provisioning. Because the functionality you'll want and need is already baked into the switch, you'll find that system presets and configurations go a long way in helping you to optimize audio, video, and data over IP even as your network grows in size and complexity.

If you're moving toward IP-based media workflows, you really can't afford not to consider a PTP-aware managed switch. Don't learn the hard way. Take our word for it! We're always happy to provide a demo or to chat with you about your switch requirements and options.

Optimal Handling of PTP

Now we've established how important it is to use a PTP-aware switch when implementing IP-based media workflows, we'll focus more on IEEE 1588 Precision Time Protocol (PTP) itself and how to handle PTP configuration and data correctly.

Accurate timing is critical to IP video/audio signal transport using SMPTE ST 2110 and AES67. Enabling synchronization of devices, including timing servers and clients, to one shared clock across packet-based networks, PTP provides a stable timebase for signals across a media operation or workflow and supports use of time code to align different media sources.

Set Consistent PTP Profiles and BMCA Settings

All devices that could potentially become the PTP timing server must execute a common Best Master Clock Algorithm (defined in IEEE 1588) to identify the best choice. Once a particular clock is selected, all other potential clocks receive timing from it. If something happens on the network that bumps that preferred or best clock, and it's no longer available, then the next best clock on the network automatically becomes the reference.

Every clock on the network, whether a boundary clock or end device (client), has a user-provisioned priority field that defines their place in the hierarchy. The lower the number, the higher the priority, with high priority indicating that a particular clock is the best one to use and low priority essentially saying, "Don't use this clock." If two or more clocks share the same high priority value, then other criteria related to clock quality are compared to determine which clock should be selected as the PTP reference.

To keep this model running smoothly, it's important to set consistent PTP profiles as you provision variables for each element participating on the PTP network. Because different industries have different PTP profiles for the different variables, you want to make sure that all the elements have the same profile — that they are all speaking the same language — and that the same (and correct) selection criteria are applied across all clocks so that every endpoint (client) will shift to the appropriate common source for PTP clock signals in the event of an interruption.

Leverage Your PTP Hierarchy Correctly

While both boundary and transparent clocks deliver the signal from the timing server to various clients, or end devices, the two generally serve different purposes.

Transparent clocks are easy to use. They simply allow PTP packets to flow through a device, adding a correction factor to each packet to account for the time it took the packet to flow through the device. With this information, the destination device can calculate the network delay and determine an accurate PTP time.

For small or mid-sized networks with only a handful of devices, a transparent clock is likely sufficient. For larger, more complex networks, a mix of transparent and boundary clocks helps to manage the larger PTP message load effectively. (More endpoints = more PTP messages.)

Because a boundary clock can act as a client device to an upstream PTP timing server and as a timing server for downstream clients/devices, it can be used to take some of the load off of that upstream timing server. The output of the upstream timing server thus can be shared with multiple downstream devices, preventing constant timing requests from various clients from causing an overload.

You can use boundary clocks to partition traffic — to create boundaries that not only keep the flows of PTP messages manageable, but also prevent an error on one endpoint from propagating and causing havoc across the network. It's an extreme case, but not impossible, and could result from a simple mistake that causes an overload on the timing server, or a more malicious denial of service attack. In either case, you can use boundary clocks to contain this kind of threat and add a layer of security to the network.

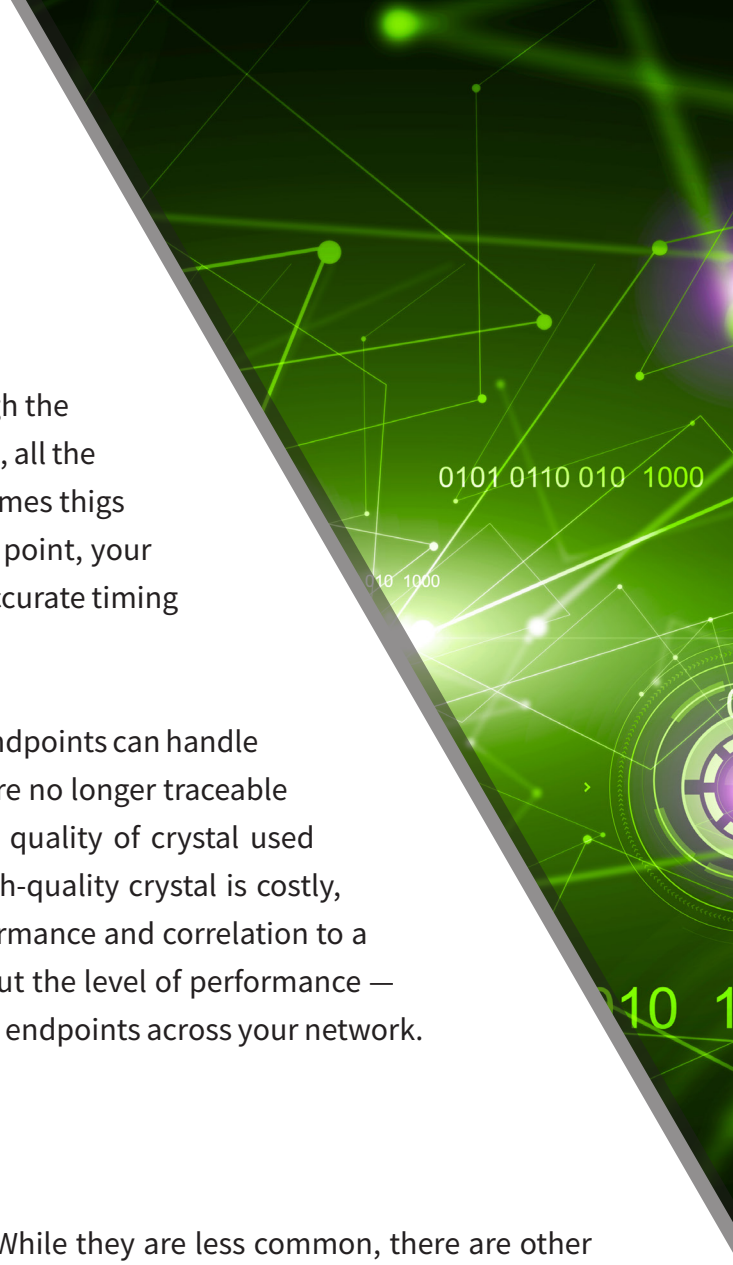

Correlate Clock Accuracy to the Applications

When everything is going well, you can trace timing through the PTP timing server, which presumably has a very good clock, all the way to the ultimate reference, an atomic clock. But sometimes thigs go wrong, and there is a break in the timing chain. At this point, your switch goes into holdover, capturing the latest and most accurate timing data and maintaining it as long as possible.

To prepare for this moment, you need to know what your endpoints can handle in terms of drift, which inevitably will occur when clocks are no longer traceable through a GPS or similar time source. Depending on the quality of crystal used in your clock, the rate of drift will be slower or faster. High-quality crystal is costly, and so vendors have to make a choice about cost of performance and correlation to a particular application. You too have to make a choice about the level of performance — and clock accuracy — required for your application and the endpoints across your network.

Other Timing Reference Options

PTP messages are just one source of timing information. While they are less common, there are other ways to communicate this data. Physical connections, for example, can carry timing in a different form, such as a waveform. Depending on the industry and application, you may find that you're working with endpoints or other devices that don't know about PTP. In this case, you want to be sure that you have the means to synchronize with other timing reference modes.

Full Support for IEEE 1588

It's always good to know that your systems support all the variables that are specified in key industry standards and specifications, and in the broadcast realm, IEEE 1588 is one of those specifications. You'll likely need to connecting end points that come from different vendors, each with their own interpretation of PTP and how to work with it. Full support for IEEE 1588 will ensure that you have the flexibility to accommodate deviations that might still be within the specification. With robust support for PTP, you can be confident that every aspect of PTP-based synchronization translates well across your network and devices.

Networking and Transport

Now we understand the importance of PTP-aware switches in an IP environment, as well as correct handling of Precision Time Protocol (PTP) data to ensure proper timing and synchronization, we look at networking and transport — the actual movement of video, audio, data, control, monitoring, intercom, and other signals over the IP network.

Know Your Needs

Before you move your operations into the IP realm, it pays to take a close look at your requirements in terms of moving media and data over the network. How many sources will the network need to support, and what kind of media do you plan to move and in what format? Will media be compressed or uncompressed, and how will that affect your bandwidth requirements?

Will you be using the network to distribute intercom, monitoring, and control signals? What kinds of external sources or signals will you need to accommodate? What about general data on the network? The better you're able to specify these details at the outset, the easier it will be for you to design or extend an IP network suited to your application and your operations.

Take a look at your existing facility and network too. What's already installed in the way of IP transport? What about dark fiber? Do you have or need connections to any other facilities? Bring these factors into your network planning discussions as well.

Although your network requirements will undoubtedly evolve over time, careful planning on the front end can help you make a smooth launch and adapt more readily as your needs change.

Manage Settings to Optimize Transport

Your IP network will support more than just media, so it makes sense to organize and prioritize transport for different payloads. One way to organize various types of traffic is to establish and use VLANs to create a logical separation between them. When you apply the appropriate QoS settings on your PTP-aware switch, you can make sure that the most essential and time-sensitive flows are delivered with the highest priority.

As you know, PTP data plays a critical role in enabling effective media-over-IP transport. In fact, as you configure your network, you'll want to set PTP as the highest-priority flow. Media would come next, and then control, intercom, and monitoring. With the right settings, you can be sure that even when traffic ramps up, your most important flows are protected.

Configure and Test With Care

With IP-based transport, you gain a great deal of flexibility in connecting different devices and managing various flows. But with that flexibility comes greater opportunity for disconnects. The details matter, and there are more of them when you move from the simpler world of SDI to the more complex IP domain. Without an understanding of your requirements (and a plan to address those requirements), you may have problems right from the start, or you might see smaller issues pop up — and sometimes turn into big issues.

In addition to planning carefully, you might consider staging a similar network prior to deployment. This is fairly common for a greenfield site, with the network and new components tested at a remote site for interoperability and performance. For a facility update or upgrade, where you're installing in place, you'll likely want to undergo a gradual deployment process that allows you to configure and test prior to going live. The time and effort you dedicate to this process can save you time, money, and headaches down the road.

Ensuring Interoperability

So we know the benefits of PTP-aware switches in an IP environment, correct handing of PTP data, and understand the actual movement of signals over the IP network. Next we shift gears just a bit to look more closely at interoperability — or, making sure everything works!

IP Is Different

Much of the broadcast industry comes from a world in which you can be confident that once you connect a cable, media or data will begin to flow through it. You don't need to worry about it. IP is different. While it introduces all kinds of power and flexibility, IP also has many more knobs that need turning to get all connected devices speaking the same language and working smoothly together. Incorrect settings or incompatible implementations of different broadcast standards can lead to little issues that turn into larger problems and even failure. This is why ensuring interoperability is among our key lessons learned.

Staging Your IP Solution

Taking the time to test a new IP deployment does take time and resources, but failure to do so can wind up being even more costly. Simply looking at a spec sheet isn't going to tell you how one product will behave as an endpoint on the network. It's not until you put devices onto the network in a dynamic setting and assess how they respond to different messages that you can start to identify any bugs that might compromise performance.

For a greenfield build, you might have the option of setting up and testing systems in place, but for migration to IP at an existing facility, you likely won't have that luxury. If money is no object, you can set up your own interop lab and ship everything there for testing, then ship it all to the facility for installation and commissioning.

A more economical approach is to work with your vendors or integrator to perform interoperability testing in a staging environment, prior to taking the IP environment live. With this approach, you give yourself time to "soak" the network, adjust it, and exercise the solution extensively before putting it to work for live day-to-day operations. You can also take advantage of your vendors' experience and expertise in specifying and supporting IP solutions for media delivery.

Get Vendor Support

In fact, as you go to select systems for your IP migration, it's worth examining different vendors' history of ensuring interoperability with other leading technologies and products. Do they have a history of working with other vendors and with standards bodies to make sure their products reflect a best practices approach? Are they active in alliances or associations that promote standards- and protocol-based interoperability and the industry's progress in implementing IP-based infrastructure and media workflows?

Working with the right partners and products, you can reduce the time and effort required to perform robust interoperability testing — and to prevent any unpleasant surprises once you take your new systems live.

Why Diagnostics Are Key to IP

For our final "lessons learned" from real-world deployments, we look at diagnostics. If you're deploying PTP-aware switches in an IP environment to support reliable transport and timing for video, audio, data, control, monitoring, intercom, and other signals, you'll want proper diagnostics during staging to get everything humming along nicely. Down the road, of course, you'll want diagnostic tools to accelerate troubleshooting, or to prevent small issues from becoming big problems.

Diagnostic tools will tell you if an endpoint or device is compliant with standards in the same way as other endpoints. They can log and capture activity, such as messaged delivered or dropped, across the network and segregate them to make analysis easier. For example, with some experience using diagnostic tools to look at information such as PTP messages, you'll be able to see when a particular endpoint isn't responding to delay requests from the timing server. You'll know that messages aren't reaching that endpoint, or that they're being dropped by the device itself.

With the right tools and experience using them, you'll gain the ability to recognize the "signatures" of different issues. Each signature points to a specific problem; it's a symptom that points directly to the underlying cause of that problem. You'll look at reports from devices and switches and think, "No wonder!" So, while IP is more complex than SDI, you will become increasingly familiar with these symptoms. You'll recognize patterns in diagnostic reports, you'll be able to see where ports and endpoints are behaving irregularly, and you'll know where to look to address the issue.

Oftentimes these issues stem from errors in provisioning the system. You might see an audio signal from a speaker joining a video flow and realize things weren't set up quite right. Maybe a port is getting flooded with messages, and packets are getting dropped. The problem may be simple, but finding the cause isn't always so. Diagnostics make it much easier to pinpoint and resolve these issues before they cause a major headache.

Fortunately, the industry's shift toward IP and embrace of IT-based systems brings a critical benefit: Networking is not new, and we already have access to sophisticated diagnostics and troubleshooting tools for IP networks. At the same time, vendors whose IP-based technologies target professional audio and video applications are growing more adept at building analytic and diagnostic tools into their products. By leveraging artificial intelligence (AI), technology suppliers are beginning to offer even more powerful tools capable of delivering even greater insights into system performance.

With tools such as these in hand, engineers can focus more time and energy on the production side of their work — creative aspects, such as where to situate microphones to get the right effect — rather than on IP infrastructure. They can focus more on the art and let diagnostics do the science necessary to support proactive maintenance and rapid troubleshooting of IP systems and endpoints.

Summary

We hope that you have found this eBook useful with your understanding of PTP and IP infrastructures and it will help you navigate the various pitfalls when moving your workflows to an IP based model. If you have questions about working with PTP in your own operations, just give us a shout. We'd love to help!