White paper
Cisco Public

Thinking About a Converged Interconnect Network?

Considerations for planning, designing, and building a network that preserves your investments

Modernizing the access network

A number of network milestones shaped today's hybrid-fiber coax (HFC), multi-service infrastructure: HFC itself, two-way activation, and digital triple-play service loads. Now a next step is at hand. This step will advance the last-mile plant to its ultimate state: Ethernet/IP. A body of work is underway to modernize the access network alongside the industrial momentum toward a distributed access architecture (DAA) and its related components, like Remote PHY devices (RPDs). It will run from point-to-point, analog and linear, to packet-switched, routed, and dynamic. The converged interconnect network (CIN) portends a massive rejuvenation of the last mile that specifically preserves the coaxial portion.

Contents

Evaluate your options

The CIN is designed for cable providers as a means to add appreciable capacity, agility, and scale to the existing HFC plant in a highly automated way. As a direct result, CINs enable service providers to strategically align with new, multi-Gigabit business opportunities, such as with adjacent, peered networks like mobile or 5G. Like any major network initiative, you need to establish long-term goals and plan, design, and build your CIN accordingly.

How to meet increasing demand

If you're a progressive cable provider, you're probably extending fiber deeper into neighborhood service groups, while virtualizing and distributing traditional headend functions. This work is underway as a means to reshape the economics of your service offerings while staying ahead of increasing demands for bandwidth. Because even though the global take-up rate for full Gigabit-grade connections (particularly in residential offerings) remains comparatively marginal, Gigabit speeds are expected in today's broadband speed wars.

It follows that future-minded cable providers are doing the research to examine available optical options and the engineering tradeoffs among them. This analysis usually happens shortly after running the numbers and shows that from a capacity and scale perspective HFC equipment will at some point fail to be economically viable.

The available optical options include a passive optical network (PON) such as gigabit passive optical networks (GPON) or Ethernet passive optical network (EPON). Another option is a CIN with Ethernet/IP. Or you can opt for full fiber, which may be any combination of optical equipment that is not HFC. The options are always good as is the ability to pivot, network-wise among the available options, depending on your engineering and marketplace needs.

Although a PON is a viable technology, it requires a complete rip and replacement of the infrastructure, which is one reason we believe a CIN is a more plausible short-, medium-, and long-term path for HFC-based service providers. A CIN is a better option because it contributes to network automation, agility, scalability, and security. Also, CINs deliberately preserve the massive investments providers have already made and continue to make into the HFC plant without the risk of vendor lock-in. In a CIN-powered DAA implementation, the HFC last mile will be poised to link with all IP/Ethernet-based networks, whether they are mobile/5G or PONs. They also will be able to link to whatever else comes along that is IP-based and needs fast, reliable, secure transit.

The application of a CIN gives a service provider the flexibility to apply all IP-based networks, known, and unknown. In that sense, moving to a CIN is part strategic positioning and part insurance policy.

Ethernet is a constant in the Internet and IP timeline. It surpassed Token Ring, Fiber Distributed Data Interface (FDDI), and numerous other LAN/WAN contenders. Our research indicates that Ethernet has won the wirelines, by growing ten times every seven years since the mid-1970s.

What is a converged interconnect network?

A CIN speaks a common language, end-to-end, and on IP-over-Ethernet. Because it also performs old and new protocol translation, it can act a single platform that is capable of hosting multiple networks, from HFC to DSL to PON. Multiple networks are an increasing reality for cable service providers because many of them grew their network portfolios through consolidation. A CIN that converges multiple networks avoids the cost and disruption associated with rip and replace strategies. For cable providers, the CIN is an economically rational and strategically maneuverable way to modernize up to and through the last-mile HFC infrastructure. You can make the move in parallel with existing or planned transitions to take fiber deeper using various DAA methodologies.

Conceptually, you can think of a CIN from the perspective of "before and after," or "now and next." In the before column, you have the existing HFC infrastructure. It's analog and point-to-point. HFC infrastructures were built before the availability of device-to-device knowledgeability and the automation that type of knowledgeability empowers. Without automation, most activities that touch the plant are manual and service provisioning can be costly. Transmission devices on the network aren't aware of each other, so mapping them, for purposes of fine-grain load balancing, isn't possible.

In the "after and next" column, the HFC infrastructure is modernized with a CIN. It's digital and switched and routed. Because it's a packet-based network, devices can dynamically communicate with one another. Additionally, because it's packet-switched and not point-to-point, you can apply techniques like statistical multiplexing. This technique is conceptually akin to video compression; it's used to more dynamically match bandwidth demand with bandwidth supply.

To visualize how this dynamic communication works, imagine a continuously running train. Each of its cars has a pre-assigned slot for a particular customer's data flow. In optical networking terms, the car is the spectral orthogonal frequency-division multiplexing (OFDM) slot. Whether or not it's carrying data, it's continuously linked into that moving train.

Clearly, inefficiencies exist in the before situation. The after scenario that features the packet switching and statistical multiplexing of a CIN removes those inefficiencies. It solves the no empty slots problem. As a direct result, operators are better able to right-size capacity, which yields three benefits:

- A massive increase in capacity and efficiency
- Preparedness for multi-service delivery (for example, peering with adjacent IP/Ethernet networks, such as mobile and 5G
- Preparedness to scale rapidly using existing infrastructure components with the proactive application of an automation suite that can grow over time. All in an open-sourced, non-proprietary, secure manner

What's converged about a CIN?

The converged aspect of a CIN is the traditional HFC access network, which acts as the last mile into the connectivity fabric of Ethernet/IP. Using a CIN makes it possible to unify the existing and increasingly disparate networks that cable operators maintain. This flexibility is desirable for cable service providers that seek faster and more modernized ways to link their long-coveted last mile infrastructure. They can connect home and business networks with services like mobile, 5G, and the PON family.

These cable operators are looking for faster, secure, and modernized ways to link with adjacent IP/Ethernet networks because the cable last-mile is an undisputed treasure. Only three wires enter a typical home: power, phone, and cable (coax). Network peering can lead to solid and lucrative business partnerships because most mobile network operators need more xHaul capacity. Most, if not all, public and private broadband networks are experiencing perpetual strain to keep up with demand.

Cable providers have excelled at capacity for more than six decades. A CIN offers yet another opportunity to lay down a technological milestone that will better position them with scale, bandwidth, and competitive and strategic maneuverability for decades to come.

Why a CIN represents a new inflection point in HFC evolution

Arguably, DOCSIS was the last major inflection point for cable broadband providers more than 20 years ago. It offered a way to reduce equipment costs through multivendor interoperability and build the basis of what is now widely known as broadband. DOCSIS underlies an entire industry sector that is valued in trillions of dollars of market capitalization. Over-the-top (OTT) and broadband native entities such as Netflix to Etsy wouldn't exist without DOCSIS.

To underscore the value of the connectivity that the cable industry provides, it's worth pointing out that broadband services reach more than 90 percent of U.S. homes. No other service has that kind of reach. Without the coaxial portion of HFC, we would still be figuring out how to eke more capacity out of twisted copper-pair telephone wires or how to get fiber directly to homes.

As momentous as it was and as reliable as it remains as a means to stay ahead of broadband demand, since that initial DOCSIS introduction, capacity gains have been largely incremental (see Figure 1).

Designing for capacity alone is no longer sufficient because capacity is only one component of scale. For progressive operators, Scale is the desired end game. To meet demand using today's processes and procedures, cable operators need to drastically increase the number of network devices they operate. Over time, and without advancements in automation, this approach will have economic implications. Too many devices, too many processes, too many people, all trying to provide the same service at greater scale without greater revenue. The business equation doesn't make sense.

A CIN represents a network inflection point because it changes the dynamics of an operator's addressable market. It opens up considerable adjacent network opportunities, enables scale through automation, and simplifies processesA CIN allows for massive device increases without the overhead associated with maintaining outdated analog equipment.

CINs are happening at an advantageous moment in time. They coincide with the coming of age of interconnects and the maturation of traditional HFCs. As part of the overall modernization of the plant, using DAA, service providers are positioned to upgrade in parallel to a CIN. The CIN enables operators to determine how they can evolve. Choosing how to build a CIN and on what technology can dictate their future.

Figure 1. DOCSIS was the last major capacity inflection point; CIN is the next.

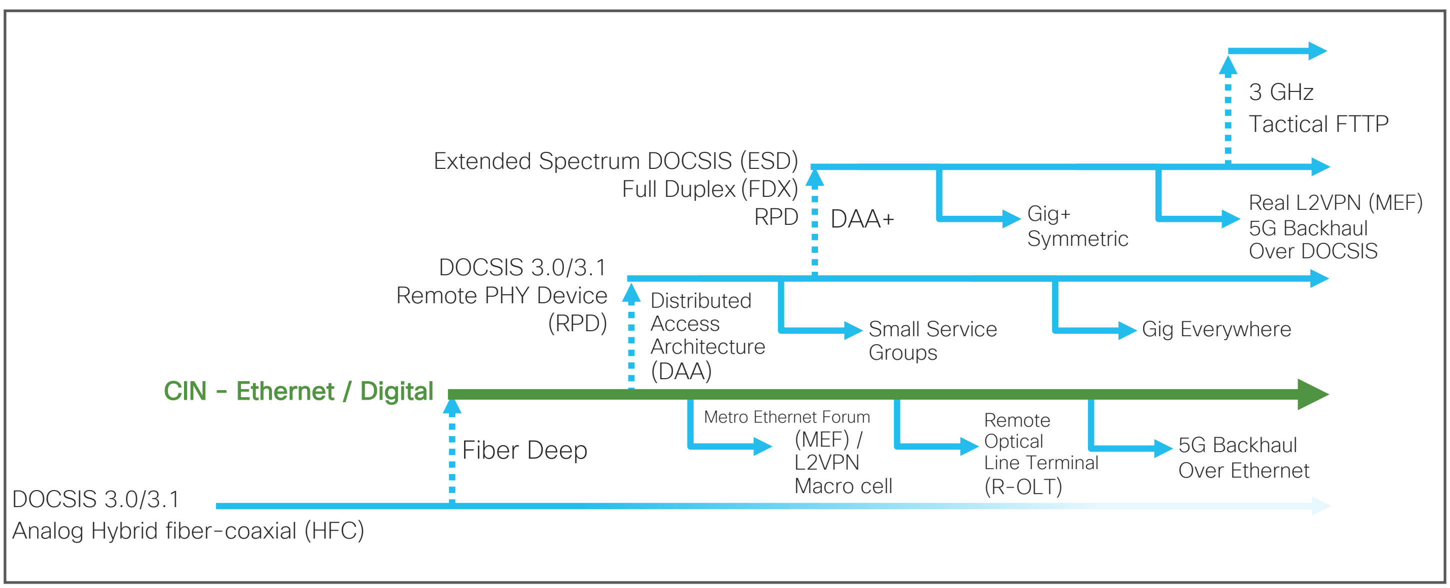

If DOCSIS enabled triple-play, CIN enables multi-tenant play

DOCSIS created a way for operators to offer more than just linear cable services. They could include voice and broadband services, over IP. It was the triple play, back when that idea was new more than 15 years ago. With a CIN, operators can widen the meaning of subscriber to include new business opportunities, such as multi-tenant networking to on board mobile and 5G partners. They can expand opportunities while maintaining their existing systems with more precision and agility. Figure 2 depicts the types of multi-tenant services that can run on a CIN.

The building blocks of a CIN

Several service providers throughout the United States and Canada are active with CIN modeling, development, and deployments. They are using their prescribed evolution to DAA and the addition of RPDs into the last-mile infrastructure as the fulcrum to introduce an over-arching CIN framework. In terms of multi-tenant expansion, this framework is inherently IP and ready for the future. Refer to Figure 3 for their core engineering considerations.

From a capacity perspective, designing a CIN means advancing to discrete, dense optics, with speeds from 10 Gigabits-per-second (Gbps) to 25G, 40G and 100G. Discrete, in this sense, means many ports.

From a powering perspective, CINs are deliberately optimized to make the most efficient use of available power. One of the fallacies of DAA implementations is that they will automatically decrease energy requirements. The idea is that the extension of fiber deeper into serving areas involves an exponential increase in network devices like RPDs, which need power. Although it's less expensive to transmit, say, 1G of data over Ethernet-IP, versus analog optics, overall power efficiency is nonetheless a vital CIN design goal.

An enhancement to the CIN is called a Field Aggregation Router (FAR). It's in development now and is typically located between the cable modem termination system (CMTS)/converged cable access platform (CCAP) core and the RPD. Ultimately, a FAR is envisioned as a temperature-hardened device with a reduced physical footprint that will work with existing node enclosures.

Among service providers, an overall design goal is ever-lower network latency, so the packet-switched nature of a CIN is desirable. Generally speaking, no technology or topology can lower latency beyond the physics of the speed of light. But what a CIN can do is to optimize the topology of the network to be logically shorter and lossless. Optimization truncates the many known and unknown network interlopers that inevitably slow bit transit or contribute to network and cloud lag.

Figure 2. A depiction of multi-tenant services running on a CIN.

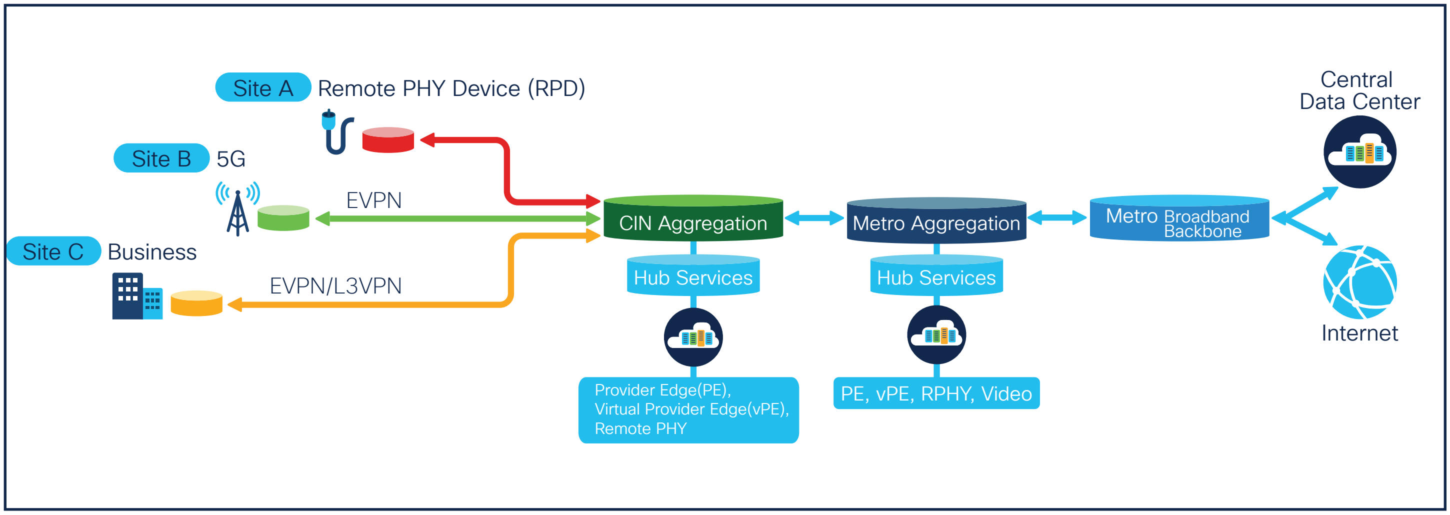

From a security standpoint, going to a CIN takes operators above and beyond existing capabilities to 802.1X and MACSec. The other tools and building blocks that are notable in CIN development are listed in Figure 3.

Figure 3. An overview of CIN requirements.

High capacity Low power	Discrete 10/25, 40/100 options Maximum power efficiency
Long reach	10G DWDM and Extended Reach (ZR) Client Optics
Timing	SyncE, 1588v2, Boundary Clock (BC) support
Remote locations	Temperature hardened systems Reduced physical footprint
Security	802.1X, MACSec
Low latency	Low latency and lossless IP transport
Multicast	IPv6, scalable, non-blocking

How a CIN-driven DAA plan is strategically powerful

Another before/now and after/next scenario helps illustrate how a CIN can make a DAA plan more powerful. In the before scenario, a typical DAA plan of record has RPDs connect to a CCAP core over dedicated, point-to-point 10G links, using colored optics (WDM, DWDM) as shown in Figure 4. This method is inefficient, relative to packet switching over IP/Ethernet. It limits both port density and the ability to scale with demand.

Up to and including its current evolution to DAA evolution, HFC is a stable, reliable, and predictable access network. Over decades, HFC has proven itself to be the workhorse that reliably kept cable broadband providers reasonably ahead of bandwidth demand. At the same time, it has been able to satisfy other needs such as video, voice, and business services.

However, HFC alone is no longer enough for a few reasons. One, it consumes a high number of dense wavelength division multiplexing (DWDM) 10G optical components, which are expensive relative to the "gray optics" and small form factor pluggable (SFP) modules used in a CIN. Two, HFC lacks sufficient hooks and components for a strategic network automation plan. Three, without the benefits of statmuxed packet switching, in the long term, HFC is an inefficient use of fiber assets.

In other words, the operator who simply replaces the point-to-point analog laser with point-to-point Ethernet (aka digital optics) is missing a big opportunity. The opportunity is to take advantage of the inherent benefits and economics of packet switching and routing, so you can evolve today's triple-play consumer bundle into tomorrow's multi-tenant play enterprise bundle.

With a CIN, service providers can continue to operate their networks as they do today. However, eventually they will need to respond to scaling needs. With a CIN in place, they can do so without disruption. A CIN can be a part of network modernization and upgrades without requiring operators to be instant experts in optical networking.

Figure 4. A common DAA layout.

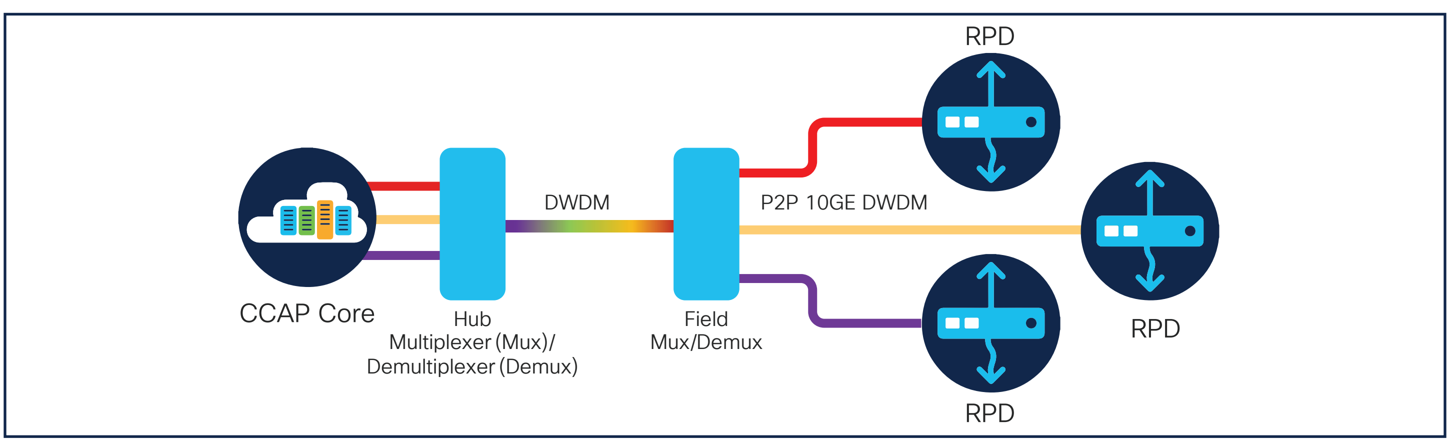

CIN design principles

After you've decided to pursue or model a CIN implementation, you need to think about design. Service providers have a plethora of options, even as they take the step to DAA. No real consensus exists among service providers and the supplier community on to how to build a CIN, but certain design principles do apply.

Timing and DAA project concurrency

The timeline or adoption rate of a CIN works best when it correlates with DAA deployments and specifically RPDs. One of the advantages of RPD deployments is that they can be gradually and gracefully deployed. Deployments can be made in step with marketplace needs, such as strained available capacity and competitive pressures. For most operators, deploying a CIN before launching RPDs isn't a priority, but it often follows quickly and can be strategically relevant to longer-term business opportunities.

Topology

Designing a CIN involves a re-evaluation of some core tenets of traditional HFC network design, starting with topology. HFC is largely a hub-and-spoke design. The HFC segment of any end-to-end, IP-based connection is considered, in essence, a long spoke. One end extends north of the CMTS, to the Internet core, and the other end goes to homes and businesses. For the type of services and data it needed to handle, the design worked, but it's insufficient for a modernized, automated, multi-tenant future.

In contrast, the shift during the DAA transition to a CIN involves a topological shift from hub-and-spoke, to leaf-and-spine. Leaf-spine topologies are common among web-scale cloud providers. They are a way to perform massive data transfers with instantaneous data transit. (They're also sometimes called "Clos" networks for Charles Clos, who formalized the design in 1952.) Leaf-spine architectures are widely used by web-scale providers in data centers where networking efficiencies have proven to be effective. Planning for a leaf-spine architecture is important, especially for carrier and cable networks, which tend to have physical challenges related to being heavily distributed. Because a perfect leaf-spine topology may not be immediately evident, planning will require unique considerations per deployment.

Automation

Certain aspects of network reliability and performance can't be added as an afterthought at the end of the design phase. Security is a prime example. Security needs to be an intrinsic mindset in network and product design. Automation is another example. It serves you best if you equate it with other key design parameters from the outset.

Back in the early days of HFC, and even until recently, automation wasn't an option because it's difficult to automate what you can't see. You can't see things that can be automated when you're using analog optics. In the interim, automation mechanisms flourished and are mature. Now, with modernized packet switching and routing and advanced silicon and optical components, network automation represents a powerful means to achieve scale and efficiency.

Network device automation will become more vital, as node sizes get smaller and more RPDs are scattered throughout the plant. All the RPDs are digital and can see each other, so they need is leaf/spine topology and digital configuration that enables them to work collaboratively. This combination enables right sizing of capacity. More bandwidth is sent where and when it's needed, and bandwidth is conserved when it isn't needed.

CIN engineering

Engineering a CIN involves aggregating the RPDs to the CCAP core in RPHY. It follows an overlay and underlay model. The overlay comprises everything running on the network to make services work, which usually includes the virtualization components. It uses L2TPv3 IP tunnels, so by using an Ethernet-IP CIN, you can maintain all services and add new peering points by creating DEPI/UEPI pseudowires within DOCSIS.

The underlay is, for the most part, the physical infrastructure. Bear in mind that the underlay is IP and RPHY demands not just IP paths, but IP circuits. IP circuits enable path symmetry, latency stability, and contention-free, high caliber broadband connectivity.

Other engineering inclusions for a CIN-based DAA design include 1588 clocking, also known as precision time control, or PTP. PTP is used to establish time throughout the network, which assures that data arrives where it needs to be, when it needs to be there, and in the right order. When you hear the terms 1588 clocking and PTP, you may hear terminology like boundary clocks, grandmaster clocks, and timing servers, which are all vitally important to enter or service the mobile marketplace. Figure 5 shows a simplified CIN design

Most early CIN deployments are for RPD aggregation

North American cable providers active with CIN deployments are starting the journey with RPD aggregation, treating that work as the first leaf off the CIN spine. For instance, imagine 30 RPDs or nodes that are deployed throughout a leg of the HFC network. Rather than each RPD linking individually back to the CCAP core, as in Figure 5, the 30 RPDs link into a CIN leaf switch/router. This design helps to optimize efficiency and network management. It also offers increased control and insights from the deeper fiber connectivity and the digitization aspects of Remote PHY.

In any conversation about network transitions and the technologies involved industrial consensus often doesn't exist. Because of an active and unsurprisingly competitive vendor environment, the DAA transition is a good example. Predictably, most vendors advocate for the technologies that reflect their respective strengths, from digital optics to multiple other iterations.

Few suppliers come naturally from a place of IP, Ethernet, packet-switching/routing, and the product breadth that comes with those strengths. Opting to simply replace a point-to-point analog laser with its

Figure 5. A pre-CIN DAA implementation and post-CIN DAA implementations.

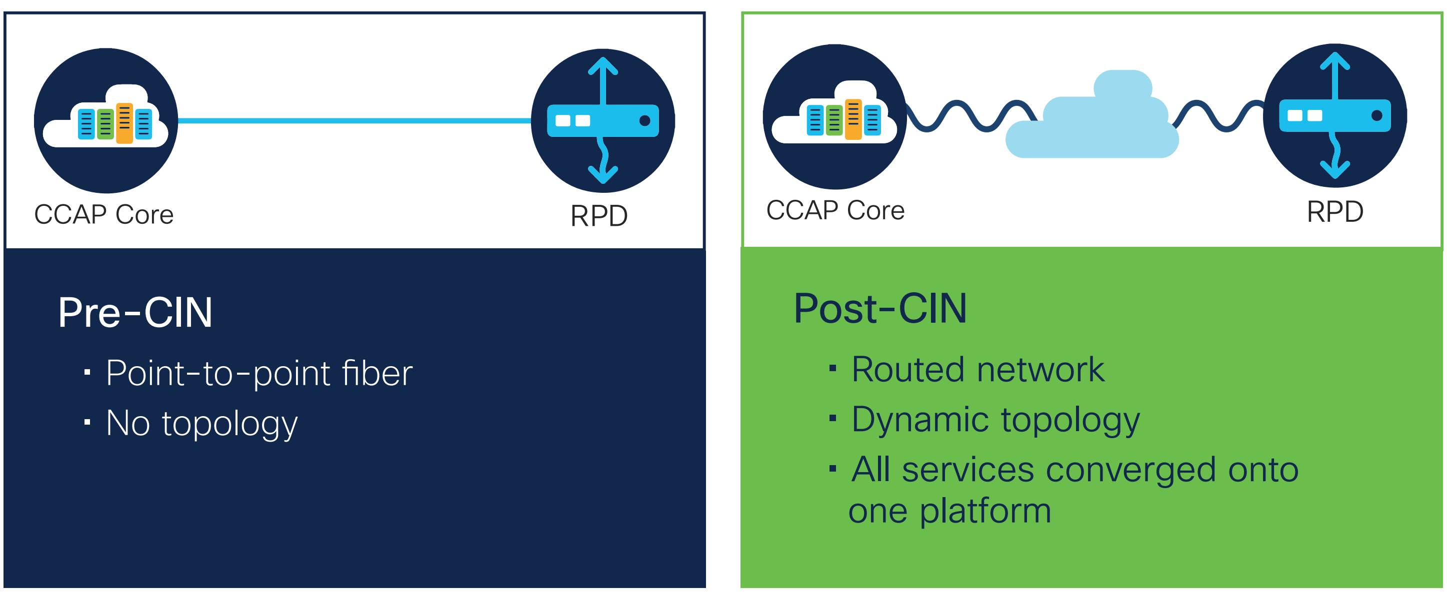

point-to-point digital equivalent leads to an obvious question. What couldn't I do before that I can do now? A lower cost-per-bit metric may be temporarily achievable, but without deep-network routing and switching, there's no statmuxing and automation. Yet both statmuxing and automation are fundamental to long-term scale and strategic agility.

CIN software for telemetry and device health monitoring

Once the network is fully digitized and connected devices like RPDs gain awareness of each other, the gathering and interpretation of telemetry data offers another advantage. The Cisco CIN framework consists in part of the NCS 5500 Series, NCS 500 Series (NCS 540), RPHY components, and IOS XR software. The data collected from RPDs, as one example, is applied to both network automation and the continuous assessment of device health.

Extracting device data as a way to assess various conditions, including health, isn't a new concept. The first SNMP data models were defined around 30 years ago. What differs in today's telemetry landscape is data collection speed. Data collection using the application of model driven telemetry (MDT) is a core component of network operating software like IOS XR.

Here is another before-and-after example. Today's device collection intervals that use SNMP typically occur in increments of five minutes. In our scale tests, collecting interface stats for a device with 576 devices took as much as five minutes. Using MDT, the collection happened every five seconds, although collection intervals of 30 seconds are more plausible, in a field environment.

MDTs use pre-defined models that can be targeted at deriving health insights about RPDs and other devices. As a component of IOS XR and as part of a CIN-based DAA, it's a modernization angle that didn't exist previously, at least not to this extent. It also enables operators to learn and grow with the data their network generates. This growth in automation and analytics-based insights can inform operational activities like proactive network maintenance.

Ultimately, a CIN based on the NCS 5500 Series, with IOS XR, can establish a common network operating system, from the core through the last-mile access plant, which operators can readily build upon with automation, insights, and analytics.

Taking the next step

The emergence of the CIN is simultaneously a big story and an old story. Networking has long had a tug-of-war between circuit-switched and packet-switched fabrics. In the long run, packet switching has proven itself repeatedly. Ethernet is forever.

The next advancement for cable providers seeking penultimate scale and strategic, multi-tenant network growth is to coordinate CIN deployments in parallel with DAA and consequent RPD deployments. All the knowledge and expertise to build a CIN already exists within the industry, particularly among engineers who work on the Internet-facing side of the CMTS. The next step is to apply that knowledge to extend the advantages of packet switching into the last-mile plant, which is the graceful, pay-as-you-go build strategy favored by services providers.

Learn more

To advance your DAA intentions with a CIN, contact your cable account team for more details.

Visit www.cisco.com/go/cable

In many ways, the emergence of CINs is reminiscent of earlier milestone chapters in cable, like the early days of the consumer triple-play of voice, video and data, over IP. Back then, consensus was difficult to find. Should you converge the services, or keep them separate? Offer some over quadrature amplitude modulation (QAM) and others over IP, or all three over IP? Invest in ATM, or not? At the time, the number of choices was daunting.

The term triple play was deserved then and now. In baseball, a triple play is both difficult and rare. A converged triple-play of voice, video, and data, over IP was also difficult. But now, decades later, nobody would doubt it was the right choice. In that sense, the cable industry has successfully gone through a similar evolution.

Building a multiservice, converged IP network is the next great chapter in the cable capacity and growth story. It's an opportunity to put in place, today, the foundation for what could be the next story of multiplay service and economic growth. It's an open pathway into the future of the network provider business. Without proprietary, technological or vendor lock-in, a CIN offers a framework for operators to plot their future – a future with services, operations, and capacity that massively scale.