



Services and Applications Perspective

Five Essential Elements of Carrier Class NTP

Telecom operators must bring to managed IP the same high levels of quality assurance consumers expect in the circuit switched environment. That calls for a better class of NTP.

Table of Contents

Overview	1
Three Ways to Synchronize Time	2
Legacy Methods	2
Enterprise NTP	3
Carrier Class NTP	4
Element 1: High Precision	5
Reference clock accuracy	5
Timing packet delay variation.	5
Capacity throughput.	6
Element 2: High Availability	7
Component redundancy	7
Blade redundancy	7
End-to-end redundancy	7
Element 3: Security	7
Network isolation	7
Clock isolation	8
Time isolation	8
Element 4: Robust Management	8
Element 5: Easy SSU Integration	9
It All Comes Back to Service Quality	9

Overview

Never before have telcos, cable operators, equipment makers, and content providers had so many ways to engage the consumer – creating a cacophony of complex provider, service, content, and device scenarios. Operators find themselves caught in a convergence of multiple new-media technology plays. Consumers meanwhile struggle to make brand choices as they switch from “watching what’s on,” i.e., linear consumption models, to more on-demand models like IPTV, online social networking, games, and downloads.

Telcos, however, have historically enjoyed a unique advantage versus the other players: superb quality assurance. For decades people have expected to hear a dial tone when they pick up the phone. Clearly, a winning strategy will be to bring that same level of customer confidence to other arenas as well.

That means bringing the same quality assurance strengths into play. One of those is the ability to synchronize the network elements that must interoperate smoothly to enable end-to-end sessions. But rather than just phone calls, today’s sessions might also be games, on-demand movies, podcasts, and other interesting interactive applications. And just as POTS (plain old telephone service) is no longer a sufficient value proposition for telcos, plain old *time* service is no longer sufficient for synchronizing the network elements on which fresh value propositions depend.

The issue with legacy synchronization services is that they are designed to work on circuit switched networks, whereas telecom service provider’s new consumer offerings are mostly deployed on IP networks. Synchronization services that *are* designed for IP networks are intended for enterprise environments. Such services are generally based on NTP (Network Time Protocol), with lower performance standards than those required by carrier networks. Ideally operators should have the best of both worlds. On the one hand, they should be able to leverage the enterprise interoperability and huge technology base of NTP, on the other, they should also be able to satisfy their own higher performance standards and leverage their own existing infrastructure and knowledge investments to reduce costs.

That means *not* inventing NTP all over again. It also means implementing NTP differently in carriers than in enterprises. How do you know if your implementation of NTP is carrier class? We suggest five essential elements:

- High precision
- High availability
- Security
- Robust management
- Easy infrastructure integration

These elements are important goals; but so is how you reach them. One approach is clearly wrong: just adding enterprise NTP to a carrier’s managed IP network. Here’s why:

Three Ways to Synchronize Time

Until the recent introduction of carrier class NTP, operators could only synchronize time in a managed IP network one of two ways – neither of which were designed for that purpose – enterprise NTP and legacy methods.

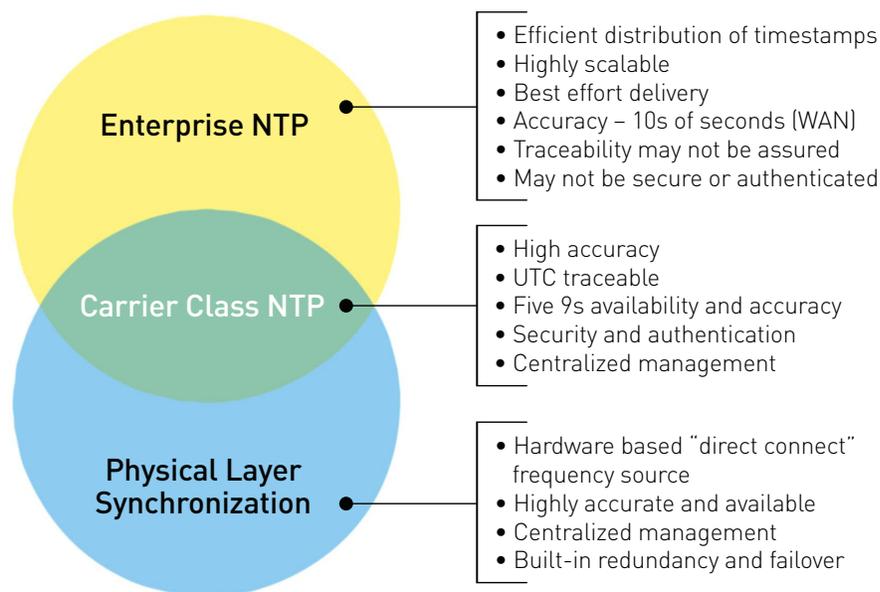


FIG. 1 Carrier class NTP is the overlap of enterprise NTP and SSU-based physical layer synchronization methods based on their respective areas of strength.

Carrier class NTP fully leverages existing network synchronization infrastructure

Physical Layer Synchronization Methods

These employ SSUs (synchronization supply units), also known as BITS (building integration timing supplies). In the SSU circuit switched world, synchronization is hardwired. Elements, such as a SONET (synchronous optical network) switch, receive timing as analog signals of very precise frequency on dedicated ports over direct connections from the SSU. All elements stay in synch because all SSUs reference a trusted time source – usually a GPS receiver – via an endpoint-to-endpoint circuit connection. If the GPS reference is lost, the SSU internal atomic clock maintains accurate time within qualified limits until the connection is regained. SSU architectures include card-based redundancy to assure high availability. SSUs also monitor the timekeeping performance of the network elements they supply, reporting back both their own status and the status of these elements to centralized management consoles in the NOC (network operations center).

The advantage of telecom legacy timekeeping is how deterministic it is. In a well-designed network infrastructure, all elements receive timing from a trusted source over a physical circuit with predetermined properties and therefore with predictable results.

Enterprise NTP

Contrast the carrier environment with enterprise IP networks, including the Internet. The world NTP was designed to handle is *non*-deterministic. Here time is distributed as time stamped packets within the same IP traffic flows as other packets, and is subject

to the same variability. That can include backed up router queues, link congestion, and an indeterminate number of hops between a sender's local loop and a receiver's. In the case of NTP, the senders are NTP servers that receive time from either an external source (like GPS) or from another NTP server. Receivers are clients that read the timestamps and reset their internal clocks accordingly. In an enterprise, clients can reside on any system that uses time, from desktop PCs, to datacenter mainframes, network routers, and firewalls.

Applications cover a wide range, but in general fulfill two missions: 1) Event synchronization, i.e., to enable events to occur at the proper time; and 2) Computer forensics, i.e., to provide proof when events did, or did not, occur.

Organizations employ event synchronization to accomplish one or both of the following objectives: 1) To schedule a process – i.e., to ensure that it starts or stops on time or runs for a specified period regardless of when it starts or stops; and 2) to ensure that cooperating processes can interoperate correctly, so that if one process hands a task off to a second process, the latter will in fact be ready to accept the handoff. In forensics, organizations use timestamps typically to discover the order of events that led to another event – such as emails leading up to a corporate fraud, or system log entries just before a crash.

Table 1 summarizes some typical examples of both missions:

The advantage of enterprise NTP is the efficiency and scalability with which it can distribute time throughout a packet switched network. Both properties are critical considering how large and widely distributed are the processes that use time within an enterprise. Hardwiring a trusted time reference to each one would not be practical.

Enterprise NTP is typically a “best effort” service used for event synchronization and computer forensics

Mission Type	When Employed	Artifact(s)	Purpose	Example Apps
Event Synchronization	During events	Application messages, Flags	Ensure events occur on time, in correct sequence	Transaction processing, Process control, Authentication
Computer Forensics	After events	Time stamps	Determine when events occurred and in what sequence	Digital signatures, Crime investigation, Fault diagnosis

TABLE 1 Why Organizations Employ Enterprise NTP

A well-architected enterprise NTP deployment can deliver time from a single NTP server to thousands of NTP clients within the millisecond accuracy most enterprise applications require. Accuracy can degrade, however, to 10s of seconds if networks become congested – say, if links go down or if IP routers momentarily send packets through more hops than network architects expect. Even if the network performs well, variability is still a factor because of how the NTP application itself acquires time in the server. It calls the operating system to get time from the clock on the server. How fast the application gets back a response from the OS depends on a number of factors, including:

- CPU context switching
- Saving and restoring states
- Flushing caches around the invocation of the interrupt handler
- Queuing as the packet makes its way through the protocol layers to the application

Carrier class NTP is required for QoS assurance in advanced telecom applications

Ideally, you would want the time in the timestamp to be the time the packet leaves the server – which would require a hardware timestamp at the instant the packet is encoded on the wire.

Carrier Class NTP

Excessive variability makes enterprise NTP unsuitable for most carrier-class applications such as those outlined in Table 2. IPTV and wireless, for example, represent this class well as their tolerance for variability is very low. If cells in a cellular network are not running off the same time index, calls will drop as users travel from one cell to another. Another issue for mobile wireless carriers is intra-provider call data record (CDR) reconciliation. If timestamps don't agree then billing may be inaccurate, potentially resulting in revenue loss, or require costly manual mediation.

The best of both worlds is for the legacy and enterprise solutions to overlap in their respective areas of strength. That is essentially the idea behind carrier class NTP. Even if variability is an issue, NTP has won some support among carriers for such applications as coarse time stamping of billing records and correlation of alarm events.

Application of Service	Carrier-Class Synchronization Needed For	Business Impact
SLA monitoring	Measurements	Revenue realization Customer satisfaction
IP performance monitoring	Measurements	Network uptime, alarms, diagnostics, Traffic management, Measurement accuracy
Video services – IPTV	Video encoding/decoding, Conditional access, Network security, Content protection	Service availability, Picture quality Revenue accounting for content QoS, QoE, Ad revenue
Circuit emulation	Measurements, SLAs Performance monitoring	QoS, Service outages
Wireless networks	AAA/content billing/ CDR reconciliation	Revenue assurance QoS Billing accuracy
IP and Ethernet services	Policy / QoS management	SLA management, Billing
Business VoIP, VoIP Centrex	Call logs, Event logs, CDRs, Quality measurement: Jitter, Delay, Packet loss	Billing, SLA assurance Event tracking Fault resolution
Network security	Intrusion defense Accurate event logging, AAA	Theft-of-service assurance Service disruption Revenue loss
E911 services	Call logs	Regulatory compliance

TABLE 2 Why Telecom Operators Employ Carrier Class NTP

Operators like the fact that NTP is easy to get. Most operating systems have NTP clients. And free public NTP servers exist on the Internet. They also like the scalability and efficiency of NTP. Then there is the appeal of running a single technology stack. Most corporate IT managers already use NTP for their enterprise systems.

So NTP certainly has value to offer. But if NTP is to deliver the QoS assurance that new value propositions require, then NTP implementations must fully support carrier-class requirements. That is why carrier class NTP should include the following key elements:

Element 1: High Precision

How accurate time needs to be depends on the applications and operations performed. Most operations impacting service assurance (e.g., session set-up, content play-out, QoS measurement) require time measured in the 10s to 100s of microseconds. These are applications that make real-time handoffs (or monitor handoffs) between points at opposite ends of an IP connection. Less time sensitive are policy management and event logging applications, such as E911 service monitoring or digitally signing expiration timestamps on tickets to particular content. Requirements here typically range in the milliseconds.

Network engineers, of course, must design their networks to meet the timing needs of the most stringent applications in their service infrastructures. For them the issue is the best way to deploy time using NTP that consistently meets their accuracy requirements.

An NTP service that is inherently accurate meets three key criteria:

- The NTP server has a precise clock – i.e., an accurate way to acquire and keep time
- A low latency path exists between NTP server and client
- Sufficient processing capacity is in place to meet demand load

Reference clock accuracy

The world standard for accurate time is GPS. The most accurate way to keep time without a direct GPS reference is an atomic clock. A precise NTP server clock is an atomic clock that also has the ability to acquire time from a GPS receiver. If the GPS reference is lost, the accuracy of the clock will stay within qualified limits until GPS reception is restored. The accuracy of all clocks, however, degrades over time, which means they must periodically be reset. The speed at which clock accuracy degrades when an external time reference is absent is called its “holdover rate.” Operators need to ask about the holdover rate of the NTP server, in addition to whether it uses an atomic clock and GPS reference.

Timing packet delay variation

Ultimately, it is not the accuracy of the time at the NTP server that counts, but the accuracy of the time employed by the client application. This is why there needs to be a low latency path between the NTP server and client. Inevitably, there is a delay between when the server timestamps an outgoing packet and the time the client receives that packet. The greater the network delay variation and asymmetry, the less accurate the timestamp in the packet.

Potentially, there are three key contributors to packet delay:

- The NTP server operating system
- Network hops between server and client
- Other NTP servers between a time reference and a client

Carrier class networks are engineered to meet accuracy requirements of their most stringent application under all operating conditions

Carrier class NTP networks are engineered to stratum level 2 and above to assure precision and traceability

As previously noted, the most accurate packet timestamp is a timestamp that has been applied in the packet at the instant the packet is encoded at the physical layer. To the extent software calls are involved, delays are incurred.

Similarly, the more network hops a packet traverses the greater the delay. Carrier networks pay close attention to hop counts and path latency as key factors in assuring quality of service. The most effective way to limit hop count and delay variation impact on timing accuracy is to move from a centralized NTP model to a distributed NTP model where NTP servers are placed at key nodes and offices in the network. This mirrors similar engineering guidelines for placement of SSUs in the network for physical layer synchronization requirements. NTP servers integrated with central office SSUs take full advantage of existing synchronization equipment and practices.

The third way packets get delayed is by increasing the number of NTP servers between the time source and the client. In an enterprise NTP servers often work in a daisy chain where upstream servers feed time to downstream servers that then feed time to clients or to other servers. Each level in the hierarchy is called a stratum, with the time source (e.g., GPS) at stratum 0. Stratum-1 servers reference the time source and may in turn provide time to stratum-2 servers and so on. Stratum-2 service is less accurate than stratum-1, because of the delay involved in reading the stratum-1 timestamp and sending out a new timestamp. Carrier class NTP is generally deployed in a very flat hierarchy with servers rarely deployed beyond stratum level 2 (stratum-1 only is preferred).

Capacity throughput

Capacity is measured in terms of the number of requests for timestamps that can be served within a given period with a given level of accuracy. For example, a server able to deliver 2000 requests per second with 10-microsecond accuracy would be considered carrier-class. Another consideration is the ability to scale up capacity within the SSU. If servers were implemented as blades, for example, you might simply add more blades to a SSU shelf to double or triple the number of clients and client requests served without re-architecting the entire server configuration or taking up more floor space.

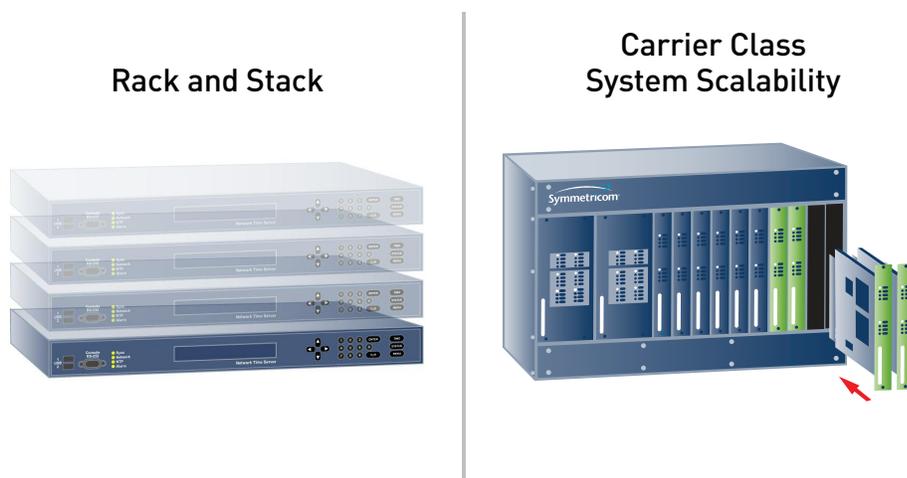


FIG. 2 Carrier class systems with card-based NTP servers are designed to scale as the network grows.

Element 2: High Availability

A blade based architecture not only promotes scalability, but also high availability as well. High availability means that clients have access to high-precision NTP service (as just described) without interruption. Carrier class five 9s availability, which means the service is available 99.999% of the time, is consistent with other key service infrastructure components. There are a number of strategies to employ to achieve this level, and using them all at the same time is the best strategy of all. They include:

Component redundancy

If an NTP server is a blade on the shelf of an SSU rack, then one way to make NTP highly available is to build the most vulnerable shelf components such as the power supply or the port, into a redundant system. If one power supply fails (or is faltering), for example, the other one should automatically sense the condition and keep the shelf working.

Blade redundancy

Making the blade itself redundant provides another layer of protection. Servers can be implemented as pairs of redundant cards, so that if one member of the pair fails, the other picks up the transaction load of the failed blade.

End-to-end redundancy

If an entire blade complex (e.g., both cards in a redundant pair) fails then what happens next? In that case, other blades should automatically serve the clients previously being served by the failed blade. This can happen in at least two ways: at the transport level or at the network level. Transport level redundancy means that blades will be configured so that if one blade fails, its inbound traffic will automatically be switched by hardware to a nearby blade in the rack. Network level redundancy means that IP packets headed for the failed NTP server blades will automatically be routed to back-up blades in the same rack or in other racks. In addition, to further support network redundancy, the router making those failover packet routing decisions can itself be redundant.

Element 3: Security

The same architecture that lends itself naturally to scalability and high-availability also lends itself well to security. Just as NTP service functionality may be discretely partitioned across distributed components, those partitions themselves may be employed as natural barriers against potential security breaches. Here are some examples:

Network isolation

NTP blades may operate in several networks simultaneously; the managed IP network, the public Internet, and the network over which the operator monitors and manages its time distribution system. All three networks should be isolated from each other, such that if one of the networks is compromised, the others are not (at least not by a NTP server blade).

How might a network be compromised? One way is a denial of service (DoS) attack in which software agents send a blizzard of time requests against a particular NTP port. Without isolation, the NTP server on that port may fail under the load, potentially impacting timing dependant services in the network. Achieving isolation can be done by hardware at the ports and also by firewalls at each port.

NTP server cards leverage existing synchronization shelves to provide redundancy for high availability applications

Not only can the three types of networks be separated from each other, but so also can the different network segments. For example, NTP packet traffic going back and forth to central office equipment can be port-isolated from packet traffic associated with street equipment cabinets, and those individual "street traffic" flows can also be port-isolated from each other.

Clock isolation

Another point of demarcation is the boundary between clock frequency source and the IP network. Isolating the clock management ports from the NTP traffic ports prevents the IP network from being used as an unauthorized entry point to clock management and configuration.

Time isolation

Physical isolation of the clock is one thing; virtual isolation of timestamps is another. Adding malicious content to timestamp packets is a well-known way to hack systems. For protection, measures exist in the NTP standard to trace time back to its intended source and ultimately to UTC¹. These measures employ digital signatures with a checksum authentication method called MD5. Servers can also employ MD5 to ensure only requests from valid clients are received (and responded to). Authenticated traceability is also a prerequisite for proving timestamps were accurate when received and not subsequently altered should timestamps ever be called into question such as when there is a billing dispute.

Element 4: Robust Management

Since QoS assurance of the carrier service infrastructure relies heavily on time synchronization, it must also rely on QoS assurance of that synchronization infrastructure. That is equally true on both the circuit switched side and on the IP side. It follows then that QoS assurance *management* is of equal importance on both sides and that their two management domains should be tightly integrated.

QoS assurance relies heavily on time synchronization

Network managers already have robust tools in place to monitor and measure timing QoS. Those should be extended, not reinvented, for IP. That implies TL1 (translation language 1) awareness on the part of management software interfaces. TL1 is a management protocol used to communicate network element status back to the NOC (network operations center). Moreover, those interfaces should be extended so that operators have a familiar look-and-feel and similar productivity tools for identifying, analyzing, and reporting synchronization performance issues of IP devices and applications. This includes real time alarms when preset performance thresholds are violated.

Time traceability, an important security element, is also important for management. Fault detection, isolation, and investigation require event logs based on known time references, which can only happen if time is traceable. Ideally, all NTP blade servers should reference the same time source, which is an additional reason for the time source to be secure. A single reference by definition ensures absolute consistency among all primary time references throughout the carrier's infrastructure. Consistency means that management and analysis are all based on the same sources at the same time – so decisions and interpretations reflect what's actually occurring.

Element 5: Easy SSU Integration

Integration of management across circuit and packet-based synchronization domains is most efficiently accomplished if NTP blades plug into existing SSUs. That means the blades have the right mechanical and electrical interfaces and are equally visible to management software as other SSU components. A common facility – both physical and virtual – allows operators to manage synchronization as a single asset regardless of whether the underlying delivery vehicle is a circuit or a packet. And, as previously noted, it offers benefits in terms of scalability, a smaller footprint, and return on training investment.

SSU Integration, if correctly implemented, can also benefit operators by pre-positioning them to leverage future technologies. In fact one example of that might be how NTP management integrates *legacy* SSU technologies like TL1. Rather than hardwire multiple specific protocols into a management layer, e.g. one for circuit, one for NTP, one potentially for something else, a better approach is to build off an extensible meta language base that offers the flexibility to employ *both* legacy and future protocols in a common context.

One new technology will be IEEE 1588 PTP (Precision Time Protocol). This is an emerging standard designed to meet stringent synchronization needs on an Ethernet link, such as synchronizing wireless base stations. PTP builds off many of the same ideas as NTP, but adds new concepts that increase precision. One is hardware time stamping as packets are encoded onto the physical layer. Another is software logic that characterizes and compensates for packet delays. Hardware time stamping, however, is already available today in advanced NTP blade servers, giving operators a head start in meeting PTP requirements.

It All Comes Back to Service Quality

The five essential elements constitute a solid base for future business. Telecom companies have trained customers to rely on them for one clear and compelling reason: quality of service. In today's crowded, noisy, and rapidly changing marketplace, that message still resonates. But to build on a position of strength, telecom operators will need to leverage the existing quality assurance technologies. That means building on existing SSU and NTP foundations. The key is how to blend these technologies within an existing quality assurance framework that already works, while meeting new and demanding requirements. This can be accomplished best by deploying SSUs and NTP services that incorporate these five elements.

Operators can leverage existing synchronization infrastructure (SSU/BITS clocks) to meet the five essential elements of carrier class NTP

¹ UTC (international abbreviation for Coordinated Universal Time) is the international time standard derived by agreement among several national standards bodies, including NIST (National Institute of Standards and Technology) in the U.S. GPS is considered a trusted UTC source.



SYMMETRICOM, INC.
2300 Orchard Parkway
San Jose, California
95131-1017
tel: 408.433.0910
fax: 408.428.7896
info@symmetricom.com
www.symmetricom.com