



Timing Requirements for Distributed Systems White Paper

As data storage migrates to the network edge to limit latency, and applications move to distributed system models to increase capacity and speed, definitive synchronization requirements have become an integral data center design element. Large, distributed databases are now in the petabyte and exabyte scale, execute thousands of transactions per second (TPS) [1], and must have multiple replications of data in multiple global locations. Managing data consistency between replications drives the need for synchronization between servers in a data center and also between data centers that may be separated by thousands of miles. The level of timing accuracy required for a specific distributed system, application, or database is dependent on several factors that must be considered when planning a timing architecture. This white paper explores several of the factors, introduces the concept of time envelopes per second to quantify the clock requirements, and discusses various solutions using distributed externally synchronized clocks.

Data Consistency and the Requirement for Precision Time

Data consistency is the need for a record, at a fixed point in time, to have the same value across all replications of a database or application. There are many alternate terms with various nuances (causal consistency, local consistency, global consistency, strong consistency, eventual consistency) [2], but they all have the same basic principle: the value is consistent across the database at a timestamped point—even if this was only achieved by altering the value in some replications through post-execution reconciliation.

Achieving this is a trivial task if the entire database or system is stored in one unreplicated instance on one machine. For globally distributed systems with multiple replications, this would still be a trivial task in a world without communication delays and if each server had access to absolute time with no uncertainty. However, communications delays exist, and absolute time without uncertainty doesn't exist. Issues such as physical distance, the communication path technology, asymmetry, and clock uncertainty at each location all create challenges that must be mitigated.

By globally distributing databases and applications at the network edge, the communications delays—or latency—between users and the databases are minimized. As an example, consider a database with all its replications located in a data center in Ashburn, VA. A user in San Jose, CA needs to communicate with the server. The theoretical fastest transfer of information between the two sites is limited by the speed of light at 12.7 ms. Something as simple as reading a value from the database requires a round trip message, meaning that the latency would be ~25 ms. Even with a dedicated fiber optic link between the two facilities, factors including the refractive index of the fiber, the actual fiber route, and other delays in the system will further limit the performance to roughly 75 ms to 100 ms per round trip.

To improve the latency, a replication of the database can be located at a server in San Jose. However, this now complicates the consistency management. If the same or related records are updated in both locations in a similar time frame, there needs to be a method to reconcile and determine which truly occurred first. Just applying a record lock will not resolve the issue because if Ashburn recognizes the need to lock the record, it will be at least 12.7 ms before San Jose would possibly know this record was locked. To minimize the consistency issues, if both data centers are synchronized to an external clock, transactions (Note 1) can be timestamped to that clock and any inconsistency can be resolved through post-execution reconciliation algorithms.

The theoretical best time granularity that the algorithms can order the transactions with absolute certainty is limited to twice the level of the clock uncertainty (CU) at the two locations. If a system can timestamp with a CU of ± 100 ms, the system can only unambiguously determine the correct sequential order of two events timestamped 200 ms or greater apart.

This spacing is referred to in this paper as a time envelope. In the first graph of Figure 1, the actual sequencing of events can be determined as the envelopes do not overlap, but in the second graph an unresolvable inconsistency occurs and there is no way to determine with certainty which transaction occurred first because of the overlapping envelopes. Other system issues may increase the size of the envelopes, but, as technology advances, the clock uncertainty will always determine the best-case minimum time envelope that can be used to determine the sequential order of events, also known as linearization.




FIGURE 1: The first graph is an example of a resolvable consistency with transactions time-stamped 300 ms apart, which is greater than the 200 ms time envelope. The second graph has an unresolvable inconsistency because the transactions are timestamped 150 ms apart, which is less than the 200 ms time envelope.

Visualize each set of related or causal records as a unique color. Envelopes of different colors may overlap, but envelopes of the same color may not. If transactions with overlapping time envelopes operate on different and unrelated records, no inconsistency will occur even if the true time order cannot be resolved. This is illustrated in Figure 2 where only the transactions on R4 and R5 are problematic. But as the duration of the envelopes increases, the probability of two identical or related records being affected in one time envelope increases, which in turn degrades system performance as transactions are rolled back. Consensus methods, probability analysis, or assigning priority to certain replications can be used to roll back only one of the transactions, but every roll back creates system inefficiencies, and the number of roll backs can reach unmanageable levels that eventually result in corruptions. A cost trade-off must be managed.

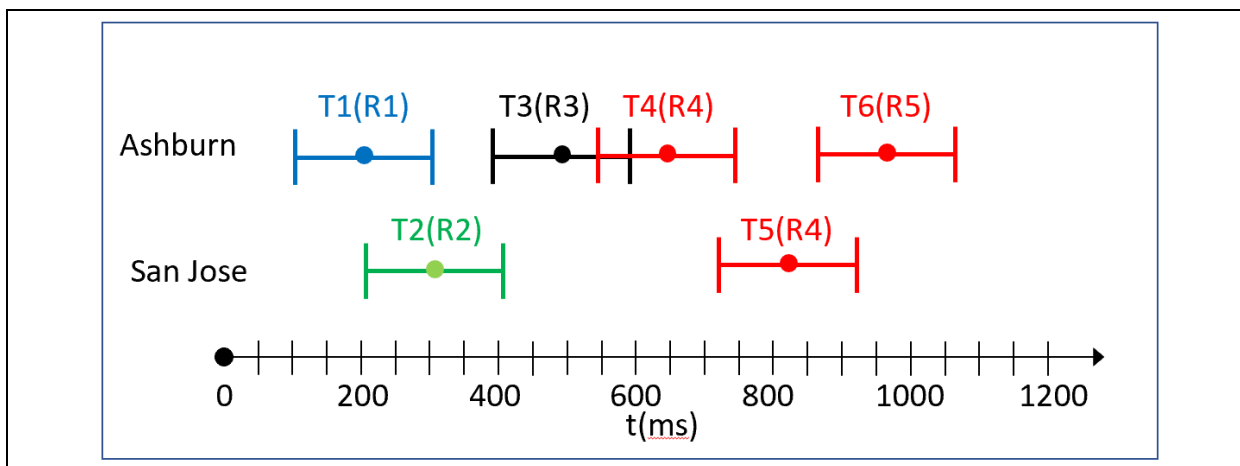


FIGURE 2: If R1, R2, R3, and R4 are completely independent records that would never interact with each other, then the ordering is not critical and the overlapping envelopes could be allowed to occur with minimal to no risk. The operation on R4 in Ashburn at 650 ms and San Jose at 825 ms are on copies of the same record, and because the time envelopes overlap this causes an issue. If R5 is a function of record R4 (for example, $R5 = 2R4 + 6$) the two transactions at 812 ms and 975 ms would be problematic as they are causal records with overlap.

Time stamping and time envelopes only enable software to reconcile inconsistencies after the fact. A system with $CU = \pm 2$ ms, with a latency between sites of 12.7 ms will only be able to achieve consistency through post-execution algorithms. In order to minimize the total burden of synchronization post-execution, portions of the databases are sometimes stored in localized shards and only records that have a high probability of demand in multiple locations are replicated.

Relationship Between System Requirements and Equipment Specifications

System architects typically specify TPS. Clock vendors specify CU. A simple equation that ensures no inconsistencies can be derived from the discussion above. For the system above with ± 100 ms, there is a maximum of five non-overlapping time envelopes that can be created, as shown in Figure 3.




FIGURE 3: For a clock with uncertainty of ± 100 ms, each time envelope for unambiguous resolution is 200 ms. This results in a maximum of five time envelopes over a one second period.

This leads to the inequality in Equation 1 that quantifies the requirement for TPS as a function of CU ensuring no irresolvable inconsistencies.

EQUATION 1:

$$TPS < \frac{1}{2CU}$$

Equation 1 is the fundamental underlying principle for determining the correct clock for a system. However, TPS and CU need to be modified to account for real world applications to avoid underspecifying the clock requirements (resulting in system degradation) or overspecifying the clock (resulting in unnecessary costs to the system).

TPS Considerations

1. **Average vs. Peak TPS.** Some databases execute write commands at a constant rate. An industrial process control database that records sensor values for a production line at specified time intervals is an example of a constant-rate system. Other applications have peak transaction rates. A financial trading application may have peak activity periods just after a market open and just before a market close. The TPS may be one hundred times higher than the average TPS during these periods. Many system specifications provide the average TPS based off the number of transactions per day divided by 86400 seconds. However, the CU needs to be designed around the peak rates, not the average daily rates.
2. **The criticality of irresolvable inconsistencies.** For many applications, the inconsistencies may have little impact to the users other than an annoyance factor. A cloud-based photography application containing metadata for photos most probably won't be significantly impacted by a couple of inconsistencies that force a handful of keywords or tags to be rolled back. However, for some applications there could be significant financial or legal costs, regulatory requirements [3,4], and even an impact on human life if the applications could affect emergency services or military decisions.
3. **The probability that two causal records could be operated upon in the same time envelope at different locations.** If a distributed database contains customer addresses and each region maintains its own customer addresses in a localized shard, the probability of someone in Tokyo updating a customer address in London at the same time as someone in New York is zero. In contrast, a military application may be the summation of multiple integrated radar images. Each location with a radar contributes to the overall air picture. If multiple radars are tracking the same fast-moving target, the systems may be reporting data essentially concurrently, and the

envelope size should be minimized to reduce the probability of simultaneous writes on the fast-moving target. Equation 1 can be modified with scaling factors for each of these considerations.

EQUATION 2:

$$S_{pk} \times S_c \times S_{pr} \times TPS < \frac{1}{2CU}$$

Where:

S_{pk} = Scaling factor for peak value vs. average value. The value can be directly determined by dividing the peak TPS by the average TPS.

S_c = Scaling factor for criticality of inconsistencies. Subjective risk value based on application—systems with no tolerance for inconsistency may opt for an S_c of higher than 1, which will require a smaller CU. A system that can tolerate inconsistencies with little impact may choose an S_c value of 0.01, relaxing the CU requirements.

S_{pr} = Scaling factor for probability of writing to identical or causal records in the same envelope. This scaling factor could be very low for certain applications where data is normally regionally isolated like the customer database listed above, which means a scaling as low as 0.001 may be used, where as the scaling factor should be set to 1 when there is a high probability of occurrence.

The weighted TPS (WTPS) can be used to consolidate the three scaling factors and the TPS into one variable.

Table 1 list several hypothetical system requirements and the process to evaluate WTPS.

TABLE 1: APPLICATIONS AND WTPS REQUIREMENTS

Use Case	Criticality of Unresolvable Inconsistency	Probability of Related Records Updated at Different Locations in Short Time Period	Avg. TPS	S_{pk}	S_c	S_{pr}	WTPS
Military Operations Database	High: Incorrectly reported data can result in loss of life.	High: Integrated battlefield operations are updated from multiple sources.	10,000	10	10	1	1,000,000
Financial Trading Company	High: Lost trades can have significant financial effects.	High: High volume of trading requests at certain times of day and stock news creates individual targets.	1,000	10	10	1	100,000
Multinational Bank Database	High: Lost transactions can lead to regulatory and legal issues.	Medium: Accounts need to be accessed at different locations.	1,000	1	10	0.1	1,000
Medical Database	High: Incorrectly reported data can result in loss of life.	Low: Individual patient data seldomly transacted at different locations simultaneously.	1,000	1	10	0.01	100
Industrial Process Control	Medium: A few missing data points can likely be averaged out.	Low: Data from sensors taken at regular intervals and specific locations.	10,000	1	1	0.01	100
Social Media Photography Application	Low: Affects customer experience, but can be resolved by reentering.	Low: Most photograph meta-data is only updated by the account holder.	10,000	10	0.01	0.01	10
Sales Contact Database	Low: Worst-case, the user has to reenter data.	Low: Updates may occur at different sites, but not very probable to happen within close time frames.	100	1	0.1	0.01	0.1

Considerations for CU

The right hand side of Equation 1, $1/2CU$, represents the number of time envelopes available based on the clock uncertainty. A more intuitive way to compare this quantity to TPS is to use the term envelopes per second (EPS). Replacing TPS with the WTPS value, and using EPS, Equation 1 can be rewritten as Equation 3:

EQUATION 3:

$$WTPS < EPS$$

This equation makes it simpler to visualize the requirements.

For the EPS of multiple sites to be consistent, they must synchronize to an external clock. UTC is typically used as the external clock source and it is distributed through various different methods before it arrives at the equipment where a transaction occurs. Figure 4 shows the two most common paths for delivering time to data center equipment.

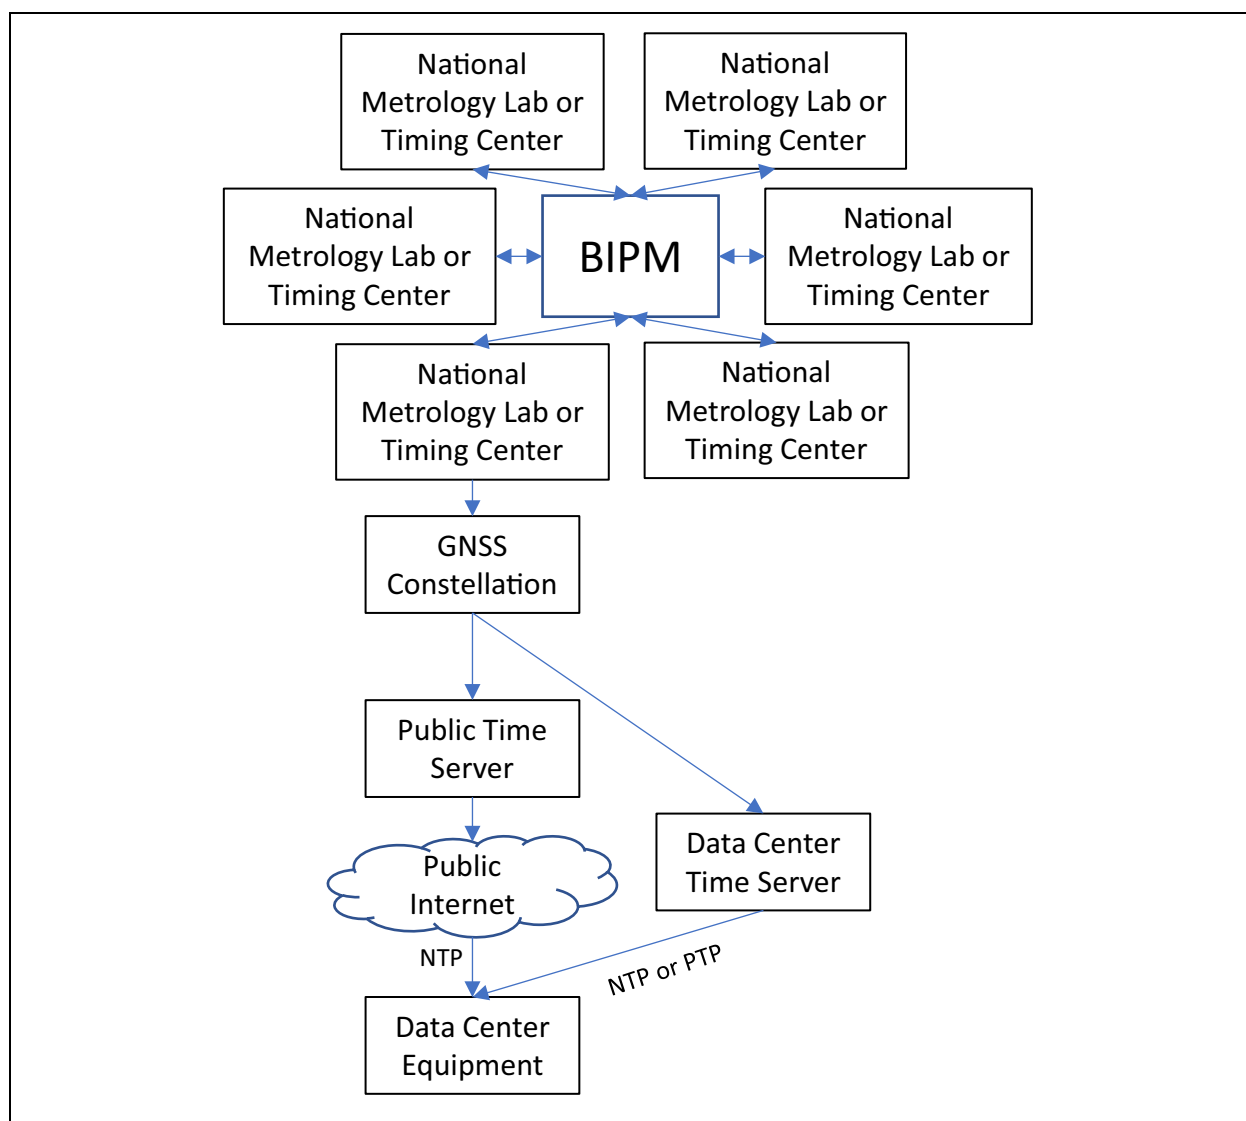


FIGURE 4: Typical paths for data center timing. In both paths time is derived from UTC published by BIPM and transmitted to a GNSS constellation. Public or private time servers then distribute the time to the data center equipment through public Internet or local area networks.

Both paths in [Figure 4](#) utilize GNSS networks to synchronize to UTC. The path on the left utilizes public network time servers for the step between the GNSS network and the data center. Public Network Time Protocol (NTP) servers distribute UTC time to the public through packet-based messages over the Internet at no fee. Published CUs for different servers vary from 100 μ s to 10 ms. This is acceptable for an external clock employed for synchronizing distributed systems that do not require high EPS (<50), security concerns are not critical, regulatory requirements for precision are relatively loose or nonexistent, and irresolvable inconsistencies caused by outlier timestamps are an acceptable risk. However, there are significant risks associated with the use of public NTP servers:

- Accuracy Risks:
 - Lack of assurance that the time is correct on Internet time servers.
 - Lack of disclosure of clock source or change in clock source for many Internet time servers.
 - Asymmetric delays in the local area network, through the firewall, and to the public time server.
- Reliability Risks:
 - Internet time servers will not send an SNMP trap or alarm when a fault has occurred.
 - The quantity of requests on an Internet time server can vary greatly and result in additional accuracy degradation during heavy traffic periods.
 - Lack of assurance of sustained responses from the Internet time server.
- Security Risks:
 - Necessary to open firewall port 123 to 2-way traffic
 - Control of public time server time source.
 - Lack of authentication methods for Internet time servers making it possible for a man-in-the-middle to manipulate the time stamps.

The path on the right in [Figure 4](#) bypasses public NTP and the Internet. This not only eliminates all the risks listed above, but also provides higher EPS values. A private time server directly acquires its time from the GNSS constellations and distributes time within the data center over a LAN as a private NTP server or Precision Time Protocol (PTP) grandmaster. Microchip offers several solutions in different form factors to address these needs. The Microchip SyncServer[®] S600 High Performance NTP Time Server [5] and the TimeProvider[®] TP4100 PTP Grandmaster [6] are both solutions that provide accurate time and distribute it to a secured local area network in a 1U rack-mount solution. These feature-rich solutions include an embedded 72 channel global navigation satellite system (GNSS) receiver that derives time from GPS, Galileo, GLONASS, Beidou, and other available constellations.




FIGURE 5: The Microchip TimeProvider TP4100 PTP Grandmaster and the SyncServer S600 High Performance NTP Time Server.

Both products provide a local time source and distribute this time through packet-based messages over a network. However, the local time source is only one contributor to the final accuracy of the timestamps at the actual equipment. Other factors that affect the timestamp accuracy include the number of routers and switches between the time server and time clients, and the hardware and software used for timestamping. When considering all factors, systems ultimately can degrade timestamp accuracy from the time source by an additional 20 ns to 5 μ s. Including system inaccuracies, the right hand side of [Equation 3](#) can be rewritten as [Equation 4](#).

EQUATION 4:

$$EPS = \frac{1}{2(CU_C + CU_N)}$$

Where:

CU_C = The clock uncertainty of the dedicated clock reference.

CU_N = The clock uncertainty of the network equipment and other hardware.

Many systems strive for final timestamp accuracy in the sub microsecond range, allowing a timestamp generated in a server in Ashton VA, San Jose CA, or Sydney, Australia and all the equipment in the distributed application to have sub microsecond accuracy not only with respect to UTC, but also with respect to each other.

Both the S600 and the TP4100 are designed to support ITU Primary Reference Time Clock (PRTC) requirements. The ITU specifications are maximum values rather than typical or RMS values. [Table 2](#) provides a comparison for the S600, TP4100, and public NTP for a state-of-the-art (SOTA) local area time network where all other sources of inaccuracy contribute 100 ns to the final timestamp accuracy at the time client. It also provides the associated EPS.

TABLE 2: COMPARISON OF PERFORMANCE OF VARIOUS EXTERNAL CLOCK SOLUTIONS

Reference Clock Type	Enabling Technology	Source Accuracy to UTC	Timestamp Accuracy SOTA Time Network	SOTA EPS	S600	TP4100
Public NTP	Packet-based time over public Internet	2.00E+06 ns	2.00E+06 ns	250	—	—
PRTC-A	Single band GNSS receiver	100 ns	200 ns	2.50E+06	X	X
PRTC-B	Single band GNSS receiver	40 ns	140 ns	3.57E+06	—	X
ePRTC	Multiband GNSS receiver and cesium atomic clock	30 ns	130 ns	3.85E+06	—	X

A common issue with specification interpretations arises when data sheets list <15 ns RMS for the source accuracy. This is an RMS value equivalent to one standard deviation, not a maximum value as shown [Table 2](#) or listed in the ITU specifications.

In addition to the rack mount solutions, Microchip offers small form factor solutions, including the Tekron brand TTM 01-G DIN-rail mounted Compact GNSS Clock [\[7\]](#) that offers NTP and PTP functionality with a 32 channel GPS and GLONASS receiver. This product, while not fully PRTCA capable, can provide typical values of less than 100 ns.



FIGURE 6: *Microchip Tekron TTM 01-G DIN Rail-Mounted Compact GNSS Clock.*

The left hand of Equation 3 (WTPS) is the system requirements. The right hand of Equation 3 (EPS) is the equipment capacity. each envelope can only contain one weighted transaction at most. [Figure 7](#) compares the WTPS of the various applications listed in [Table 1](#) and the equipment EPS values listed in [Table 2](#). When designing a system, the red bar for the specific application (WTPS) must be less than the black bar for the capacity of the system (EPS). The graph illustrates that public NTP does not provide a high enough EPS for the military, financial, and trading applications. For the other four applications, public NTP provides an adequate EPS. However, it may be an inappropriate choice for security or other reasons.




FIGURE 7: Application WTPS Requirements vs. Equipment EPS Capability.

External Precision Time Security Solutions

All the clocks mentioned above rely on GNSS receivers to provide the time traceability to UTC. Due to the very low power of GNSS signals, they are susceptible to jamming, which occurs when the GNSS signal cannot be read by the receiver due to energy from another signal. The sources of jamming signals can be either intentional by an individual with malicious intent or unintentional by signals operating in similar frequency ranges. GNSS signals can also be intentionally spoofed by transmitting a simulated GNSS signal with malicious intent. GNSS spoofing and jamming are identified as major threats to critical infrastructure by the US Cybersecurity and Infrastructure Security Agency (CISA). Executive Order 13905 [8] was issued in 2020 requiring US Federal Agencies and Critical Infrastructure operators to take actions to mitigate risks related to spoofing and jamming. Jamming and spoofing events are occurring with increasing frequency as the technology for both becomes widely available. During these events, also referred to as GNSS denial, the NTP server or PTP grandmaster will drift away from UTC, thus increasing the clock uncertainty.

If spoofing remains undetected, it becomes an even greater risk to security as third parties could intentionally manipulate time at a data center location, allowing unauthorized access to data, and possibly resequencing data.

Microchip's Bluesky™ GNSS Firewall [9] hardware solution detects spoofing and jamming and provides either a simulated GPS signal during these outages or squelches the signal from the receiver before the time is transmitted to the network. It can also be configured to pass through alternate time sources during these events.



FIGURE 8: Microchip's Bluesky™ GNSS Firewall.

In addition to the hardware solution, the TP4100 and S600 are both available with similar anti-jamming, anti-spoofing software, including the recent "Bluesky Inside" software release.

External Precision Time Resiliency Solutions

When an external precision time source is distributed securely to a data center, the EPS is dependent upon the accuracy of the external time source. However, once the external time source is lost, that data center must rely on an internally located clock or an alternate external time source. The external clock could be lost for any number of reasons: an Inter-

net outage for public NTP, a solar flare for GNSS-based time, or something as simple as an improperly grounded antenna transmitting a signal that unintentionally jams GPS signals at nearby antennas. Once the outages occur, the local timing will drift from the external time source, resulting in increased uncertainty and a decrease in EPS. This rapidly increases the probability of irresolvable inconsistencies and can quickly jeopardize an entire global database because the sites are no longer synchronized. The amount of drift will depend upon the stability of the local clock source available. Microchip offers a number of quartz- and rubidium-based embedded clock options for the S600 and TP4100, and offers external cesium atomic standards clocks including the 5071A [10] and the Time Cesium 4310B [11] that can be used as alternate time references for the time sources.



FIGURE 9: Microchip 5071A and 4310B Cesium References.

Figure 10 shows the EPS for several Microchip solutions with different clock types while locked to GNSS, along with the EPS degradation after 24 hours of GNSS denial, and after one week of GNSS denial for a SOTA time network that adds 100 ns of inaccuracy to the timestamps (Note 2). Public NTP at an accuracy of 2 ms with a local crystal oscillator back-up has also been added to the chart for reference. The chart illustrates that, while locked to GNSS, all the Microchip solutions provide three orders of magnitude of improvement in EPS compared to public NTP. The impact of the different local clocks on EPS is also clearly evident, with local TCXOs quickly dropping to public NTP level of performances when not locked, whereas a cesium reference offers almost no degradation in EPS for a full week. Data center architects must determine the impact of the decreased EPS and design in the appropriate references. In conjunction, many architects may choose to still have a public NTP or another backup external time source during GNSS denial.




FIGURE 10: EPS for various different external clock options during GNSS disciplining and during GNSS denial.

Figure 11 and Figure 12 compare the application WTPS vs. EPS in unlocked (GNSS denied) operation for 24 hours and one week. Both graphs indicate that, while unlocked, the choice of the local oscillator is critical to ensuring that the application WTPS requirements are met. After only one day of GNSS loss, the banking application will require at least an OCXO, and the trading and military applications require a rubidium. After one week, most of the applications require an atomic clock.




FIGURE 11: Application WTPS requirements vs. 24 hours unlocked equipment EPS capabilities.




FIGURE 12: Application WTPS requirements vs. 1 week unlocked equipment EPS capabilities.

Virtual Primary Reference Clocks

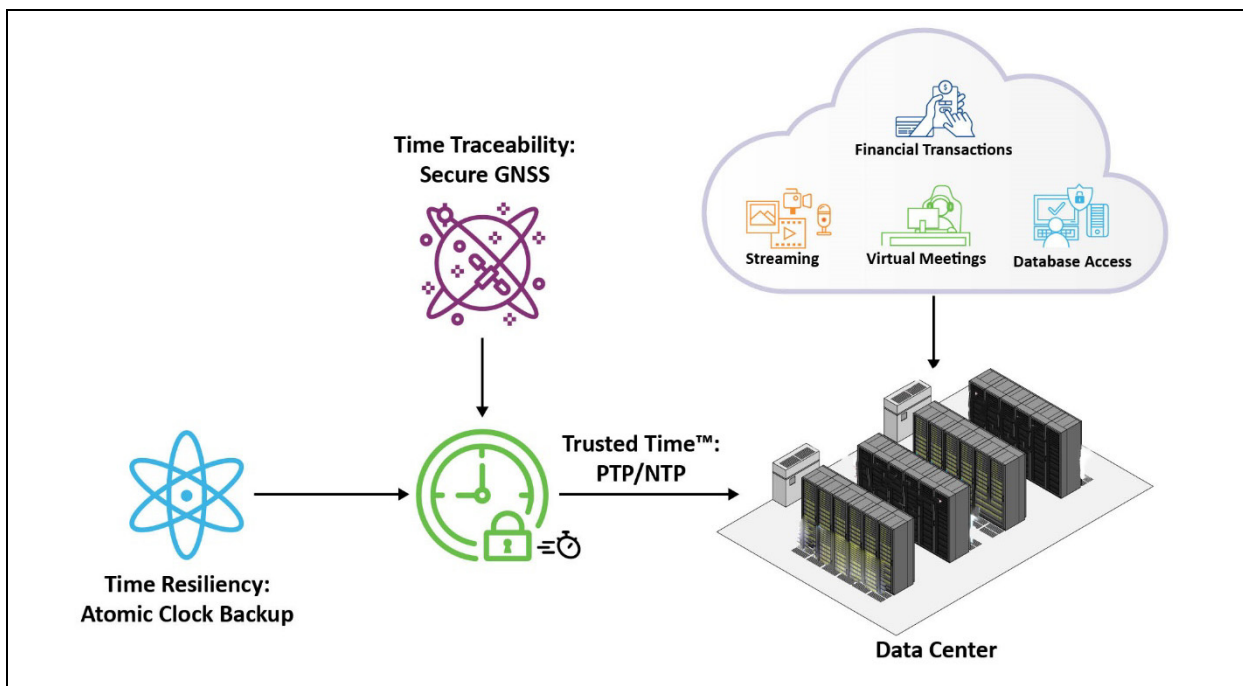


FIGURE 13: Conceptual Diagram of a Virtual Primary Reference Clock.

To meet the demands of distributed system models that require increased capacity and speed, Microchip has developed the “Virtual Primary Reference Clock” to distribute precise, secure, and resilient time in a turnkey solution [12]. By pairing the S600 or TP4100 time server with an embedded rubidium clock or an external cesium clock and securing it with Bluesky Firewall hardware or software, a high precision time source can be located in each data center to increase its overall EPS, enabling millions of TPS to occur without significant consistency risks, even if the external time source is lost for a day or even weeks.

Conclusion

Determining the correct timing architecture to support globally distributed data is dependent upon multiple factors including the transaction per second, the criticality of unresolvable inconsistencies in data across replications, and the probability of records or related records being updated in multiple locations over short time periods. Once these factors are analyzed, the WTPS can be determined and the correct equipment can then be specified to ensure the required level of consistency within the data is achieved.

Notes

- Note 1:** This paper assumes transactions consist of one operation. Most database transactions involve multiple operations. The principles in this paper would still apply, however values would need to be scaled to account for the number of operations in a transaction.
- 2:** Clock accuracies assume constant temperature. Values are either taken from published values, empirical data, or projected values from models. Unlocked values assume the clocks have stabilized and disciplined to UTC for a period of time as specified on data sheets.

References

1. <https://storage.googleapis.com/pub-tools-public-publication-data/pdf/65b514eda12d025585183a641b5a9e096a3c4be5.pdf>
2. <https://www.cs.princeton.edu/courses/archive/spr11/cos461/docs/lec24-strong.pdf>
3. <https://www.esma.europa.eu/document/guidelines-transaction-reporting-order-record-keeping-and-clock-synchronisation-under-mifid>
4. <https://www.nyse.com/publicdocs/nyse/markets/nyse/rule-interpretations/2017/NYSE%20Number%2017-02%203.8.pdf>
5. <https://www.microsemi.com/product-directory/enterprise-network-time-servers/4117-syncserver-s600>
6. <https://www.microsemi.com/product-directory/carrier-grade-ntp-ptp-ieee-1588-grand-masters/4422-timeprovider-4100>
7. <https://tektron.com/products/compact-gnss-clock-ttm01g/>
8. <https://www.federalregister.gov/documents/2020/02/18/2020-03337/strengthening-national-resilience-through-responsible-use-of-positioning-navigation-and-timing>
9. <https://www.microsemi.com/product-directory/gps-instruments/4398-bluesky-gps-firewall>
10. <https://www.microsemi.com/product-directory/cesium-frequency-references/4115-5071a-cesium-primary-frequency-standard>
11. <https://www.microsemi.com/product-directory/cesium-frequency-references/4114-csiii-model-4310b>
12. <https://www.microchip.com/en-us/products/synchronization-and-timing-systems/virtual-primary-reference-time-clock>

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods being used in attempts to breach the code protection features of the Microchip devices. We believe that these methods require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Attempts to breach these code protection features, most likely, cannot be accomplished without violating Microchip's intellectual property rights.
- Microchip is willing to work with any customer who is concerned about the integrity of its code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is "unbreakable." Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Information contained in this publication is provided for the sole purpose of designing with and using Microchip products. Information regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications.

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL OR CONSEQUENTIAL LOSS, DAMAGE, COST OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AnyRate, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maxStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PackTime, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, FlashTec, Hyper Speed Control, HyperLight Load, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, WinPath, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, Inter-Chip Connectivity, JitterBlocker, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SMART-I.S., storClad, SQI, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2021, Microchip Technology Incorporated, All Rights Reserved.

ISBN: 978-1-5224-8919-1

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS

Corporate Office
2355 West Chandler Blvd.
Chandler, AZ 85224-6199
Tel: 480-792-7200
Fax: 480-792-7277
Technical Support:
<http://www.microchip.com/support>
Web Address:
www.microchip.com

Atlanta
Duluth, GA
Tel: 678-957-9614
Fax: 678-957-1455

Austin, TX
Tel: 512-257-3370

Boston
Westborough, MA
Tel: 774-760-0087
Fax: 774-760-0088

Chicago
Itasca, IL
Tel: 630-285-0071
Fax: 630-285-0075

Dallas
Addison, TX
Tel: 972-818-7423
Fax: 972-818-2924

Detroit
Novi, MI
Tel: 248-848-4000

Houston, TX
Tel: 281-894-5983

Indianapolis
Noblesville, IN
Tel: 317-773-8323
Fax: 317-773-5453
Tel: 317-536-2380

Los Angeles
Mission Viejo, CA
Tel: 949-462-9523
Fax: 949-462-9608
Tel: 951-273-7800

Raleigh, NC
Tel: 919-844-7510

New York, NY
Tel: 631-435-6000

San Jose, CA
Tel: 408-735-9110
Tel: 408-436-4270

Canada - Toronto
Tel: 905-695-1980
Fax: 905-695-2078

ASIA/PACIFIC

Australia - Sydney
Tel: 61-2-9868-6733

China - Beijing
Tel: 86-10-8569-7000

China - Chengdu
Tel: 86-28-8665-5511

China - Chongqing
Tel: 86-23-8980-9588

China - Dongguan
Tel: 86-769-8702-9880

China - Guangzhou
Tel: 86-20-8755-8029

China - Hangzhou
Tel: 86-571-8792-8115

China - Hong Kong SAR
Tel: 852-2943-5100

China - Nanjing
Tel: 86-25-8473-2460

China - Qingdao
Tel: 86-532-8502-7355

China - Shanghai
Tel: 86-21-3326-8000

China - Shenyang
Tel: 86-24-2334-2829

China - Shenzhen
Tel: 86-755-8864-2200

China - Suzhou
Tel: 86-186-6233-1526

China - Wuhan
Tel: 86-27-5980-5300

China - Xian
Tel: 86-29-8833-7252

China - Xiamen
Tel: 86-592-2388138

China - Zhuhai
Tel: 86-756-3210040

ASIA/PACIFIC

India - Bangalore
Tel: 91-80-3090-4444

India - New Delhi
Tel: 91-11-4160-8631

India - Pune
Tel: 91-20-4121-0141

Japan - Osaka
Tel: 81-6-6152-7160

Japan - Tokyo
Tel: 81-3-6880-3770

Korea - Daegu
Tel: 82-53-744-4301

Korea - Seoul
Tel: 82-2-554-7200

Malaysia - Kuala Lumpur
Tel: 60-3-7651-7906

Malaysia - Penang
Tel: 60-4-227-8870

Philippines - Manila
Tel: 63-2-634-9065

Singapore
Tel: 65-6334-8870

Taiwan - Hsin Chu
Tel: 886-3-577-8366

Taiwan - Kaohsiung
Tel: 886-7-213-7830

Taiwan - Taipei
Tel: 886-2-2508-8600

Thailand - Bangkok
Tel: 66-2-694-1351

Vietnam - Ho Chi Minh
Tel: 84-28-5448-2100

EUROPE

Austria - Wels
Tel: 43-7242-2244-39
Fax: 43-7242-2244-393

Denmark - Copenhagen
Tel: 45-4485-5910
Fax: 45-4485-2829

Finland - Espoo
Tel: 358-9-4520-820

France - Paris
Tel: 33-1-69-53-63-20
Fax: 33-1-69-30-90-79

Germany - Garching
Tel: 49-8931-9700

Germany - Haan
Tel: 49-2129-3766400

Germany - Heilbronn
Tel: 49-7131-72400

Germany - Karlsruhe
Tel: 49-721-625370

Germany - Munich
Tel: 49-89-627-144-0
Fax: 49-89-627-144-44

Germany - Rosenheim
Tel: 49-8031-354-560

Israel - Ra'anana
Tel: 972-9-744-7705

Italy - Milan
Tel: 39-0331-742611
Fax: 39-0331-466781

Italy - Padova
Tel: 39-049-7625286

Netherlands - Drunen
Tel: 31-416-690399
Fax: 31-416-690340

Norway - Trondheim
Tel: 47-7288-4388

Poland - Warsaw
Tel: 48-22-3325737

Romania - Bucharest
Tel: 40-21-407-87-50

Spain - Madrid
Tel: 34-91-708-08-90
Fax: 34-91-708-08-91

Sweden - Gothenberg
Tel: 46-31-704-60-40

Sweden - Stockholm
Tel: 46-8-5090-4654

UK - Wokingham
Tel: 44-118-921-5800
Fax: 44-118-921-5820