

LEVERAGING THE AIR FORCE ENTERPRISE APPROACH TO TIME SYNCHRONIZATION: A PROPOSAL

Glenn Bell and Jay Scarano
The MITRE Corporation
Bedford, MA 01730, USA
E-Mail: glennb@mitre.org and jgs@mitre.org

Abstract

Net-centric warfare requires an unprecedented level of integration. Time synchronization is a key driver for that integration. The U.S. Naval Observatory (USNO) provides PTTI capability for the Department of Defense (DoD). To date, we do not have a comprehensive method for distributing time from the USNO to systems with varying precision and accuracy requirements. The Air Force, as part of its Common Integrated Infrastructure initiative, has taken on this problem. The result is an Air Force Enterprise Time architecture that encompasses the USNO provided delivery mechanisms and provides the framework for delivery of pedigreed time to end systems. We offer the Air Force approach to achieving enterprise time delivery as a starting point for the DoD. Further details and explanations of CII Services can be obtained by contacting the authors directly.

THE TIME CHALLENGE

Net-Centric Warfare (NCW) has significant implications for military operations. NCW will change “*the nature of our mission(s), the battlespace in which we operate, our adversaries’ capabilities, our ability to sense and understand the battlespace, the capability of the weapons at our disposal, and ... our ability to command and control*” [1]. Key to NCW is information superiority, defined in Joint Publication 3-13 as “*the capability to collect, process, and disseminate an uninterrupted flow of information while exploiting or denying an adversary’s ability*” [2]. Information superiority is viewed as a capability, analogous to air superiority or sea control, which increases the effectiveness of offensive and defensive operations. It enables the compression of timelines, allowing the DoD to operate more effectively.

The increased operational tempo introduced by NCW significantly increases the complexity of the supporting systems. All branches of the government must act in concert with the joint and coalition community. Accordingly, the supporting systems must also interact in complex and often unpredictable ways. Achieving this unprecedented level of cooperation requires a set of core functions common to all systems. These common functions, or Enterprise Services, provide consistency and coherence across all systems. Time synchronization is one foundational service essential for establishing a common reference for all systems and organizations within the Department of Defense.

AN APPROACH TO ENTERPRISE SERVICES

The Air Force Network Time Service was born out of a study conducted jointly between the Defense Information Services Agency (DISA) and the U.S. Air Force Electronic Systems Center (ESC). The study identified many computing capabilities for consideration as DoD enterprise services. The ability to synchronize computers on the network was one of the services proposed.

In 2002, the U.S. Air Force Electronic Systems Center, at Hanscom AFB, developed the Command and Control Enterprise Reference Architecture (C2ERA) to “supply a technical design pattern to program office, contractor and user architects and developers with the goal of guiding and constraining key system integration and interoperability decisions. This architecture document describes an enterprise architectural plan, one that subscribes to common standards, utilizes pervasive commercial technologies and is based on a computing services-oriented approach.” [3, p. V]

The C2ERA introduced an engineering approach to building highly integrated systems and services. The key components of the reference architecture are the Common Integrated Infrastructure (CII), a set of “smart, adaptive services that allow users to rapidly access, manipulate and display trusted data in a changing environment” [3, p. 15], and the C2 Node, which “provides common local services such as messaging and distributed transaction support ... reducing the effort required to integrate modernized and newly developed applications” [3, p. 35]. The CII built on the original DISA/ESC investigation.

Figure 1 illustrates the vision of an integrated infrastructure supporting the DoD and integrating into Coalition Communities as well as other government and non-government networks.

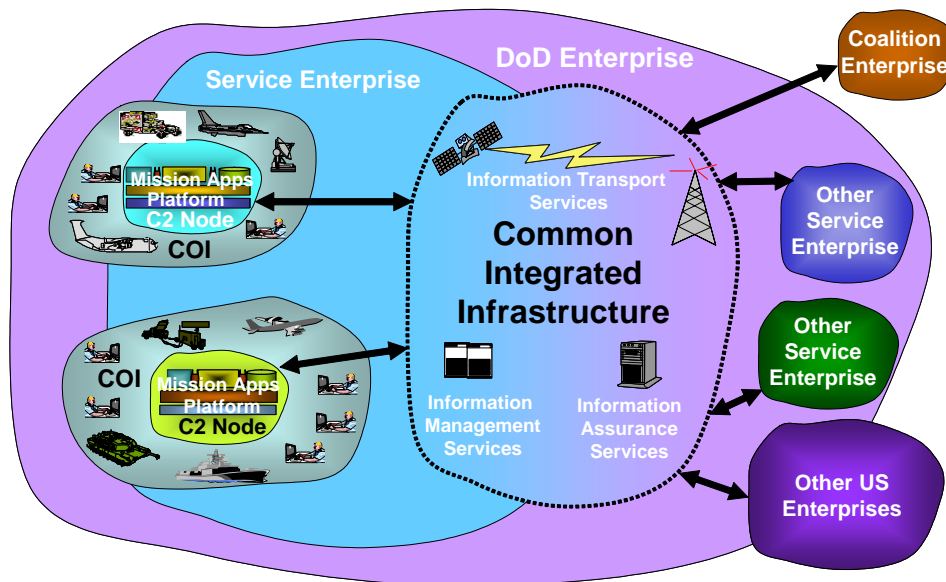


Figure 1. C2ERA Vision.

COMMON INTEGRATED INFRASTRUCTURE

The Common Integrated Infrastructure (CII) is a set of enterprise services that are used by all systems and are the building blocks upon which other services are based. The following criteria were used to identify CII services:

- Is this a common utility (“service”) essential for enabling operational capability across the Enterprise?
- Is Enterprise control of the service required to ensure successful implementation?
- Does the service scale to support the Enterprise?
- Is Enterprise content, consistency, or connection required? That is, must the service content be complete, does the service behavior need to be the same every time it’s accessed, or will the service be required to accept connections from across the enterprise?
- Is a single service specification possible (or practical)?

The set of attributes that we believe should be common to all enterprise services is:

- Usable by all systems and nodes (with development tool support)
- Available across the enterprise (millions of users in diverse locations)
- Accessible via a single service specification (a least common denominator interface available to all users)
- Well-defined quality of service measures (e.g., reliability, performance)
- Well managed (24x365 support, worldwide)
- Potential for multiple service/business models (few, if any, enterprise services are provided by a single organization).

The implications of these attributes are that CII Services are constrained to those services that provide capabilities to most, if not all, systems participating in the enterprise. Accordingly, the services must be ubiquitous and will utilize a variety of deployment, operations, and management models. As enterprise services will be used by a wide variety of applications – most unknown to the provider – the service must be highly available with published service levels and around-the-clock support.

By applying the enterprise service criteria to the original DISA/ESC list, an initial list of CII Services was developed:

- Network Time Service (NTS)
- Domain Name System (DNS)

- Communications Transport (CT)
- Privilege Management Infrastructure (PMI)
- Authentication
- Key Management Infrastructure (KMI)
- Directory Services (DS)
- Information/Service Broker
- Messaging Services
- Connecting Links to Network
- Voice and Voice over IP
- Enterprise Security/System/Network Management Services (ESSNM).

The list of services is not definitive. As the enterprise matures, some services may turn out to be inappropriate or immature, and some new services may need to be added.

The CII team suggested the following approach to implementing CII services:

- Describe the service (identifying requirements, available technologies, and potential service providers)
- Architect the service (providing technology-independent implementation guidance)
- Deploy and operate the service (leveraging existing programs where possible, and integrating capabilities across providers).

The CII Team selected the Network Time Service (NTS) as one of two pilots for exercising and evolving the enterprise services process. NTS was selected due to several factors: time-related standards were very mature; there is a responsible organization within the DoD; there are many commercial implementations available; implementation costs were expected to be negligible; and a case for time synchronization requirements could be easily made. In short, we thought that the path to an enterprise network time service would be easy. The Domain Name System, the second pilot, was chosen for similar reasons.

THE AIR FORCE NETWORK TIME SERVICE

Several challenges beset the establishment of a common time service.

DEVELOPING THE TIME SERVICE – FIRST PASS

Initial work on NTS was dual-focused: identifying enterprise time synchronization requirements; and studying time distribution mechanisms.

The requirements collection activity fell short of expectations. Many programs do not understand the need for time synchronization, many programs do not have definitive analysis of synchronization requirements, and, where requirements were provided, there were widely divergent accuracy and precision needs. Table 1 details some of the accuracy requirements presented to the team.

Table 1. Time Service Accuracy Requirements.

Function or Service	Identified Requirement	Normalized (Seconds)
Windows [™] Domain	5 minutes	300
Information Assurance (logging)	1 second	1
Communications	500 ns	500×10^{-9}
Correlation	100 ns	100×10^{-9}

DoD and Air Force regulations with respect to Precise Time and Time Interval (PTTI) are not universally understood, referenced, or enforced. Additionally, there are few organizations responsible for providing time services to edge users. These were problems the Air Force attempted to address when describing, architecting, and deploying NTS.

DELIVERING A TIME SERVICE

Initially, the Network Time Protocol (NTP) was selected as a primary delivery mechanism for the Air Force. An NTP-based solution was proposed for several reasons:

- NTP is an Internet Protocol (IP)-based service that provides for millisecond accuracy across the wide area network and micro/nano second accuracies on local area networks.
- Virtually all modern computer operating systems support the protocol.
- There are well-managed NTP servers hosted on the principal DoD networks.
- The protocol includes redundancy and failover capabilities.
- NTP is a mature, well-understood standard.

A white paper was prepared and submitted to stakeholders for comment. The paper was a mix of service description, architecture views, and implementation guidance. This original NTP proposal was met with resistance from the user and provider communities for several reasons:

- Programs were not yet ready to be identified as Enterprise Service providers. They would not be able to support additional burden within their existing funding profile.
- The NTP solution did not support high-accuracy requirements.
- NTP use across the enterprise would create security issues. Air Force policy directs that Network Control Centers “*will not allow external NTP sources through the BIP boundary due to inherent security problems*” [4].
- Programs were not ready to accept wholesale changes
 - The paper recommended that Simple Network Time Protocol (SNTP)-based systems (i.e., Microsoft[™] Windows) be upgraded to NTP. This change would impact all Microsoft-based systems, including desktops and laptops, in the enterprise.
 - Specific versions of NTP software were stipulated to address security concerns. This would have required upgrading most NTP-based systems (i.e. Solaris, HP-UX).

As an architecture and implementation guidance, the white paper was not very successful. However, as a service description, the paper made time synchronization understandable to the Air Force community. The lesson learned was that before attempting to define a generic architecture for any given enterprise service, ensure that a clear description is developed, vetted, and accepted by the stakeholders.

Based on this lesson and all the feedback received, the next step was to architect NTS.

DEFINING AN ENTERPRISE SERVICE PROFILE

The task of architecting enterprise services posed some challenges of its own. The system engineering and architecture tools available are primarily designed to develop “stovepipe” systems, systems that deliver all required functionality as a single entity, with little dependence or interoperability with others. Additionally, much of the traditional system engineering wisdom does not directly apply to developing enterprise services. The available tools and knowledge must be carefully reviewed for applicability and shortcomings.

The DoD Architectural Framework (DoDAF) was selected as the modeling tool for the architecture, not because it addressed all the “Enterprise” concerns, but rather because it provided a familiar construct for the DoD architecting community. Since the DoDAF is essentially a system-focused architecture tool, the objective of each DoDAF component was reviewed and applied to the question of defining “Enterprise.” Supplemental products were developed to satisfy any deficiencies identified. One significant shortcoming was the inability to portray dependencies between enterprise services, leading to the development of an enterprise-specific view.

A Service Profile template was developed, based upon the DoDAF, as a vehicle for describing the architecture. The format was designed to be flexible enough to change as architecture techniques were explored, refined, and, sometimes, discarded. The Service Profile is considered a living document, but

generally is designed to describe only those standards and methods necessary to ensure service consistency and cohesion across the Enterprise.

THE AIR FORCE NETWORK TIME SERVICE PROFILE

“NTS is a core Information Technology (IT) service facilitating Net Centric Warfare. NTS facilitates the integration of data exchanges by introducing standardized, synchronized time services throughout the enterprise. Additionally, NTS has Information Assurance (IA) and Information Management (IM) implications that support services in those disciplines” [3, p. 9].

A set of Key Performance Parameters (KPP) was defined and used as a basis for the architecture. The KPPs are:

- Traceability – All systems shall trace their time pedigree back to the root time services, provided by the U.S. Naval Observatory (USNO).
- Accuracy – All time service providers shall maintain the accuracy of their service to within ± 100 ms of UTC as provided by USNO (UTC (USNO)).
- Availability – All time service providers shall utilize DoD provided delivery mechanisms.

The Traceability and Accuracy KPPs are linked to the Chairman of the Joint Chiefs of Staff Instruction (CJCSI) 6130.01 Master Navigation and Timing Plan [5], which directs that *“Any information that makes reference to time must be able to provide that time in terms of the standard temporal reference defined by Coordinated Universal Time (UTC) as maintained by the US Naval Observatory (USNO) Master Clock, which is the standard for military systems.”* (p. A-1)

The ± 100 ms value delineated in the Accuracy KPP is intended to provide a discrete standard attainable by all Air Force systems without imposing an undue financial or engineering burden upon program offices. 100 ms is the minimum accuracy requirement and all delivery mechanisms detailed in the Air Force NTS Profile support accuracy significantly better than this minimum.

The Availability KPP ensures that delivery mechanisms have clear traceability to USNO and are designed to support the warfighter in times of conflict.

The NTS Profile describes multiple delivery mechanisms, each with specific accuracy, cost, and engineering characteristics that enable programs to make selections based upon the dynamics of their system(s). This is a fundamental shift from the white paper proposal. The white paper proposed an NTP-centric solution which directed unilateral software and process solutions. The NTS Profile proposed five mechanisms; the selection of which one to utilize is left to the implementing program. Additionally, the implementation details (i.e. hardware, software, processes) are left to the implementing program, so long as one of the approved delivery mechanisms are utilized and the KPPs are satisfied.

The Operational Concept for the NTS Architecture is summarized in Figure 1, which depicts the *“High Level Operational Concept (OV-1). The purpose of the OV-1 is to provide a quick, high-level description of what the architecture is supposed to do, and how it is supposed to do it.”* [6, 4-1]

At the core of the Enterprise are the UTC nodes. Initially, USNO and the National Institute of Standards and Technology (NIST) were selected, as shown in Figure 2.

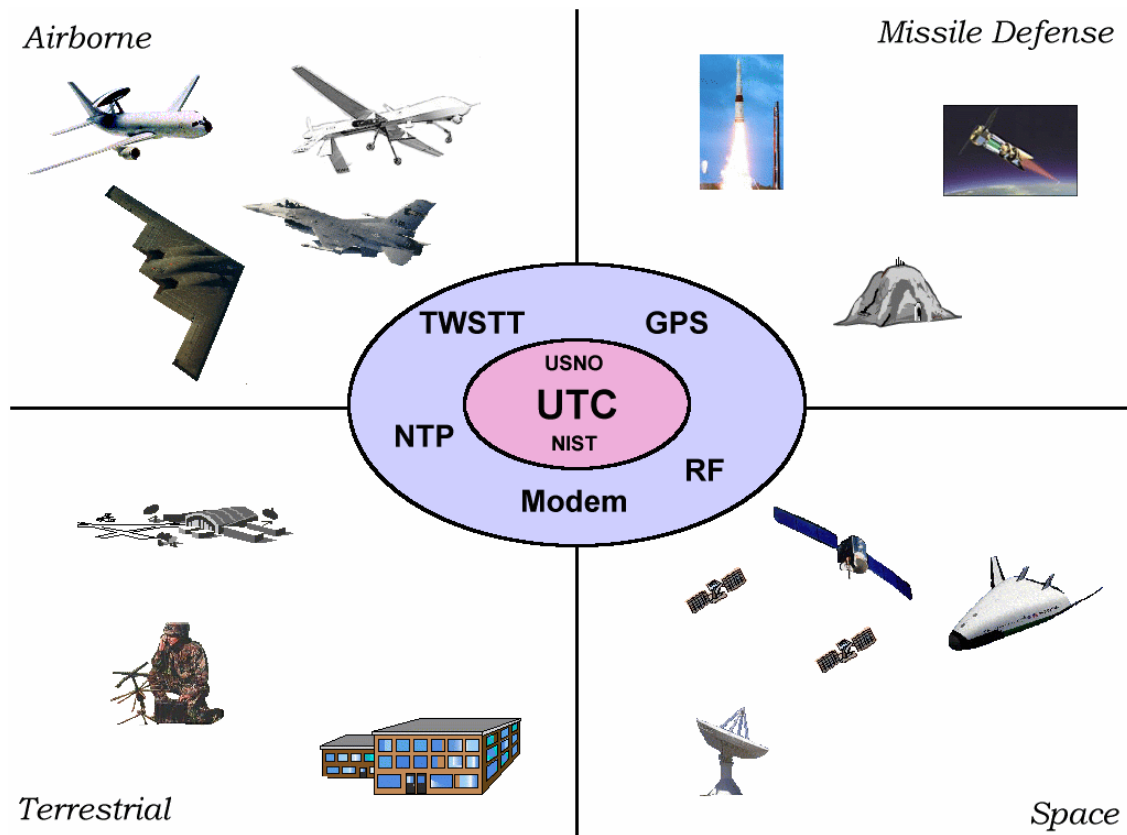


Figure 2. NTS High Level Operational Concept (OV-1).

The inclusion of UTC (NIST) is an indication of the team’s difficulty identifying pertinent regulations. NIST was included in several sections of the NTS Profile. Time services supplied by NIST violate the availability KPP, directing that only DoD-provided mechanisms may be utilized. In the initial version of the NTS Profile, NIST was indicated as a provider for RF-based, NTP, and Modem time services. It has since been determined that NIST services are not appropriate for time synchronization within the DoD. NTS Profile references to UTC (NIST) will be removed in future versions. To maintain the integrity of the profile description below, UTC (NIST) references have been preserved.

Time is delivered from UTC (USNO) using several delivery mechanisms: Two-Way Satellite Time Transfer, Global Positioning System, Radio Frequencies, Network Time Protocol, and Modem. NTS service providers also use these delivery mechanisms to furnish time to the user community. High-precision time devices, such as atomic clocks, may be employed as “flywheels” to provide accurate time during periods when access to UTC (USNO) is impeded.

Several *System (Service) Interface Descriptions (SV-1)* are provided as guides to interface mission application, Command, and Control (C2), Communities of Interest (COIs), and enterprise services with the Enterprise NTS Architecture. The SV-1 is a standard architectural view within the DoD Architectural Framework and depicts the systems and/or nodes and system interfaces. “An interface, as depicted in SV-1, is a simplified, abstract representation of one or more communications paths between systems nodes or between systems.” [6, p. 5-2]

TWO-WAY SATELLITE TIME TRANSFER

Two-Way Satellite Time Transfer (TWSTT) is the most accurate method of delivering time to a node. A dedicated satellite connection to USNO is established to provide highly accurate (up to ± 1 ns) synchronization. TWSTT employs two satellite communications channels to synchronize the reference clocks. Short-term accuracy of better than 1 ns is achievable. Long-term calibration to 1 ns is possible through the use of portable Earth stations.

The *TWSTT System Interface Description SV-1*, shown in Figure 3, shows the key nodes involved in establishing TWSTT-based time synchronization and dissemination. Time is synchronized using ground-based satellite terminals communicating through a communications satellite. The “Intermediate Time Server” identifies the system requiring synchronization. In the case of a Time Service Provider, synchronization services will be provided to supported clients, commonly via the Network Time Protocol. The “Backup Time Source” may be either an alternate time delivery mechanism or a “flywheel” atomic clock.




Figure 3. TWSTT System Interface Description.

All system components for TWSTT are commercially available, as are the satellite communications channels. This solution does not require dedicated space-based equipment, but has a high initial cost for equipment, as well as high recurring costs for the communications channels.

Organizations using TWSTT to synchronize their reference clock(s) may further distribute time using one of the approved delivery mechanisms. Additionally, service providers should employ a backup time source to ensure service availability.

GLOBAL POSITIONING SYSTEM

The Global Positioning System (GPS) is managed by the Department of Defense and provides a precise geo-positioning and highly accurate (± 200 ns) time synchronization service to any user worldwide, free of charge. GPS employs a constellation of at least 24 medium earth-orbit satellites, each synchronized with UTC (USNO). Each satellite in the constellation carries redundant atomic clocks to ensure precise timekeeping and availability.

Figure 4 illustrates the *GPS System Interface Description (SV-1)*. The USNO Master Clock node steers the time of the GPS Constellation. The Intermediate Time Server employs a GPS receiver to synchronize its time, which may be further distributed to subordinate nodes.

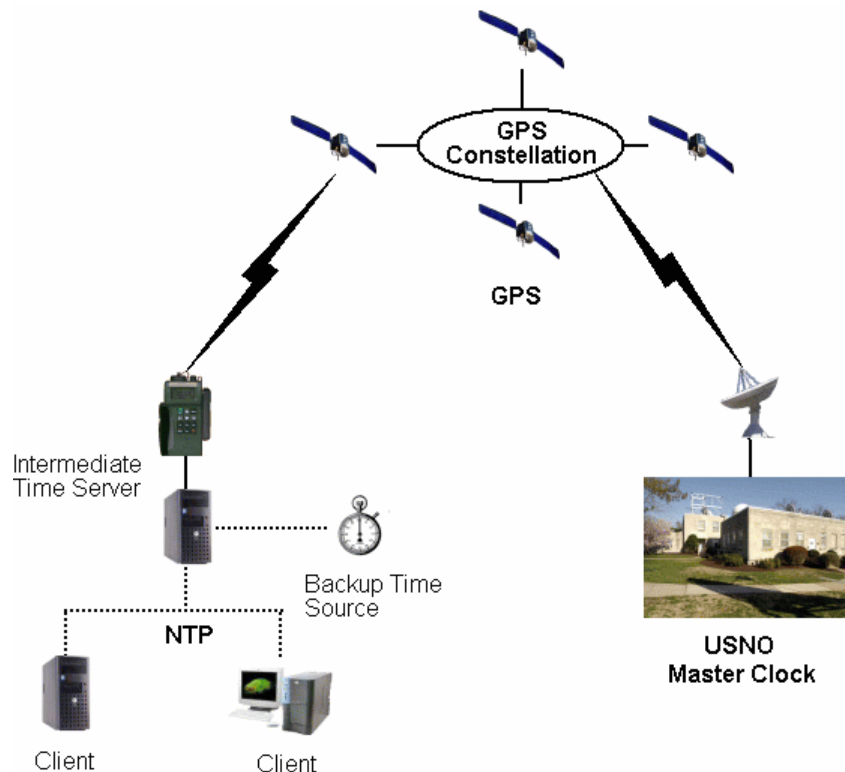


Figure 4. GPS System Interface Description.

GPS operates in two modes. The Standard Positioning System (SPS) provides time accuracies to within 340 ns, and is provided for most non-military applications. SPS may be degraded or disabled regionally to deny its use during period of conflict.

The Precise Positioning System (PPS) provides time accuracies to within 200 ns, and is provided for military applications and is protected via cryptography. “Systems used for combat, combat support, or combat service support missions must use PPS capable GPS receivers operating in the Precise Positioning System (PPS) mode.” [2, p. E-3]

Both Standard Positioning and Precise Positioning System-based equipment is commercially available and provided by multiple vendors. GPS is very flexible and available, but is subject to electronic countermeasures. As in other delivery methods, service providers should employ backup time sources and/or “flywheel” clocks to ensure availability.

RADIO FREQUENCY

A time service via radio frequencies (RF) is a classic method of dissemination and synchronization. The National Bureau of Standards (now National Institute for Standards and Technology (NIST)) established the first shortwave broadcast of time and frequency standards in 1923. Since that time, NIST has increased RF-based time services to support the continental US and Hawaii.

While NIST is not an approved master clock for DoD systems, the RF-based delivery mechanism will be retained in the NTS Profile as a guideline for delivering DoD RF-based time services.

The RF System Interface Description (SV-1), as shown in Figure 5, shows NIST as the master clock node. The reference clock, synchronized to UTC (USNO), transmits a Time Synchronization Message on a regular schedule. Any node within the broadcast region uses an RF receiver to synchronize the Intermediate Time Server node. The synchronized clock on the Intermediate Time Server may then be used as a reference clock for others. If RF is selected as a primary delivery mechanism, a backup mechanism or flywheel clock is an important consideration.

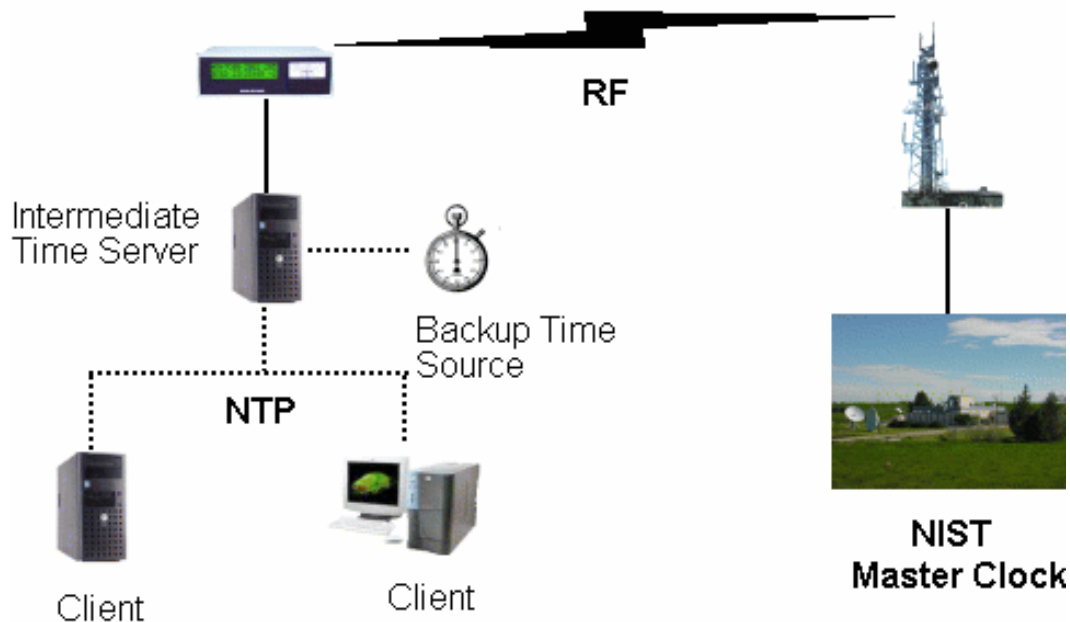


Figure 5. Radio Frequency System Interface Description (SV-1).

RF receivers are commercially available and support synchronization accuracy to ± 1 ms. RF time synchronization provides regional service only and is susceptible to electronic countermeasures as well as natural events, such as weather and solar activity.

RF is recommended as a backup delivery mechanism and as a means of delivering time synchronization to edge users with limited network connectivity.

NETWORK TIME PROTOCOL

The Network Time Protocol (NTP) is defined in IETF RFCs 1305 [7] and 2030. First proposed in 1985, NTP is designed to provide time synchronization across an IP-based network. NTP “*utilizes a... design in which a distributed subnet of time servers operating on a self-organizing, hierarchical-master-slave configuration synchronizes local clocks within a subnet and to national time standards.*” [7, p. 1]

The NTP System Interface Description (SV-1), depicted in Figure 6, identifies USNO as the Master Clock, which provides several NTP time servers, synchronized to UTC (USNO) on principal DoD networks. Intermediate Time Server nodes connect to the USNO time servers via UDP/IP on port 123. Unsynchronized servers query the USNO time server(s) approximately once per minute. As synchronization improves, the query rate decreases gradually to about once every 17 minutes when fully synchronized. Intermediate Time Server nodes may further distribute time services, normally using NTP. Network Time Service Providers, and other systems requiring high availability, should employ a backup time source to ensure continuous operations.




Figure 6. NTP System Interface Description (SV-1).

The NIST Master Clock should not be used for DoD time synchronization.

The NTP delivery mechanism is inexpensive and provides good accuracy (± 30 ms across wide area networks, < 1 ms on local area networks); however, due to security concerns its availability across DoD and AF firewalls is restricted. NTP requires IP connectivity and, accordingly, is susceptible to network-based attacks.

NTP is commonly used to distribute time, acquired using other delivery mechanisms, such as GPS, within local networks.

MODEM

Telephonic time synchronization utilizing modems is another venerable technique. USNO provides a bank of 1200 baud (8-N-1) modems. The modems are available in Washington, DC at (202) 762-1594 (DSN 762-1594), and in Colorado Springs, CO at (719) 567-6743 (DSN 560-6743).

Figure 7 illustrates the Modem System Interface Description (SV-1). Time services are provided through modem pools. The Intermediate Time Server initiates a call via commercial or DSN telephonic systems. The USNO time server transmits a time string followed by a mark signal. The mark signal is used to identify when the current time matches the time string. This method is used to address latencies in the telephone switching system. The Intermediate Time Server may further distribute the synchronized time. As with other methods, a backup mechanism or flywheel clock should be used to ensure availability.

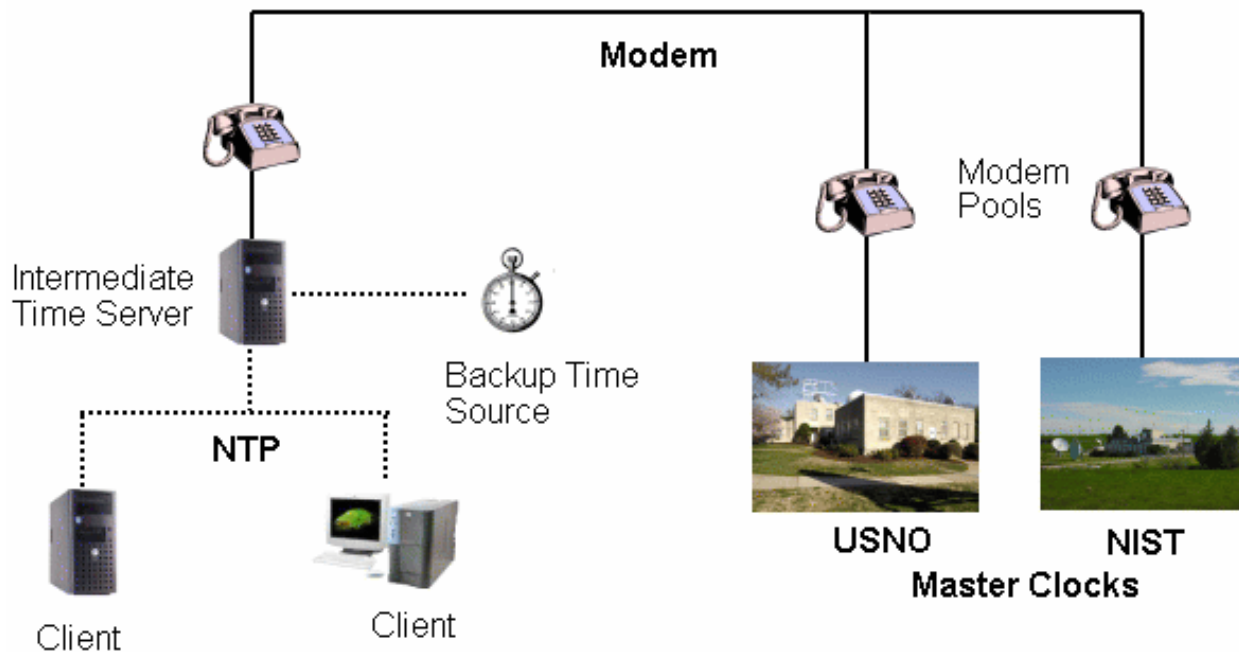


Figure 7. Modem System Interface Description (SV-1).

The NIST master clock should not be used for DoD time synchronization.

The hardware and software required for modem time synchronization is widely available and inexpensive. Telephonic support is near ubiquitous worldwide. The modem delivery method provides ± 45 ms accuracy for a single-sample time message. Using available handshaking algorithms, accuracies < 10 ms are achievable. Modem-based time synchronization is the least accurate method and is dependent upon telephone availability.

The Modem Delivery Mechanism is most appropriate as a backup method to other methods.

THE PROPOSAL

Our approach, presenting several techniques for time synchronization, was ultimately successful in the Air Force community. By offering multiple solutions and describing the implications of each, Air Force Programs can decide which method(s) is best suited to their particular situation.

We offer the Air Force NTS profile for consideration as the starting point for developing time service architectures within the DoD. We believe the strength of the profile is in the process rather than the specific solutions identified for the Air Force. Keys to the process are:

- Develop KPPs, based upon Service and DoD requirements and regulations
- Identify the requirements scope
- What is the minimum and maximum acceptable accuracy to support all programs?
- Avoid the temptation to over-specify
 - Only the standards, mechanisms and processes necessary to achieve consistency and coherence throughout the enterprise should be included
- Ensure delivery methods are compatible across the DoD. Service Time Architectures must be able to integrate into a unified DoD Enterprise.

We believe that the Air Force will participate in activities related to an enterprise architecture for time synchronization, regardless of whether the Air Force NTS profile is selected as the baseline or not.

ACKNOWLEDGMENTS

The authors would like to thank our sponsors: Mr. Tom Powis, Mr. Bert Hopkins, and Dr. Margaret Corasick. We would also like to thank the team that pioneered this activity: Mr. Barry Smith, Dr. George Huff, and Mr. Jason Mathews.

REFERENCES

- [1] David S. Alberts, John J. Garstka, and Frederick P. Stein, *Network Centric Warfare: Developing and Leveraging Information Superiority*, 2nd Edition (February 2000), DoD Command and Control Research Program, p. 53.
- [2] *Joint Publication 3-13 Joint Doctrine for Information Operations* (9 October 1998), Chairman of the Joint Chiefs of Staff (CJCS), p. I-10.
- [3] *C2 Enterprise Reference Architecture v3.0-14* (1 December 2002), U.S. Air Force Electronic Systems Center (ESC).

- [4] *U.S. Air Force Instruction 33-115 Network Management* (15 December 2002), U.S. Air Force Communications Agency (AFCA), p. 16.
- [5] *2003 CJCS Master Positioning, Navigation and Timing Plan CJCSI 6130.01C* (31 March 2003), Chairman of the Joint Chiefs of Staff (CJCS).
- [6] *DoD Architecture Framework version 1.0 Volume I: Definitions and Guidelines* (9 February 2004), DoD Architecture Framework Working Group (DoDAF-WG).
- [7] D. Mills, *Network Time Protocol (Version 3) Specification, Implementation and Analysis* (March 1992), Internet Engineering Task Force (IETF) Request For Comments (RFC) 1305.

The Common Integrated Infrastructure Network Time Service (June 2003), U.S. Air Force Electronic Systems Center (ESC).

U.S. Air Force Network Time Service Profile v1.0 (28 January 2004), U.S. Air Force Electronic Systems Center (ESC).

D. W. Hanson, T. R. Meeker, and L. C. Heishman, 1989, “*Fundamentals of Two Way Time Transfers by Satellite*,” in Proceedings of the 43rd Annual Frequency Control Symposium, 23-25 May 1990, Baltimore, Maryland, USA (IEEE 90CH2818-3), pp. 478-487.

Department of Defense Directive 4650.5 Positioning, Navigation and Timing (2 June 2003), Office of the Deputy Secretary of Defense.

Scott L. Surer, *DoD C4ISR Enterprise Services: A Joint DISA/ESC Initiative Stakeholder Analysis, Version 0.9 Draft* (21 November 2001), The MITRE Corporation.

QUESTIONS AND ANSWERS

DOUGLAS BOEHM (Defense Information Systems Agency): Is there a site we could go to get this information that you have, and then also proposed items to be posted onto the Web site?

GLENN BELL: The architecture itself is on the Air Force ITRIM site, Scott Air Force Base. That is probably not available to everybody, DISA folks.

BOEHM: I have DOT mail.

BELL: Okay, within DOT mail, it is on the Air Force ITRIM Web site. I do not have the URL right now, but I can get it to you.

MARC DAMASHEK (Department of Defense): Glenn, there is a very useful distinction to be made that will head off a lot of the confusion about Web-based versus non-Web-based. That is to figure out which things are inherently synchronous and which things are inherently asynchronous.

The bad lesson that we learned 20-30 years ago is that all the synchronous stuff in early computer code was due to the need for programmer convenience, and Apple and Xerox and so on fixed that. But the pendulum swung in the other direction, and now the assumption is that everything can be made asynchronous.

The stuff we are talking about is inherently synchronous. You cannot get away from that. And that is not a familiar concept to the user community.

BELL: That is absolutely true. I made that comment to a number of people. In defining an enterprise service, the technical part of defining enterprise is about that big (gesture). It is very doable. It is easy to address the technical problems.

The real problem is addressing the social and political climates – you know, why should I give up control to the Enterprise? I currently built it into my system; I am going to give up control and let somebody that I do not know manage this service for me? And that is a real challenge.