Workgroup: Network Time Protocol
Internet-Draft:
draft-gruessing-ntp-ntpv5-requirements-05
Published: 21 May 2022
Intended Status: Informational
Expires: 22 November 2022
Authors: J. Gruessing
Nederlandse Publieke Omroep

NTPv5 use cases and requirements

Abstract

This document describes the use cases, requirements, and considerations that should be factored in the design of a successor protocol to supersede version 4 of the NTP protocol [RFC5905] presently referred to as NTP version 5 ("NTPv5"). This document is non-exhaustive and does not in its current version represent working group consensus.

Note to Readers

RFC Editor: please remove this section before publication

Source code and issues for this draft can be found at https://github.com/fiestajetsam/draft-gruessing-ntp-ntpv5-requirements.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

NTP version 4 [RFC5905] has seen active use for over a decade, and within this time period the protocol has not only been extended to support new requirements but has also fallen victim to vulnerabilities that have been used for distributed denial of service (DDoS) amplification attacks.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Use cases and existing deployments of NTP

There are several common scenarios for existing NTPv4 deployments: publicly accessible NTP services such as the NTP Pool [ntppool] are used to offer clock synchronisation for end users and embedded devices, ISP-provided servers are used to synchronise devices such as customer-premises equipment where reduced accuracy may be tolerable. Depending on the network and path these deployments may be affected by variable latency as well as throttling or blocking by providers.

Data centres and cloud computing providers also have deployed and offer NTP services both for internal use and for customers, particularly where the network is unable to offer or does not require PTP [IEEE-1588-2008]. As these deployments are less likely to be constrained by network latency or power the potential for higher levels of accuracy and precision within the bounds of the protocol are possible.

3. Requirements

At a high level, NTPv5 should be a protocol that is capable of operating in local networks and over public internet connections where packet loss, delay, and filtering may occur. It should be able to provide enough information for both basic time information and synchronisation.

3.1. Resource management

Historically there have been many documented instances of NTP servers receiving large amounts of unauthorised traffic [ntp-misuse] and the design of NTPv5 must ensure the risk of these can be minimised.

Servers SHOULD have a new identifier that peers use as reference, this SHOULD NOT be a FQDN, an IP address, or an identifier tied to a public certificate. Servers SHOULD be able to migrate and change their identifiers as stratum topologies or network configuration changes occur.

The protocol MUST have the capability for servers to notify clients that the service is unavailable, and clients MUST have clearly defined behaviours for honouring this signalling. In addition servers SHOULD be able to communicate to clients that they should reduce their query rate when the server is under high load or has reduced capacity.

Clients SHOULD periodically re-establish connections with servers to prevent attempting to maintain connectivity to a dead host and give network operators the ability to move traffic away from hosts in a timely manner.

The protocol SHOULD have provisions for deployments where Network Address Translation occurs, and define behaviours when NAT rebinding occurs. This should also not compromise any DDoS mitigation(s) that the protocol may define.

3.2. Algorithms

The use of algorithms describing functions such as clock filtering, selection, and clustering SHOULD have agility, allowing for implementations to develop and deploy new algorithms independently. Signalling of algorithm use or preference SHOULD NOT be transmitted by servers.

The working group should consider creating a separate informational document to describe an algorithm to assist with implementation, and consider adopting future documents which describe new algorithms as they are developed. Specifying client algorithms separately from the protocol will allow NTPv5 to meet the needs of applications with a variety of network properties and performance requirements.

3.3. Timescales

The protocol SHOULD adopt a linear, monotonic timescale as the basis for communicating time. The format should provide sufficient scale, precision, and resolution to meet or exceed NTPv4's capabilties, and have a rollover date sufficiently far into the future that the protocol's complete obsolescence is likely to occur first.

The timescale, in addition to any other time-sensitive information, MUST be sufficient to calculate representations of both UTC and TAI. Through extensions the protocol SHOULD support additional timescale representations outside of the main specification, and all transmissions of time data SHALL indicate the timescale in use.

3.4. Leap seconds

Tranmission of UTC leap second information MUST be included in the protocol in order for clients to generate a UTC representation, but must be transmitted as separate information to the timescale. The specification SHOULD be capable of transmitting upcoming leap seconds greater than 1 calendar day in advance.

Leap second smearing SHOULD NOT be applied to timestamps transmitted by the server, however this should not prevent implementers from applying leap second smearing between the client and any clock it is training.

3.5. Backwards compatibility with NTS and NTPv4

The desire for compatibility with older protocols should not prevent addressing deployment issues or cause ossification of the protocol.

The model for backward compatibility is: servers that support multiple versions of NTP must send a response in the same version as the request. This does not preclude servers from acting as a client in one version of NTP and a server in another.

Protocol ossification MUST be addressed to prevent existing NTPv4 deployments which respond incorrectly to clients posing as NTPv5 from causing issues. Forward prevention of ossification (for a potential NTPv6 protocol in the future) should also be taken into consideration.

3.5.1. Dependent Specifications

Many other documents make use of NTP's data formats ([RFC5905] Section 6) for representing time, notably for media and packet timestamp measurements. Any changes to the data formats should consider the potential implementation complexity that may be incurred.

3.6. Extensibility

The protocol MUST have the capability to be extended; implementations MUST ignore unknown extensions. Unknown extensions received by a server from a lower stratum server SHALL not be added to response messages sent by the server receiving these extensions.

3.7. Security

Data authentication and optional data confidentiality MUST be integrated into the protocol, and downgrade attacks by an in-path attacker must be mitigated.

Cryptographic agility must be supported, allowing for more secure cryptographic primitives to be incorporated as they are developed and as attacks and vulnerabilities with incumbent primitives are discovered.

Intermediate devices such as hardware capable of performing timestamping of packets SHOULD be able to add information to packets in flight without requiring modification or removal of authentication or confidentiality on the packet.

Consideration must be given to how this will be incorporated into any applicable trust model. Downgrading attacks that could lead to an adversary disabling or removing encryption or authentication MUST NOT be possible in the design of the protocol.

4. Non-requirements

This section covers topics that are explicitly out of scope.

4.1. Server malfeasence detection

Detection and reporting of server malfeasance should remain out of scope as [I-D.ietf-ntp-roughtime] already provides this capability as a core functionality of the protocol.

5. Threat model

The assumptions that apply to all of the threats and risks within this section are based on observations of the use cases defined earlier in this document, and focus on external threats outside of the trust boundaries which may be in place within a network. Internal threats and risks such as a trusted operator are out of scope.

5.1. Delay-based attacks

The risk that an on-path attacker can delay packets between a client and server exists in all time protocols operating on insecure networks and its mitigations within the protocol are limited for a clock which is not yet synchronised. Increased path diversity and protocol support for synchronisation across multiple heterogeneous sources are likely the most effective mitigations.

5.2. Payload manipulation

Conversely, on-path attackers who can manipulate timestamps could also speed up a client's clock, resulting in drift-related malfunctions and errors such as premature expiration of certificates on affected hosts. An attacker may also manipulate other data in flight to disrupt service and cause de-synchronisation. Message authentication with regular key rotation should mitigate both of these cases; however consideration should also be made for hardware-based timestamping.

5.3. Denial of Service and Amplification

NTPv4 has previously suffered from DDoS amplification attacks using a combination of IP address spoofing and private mode commands used in many NTP implementations, leading to an attacker being able to direct very large volumes of traffic to a victim IP address. Current mitigations are disabling private mode commands and encouraging network operators to implement BCP 38 [RFC2827]. The NTPv5 protocol specification should reduce the amplification factor in request/response payload sizes [drdos-amplification] through the use of padding and consideration of payload data.

6. IANA Considerations

This document makes no requests of IANA.

7. Security Considerations

As this document is intended to create discussion and consensus, it introduces no security considerations of its own.

8. References

8.1. Normative References

[I-D.ietf-ntp-roughtime] Malhotra, A., Langley, A., Ladd, W., and M. Dansarie, "Roughtime", Work in Progress, Internet-Draft, draft-ietf-ntp-roughtime-05, 24 May 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-ntp-roughtime-05>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/rfc/rfc2827>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/rfc/rfc5905>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

8.2. Informative References

[drdos-amplification] "Amplification and DRDoS Attack Defense -- A Survey and New Perspectives", n.d., <https://arxiv.org/abs/1505.07892>.

[IEEE-1588-2008] "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", n.d..

[ntp-misuse] "NTP server misuse and abuse", n.d., <https://en.wikipedia.org/wiki/NTP_server_misuse_and_abuse>.

[ntppool] "pool.ntp.org: the internet cluster of ntp servers", n.d., <https://www.ntppool.org>.

Appendix A. Acknowledgements

The author would like to thank Doug Arnold, Hal Murray, and Paul Gear for contributions to this document, and would like to acknowledge Daniel Franke, Watson Ladd, Miroslav Lichvar for their existing documents and ideas. The author would also like to thank Angelo Moriondo, Franz Karl Achard, and Malcom McLean for providing the author with motivation.

Author's Address

James Gruessing
Nederlandse Publieke Omroep
Postbus 26444
1202 JJ Hilversum
Netherlands

Email: james.ietf@gmail.com