The Past, Present, & Future of NTP Operations

Paul Gear
AusNOG 2017
Melbourne

Thanks

AusNOG

CANONICAL

ubuntu

Caveats

- My opinions aren't Network Time Foundation's
- My opinions aren't my employer's
- “I'm no expert, I just try my best not to be a total screw-up.” – Sarah White, NTP Pool mailing list
- Draft proposals

The Plan

- Introduction & historical recap
- The Past: Operational issues faced by NTP
- The Future: Recent RFC drafts
- The Present: Best Current Practices

What is NTP?

- The preferred way to synchronise clocks over TCP/IP networks since the 1980s
- More than 10 RFCs, spanning the years 1985-2016, and covering 5 protocol versions
- NTP reference implementation from NTF

Why care about NTP?

- Time is fundamental
  - Maybe your data contains timestamps
  - Maybe you run distributed systems
- Focus on security and Open Source infrastructure in the post-Heartbleed/Shellshock era
- Learning the tools well

NTP Fundamentals

“The goal of the NTP algorithms is to minimize both the time difference and frequency difference between UTC and the system clock. When these differences have been reduced below nominal tolerances, the system clock is said to be synchronized to UTC.” [RFC5905, sect. 4]

NTP Fundamentals

On-wire protocol:

- 32 bits of version & control
- Two 32-bit timestamps
- 32-bit reference identifier
- Four 64-bit timestamps
- Optional extension fields
- Optional MAC

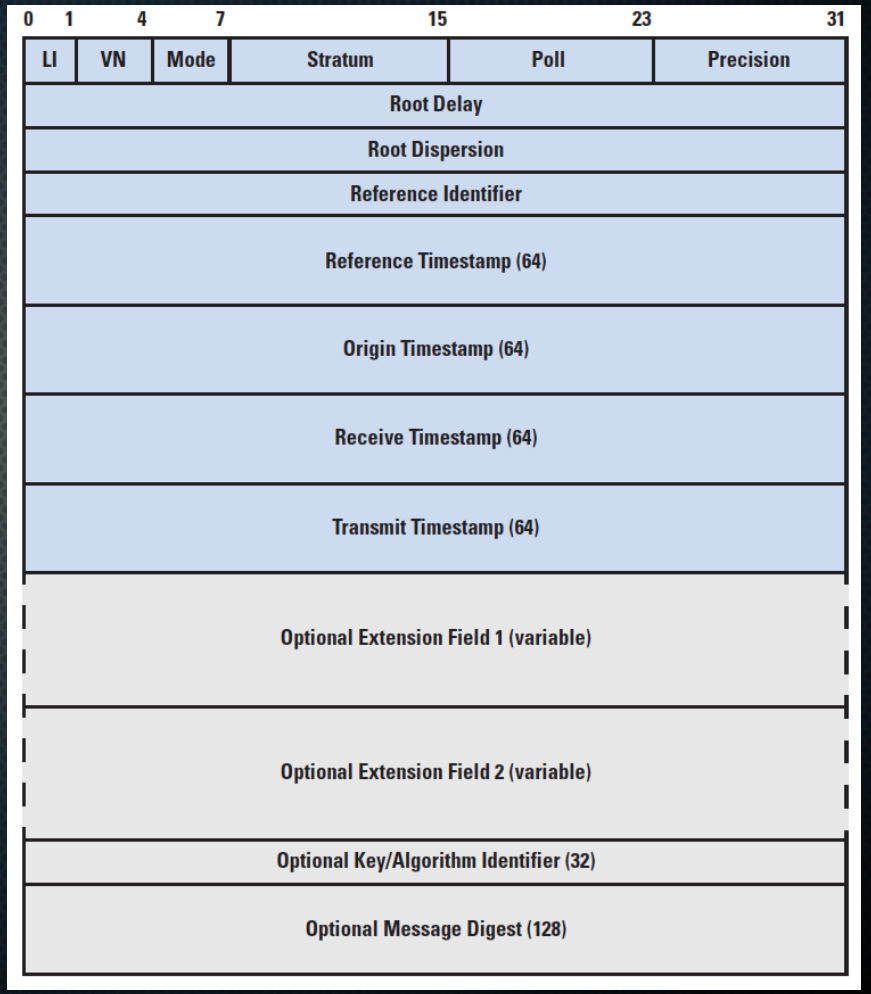

The Past: Operational Issues

Leap seconds

- Solar (UTC) & atomic (TAI) time differ due to orbital degradation
- 28 leap seconds since 1972; negative leap seconds possible
- No leap seconds in POSIX; each day is 86,400 seconds
- Leap seconds have triggered Linux kernel & application bugs
- Leap smearing is one solution: the extra second is spread out over a longer period (usually ½ a day on either side)

Reflective DDoS

- Late 2013/early 2014 saw reflective DDoS attacks
- Used on NTP **`monlist`** command
- DDoS bandwidth: 90 – 400 Gbps
- Amplification factor: 5x – 4670x?
- Causes:
  - poor default configuration in **`ntpd`** (CVE-2013-5211)
  - failure to implement spoofed IP egress filtering

Reflective DDoS

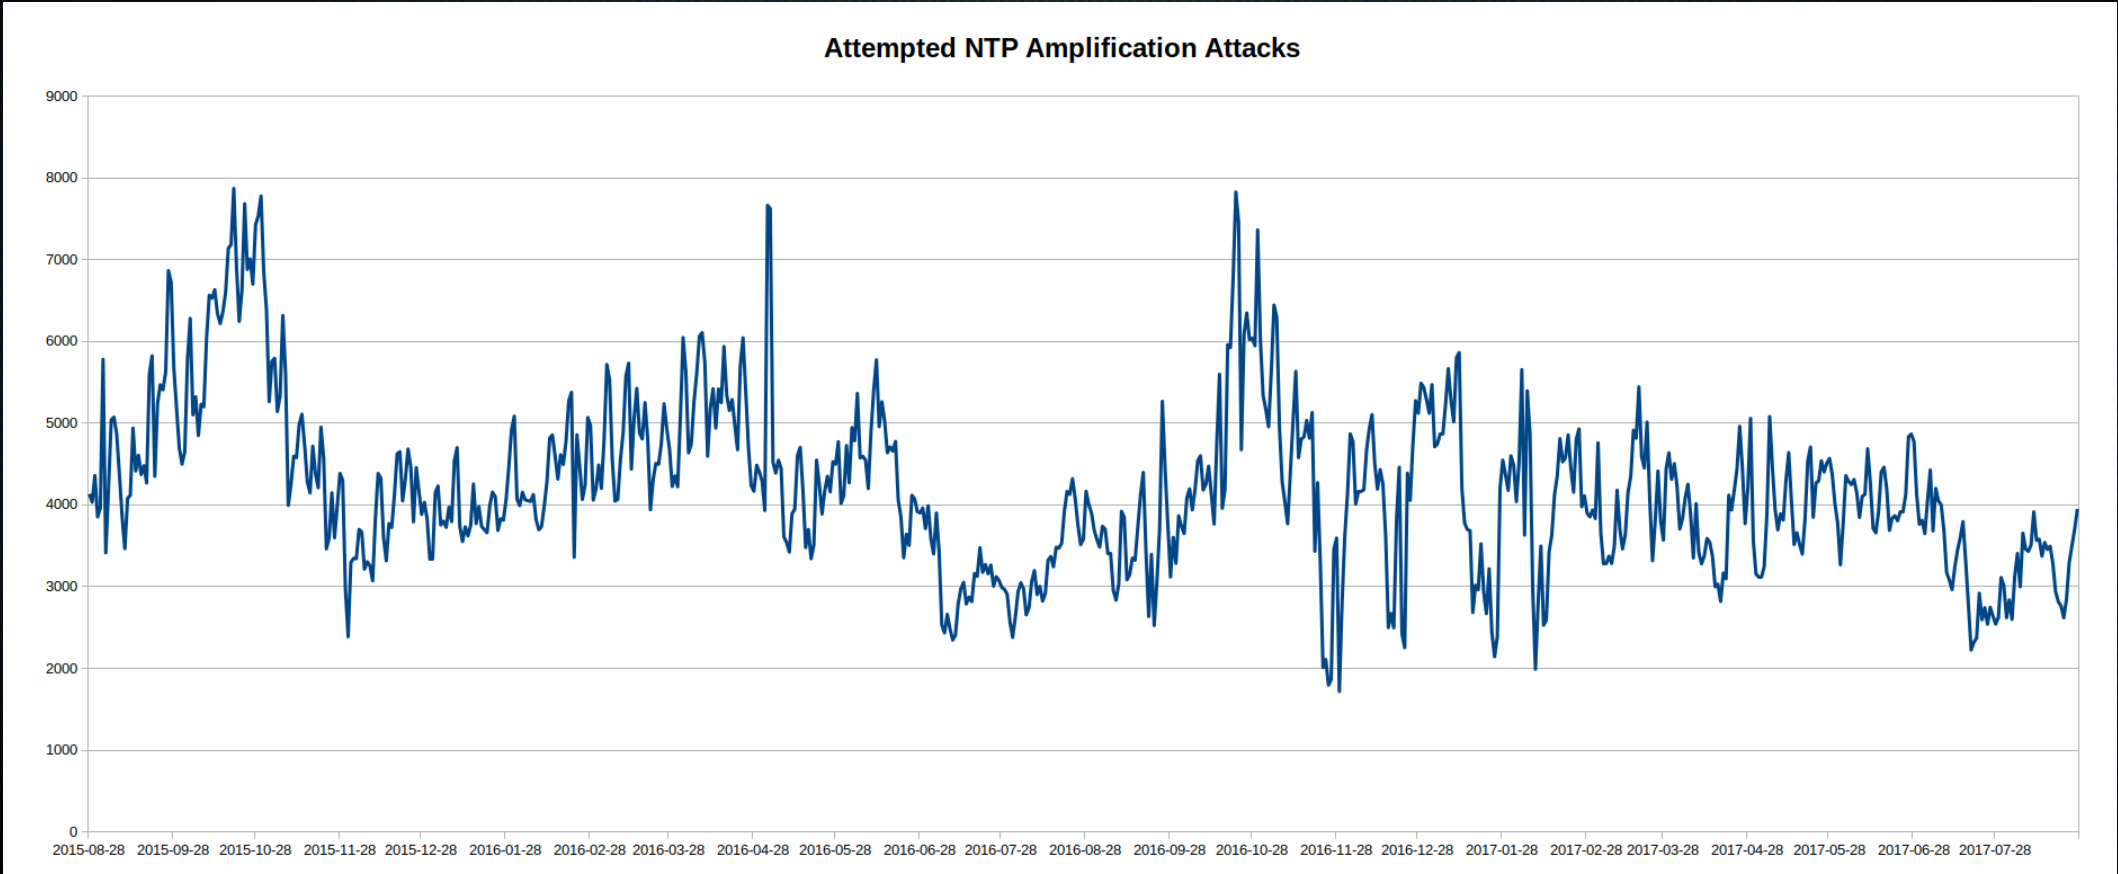

Authentication & Privacy

- NTP was originally designed with no security-specific features. However:
  - The protocol itself is designed to be resilient to attack
  - Most fields aren't security-sensitive
  - Shared-key MD5 MAC available
- But crypto is a moving target:
  - Autokey and extension fields were proposed to provide strong authentication, but have been demonstrated insecure
  - Growing concern regarding client fingerprinting
  - MD5 now considered broken

Netgear vs. University of Wisconsin

- Some Netgear ADSL & cable modem models shipped with:
  - University of Wisconsin's NTP server IP address hard-coded
  - 1-second polling intervals (until a valid NTP response was received)
- Flawed Routers Flood U. Wisconsin Internet Time Server
  - Dave Plonka, August 2003; also LISA & NANOG talks
- NTP server misuse and abuse – Wikipedia

Snapchat NTP Surge 2016-12-13

- Snapchat iOS client included an NTP library which queried 35-60 NTP servers every time a user opened the app
- My pool server saw:
  - 2x unique IPs/day, 40x unique IPs/hour
  - 7x peak packet count, 6x peak byte count

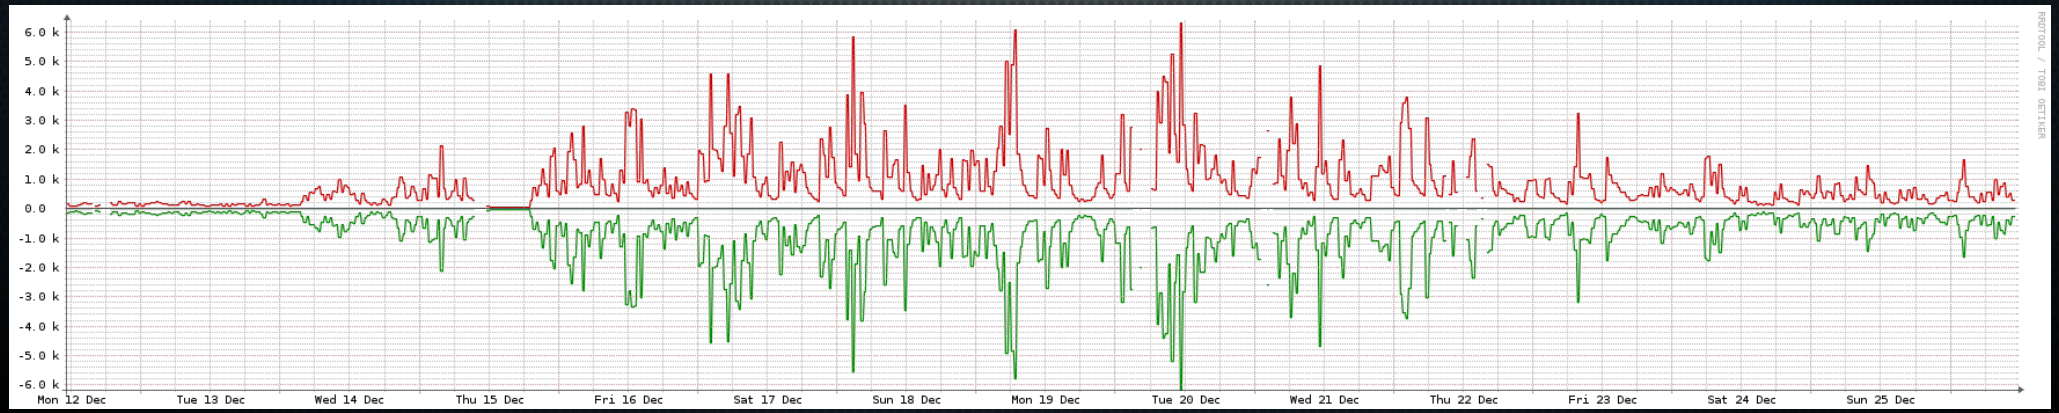

Snapchat NTP Surge 2016-12-13

- IPv6 was unaffected, large servers in US & Europe only a little
- Africa, Australia, and South America hard-hit, with many servers forced out of the pool due to load

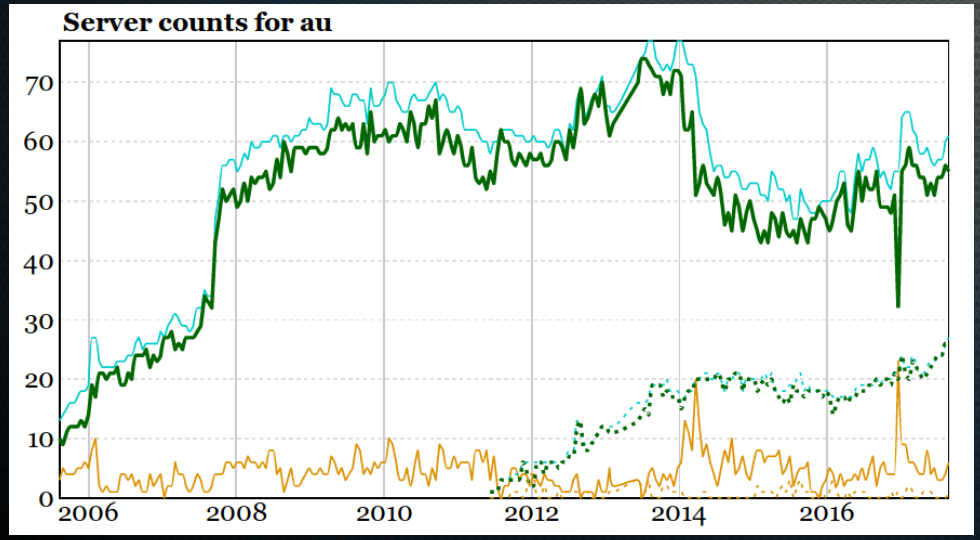

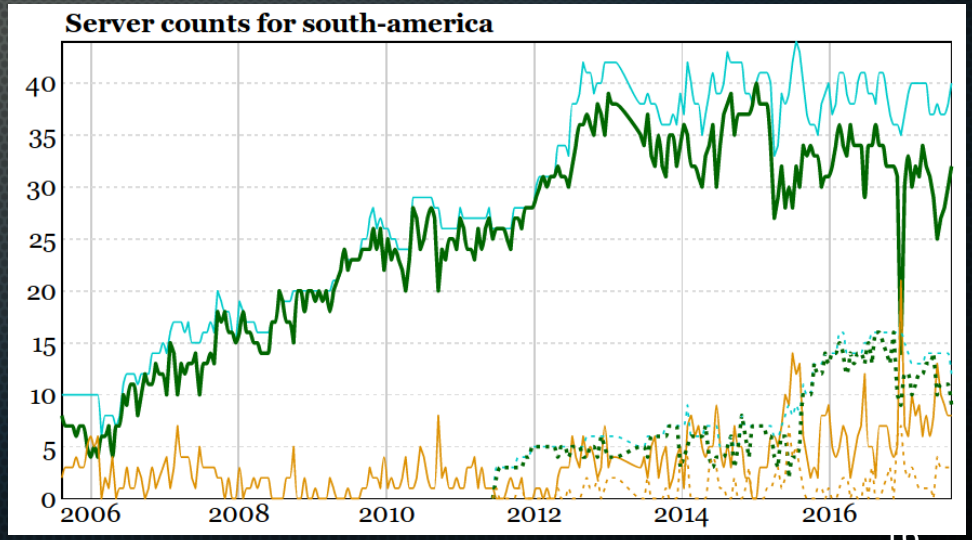

The Future: RFC drafts

Message Authentication Code for NTP

- Deprecates 128-bit MD5 MAC field
- Replaces with 128-bit AES-CMAC
- No protocol changes required

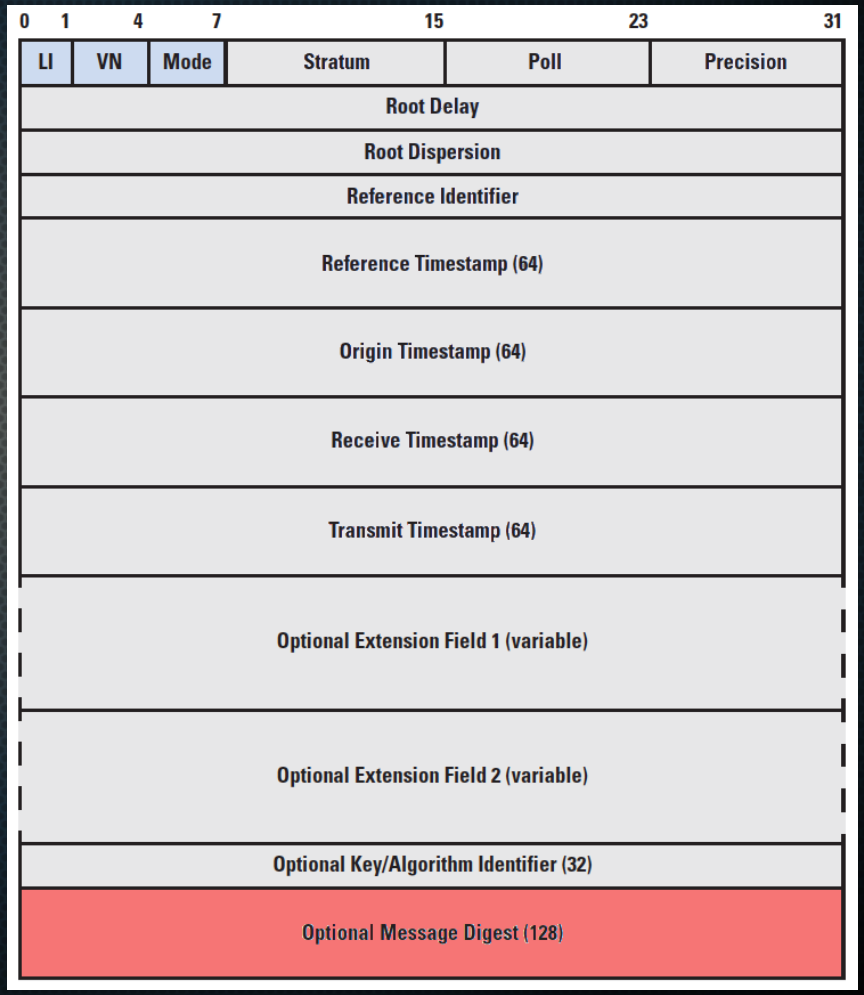

NTP Client Data Minimization

- Aims to reduce the amount of state revealed by NTP clients
- Polling interval set to the real polling interval, or zero
- Transmit timestamp: random
- All other time fields: zero
- No protocol changes required

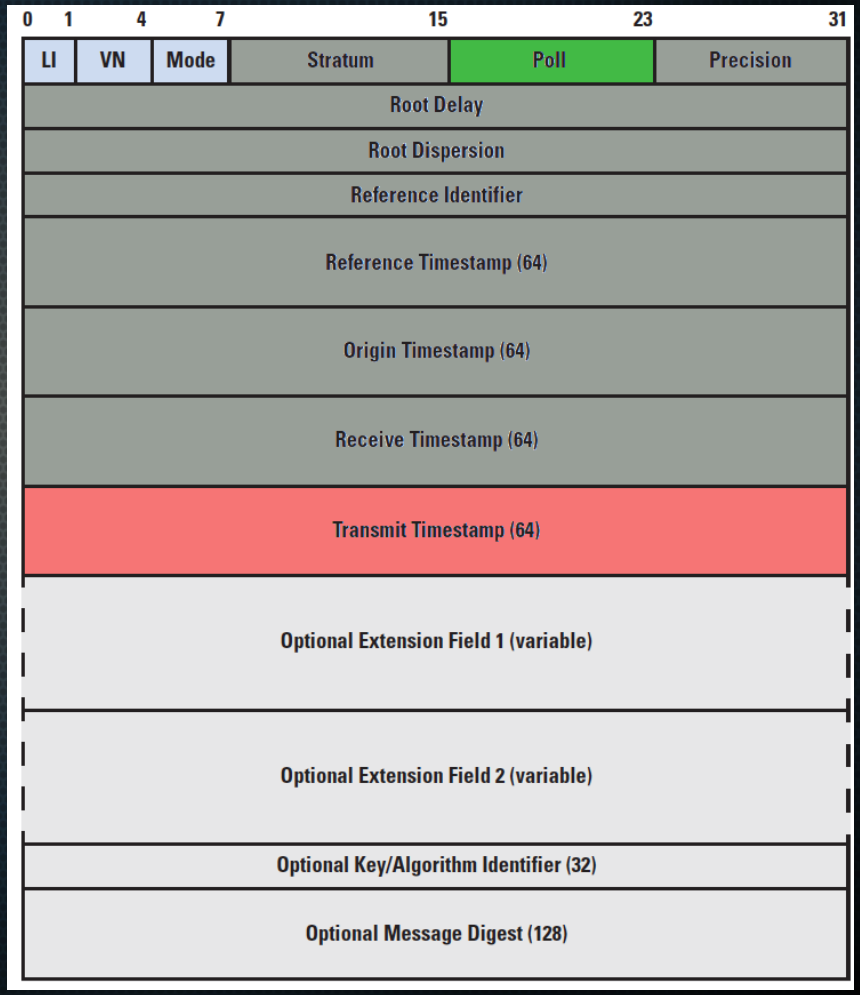

Network Time Security (NTS) for NTP

- Replacement for Autokey
- NTF contracted to create proof-of-concept
- 15 revisions of original NTS draft; 9 revisions of the follow-on NTS for NTP draft
- Two different modes of operation

- Aims as stated in draft:
  - (D)TLS-based authentication
  - encrypt some/all of packet
  - increase protection against replays & spoofing
  - avoid contributing to DDoS amplification attacks
  - maintain high scalability

NTS for NTP: Symmetric & Control Modes

- Uses DTLS (TLS over UDP)
- Encrypts entire NTP packet as DTLS payload
- Requires TLS1.2, prohibits RC4
- Questions remain about:
  - packet size
  - certificate authority infrastructure
  - impact on timing accuracy of encapsulation vs. extension fields
  - compatibility with hardware timestamping

NTS for NTP: Cryptosystem

- Authenticated Encryption with Associated Data (AEAD)
  - RFC 5116: An Interface and Algorithms for Authenticated Encryption
- Aims to standardise & abstract crypto to increase adoption
- 128- & 256-bit AES, two block cipher modes: GCM & CCM
- Simultaneously authenticates & decrypts

NTS for NTP: Cryptosystem

Operation	Inputs	Outputs
Encryption	secret key (K) unique nonce (N) [authenticated data (A)] plaintext (P)	ciphertext (C)
Decryption	secret key (K) unique nonce (N) [authenticated data (A)] ciphertext (C)	plaintext (P) or FAIL

NTS for NTP: Client & Server Modes

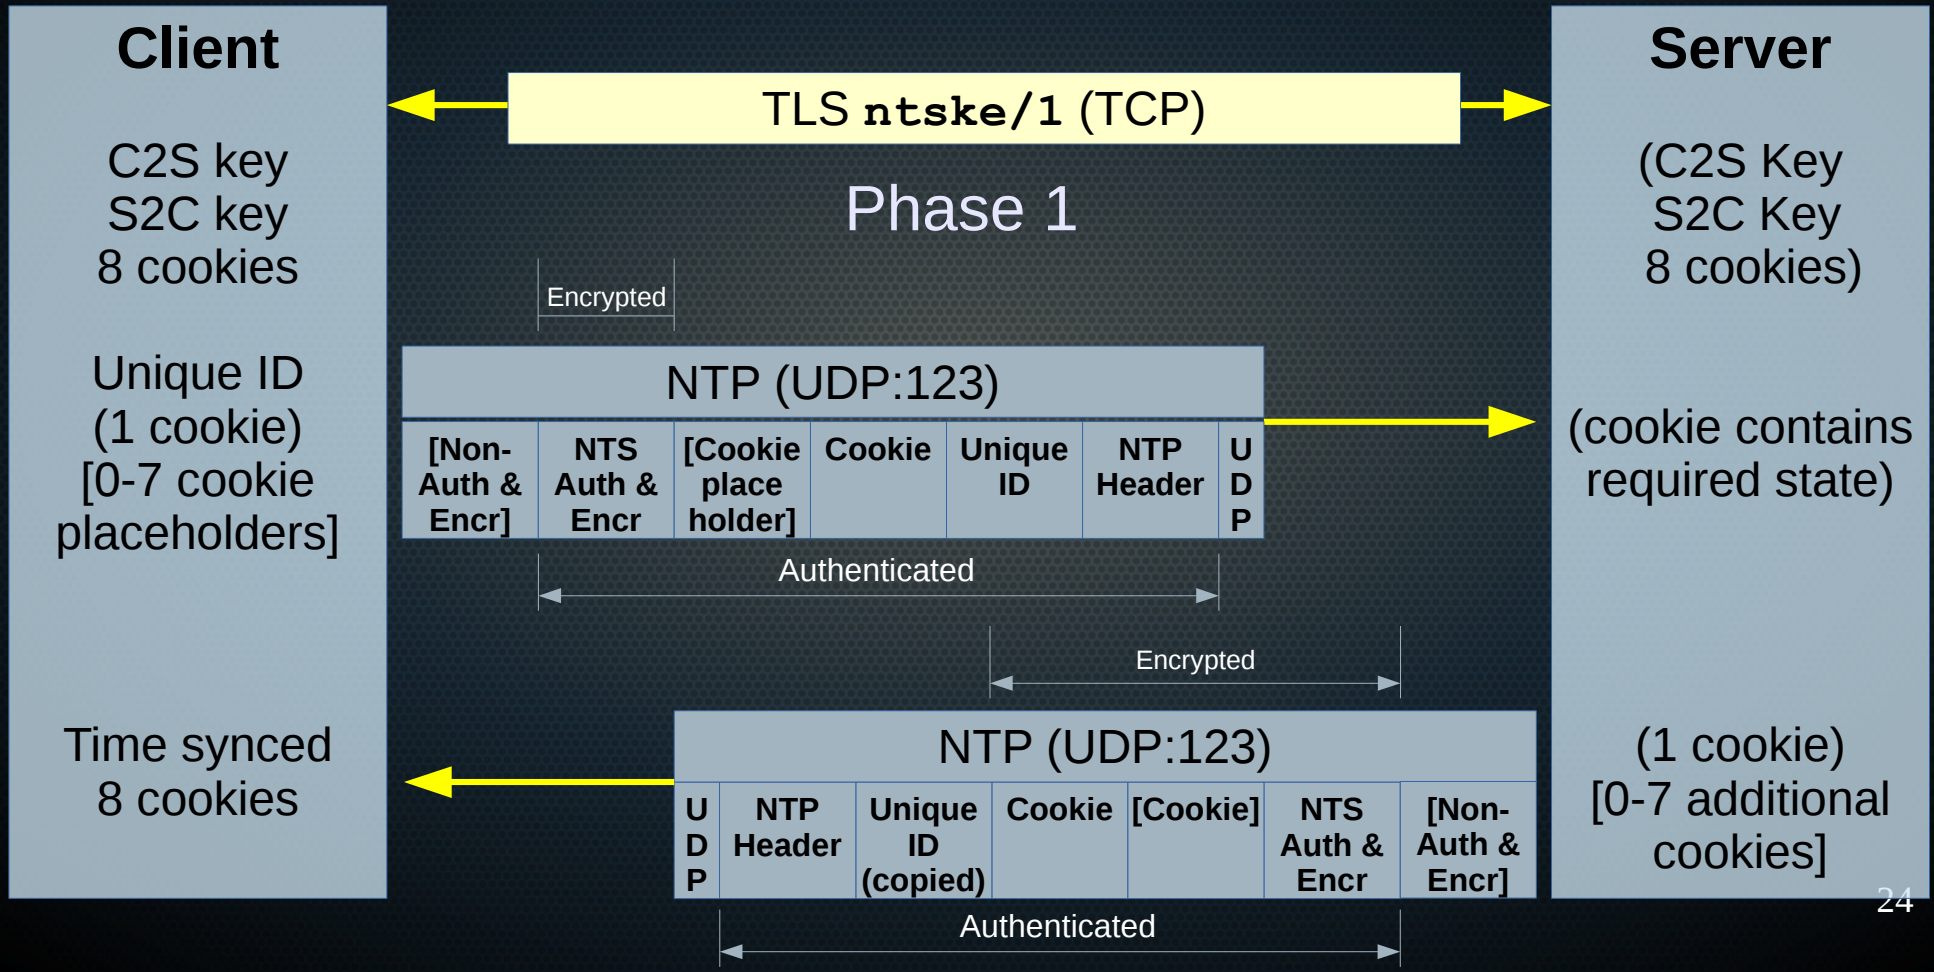

NTS for NTP: Client & Server Modes

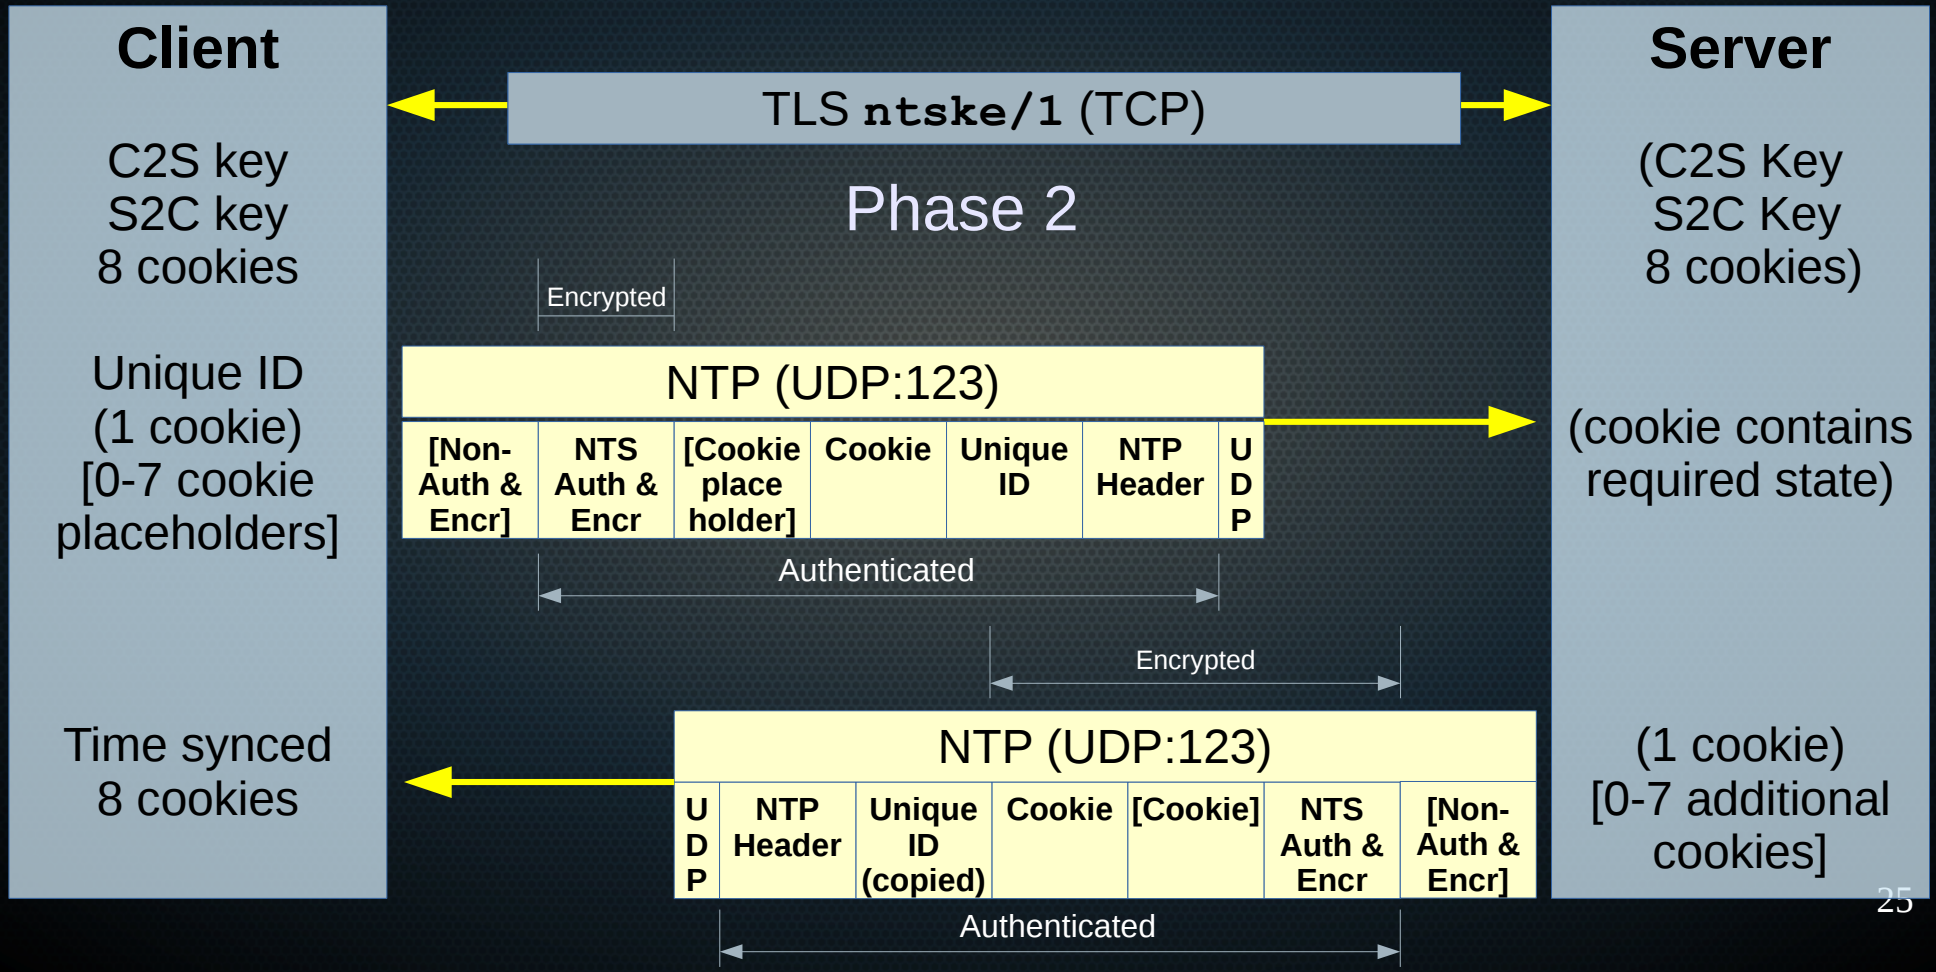

NTS for NTP: closing thoughts

- Ambitious draft, trying to solve multiple problems
- Elegant cryptosystem – allows future changes
- Questions remain about:
  - why the Unique Identifier is necessary
  - compatibility with hardware timestamping
  - certificate authority infrastructure & use with NTP pool
  - entropy & encryption demands on high-volume servers

Other recent drafts

- Yang Data Model
- REFID updates
- Control Messages Protocol for Use with NTPv4

```
     remote           refid      st t when
==============================================
o127.127.20.0    .PPS.            0 l    2
 time.apple.com  .POOL.          16 p    -
 time-ios.apple. .POOL.          16 p    -
 172.22.254.2    172.22.254.53    2 s   11
 172.22.254.1    172.22.254.53    2 s   39
-17.253.82.125   .GPSs.           1 u   62
+17.253.66.125   .GPSs.           1 u   62
+17.253.66.253   .GPSs.           1 u   44
```

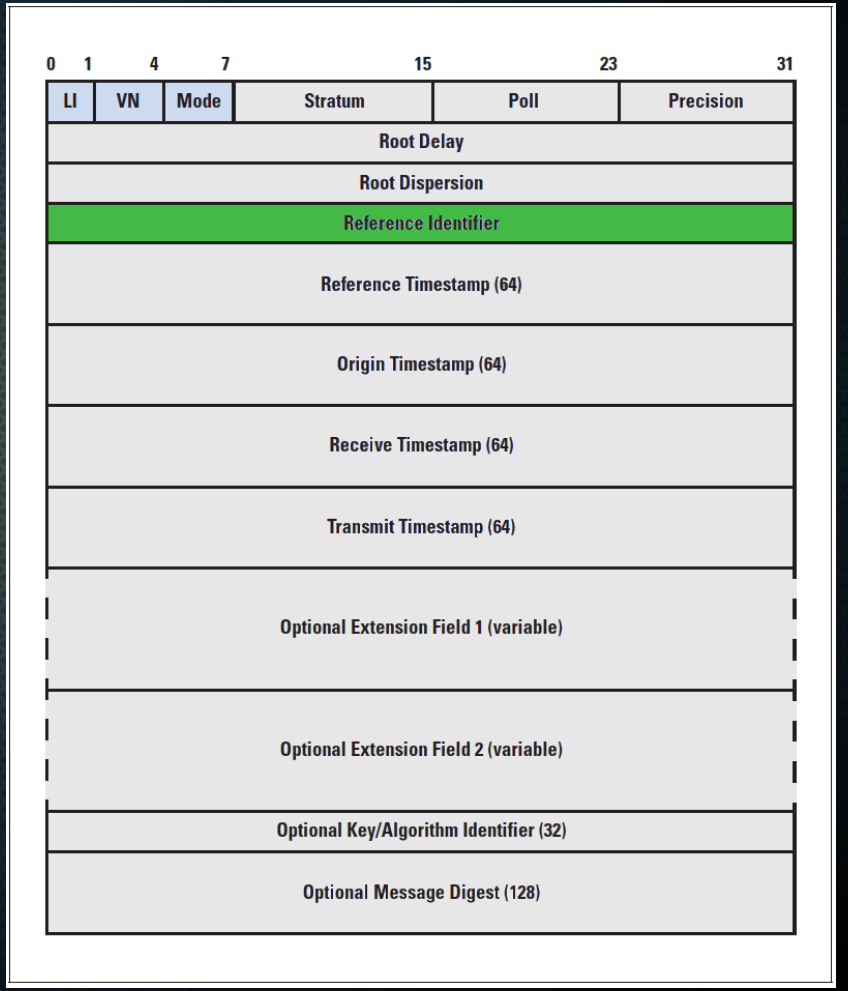

Extension Field Proposals

- Originally EFs were only used by Autokey, but now NTS and:
  - Suggested REFID: allow identifying first-degree timing loops without leaking information about the upstream server
  - I-Do: allow NTP instances to exchange information about which EFs they support
  - Distributing the UTC/TAI difference
  - Indicating the use (and possibly method) of leap smearing
  - Legacy MAC

Upcoming NTP releases

- ntp-4.2.8p11 probably last 4.2.x release
- ntp-4.4.0 should have AES-128-CMAC support & new REFIDs
- NTS may arrive in 4.4.0, or 4.6.0 if the draft is not finalised
- Moving to simplified versioning:
  - first digit: protocol version
  - second digit: major release
  - third digit: patch release
  - even major releases: stable, odd releases: development

The Present:
Best Current Practices

Keep up-to-date

- Update your NTP software
  - new versions released frequently
  - router/switch vendors can be quite slow to patch
- Update your NTP config
  - behaviours, options, and default settings can change

Implement BCP 38

Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38
Implement BCP 38 Implement BCP 38 Implement BCP 38

Number of time sources

- Most common NTP configuration error
- Not just for redundancy – NTP requires multiple sources for **accuracy**
- Use “enough” time sources
  - how many is enough?
  - typical recommendation has been a **minimum** of 4 sources
  - use the `pool` directive to ensure “enough” quality sources

Diversify reference clocks

- Reliance on any given vendor, model, chipset, or firmware version means that systemic problems have a larger impact
- Ongoing cost reduction of GPS-based NTP servers → over-reliance on GPS?
  - GPS affected by EMP, solar flares, & hardware failures
  - US & Europe have the additional option of radio reference clocks
  - Lower-cost atomic clocks, please!

Diversify reference clocks

There's a stratum 1 server to suit every budget

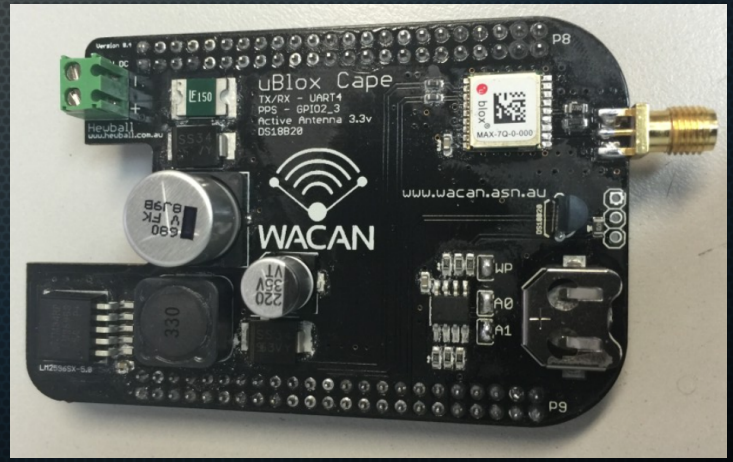

Disable or secure remote control

- Ensure mode 6 and mode 7 are disabled
- This is the default for all new versions

```
restrict default -4 nomodify notrap nopeer noquery
restrict default -6 nomodify notrap nopeer noquery
restrict source nomodify notrap noquery
```

Monitor NTP

- configured & working sources
- quality of synchronisation with sources
- system logs – `ntpd` provides some attack indications
- https://launchpad.net/ntpmon

Use the pool responsibly

- Software vendors & Open Source projects: vendor subdomain
- Millions of external clients: provide public servers
  - examples: Apple, Microsoft, Ubuntu
- Large private networks: provide private servers
  - reduces chances of DDoS, simplifies firewall rules
- Spare server? Please consider participating
  - bandwidth requirements need not be high
  - keep GeoIP data on your IP space up-to-date

Handle leap seconds appropriately

- Reference implementation includes leap second support
  - ntp-4.2.6 changed behaviour from any upstream source to the majority clique
- Some servers report erroneous leap seconds
  - can be mitigated with a local leap second file
- Leap smearing is available in the reference implementation
  - not recommended for public servers
  - check compliance and legal implications

Security issues

- Autokey deprecated; MD5 MAC to follow
  - NTS and AES-CMAC are expected replacements
- Leaf nodes should not respond to unsolicited NTP requests
- Daemon restart attacks
- Use broadcast mode only on trusted networks
- Use symmetric mode only with trusted peers

Embedded device manufacturers

- Ensure you have a rock-solid update mechanism
  - both code and configuration
- Get a vendor subdomain from **`pool.ntp.org`**
  - and/or provide your own servers

Anycast precautions

- Use only when a higher-than-average level of time variation is acceptable
  - Anycast routes switching between endpoints increases jitter
  - Load balancing unnecessary and unhelpful
- DNS-based pools with 4+ entries are preferred over anycast

NTP and NTF want you!

- Contribute operational expertise to IETF discussions
- Test new versions & features
- Participate in the pool
- Become a supporter of NTF

Thanks for listening!

- NTP docs
- Geoff Huston's Protocol Basics – The Network Time Protocol
- My talk at Linux.conf.au 2017 sysadmin miniconf
- My blog series
- Any questions?