


Postfix: recent adoption of new technologies



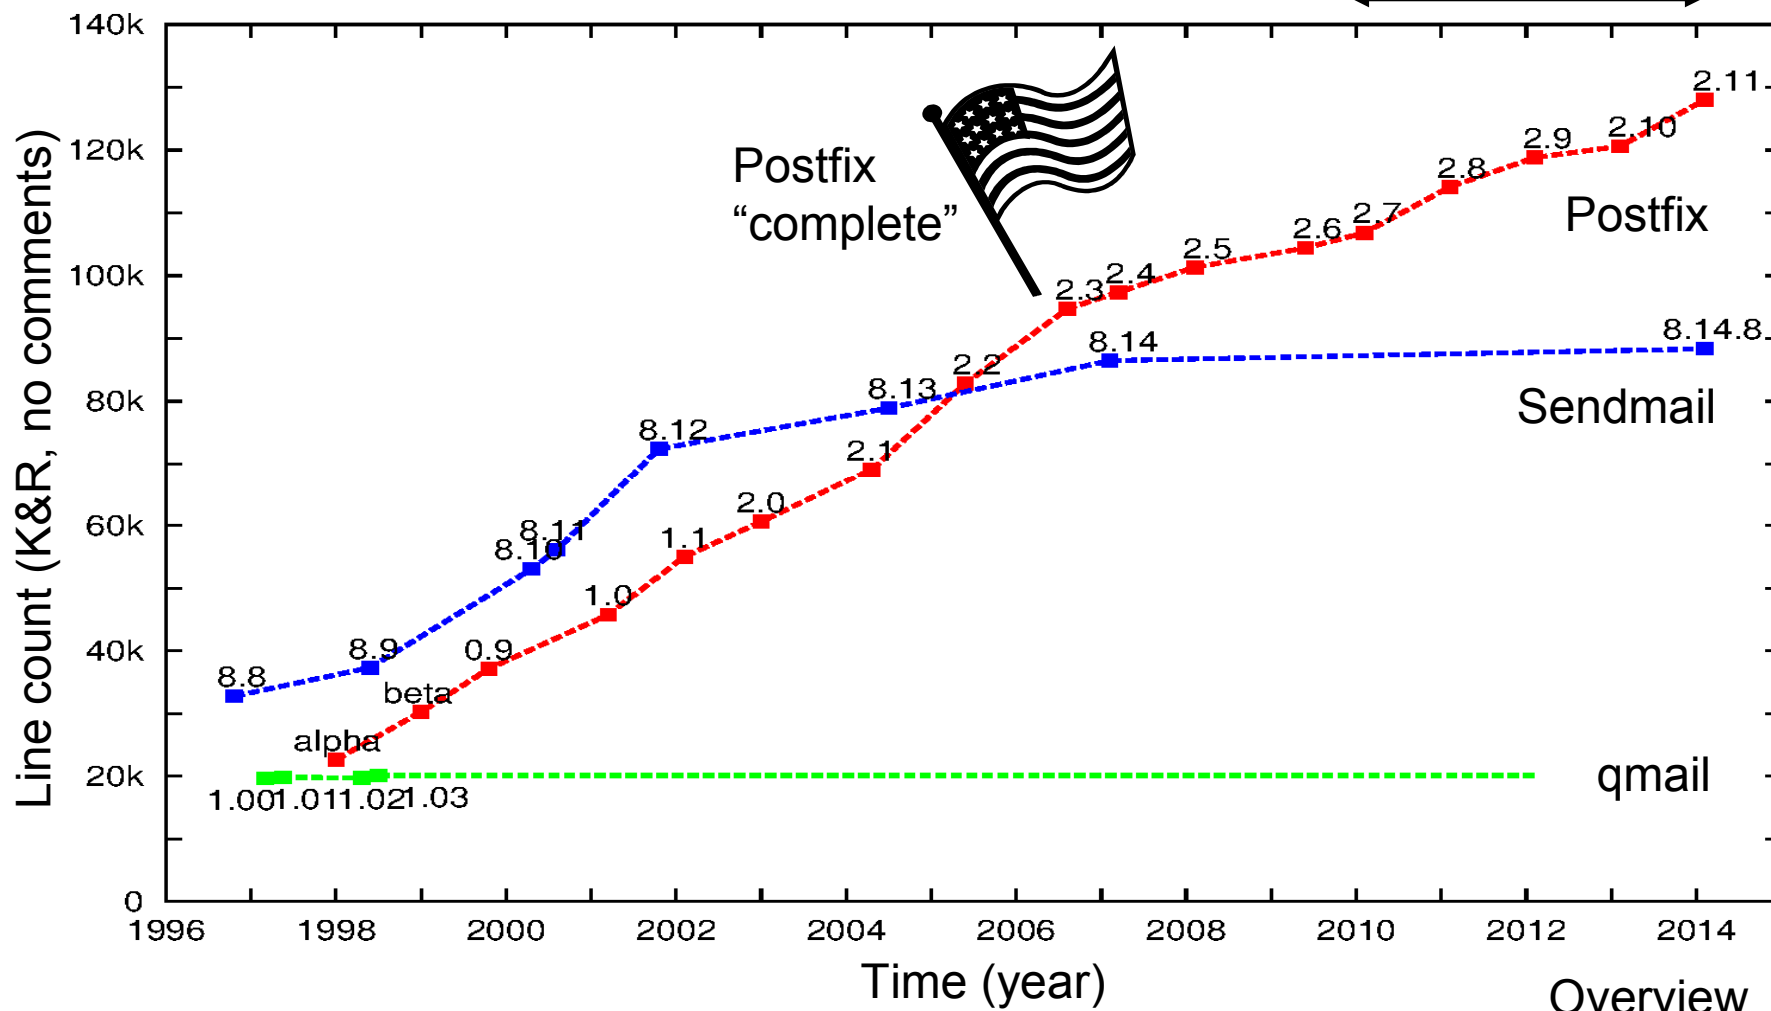
Overview

- Overview.
- Motivation and architecture.
- Spam around the clock.
- Scalable defense (postscreen zombie blocker).
- New: miscellaneous improvements.
- New: security without global PKI (DANE).
- New: replacing Berkeley DB with LMDB.
- Conclusion.

Postfix timeline

Larger is not necessarily better

Most of this presentation



- Postfix motivation and architecture

Why (not) write another UNIX mail system

CERT/CC advisories for Sendmail

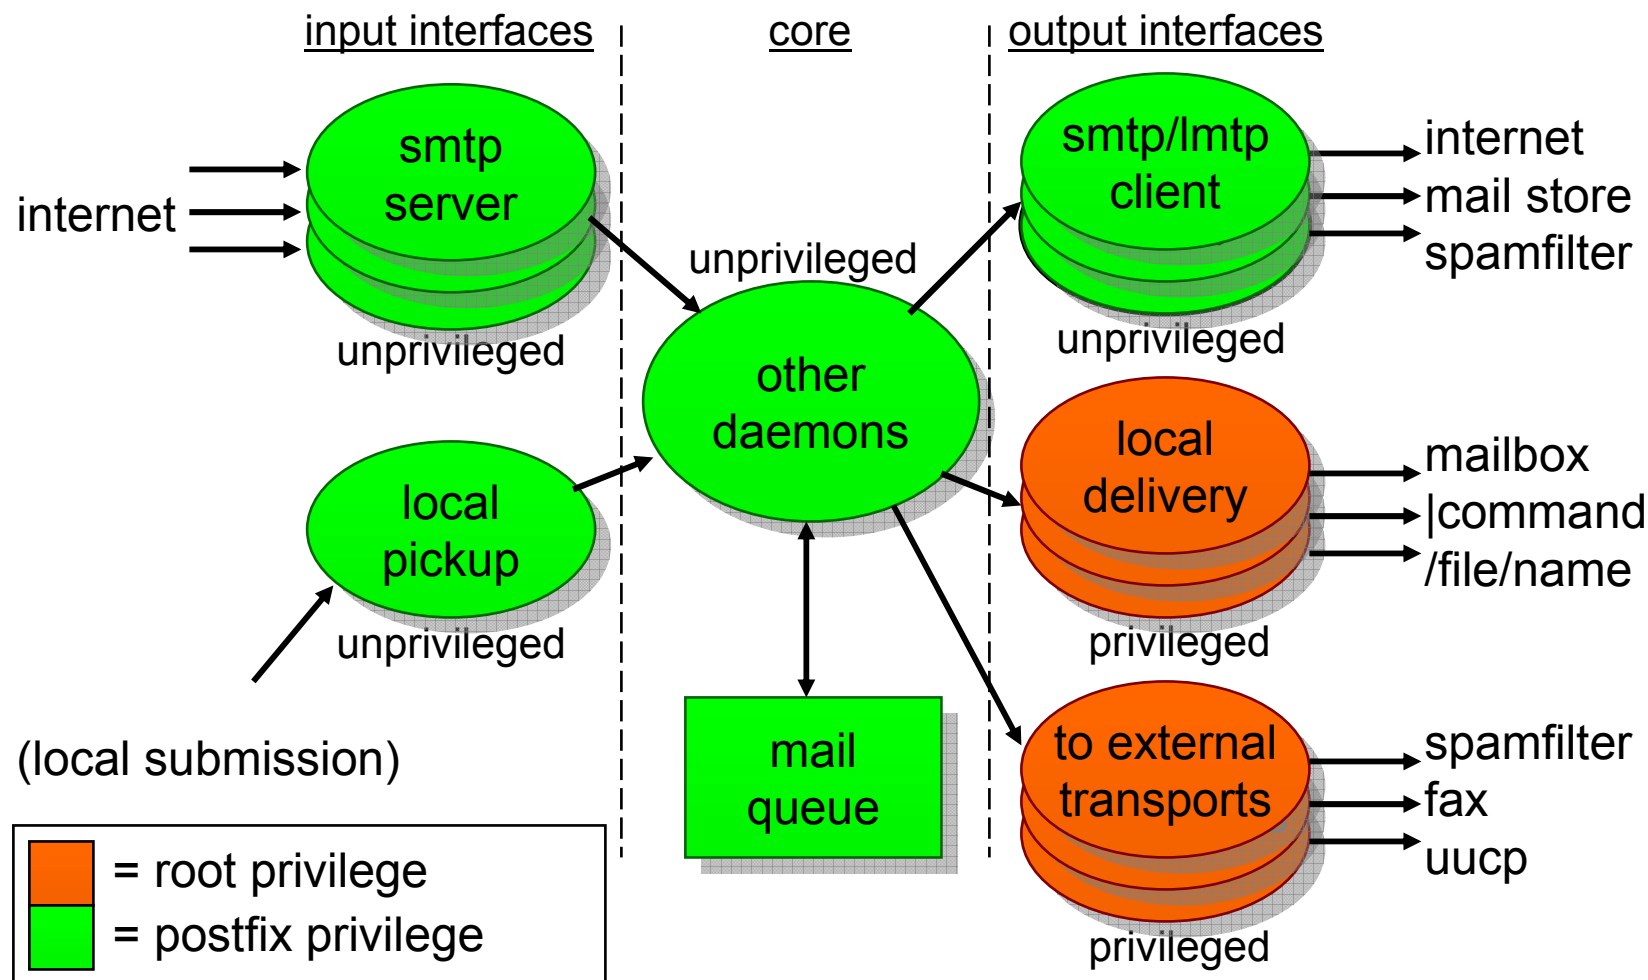
The initial threat model: mail server attacks

Advisory	Version	Impact
CA-1988-01	5.58	Unprivileged access
CA-1993-16	8.6.3	Unprivileged access
CA-1994-12	8.6.7	Full system privilege
CA-1995-05	8.6.9	Full system privilege
CA-1995-13	8.7.0	Full system privilege
CA-1996-04	8.7.3	Full system privilege
CA-1996-20	8.7.5	Full system privilege
CA-1996-24	8.8.2	Full system privilege
CA-1996-25	8.8.3	Group privileges
CA-1997-05	8.8.4	Full system privilege
CA-2003-07	8.12.7	Full system privilege
CA-2003-12	8.12.8	Full system privilege
CA-2003-25	8.12.9	Full system privilege

Motivation & architecture

Postfix low-privilege architecture


(omitted: non-daemon programs for submission and management)



Motivation & architecture

- Spam around the clock

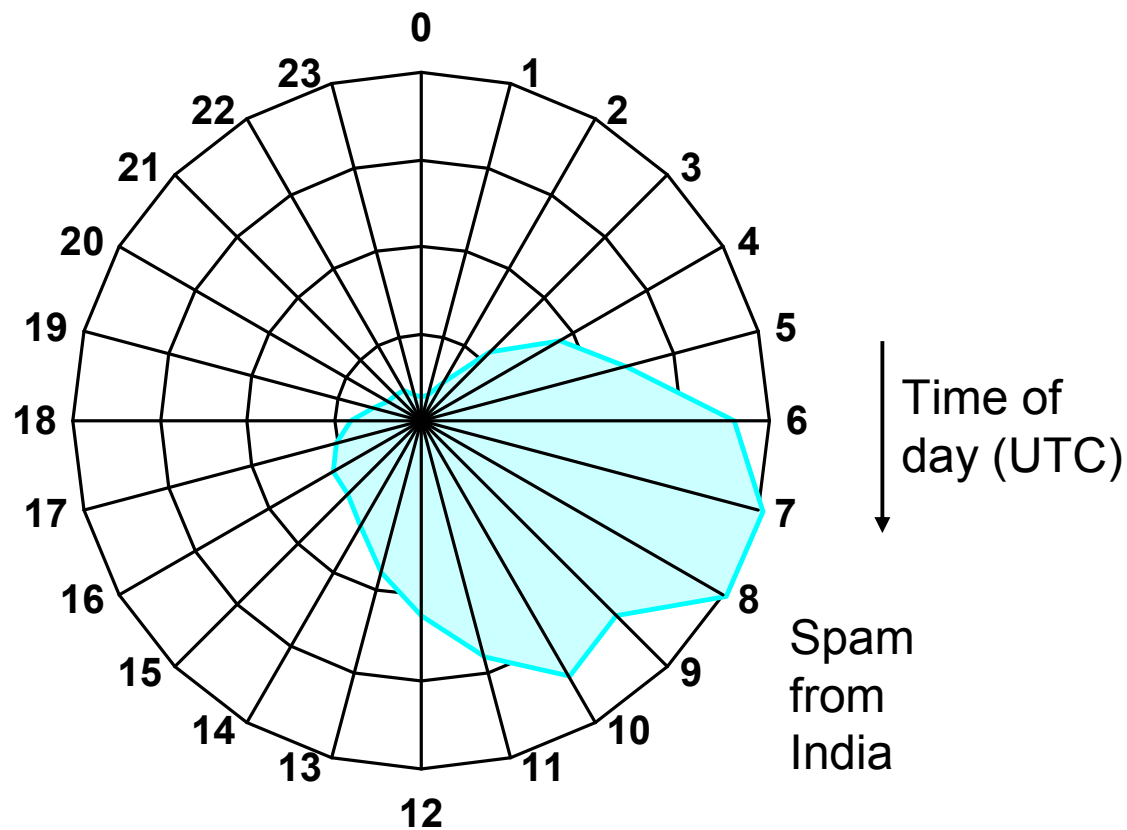
SPAM is a 24-hour operation ...



- Spam connections to charite.de (Berlin, Germany) Oct 29 – Jan 23, 2014, from IP addresses blacklisted at zen.spamhaus.org.

Spam around the clock

... but many spambots are not









- Spam connections to charite.de (Berlin, Germany) Oct 29 – Jan 23, 2014, from IP addresses blacklisted at zen.spamhaus.org.

Spam around the clock

Spam connections/hour at charite.de (time in UTC)

From IP addresses blacklisted at zen.spamhaus.org, Oct 29 – Jan 23, 2014



Spam around the clock

- **Zombies suck the life out of the mail server**

Adapting to changing threats

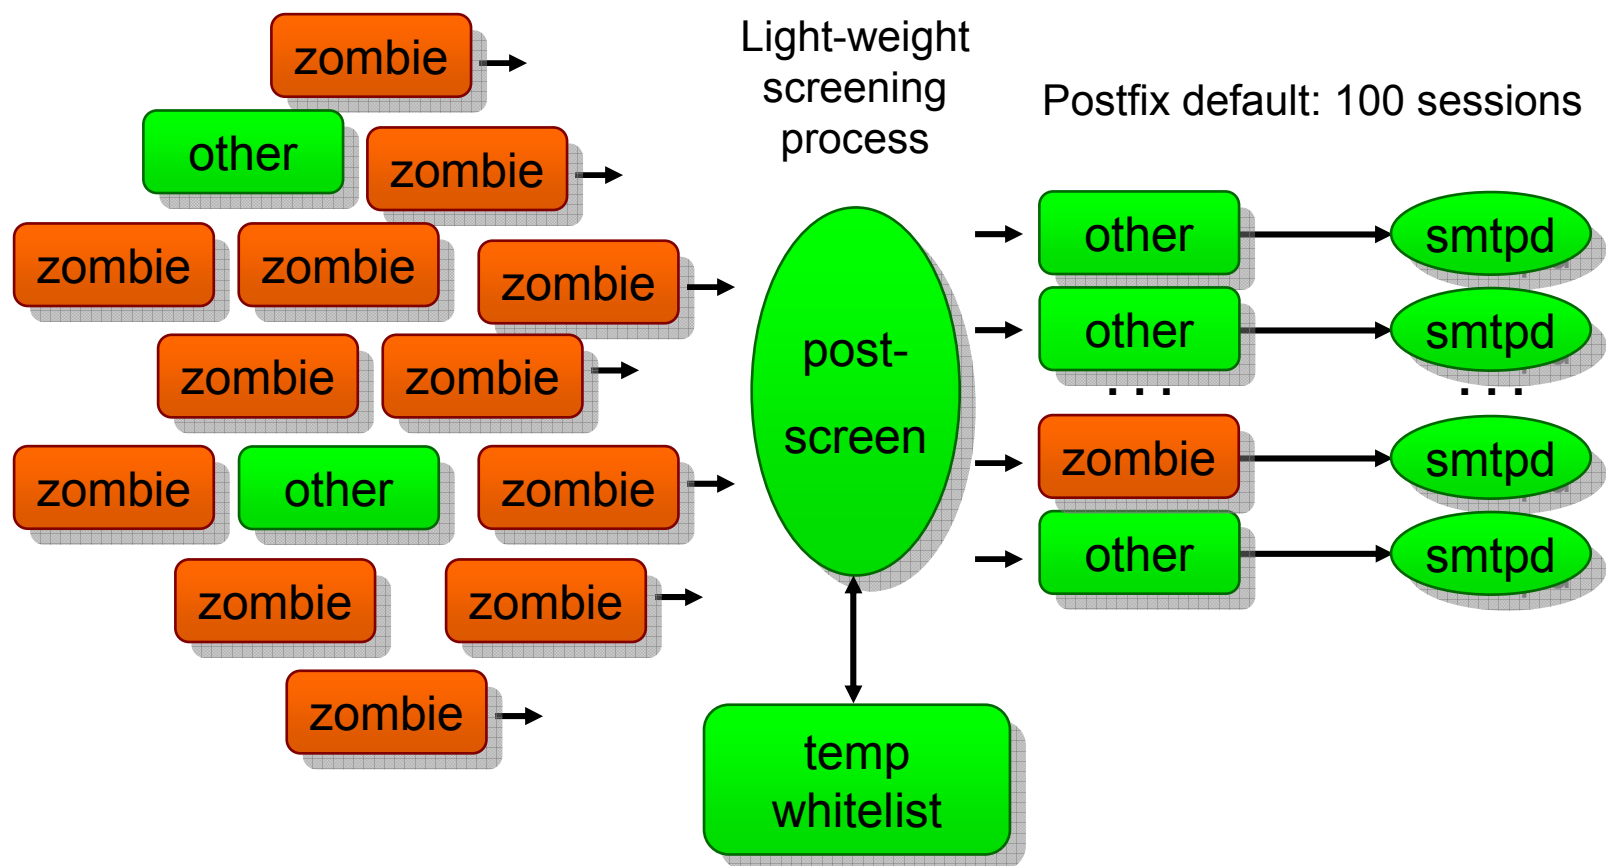
Email spam percentage over time (Symantec)

August 2010: 92% Of email is spam, 95% of spam is from botnets



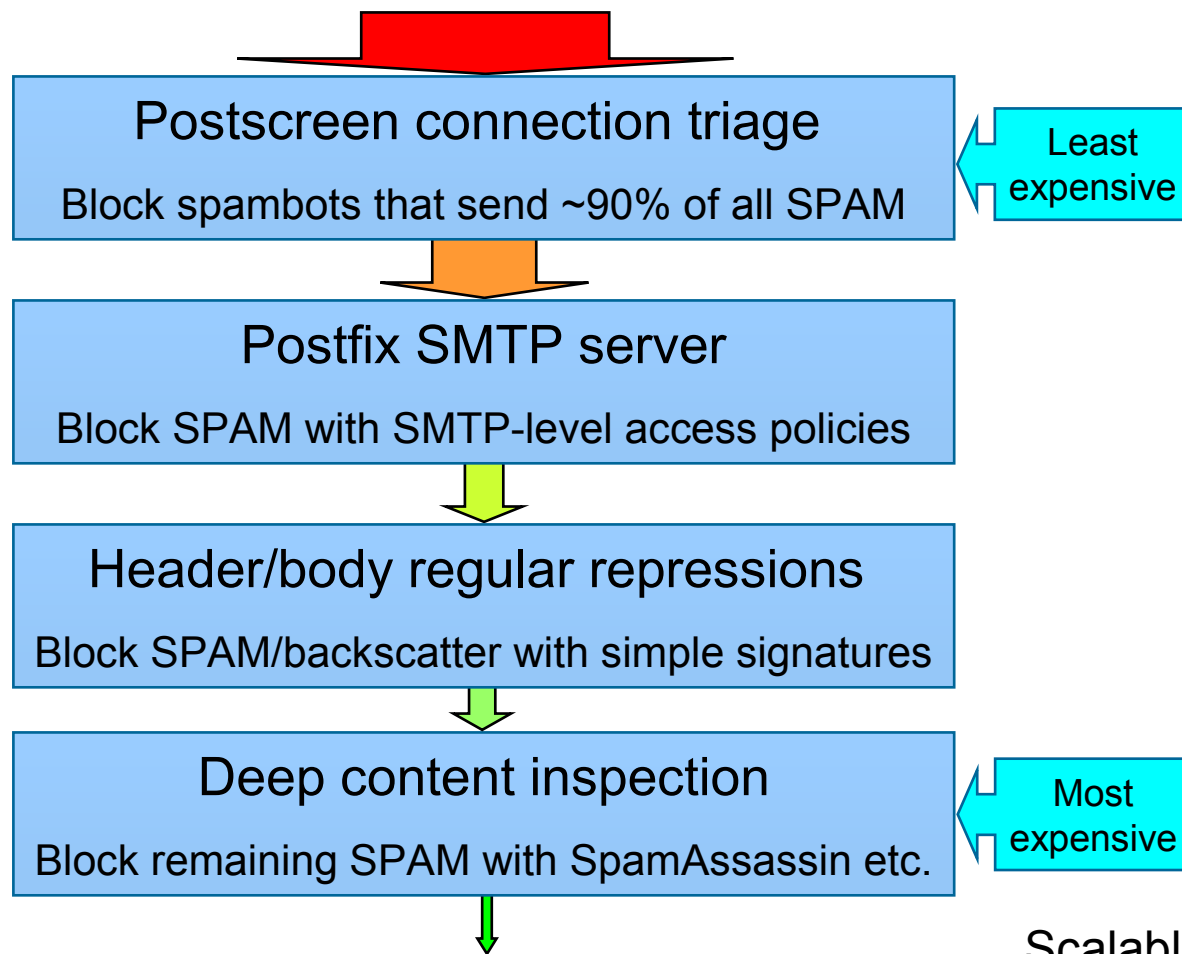
postscreen zombie blocker

Prior work: OpenBSD spamd, MailChannels TrafficControl, M.Tokarev



Scalable defense

postscreen – one step in a four-layer defense

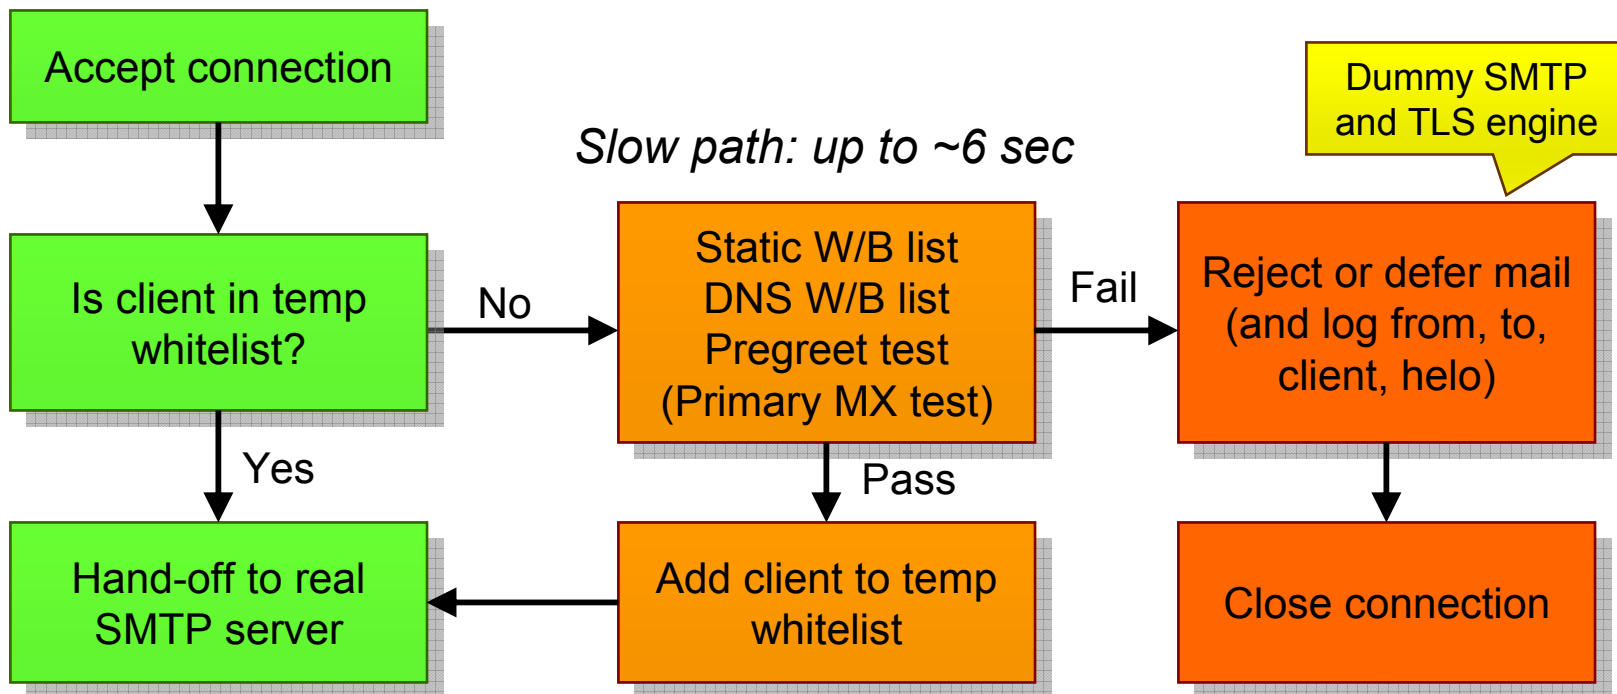


Scalable defense

postscreen workflow - tests before SMTP handshake

One daemon screens multiple connections simultaneously

Fast path: ~0.1 ms



Scalable defense

postscreen DNSBL/DNSWL support

Parallel DNS lookups

- Weight factors (to *whitelist*, use *negative* numbers).

```
postscreen_dnsbl_sites = zen.spamhaus.org*2, bl.spamcop.net*1,  
                        b.barracudacentral.org*1  
postscreen_dnsbl_threshold = 2
```

- Reply filters.

```
postscreen_dnsbl_sites = zen.spamhaus.org=127.0.0.4 ...  
postscreen_dnsbl_sites = zen.spamhaus.org=127.0.0.[1..11] ...
```

- Allow “good” clients to skip tests (Postfix 2.11).

```
postscreen_dnsbl_sites = list.dnswl.org=127.0.[0..255].[1..3]*-2 ...  
postscreen_dnsbl_whitelist_threshold = -2
```

Scalable defense

Making zombies bark - multi-line greeting trap



- Good clients wait for the full multi-line server greeting:

```
postscreen: 220-server.example.com ESMTP Postfix<CR><LF>  
smtp server: 220 server.example.com ESMTP Postfix<CR><LF>  
good client: HELO client.example.org<CR><LF>
```


- Many spambots talk immediately after the first line of the multi-line server greeting:

```
postscreen: 220-server.example.com ESMTP Postfix<CR><LF>  
postscreen: (wait a few seconds)  
spambot: HELO i-am-a-bot<CR><LF>
```

Scalable defense

Over 60% of bots pregreet (8% not on DNSBL)

mail.charite.de, Berlin, Aug 26 – Sep 29, 2010



Scalable defense

- New developments: miscellaneous improvements

Miscellaneous Postfix 2.11 improvements

- Documentation: “Perfect” Forward Secrecy.
 - Network communication remains confidential when server private key is exposed *later* (heartbleed-class bug, government agency, hacking).
- TLS-encrypted MySQL database connections.
 - Useful in multi-tenant network or cloud.
- Both “user+suffix@example” and “user-suffix@example”.
 - In main.cf: `recipient_delimiter = +-`
 - Even small things can make a difference.

Miscellaneous Postfix 2.11 improvements

- Managing master.cf files without text editor (just like main.cf).
 - Primary target: third-party management tools.
 - Basic idea: “*service-name/service-type/attribute = value*”.
 - Modify/add/delete parameter setting for specific Postfix service:
 - `postconf -P "smtp/unix/smtp_bind_address = 192.0.2.1"`
 - Modify specific service field for all Postfix services:
 - `postconf -F "**/*/chroot = n"` (turn off chroot globally)
 - Similar syntax to “query” or “list” service attributes or parameters.

Miscellaneous Postfix 2.12 improvements

- Delivery status filter for Postfix delivery agents (March 2014).
 - Delivery status format: “X.Y.Z Descriptive text...”
 - Replace descriptive text to hide file names or command details.
 - Replace X.Y.Z delivery status code, for example to **change soft errors into hard errors** (can't change non-errors ↔ errors).

```
main.cf:
    smtp_delivery_status_filter = pcre:smtp_dsn_filter

smtp_dsn_filter:
    /^4(\.\d+\.\d+ TLS is required, but host \S+ refused+)/ 5$1
    /^4(\.\d+\.\d+ TLS is required, but was not offered .+)/ 5$1
```

- SMTP client delays soft→hard mapping until “final” MX host.

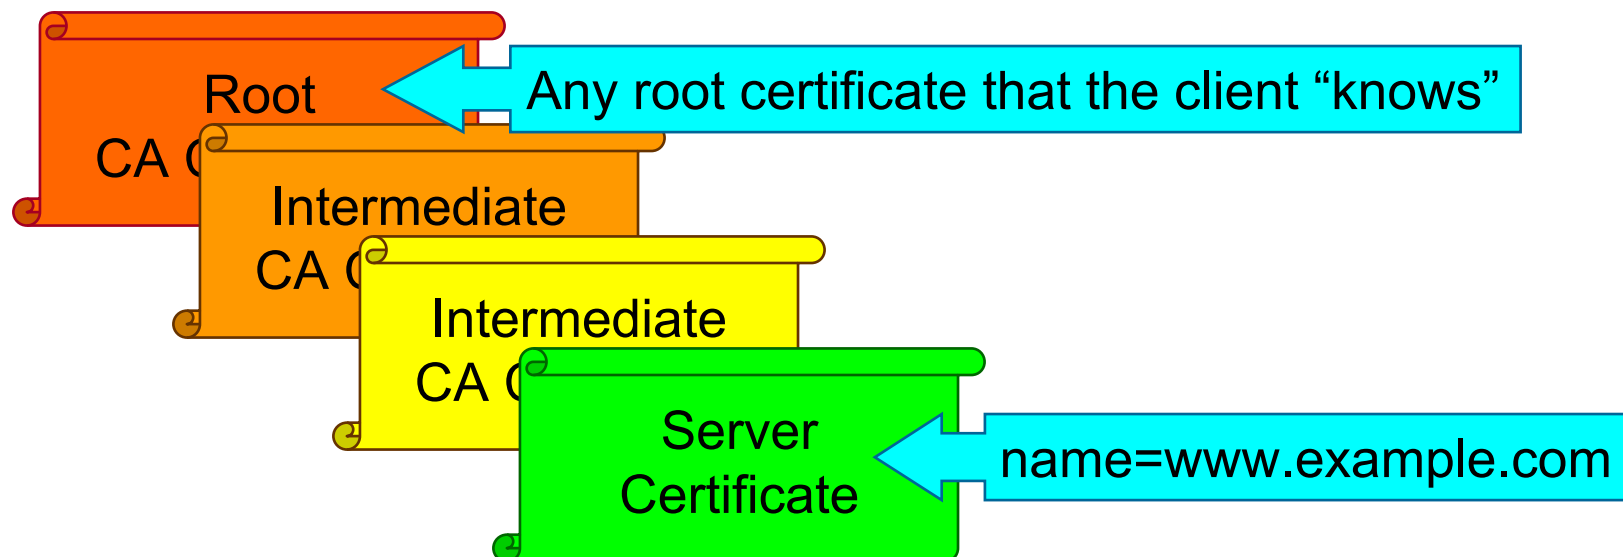
Miscellaneous

- New developments: security without global PKI

DNS-based authentication of named entities (DANE)

Global PKI, in theory

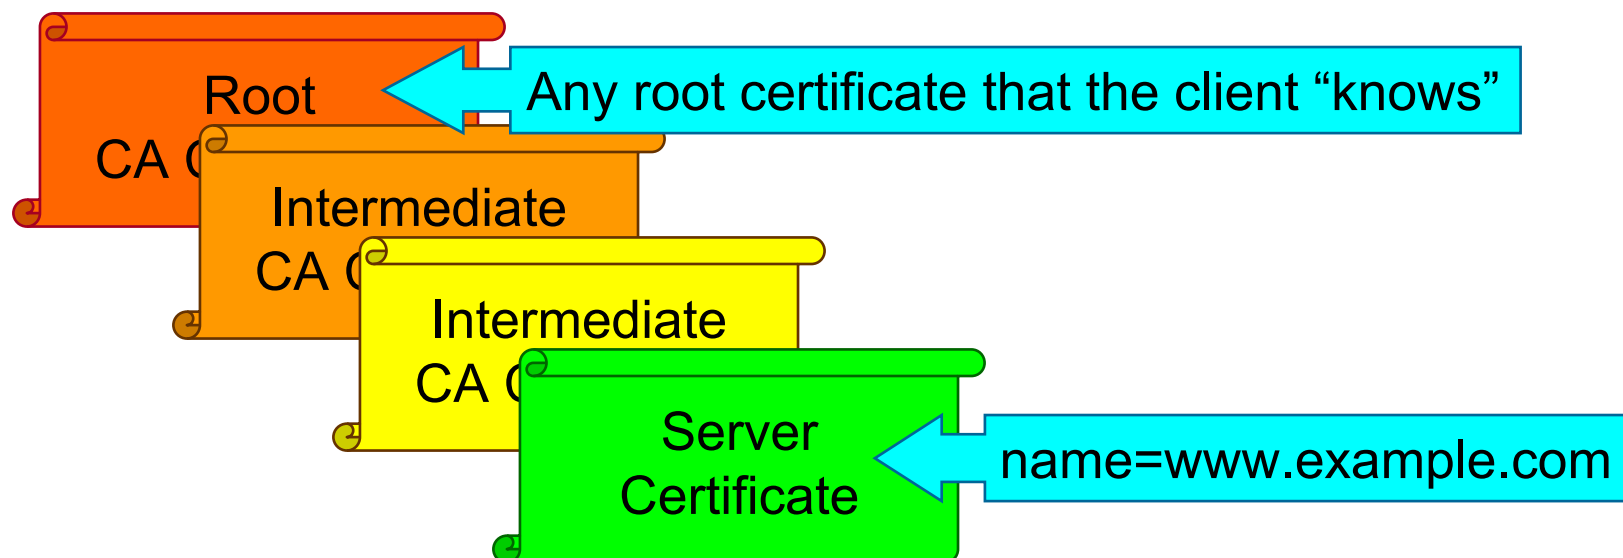
- Example: server authentication without trusting IP or DNS.
 - Client connects to server.
 - Client verifies name in server certificate.



Security without global PKI

In practice, PKI violates the principle of least privilege

- Hundreds of root CA certificates (Windows ~350, IOS ~200).
 - Owned by ~100 distinct organizations world-wide.
 - Hundreds (or more) registration authorities (RAs) world-wide.



Security without global PKI

SMTP over TLS – no server certificate verification

RFC 3207, published 2002

- Problem: RFC does not require certificate *name* verification.
 - Why not the recipient domain name (*example.com* below)?
 - One mail server may host many domains (RFC predates SNI).
 - Why not the mail server hostname (*spamfilter.example* below)?
 - The mail server hostname is looked up with insecure DNS.

Simplified connection setup procedure

DNS query: example.com MX?

DNS reply: example.com MX spamfilter.example

DNS query: spamfilter.example A?

DNS reply: spamfilter.example A 192.0.2.1

Negotiate TLS with host = 192.0.2.1, port = 25

Security without global PKI

SMTP over TLS – downgrade vulnerability

- Problem: the client doesn't know that it should use TLS.
 - <user@example.com>, not <smtps://user@example.com>.
 - Plaintext is the default, TLS is optional.

Without downgrade attack

S:	220 server.example.com
C:	EHLO client.example.org
S:	250-server.example.com 250 STARTTLS
C:	STARTTLS
S:	220 Ready to start TLS
<i>No plaintext from here on</i>	

With man-in-the-middle downgrade attack

S:	220 server.example.com
C:	EHLO client.example.org
S:	250 server.example.com (No STARTTLS announcement)
C:	MAIL FROM:<user@example.org>
S:	250 Sender address accepted
<i>Plaintext throughout the entire session</i>	

Security without global PKI

RFC 6698 (DANE) to the rescue

DNS-Based Authentication of Named Entities



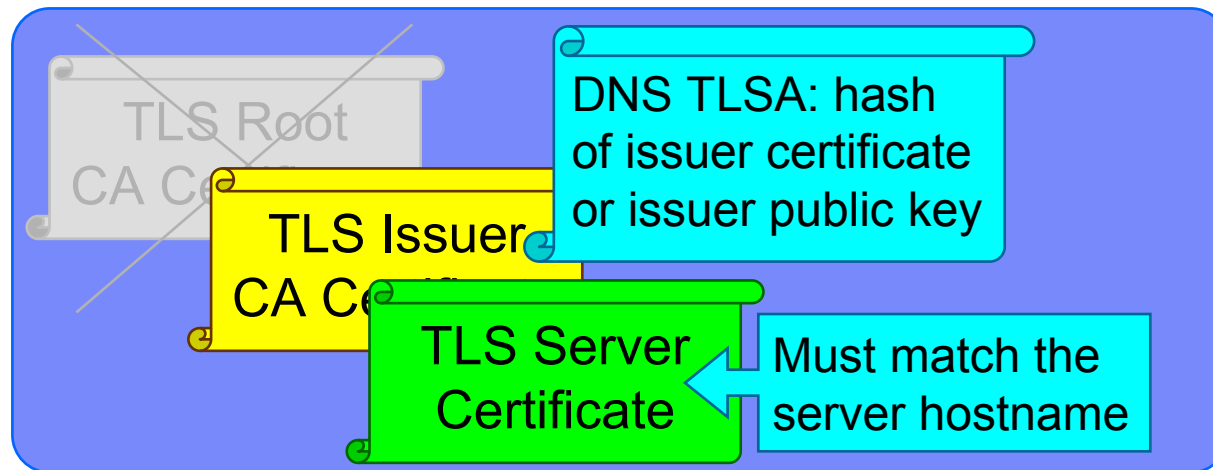
- Introduces TLSA¹ DNS records with:
 - Expected server (or issuer) certificate or public key.
 - Or better: their SHA-256 or SHA-512 hash.
- Requires secure DNS (DNSSec).
 - Unavoidable when using DNS for secure authentication.

¹RFC 6698: "TLSA does not stand for anything".

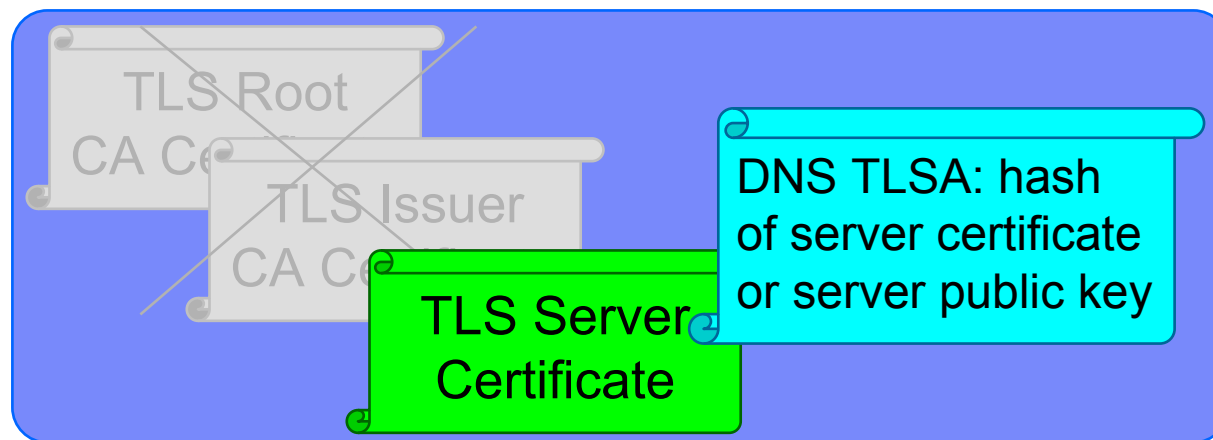
Two preferred (SMTP) DANE deployments

(the other two modes combine DANE with conventional PKI)

- Mini PKI



- No PKI



Security without global PKI

Simplified example with debian.org

- Look up the debian.org mail server names:

Reply: `debian.org MX 0 mailly.debian.org`
`debian.org MX 0 muffat.debian.org`

- Look up mailly A records:

Reply: `mailly.debian.org A 82.195.75.114`

- Connect to 82.195.75.114 port 25.

- Look up mailly TLSA records:

Reply: `_25._tcp.mailly.debian.org TLSA 3 1 1 [SHA-256 of TLS server public key]`

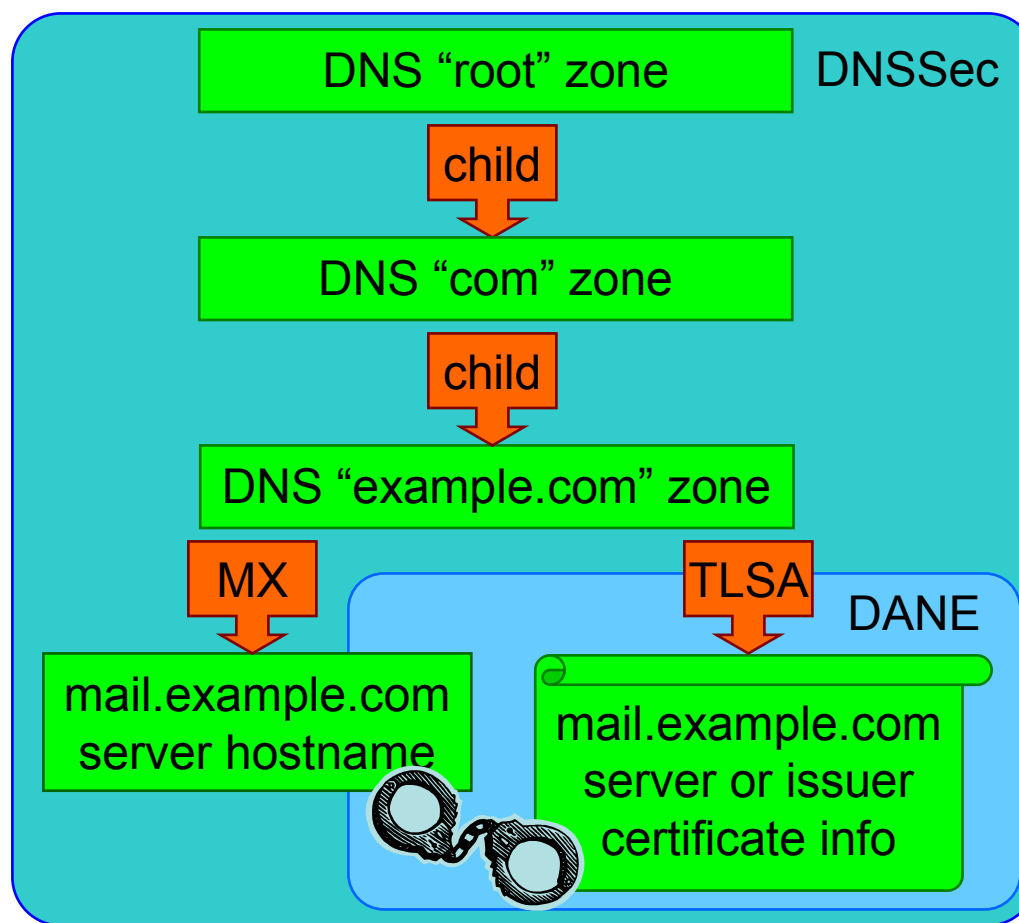
- Match TLSA record with SHA-256 of TLS server public key.

Not showing:
DNSSEC signatures
(RRSIG records)

Security without global PKI

Securing SMTP with DNSSec and DANE

- Minimized trust.
 - Not: 100s of RAs.
 - Secure copy of root zone public keys.
 - DNS target zone plus its ancestors.
 - Maybe: issuer cert.
- No downgrade attack.
 - Use TLS when DNS TLSA record exists.



Security without global PKI

Postfix 2.11 – two new DANE TLS profiles

none	No TLS (always use plaintext)
may	Opportunistic TLS (use TLS if available, otherwise plaintext)
encrypt	Mandatory TLS (no plaintext)
verify	Mandatory server certificate signature & name verification
secure	Verify, with pre-determined name (prevent DNS spoofing)
dane	Opportunistic DANE TLS (use TLSA record if available)
dane-only	Mandatory DANE TLS (require TLSA record is available)
fingerprint	Certificate or public-key fingerprint verification (no PKI)

Security without global PKI

DANE support in Postfix 2.11 stable release

- Aug 2012: RFC 6698 is published.
- Q1 2013: Start of Postfix implementation (Viktor Dukhovni).
- Jan 2014: DANE support in Postfix stable release.
 - Requires DNSSEC validating resolver (e.g., BIND or unbound).
- Please try DANE support, but be prepared.
 - A few DNS servers mis-handle TLSA queries.
 - Use “opportunistic DANE” as default.
 - Use “non-DANE profile” SMTP TLS policy for problem sites.
 - See TLS_README (or the upcoming DANE_README).

Security without global PKI

- New developments: LMDB database support

Unintended consequences of adopting AGPL

June 2013: Oracle updates Berkeley DB 6.0 license

Popular open-source key-value store

- Berkeley DB v5: two licenses, copyleft and commercial.

Copyleft: make all source code available if you *distribute work* that uses Berkeley DB.

- Berkeley DB v6: two licenses, AGPLv3¹ and commercial.

AGPL: *also* make all source code available if you *provide network service* that uses Berkeley DB.

- Problem: cannot combine GPLv2 with AGPLv3 code, without relicensing all the GPLv2 code (GPLv3 would be OK).

- ¹GNU Affero General Public License version 3.0. Pronunciation: /'af.fe.ro:/.

Replacing Berkeley DB


LMDB - Lightning Memory-Mapped Database

Author: Howard Chu

- Described by some as a Berkeley DB replacement.
- OpenLDAP Public License.
- Memory-mapped, max size limited by memory address range (typically ~31 bits on i386, ~47 bits on x86_64 or ~128 TB).
- Copy-on-write, zero-copy, MVCC, multi-reader, single-writer.
- Ported by its author to dozen+ other open source projects.
- Postfix integration took 5 iterations with changes to both Postfix and LMDB.


Replacing Berkeley DB

Challenges integrating LMDB into Postfix

- Hard database size limit, specified when database is opened.
 - Postfix processes fail unexpectedly if size limit is set too low.
 - LMDB 0.9.8 allows Postfix to resize database on the fly.
- LMDB lockfile must be writable by readers. Hard limit on number of readers, specified when database is opened.
 - World-writable lock files, for example under /etc/postfix. 
 - Postfix process fail unexpectedly if reader limit is set too low.
 - LMDB 0.9.9 allows Postfix to use external (fcntl()-based) locks.

Replacing Berkeley DB

Challenges integrating LMDB into Postfix

- Information leak: writing ~4kbyte chunks of uninitialized heap memory to the LMDB database file.
 - Contains traces from past activity in the same process, not necessarily meant to persisted or shared. 
 - LMDB 0.9.10 initializes malloc()ed memory by default.
- LMDB library functions rely on assert() extensively.
 - Writes message to stderr and aborts the program immediately.
 - Postfix daemons fall out of the sky without logging any error.
 - LMDB 0.9.11 allows Postfix to log an error message.

Replacing Berkeley DB

LMDB support in Postfix 2.11 stable release

- First persistent Postfix database that safely supports multiple writers such as `postscreen(8)` or `verify(8)`.
- *Not exactly* a Berkeley DB drop-in replacement – requires additional Postfix code to recover from “hard limit” errors.
- Expect better safety than Berkeley DB, mainly due to COW.
- Not necessarily faster, due to Postfix’s small transactions.

Replacing Berkeley DB

Postfix lessons learned

- Invent sparingly: don't re-invent what works.
 - SMTP, Milter, maildir, Sendmail lookup tables.
- Build the stable protocols into Postfix.
 - SMTP, LMTP, TLS, SASL, IPv6, DSN, MIME, LDAP, SQL, CDB, memcache, LMDB, (DANE).
- Plan for change: provide safe plug-in interfaces for future proofing.
 - Anti-Spam, Anti-Virus, DKIM, SenderID, SPF, greylist.
- Optimize both worst cases and common cases.
 - On the Internet, worst cases will become common cases.
- Don't let a C prototype become the final implementation.

Conclusion