BIND 9

Making a modern DNS Server

Ondřej Surý @ ISC

ISC Internet Systems Consortium

BIND 9 - The History

- First released in y2k
- Written from scratch
- Design by Contract
    - Rather crash than overwrite memory
- First DNSSEC implementation

Photo by Mr Cup / Fabien Barral on Unsplash

BIND 9.11 ESV

- Released in 2016
- Extended Support Version
- Under Mozilla Public License
- New Features:
  - Catalog Zones
  - Addzone/Delzone Provisioning
  - DNSSEC Key Manager
  - Negative Trust Anchors
  - DNSTAP
  - DNS Cookies
  - Minimal responses to ANY

Photo by Nathan Dumlao on Unsplash

BIND 9.12

- Released in 2018
- NSEC Aggressive Use
- Serve Stale (TTL Stretching)
- Response Policy Interface
- Major Refactoring
- Speedup factor: 1.25-6
- CDS/CDNSKEY tools
- ED25519 Support
- Obsoleted by BIND 9.14

Photo by Fachry Zella Devandra on Unsplash

BIND 9.14

- Released in March 2019
- New release schedule
- Refactoring and Modernization
- New features:
  - QNAME Minimization
  - Mirror Zones
  - Plugins for Query-Response Processing

Photo by Drew Hays on Unsplash

Modernization

- On *NIX, BIND requires:
  - C99 Support in Compiler
  - POSIX Threads
    - Advanced Sockets API for IPv6
    - Standard Atomic <stdatomic.h>
- Support for lot of old systems dropped

Photo by Samuel Zeller on Unsplash

Crypto Refactoring

- PKCS#11 now used only for public-key cryptography
    - OpenSSL is mandatory
- Performance Improvements
    - Task Manager is now multithreaded
    - Socket Code has multiple event loops

Photo by Alex Jones on Unsplash

Performance Improvements

- Extra 100k or more QPS Improvement
- Due to Refactoring
    - Task Manager is now multithreaded
    - Socket Code has multiple event loops

Photo by Alex Jones on Unsplash

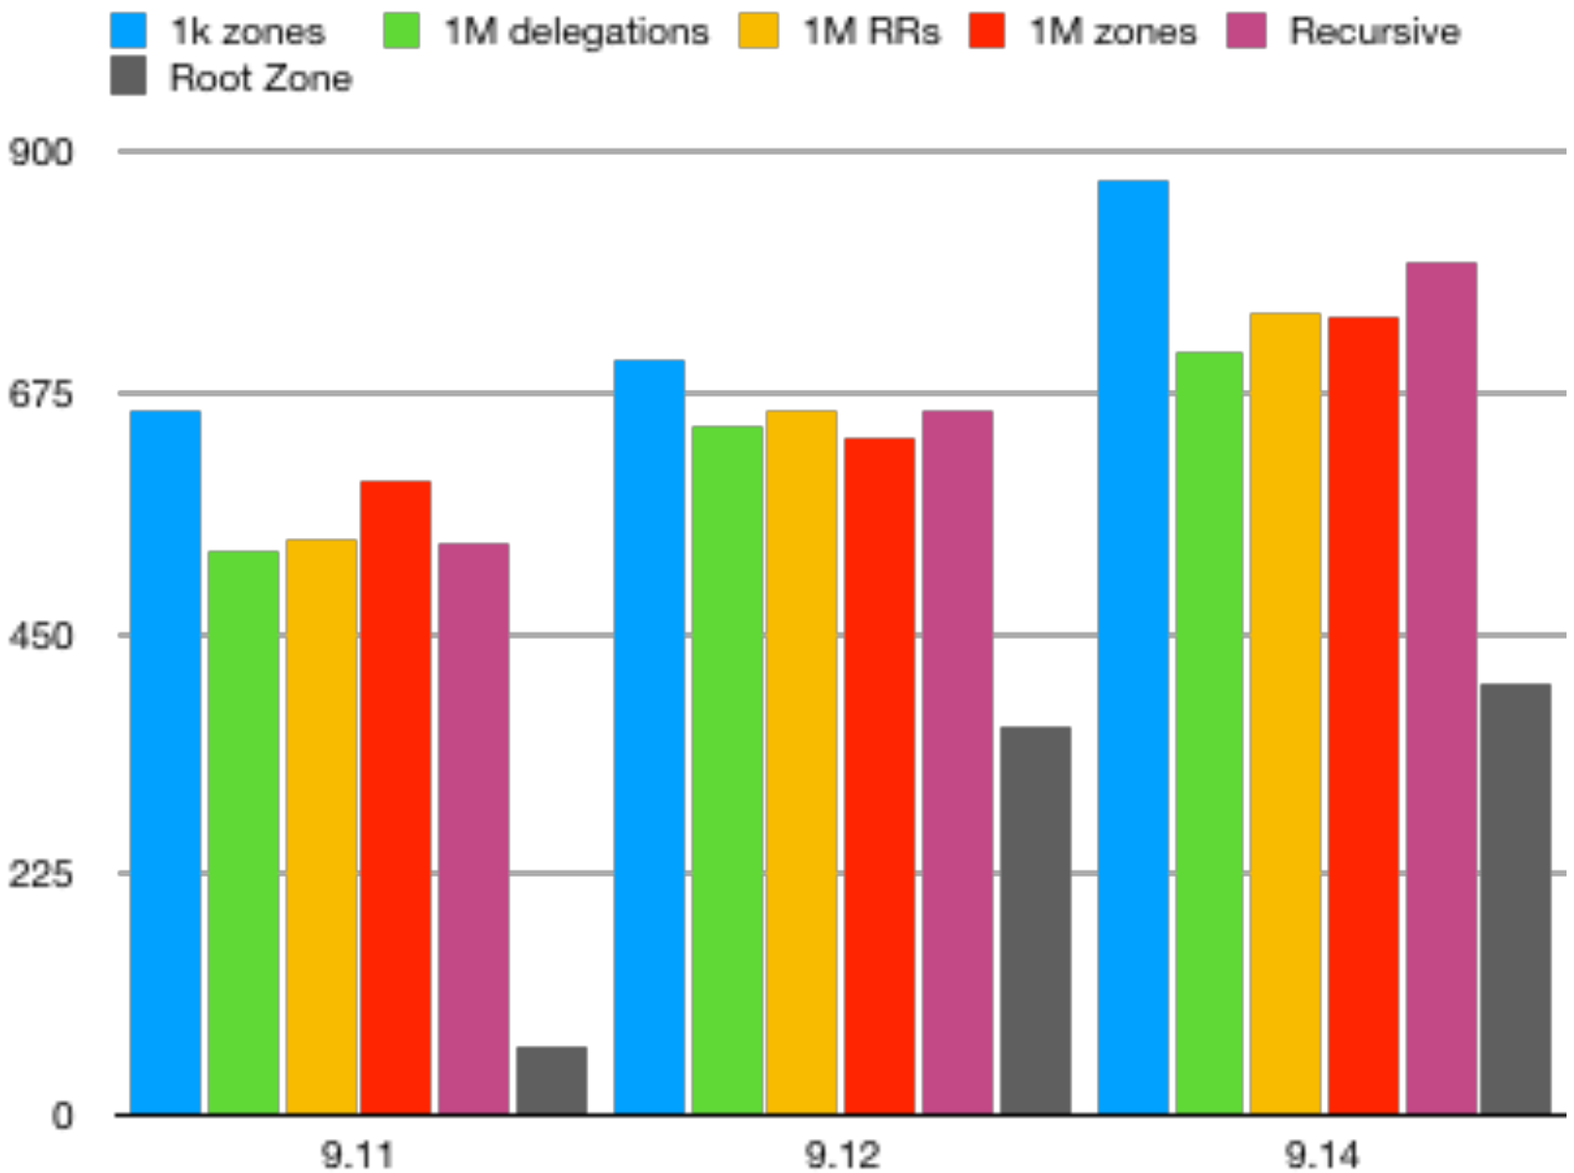

Performance

Responses Per Second

QNAME Minimization

- Defined in RFC7816
- Improves DNS Privacy
- Protects DNS transactions
    - Resolver sends only the minimal info needed to resolve the query
- Enabled by default in a relaxed mode in BIND 9.14

Photo by Josh Shaw on Unsplash

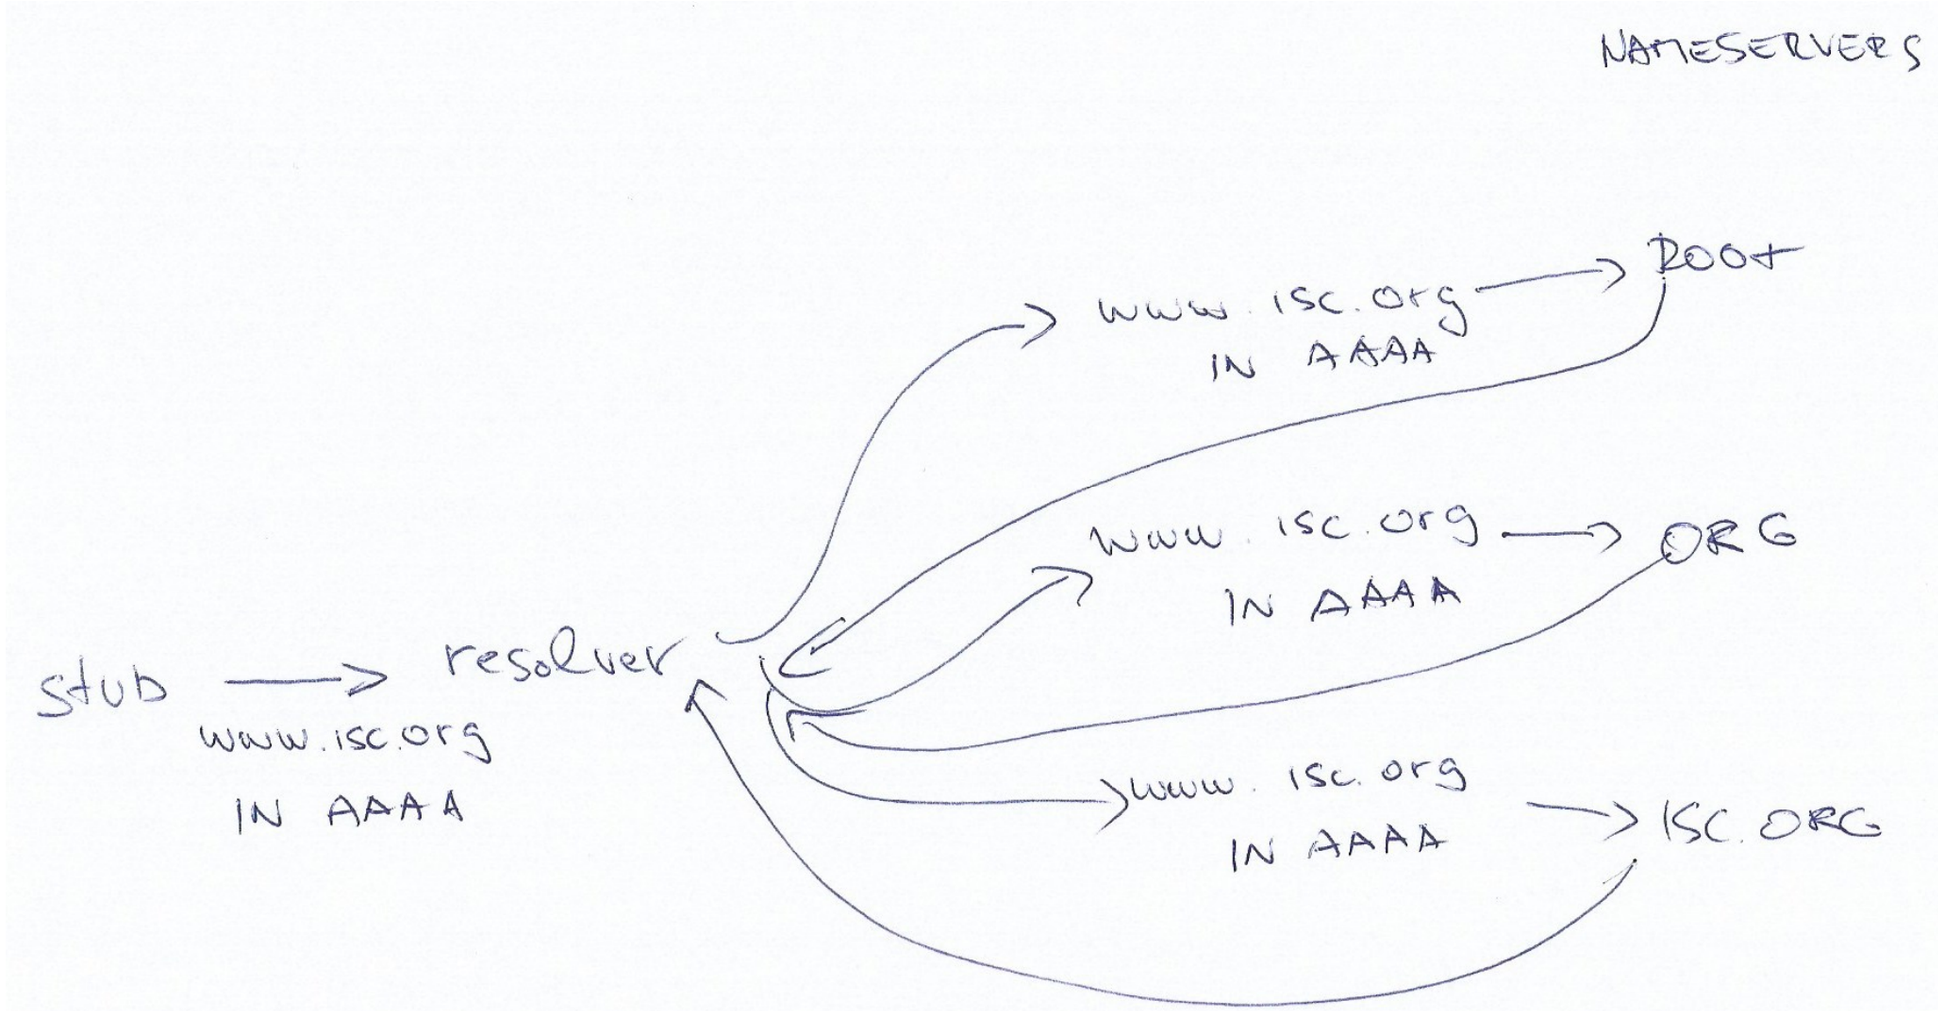

QNAME Minimization

Normal DNS Traffic

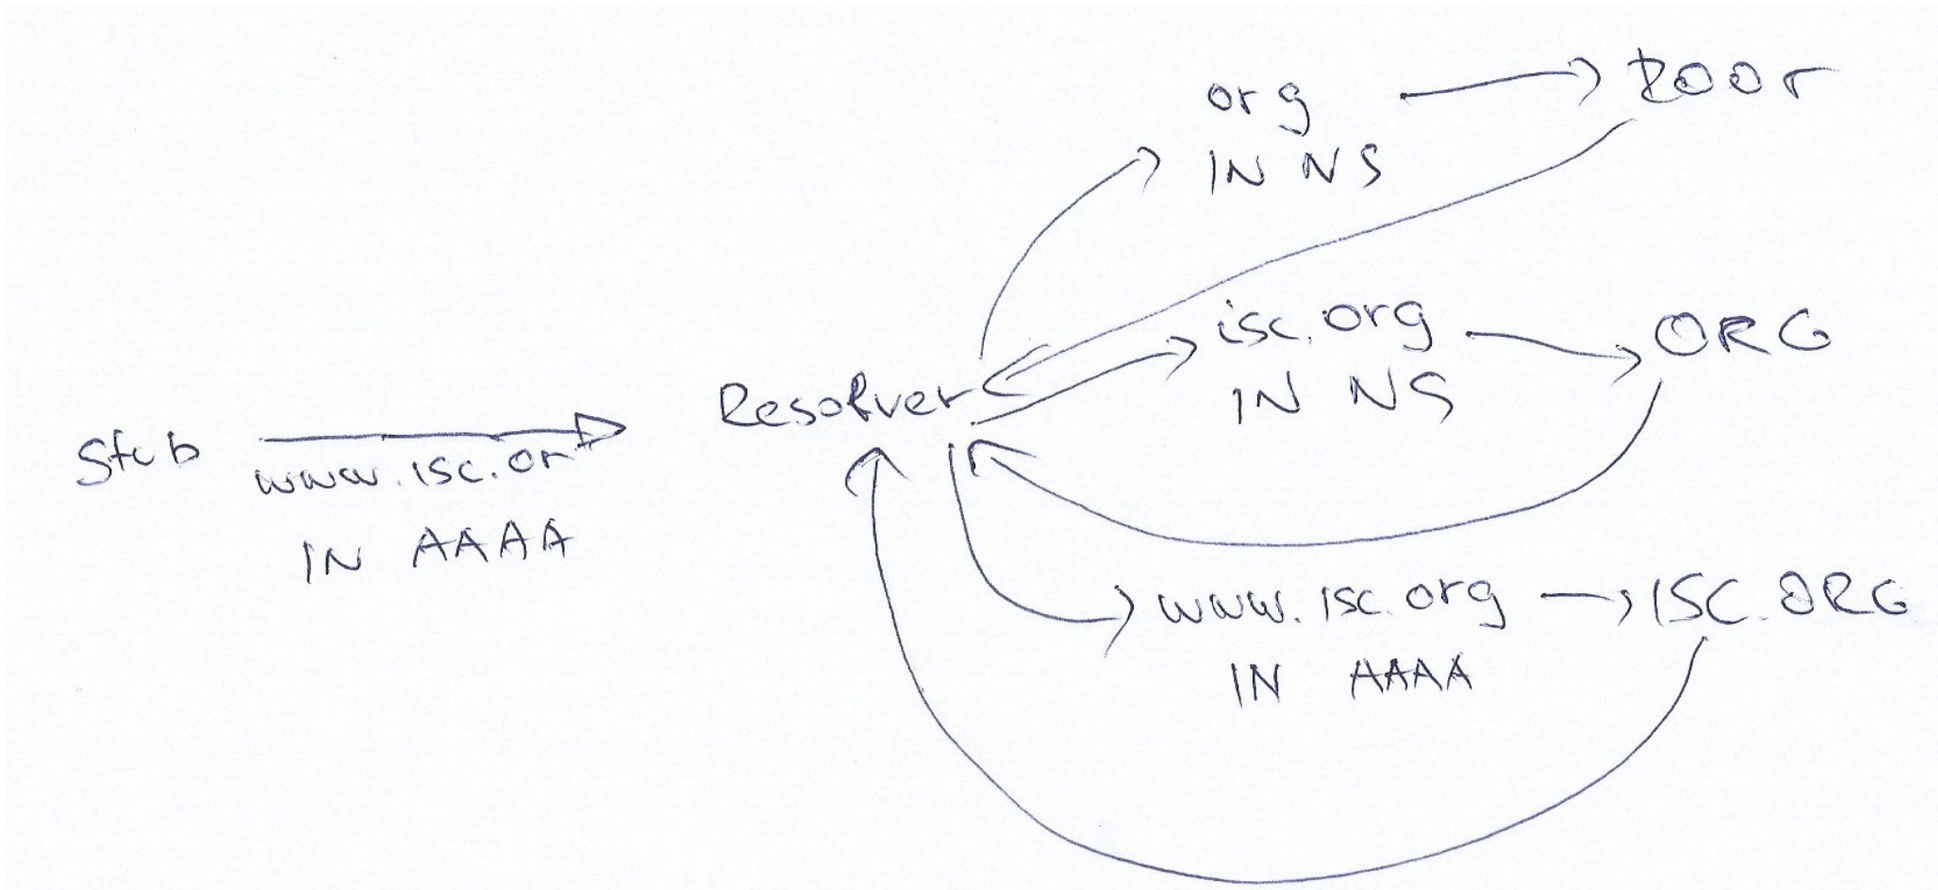

QNAME Minimization

Minimized DNS

Root Zone Local Copy

- Defined in RFC7706
- Reduced traffic to the Root
- Transfer (XFR) RZ from well-known sources
- DNSSEC Validated
- Root Zone used only internally

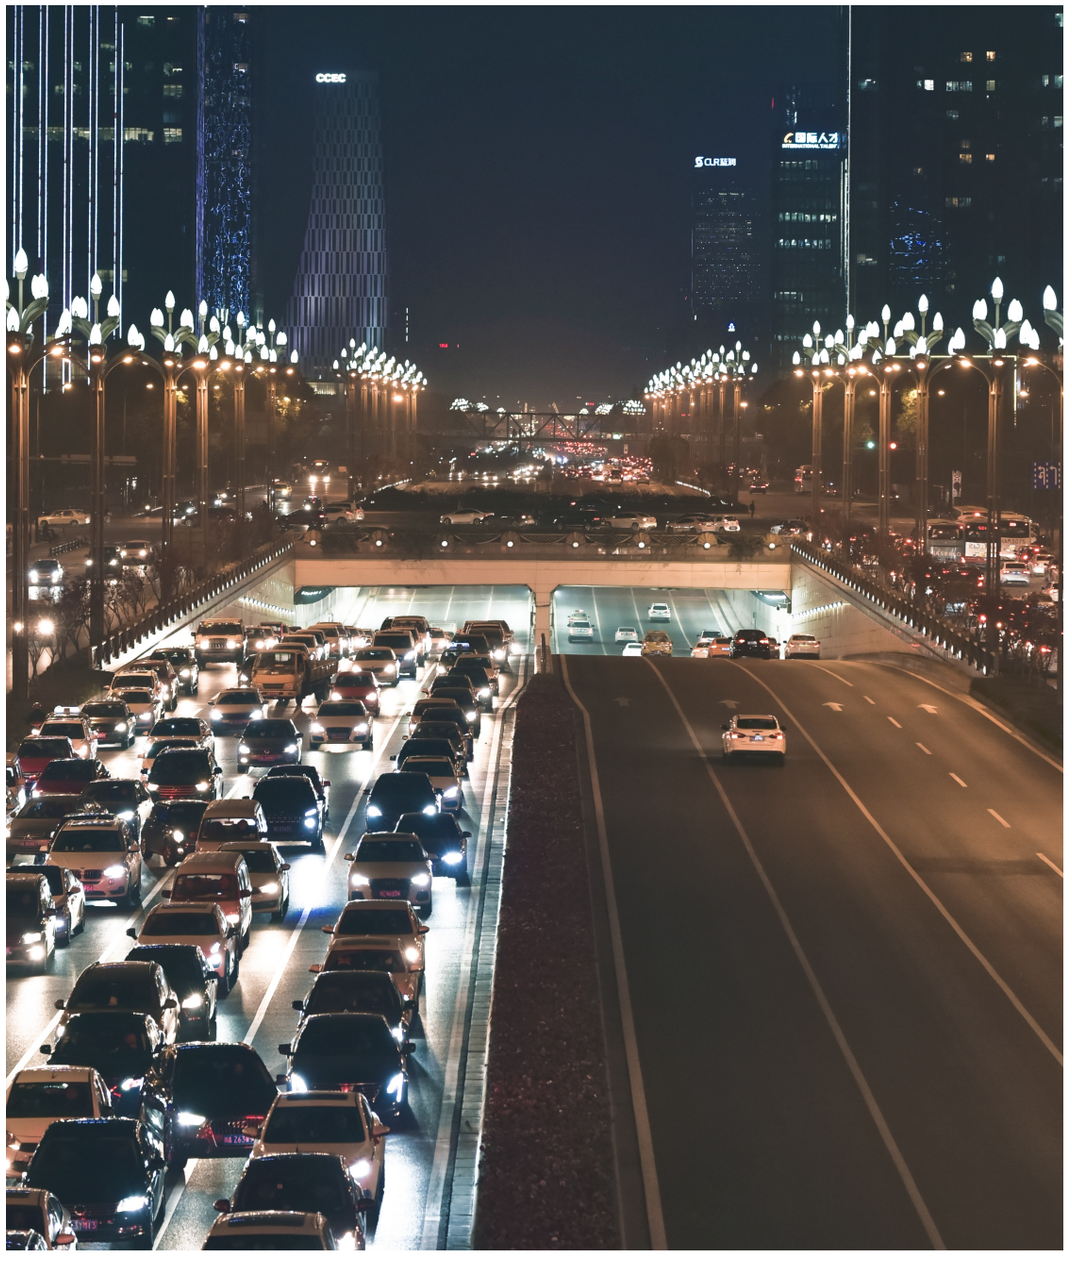

Photo by 7 SeTh on Unsplash

BIND 9.15 Plans

- Performance
- Management
- Security & Privacy
- Operations

Photo by Daniel McCullough on Unsplash

Performance

- Improve BIND performance, so you don’t have to care
- Network stack rewrite / reengineering
- Improve both UDP and TCP performance
- Using libuv
    - Full-featured event loop backed by epoll, kqueue, IOCP, event ports
    - Asynchronous TCP and UDP sockets
    - And more…
        - Thread pools
        - Signal handling
        - High resolution clock
- Using external library allows us to focus on DNS

Management

- Metrics, statistics, …
  - Too many?
  - Too few?
  - What’s missing?
  - What’s extra?
- Update Catalog Zones

Photo by Kyle Glenn on Unsplash

DNS over TLS

- Defined in RFC 7858 & RFC 8310
- Running on port 853
- Add as a module to new networking stack
- Support in clients:
  - CLI Tools: dig, delv, …
  - Forwarder
- Support in server:
  - Stub to BIND Resolver

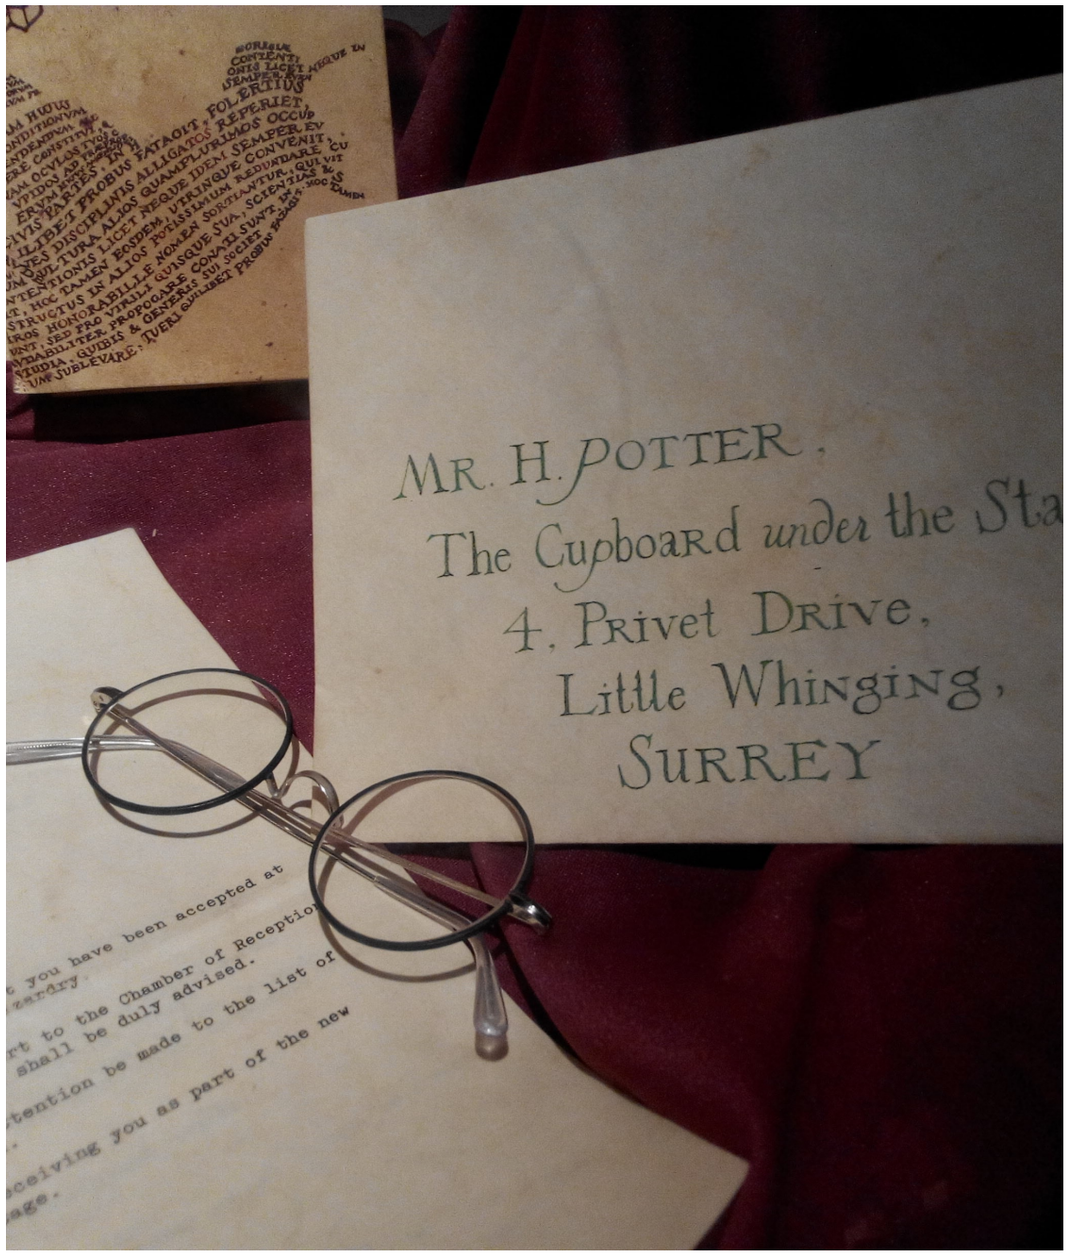

Photo by Rae Tian on Unsplash

DNS over HTTPS

- RFC 8484
- Mixed traffic with HTTP
- Add counterweight to “DNS Silos”
- Support in clients
- Support in server via proxy:
  - Add support for “proxy” protocol (acl, RPZ, …)
  - Proxy module for popular webservers (nginx, apache2) as cross-vendor project

Photo by Pawel Nolbert on Unsplash

DNSSEC Made Easy

- Go from “DNSSEC in 6 minutes” to simple “Yes”
- Keep the existing tools (dnssec-signzone, …)
- Pick reasonable defaults
    - Elliptic Curve keys
- Automate everything
    - Key creation
    - Key rotation
    - Child-to-parent (both sides)
    - Periodic signing
- Add support for:
    - Offline KSKs
    - Combined Signing Keys (CSK)

Photo by John Salvino on Unsplash

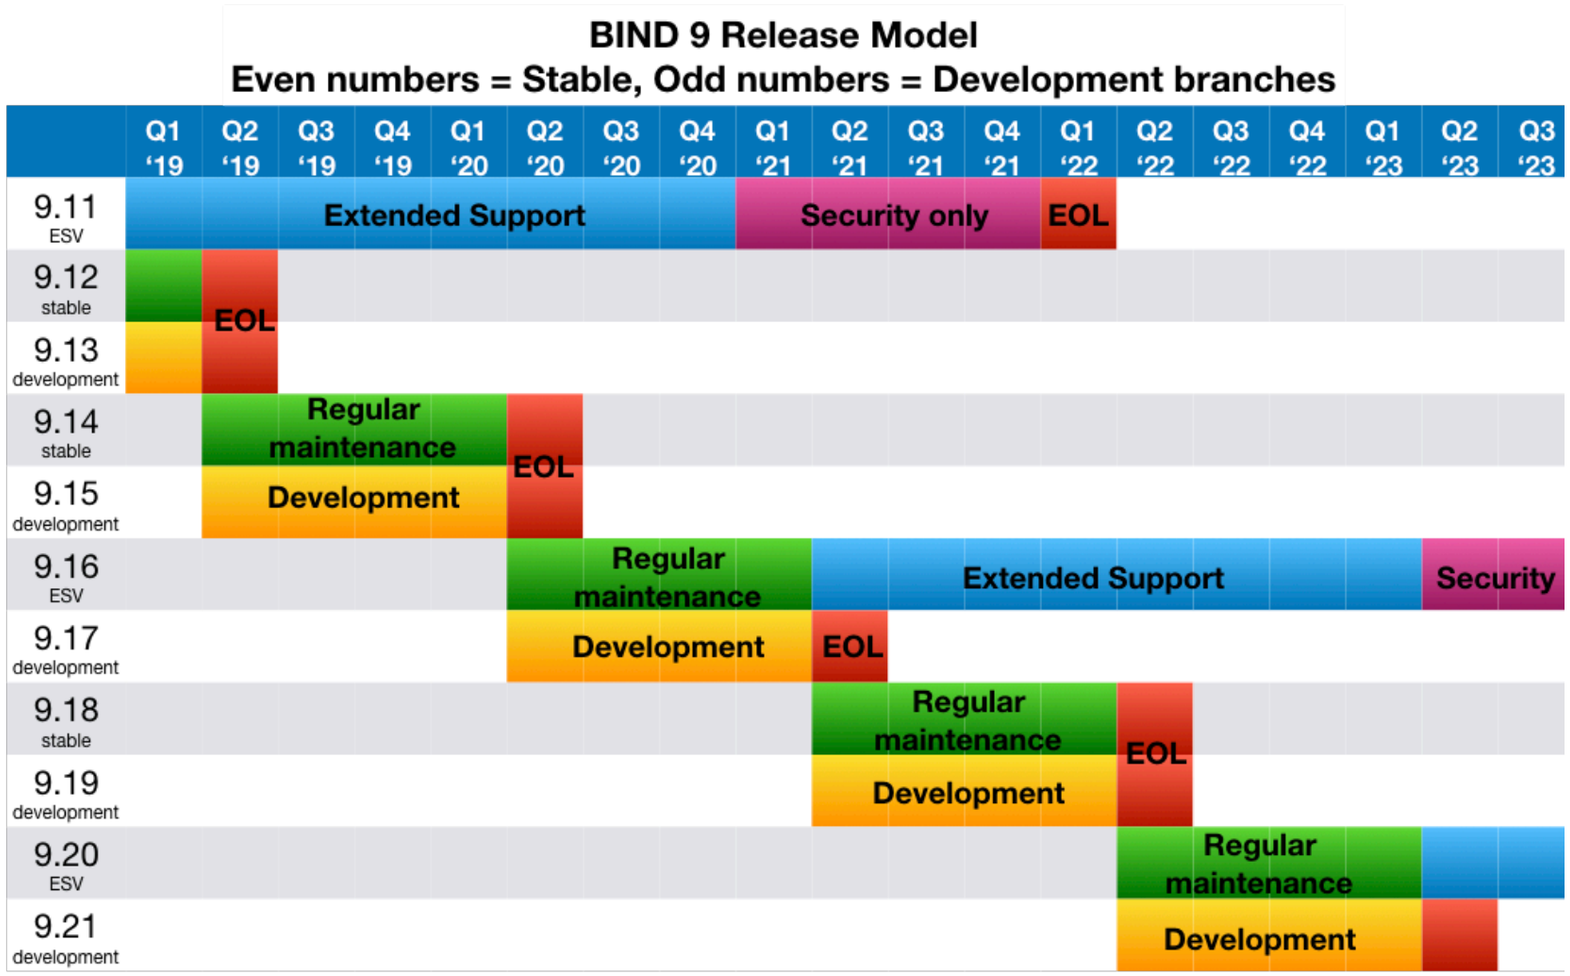

Release Schedule

Predictable (Time-Based) Releases

Photo by Evan Dennis on Unsplash

Questions?

Internet Systems Consortium