DNS Hackathon

Revisiting DNS…

David Huberman
Nicolas Antoniello
Carlos Martinez

LACNIC 38 – LACNOG 2022

03 October 2022

Once upon a time…

Names and Numbers

- Devices are identified over the Internet using IP addresses.
  - IPv4: 192.0.2.7
  - IPv6: 2001:db8::7
- Whie IP addresses are easy for machines to use, people prefer to use names.
- In the early days of the Internet, names were simple
  - No domain names yet
  - “Single-label names”, 24 characters maximum
  - Referred to as ***host names***
  - No big security concerns

Rise of the DNS !

Domain Names

- DNS database structure is an inverted tree called the ***name space*** (*think of DNS as a huge distributed warehouse to store information meant to be accessible for anyone in the Internet*).
- Every node (except the root) has a label called ***domain name.***
- That ***domain name*** is built by sequencing node labels from one specified node up to the root, separated by dots (this unambiguous identification of a domain name is often called ***fully qualified domain name (FQDN)***.

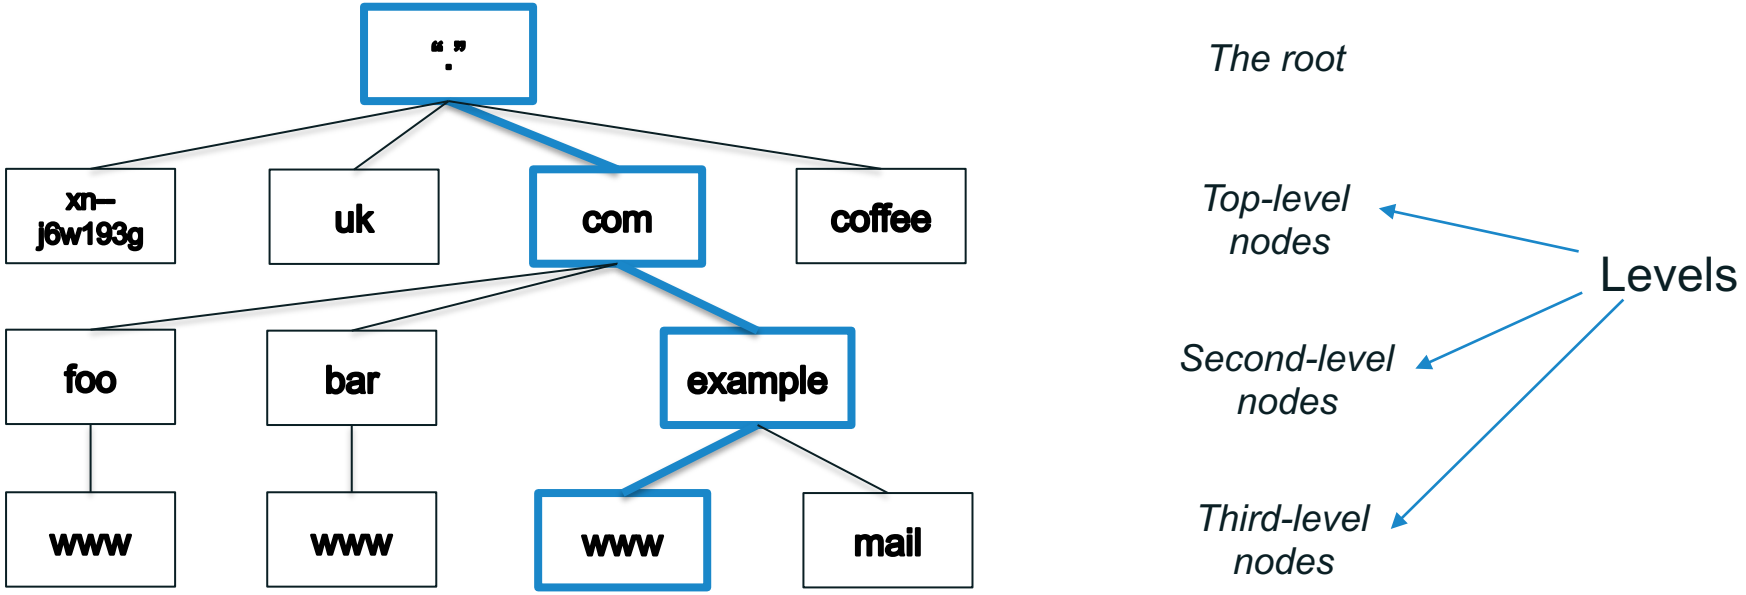

Domains and Delegation of Administration (Zones)

- Now, if we manage a huge distributed warehouse of information meant to be accessible for anyone in the Internet, we'll need some help to operate it and some standards to be able to store, manage and exchange that info.
- And there it comes the ***domain*** definition and the possibility of ***delegating*** the management of parts (***zones***) of the warehouse (DNS) to allow distributed administration (so administrative divisions are called *zones*).
- A *zone* then would be a shelf in our warehouse for storing data. While a *domain* is defined as a node and everything below it (all the subtree under a node).

Delegation creates zones:

- Delegating zone is the ***parent***
- Created zone is the ***child***

Some of the data stored in zones:

- IPv4 address for a name (A)
- IPv6 address for a name (AAAA)
- Mail server for a name (MX)

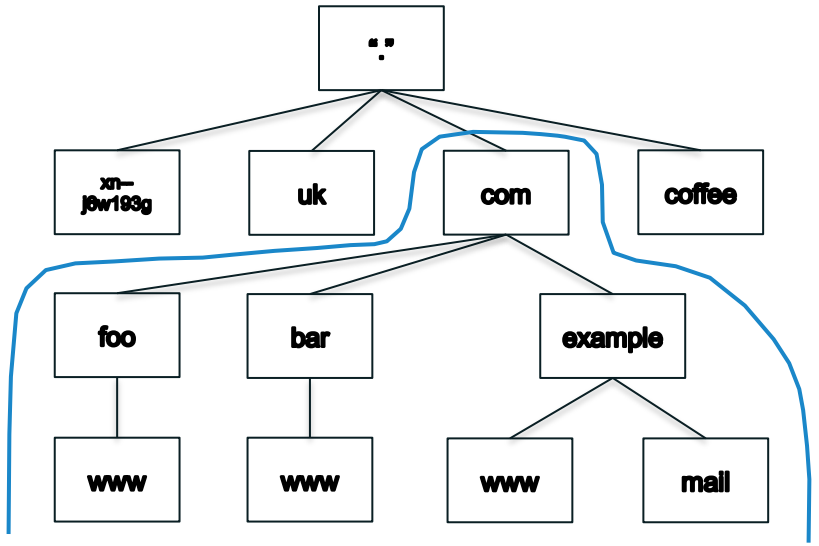

"**com**" domain

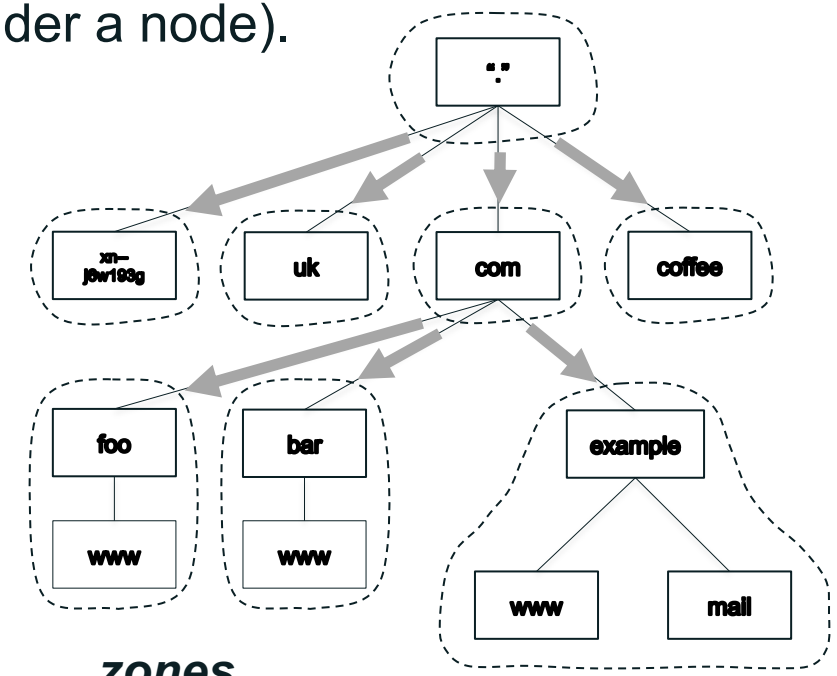

zones

Name Resolution Process in Action…

Resolution Process Definitions

- Recall DNS is a distributed database:
  - Data is maintained locally (into our huge and distributed warehouse) but available globally.
- ***Resolvers:*** send queries (this are like service providers that find the data for us, so we don't have to sneak into the warehouse shelves ourselves).
- ***Name servers:*** answer queries (the boxes or containers in our warehouse's shelves that contains all the data).
- The ***resolution process*** is the implementation of translating from an IP address to a domain name, or more general getting the answer for a specific query.

Additional System Optimization

- ***Caching***: basically, this service providers (***resolvers***) may remember the info they find in the warehouse so as they don't have to search again for it each time someone asks for the same (so remember caching happens in the *resolvers*).

- ***Replication***: it's wise to maintain several copies of our shelves (all containing the same info) so as to be able to place them as near to the service providers (resolvers) as we can. Thus providing less resolution times (info storage is closer), load balancing (as provides access different shelves instead of all the same), and of course makes the whole system more robust and resilient (using ***Anycast***).

DNS Components & Resolution Process

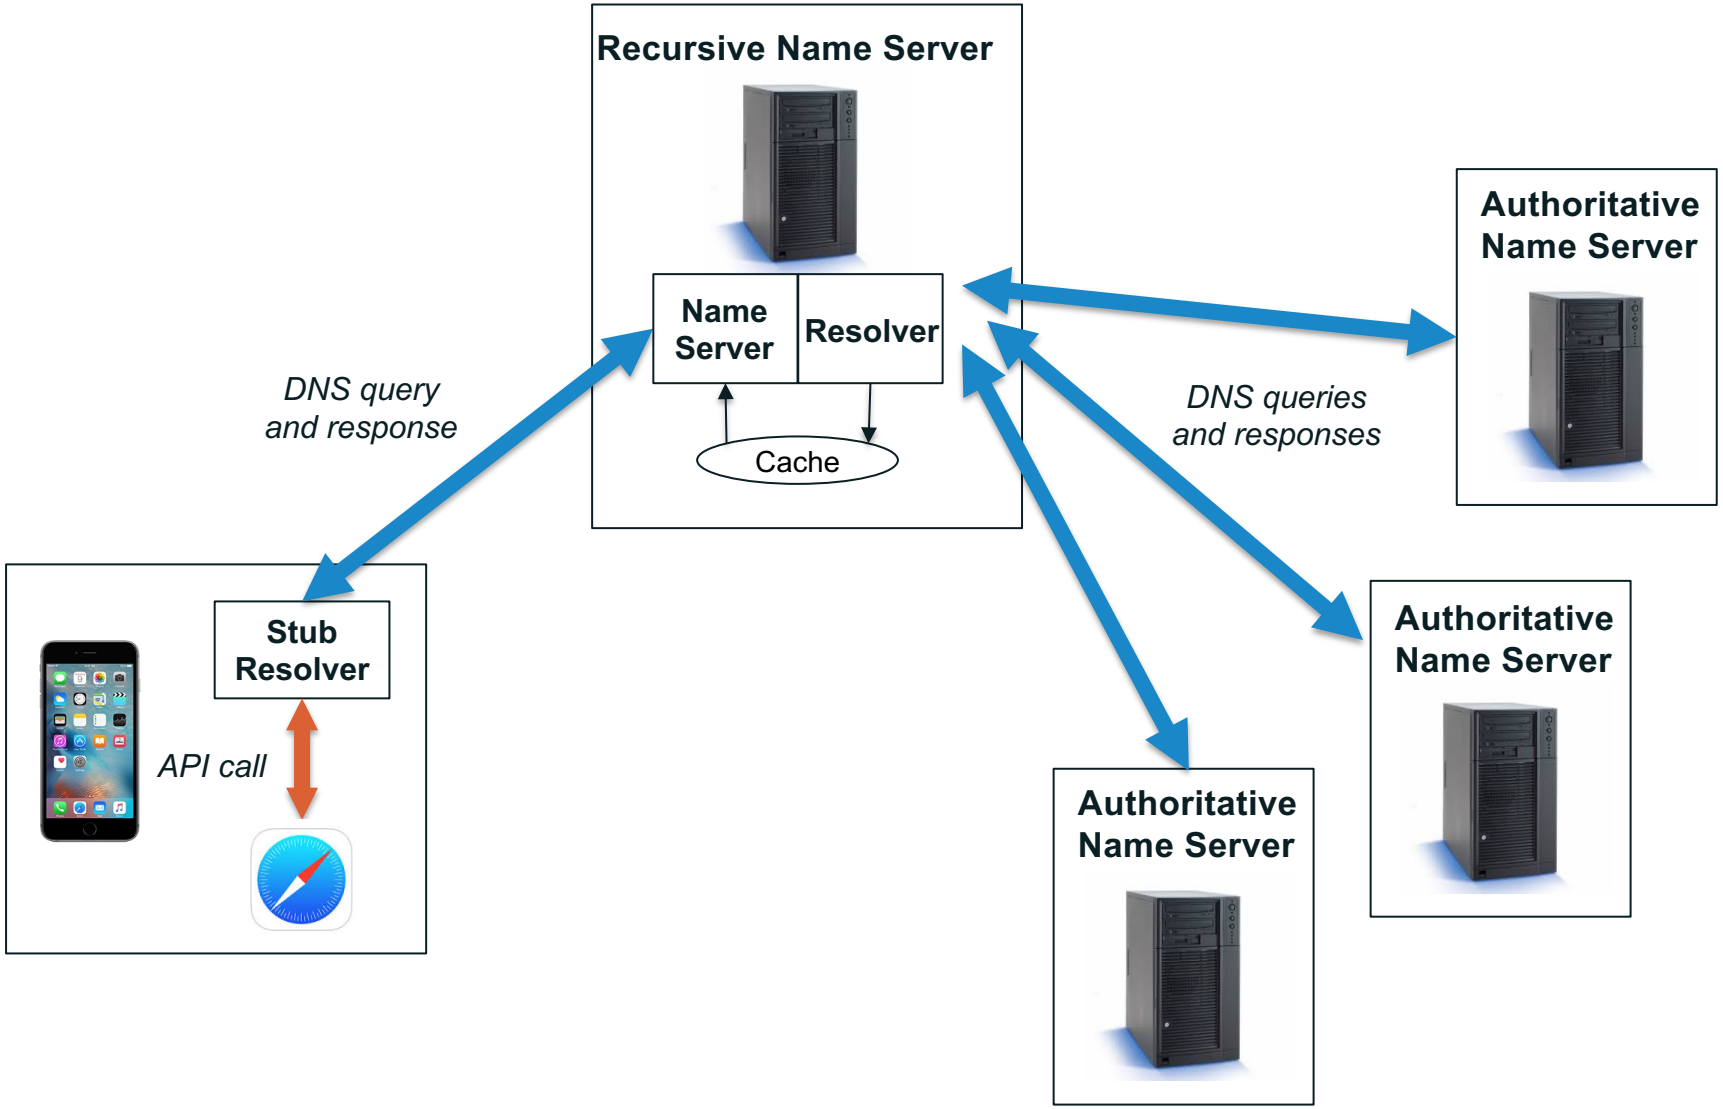

Some solution or mitigation mechanisms to consider, apply and/or deploy

Resiliency: DNS Server Copies…

Multiple NS Servers

- Zones can and should (if possible) have multiple authoritative servers:
  - Provides redundancy and resiliency
  - Distribute query load
- Zone replication is part of the DNS protocol, so this functionality is provided for in the standards and is implemented by all DNS server software.

Anycast

Anycast could be defined as a combination of IP addressing and routing scheme, where:

- the same IP address is assigned to many target devices; and
- the decision of which destination the packet will reach is decided by the network's routing mechanisms and metrics.

Anycast does not require any special configuration at the application level or at the client level. It is a transparent process for the client.

The goal is for packets to reach the closest Anycast destination based on routing metrics that the network deems important (for example, the number of hops).

Anycast for DNS servers

- Root server operators often use Anycast, distributing many instances of their root server around the world.
- Anycast is also commonly used by recursive resolution operators, distributing many instances of their recursives around the world.
- Some of the benefits of Anycast applied to DNS:
  - Provides redundancy and resiliency to the global DNS infrastructure.
  - Distributes the load of queries and responses across many servers.
  - Reduce latency by allowing more instances closer to more clients.
  - Provides more robustness, helping to mitigate events such as DoS attacks on the DNS infrastructure.
- The technique can be applied to authoritative at any level as well as recursive.
- In case of Authoritative servers, all must maintain the same information so that the response is the same regardless of which of the copies is consulted (DNS standards provides solution for this).

The Root Servers Operators

- **A** Verisign
- **B** University of Southern California Information Sciences Institute
- **C** Cogent Communications, Inc.
- **D** University of Maryland
- **E** United States National Aeronautics and Space Administration
  (NASA) Ames Research Center
- **F** Information Systems Consortium (ISC)
- **G** United States Department of Defense (US DoD)
  Defense Information Systems Agency (DISA)
- **H** United States Army (Aberdeen Proving Ground)
- **I** Netnod Internet Exchange i Sverige
- **J** Verisign
- **K** Réseaux IP Européens Network Coordination Centre (RIPE NCC)
- **L** Internet Corporation For Assigned Names and Numbers (ICANN)
- **M** WIDE Project (Widely Integrated Distributed Environment)

root-servers.org

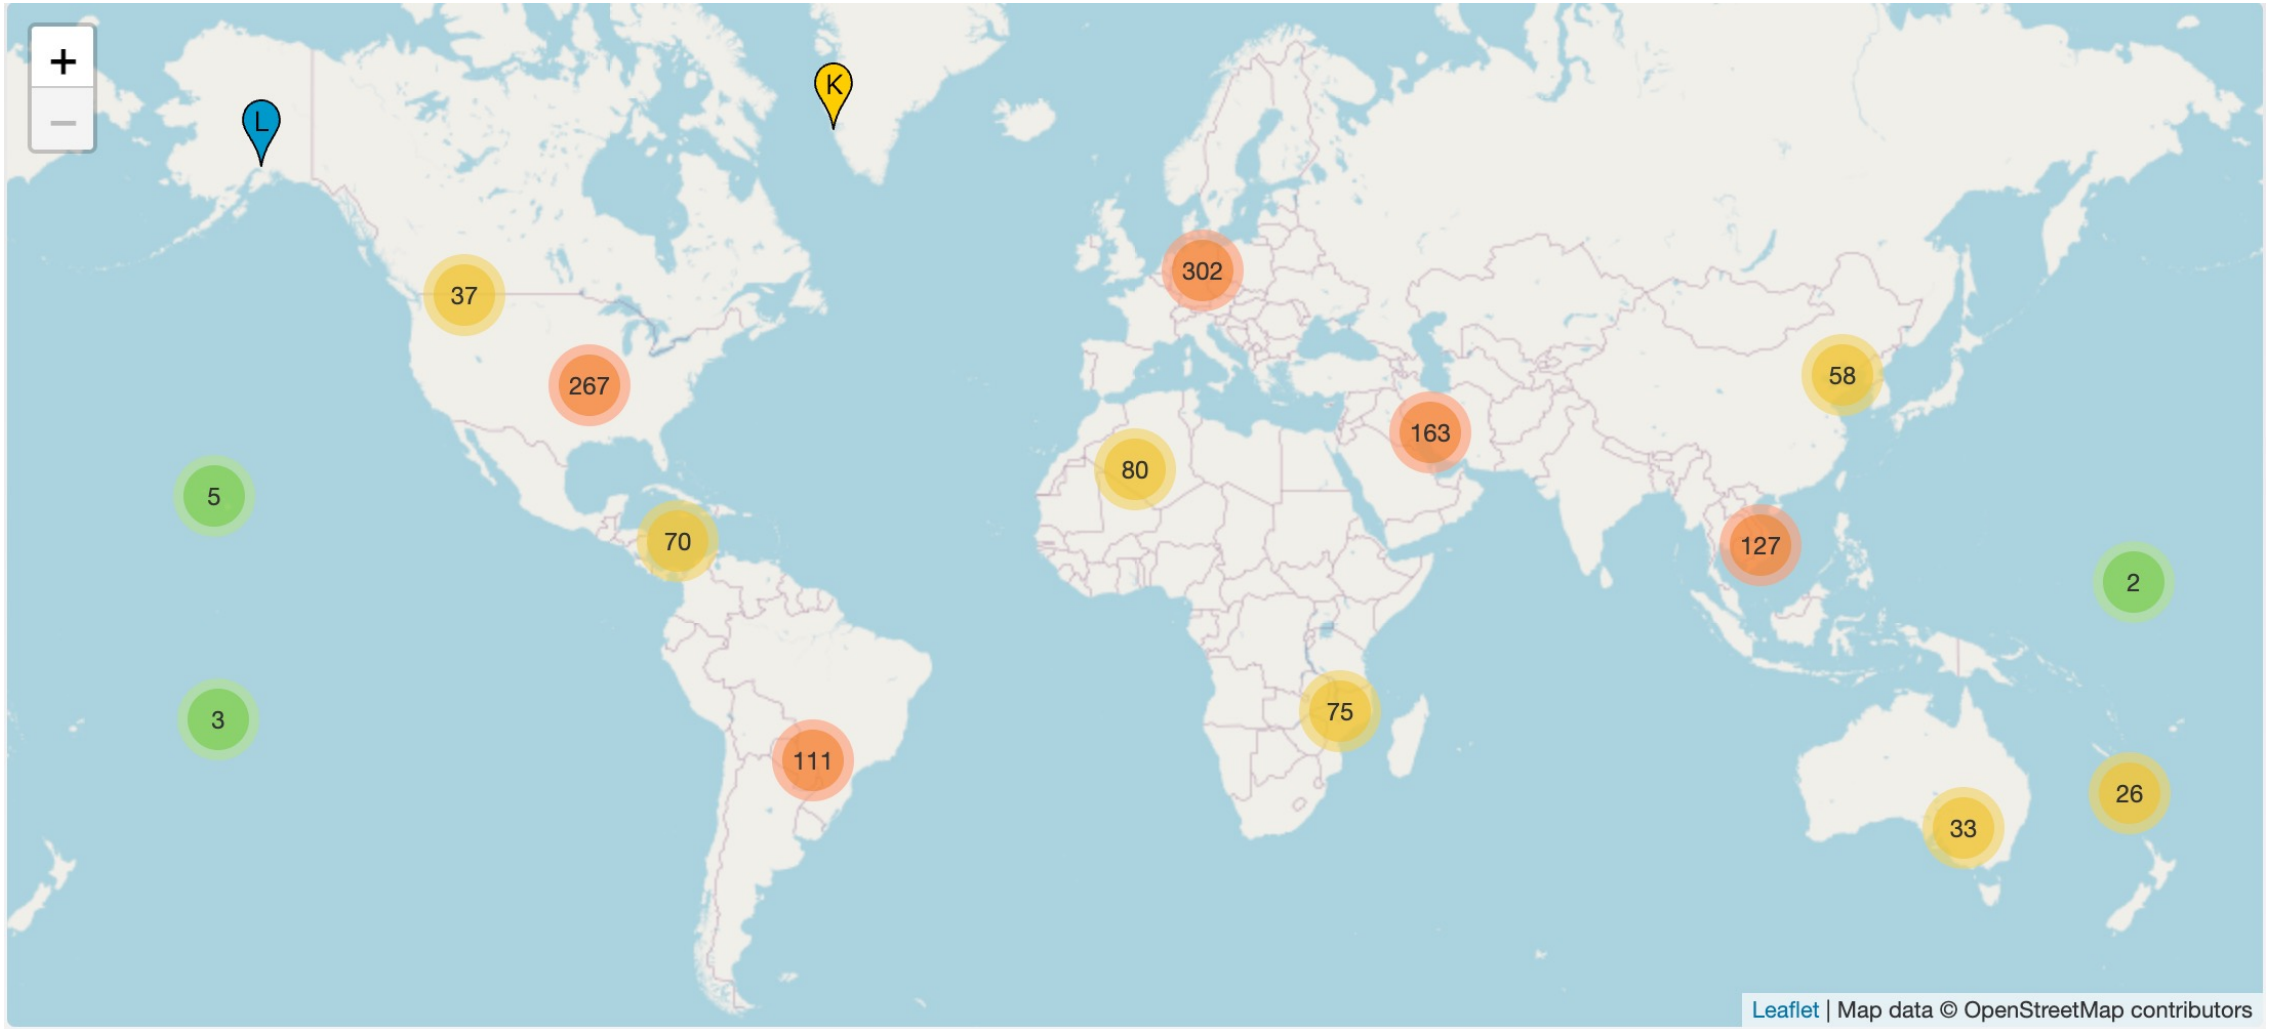

As of 10/01/2022 3:05 p.m., the root server system consists of 1551 instances operated by the 12 independent root server operators.

Security: DNSSEC …

DNS Ecosystem

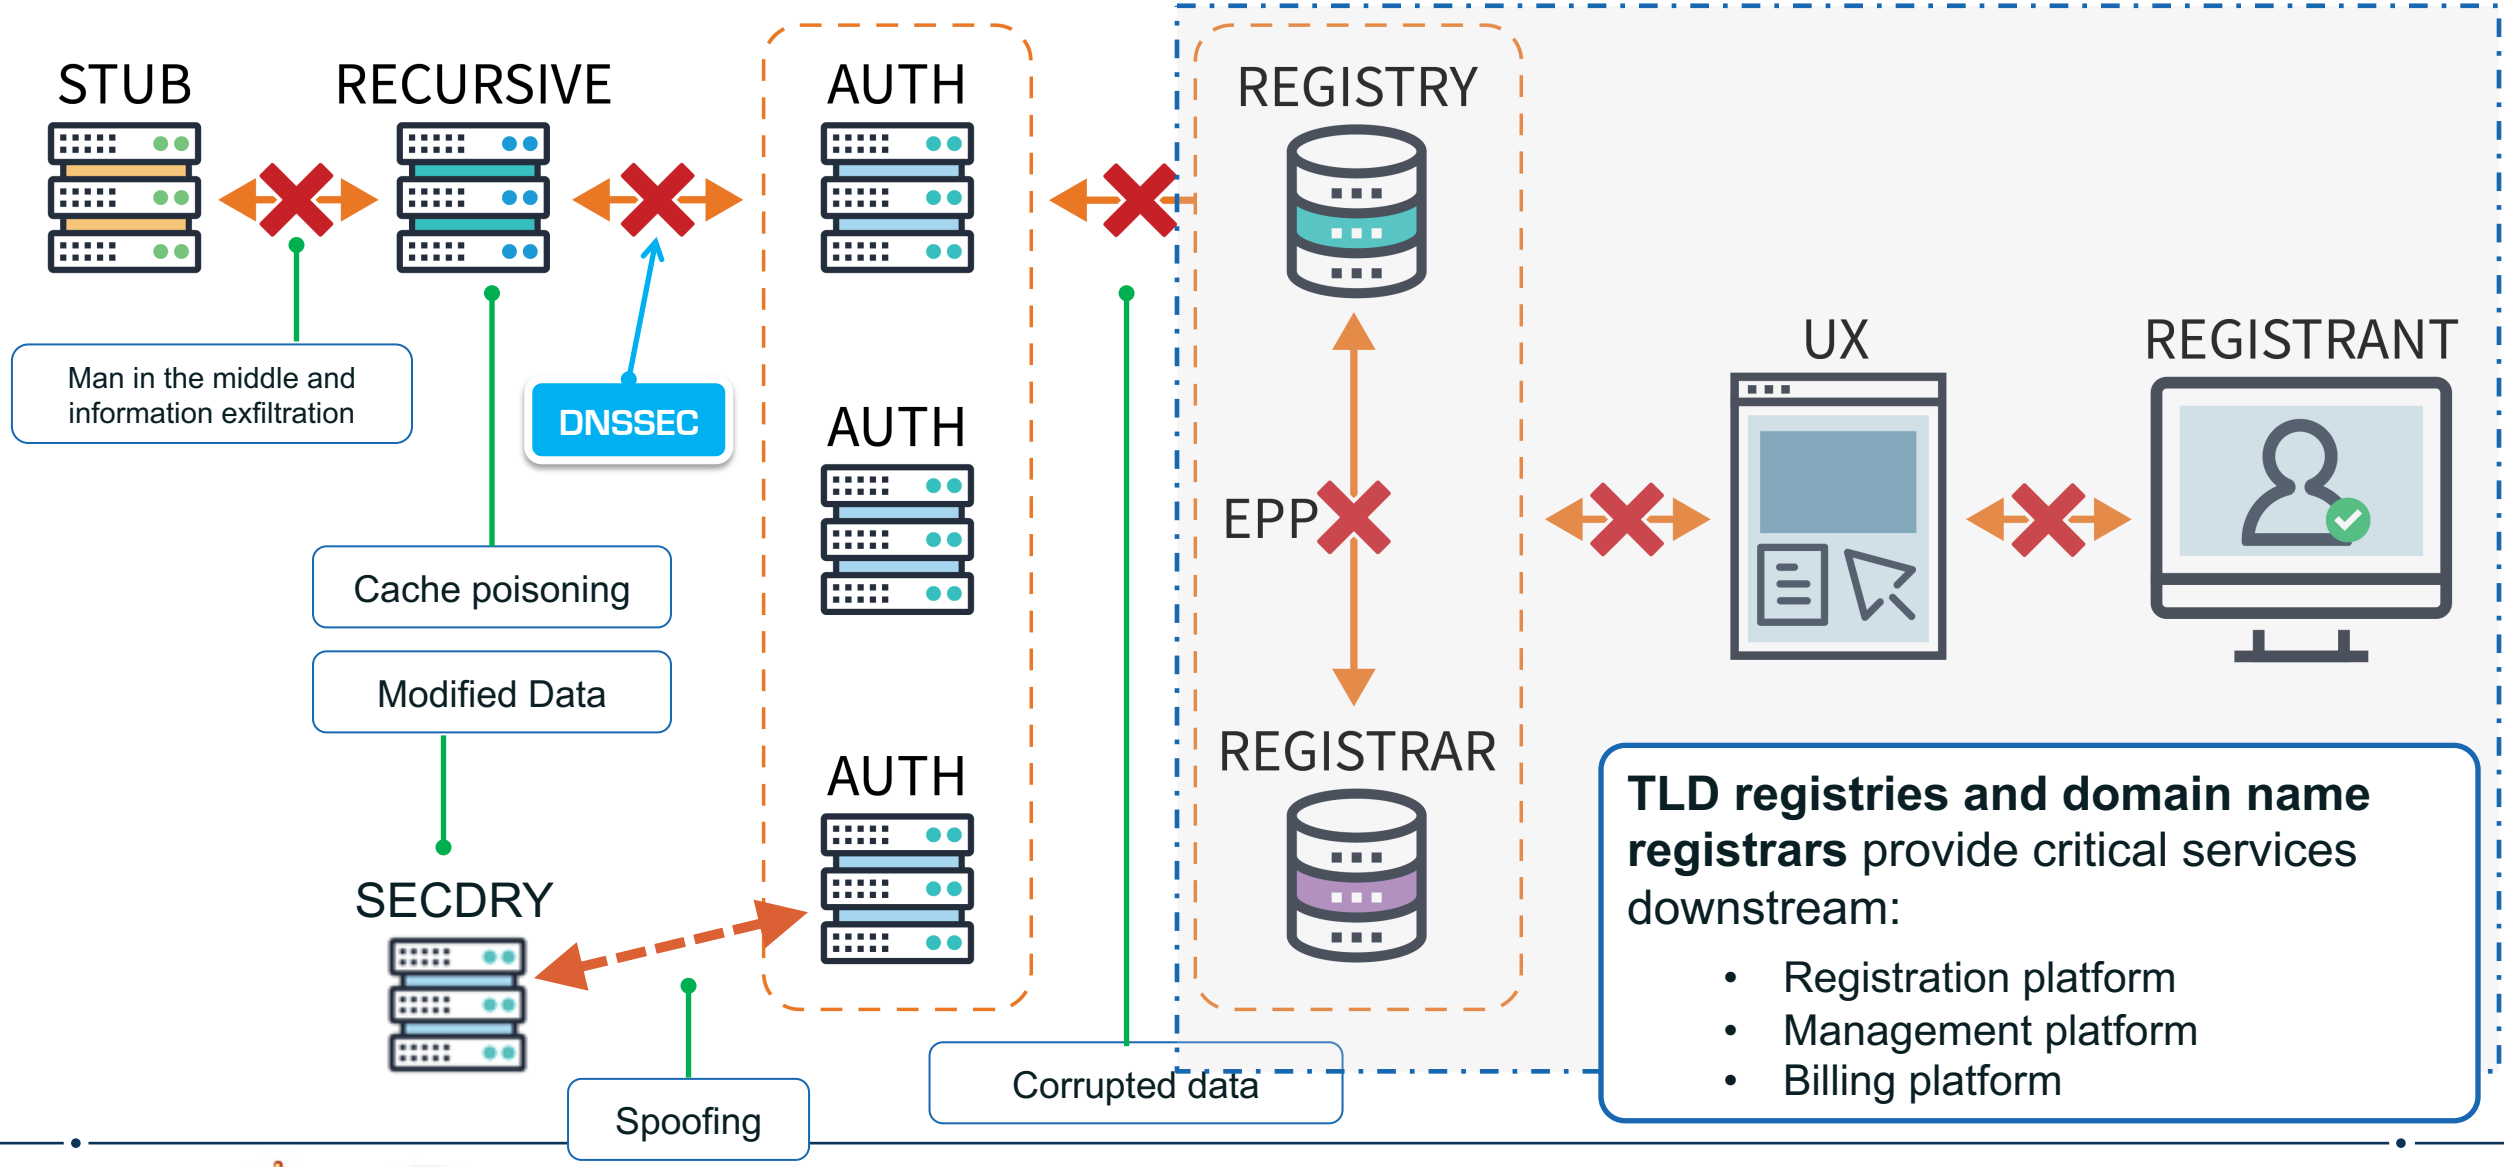

TLD registries and domain name registrars provide critical services downstream:

- Registration platform
- Management platform
- Billing platform

What DNSSEC Does

- DNSSEC uses public-key cryptography and digital signatures to provide:
  - Data origin authentication
    - "Did this response really come from the *example.com* zone authority?"
  - Data integrity
    - "Did an attacker (e.g., a man in the middle) modify the data in this response since the data was originally signed?"
- DNSSEC offers protection against spoofing of DNS data (and so, for attacks like cache-poisoning, etc.).

What DNSSEC Doesn’t Do

- **DNSSEC does not**:
  - Provide any confidentiality for DNS data
    - No encryption.
    - Transferred data will be readable for person-in-the-middle.
  - Address attacks against DNS software
    - DDoS
    - “packets of death”
    - Etc.

Signing DNS Data

- In DNSSEC, each zone has a public/private key pair
- Data in the zone is signed with the private key
  - Signing the data is usually de-coupled from serving the data
  - The design allows data to be signed ahead of time rather than “on the fly” for each response
- Important: In DNSSEC, DNS *data* is signed, not DNS *messages*
  - Signing messages is called transaction security
  - A separate protocol called TSIG handles that

Zone Key Pairs

- The zone's public key is published in the zone in a specific record.
- The zone's private key is kept safe:
  - The amount of protection required depends on how the zone owner evaluate the risks involved in case the private key is disclosed or compromised.
- Options for protecting a zone's private key:
  - Stored on-line in some encrypted form, only decrypted when needed for signing data
    - The minimum.
  - Stored offline also in some encrypted form
    - Offers more protection.
  - Stored in a hardware security module (HSM)
    - Offers the most protection but overkill (may also be costly) for many applications.

Recalling Resource Records (RR)

- Data associated with domain names is contained in Resource Records.
  - **A** IPv4 address
  - **AAAA** IPv6 address
  - **NS** Name of an authoritative name server
  - **SOA** “Start of authority”, appears at zone apex
  - **CNAME** Name of an alias to another domain name
  - **MX** Name of a “mail exchange server”
  - **PTR** IP address encoded as a domain name (for reverse mapping)

DNSSEC adds some others:

- DNSKEY
- RRSIG
- NSEC
- DS

DNSSEC Benefits

- **Technical benefits**
  - Provide authentication/origin validation.
  - Guarantee the integrity and non-manipulation of DNS data.
  - Authenticated denial of existence of DNS data (NSEC).

- **Impact on the different members of the ecosystem**
  - End user: Confidence of reaching the desired/correct website (plugin of https).
  - Registrant: Fraud mitigation and increased brand protection (country code reputation).
  - Registrar: Meet industry standards and meet registrant demands for increased security (attract and retain security- and reputation-focused registrants).
  - Registry: Comply with industry best practices and registrar demands for stronger domain security.

Privacy: DoT & DoH …

DNS Ecosystem

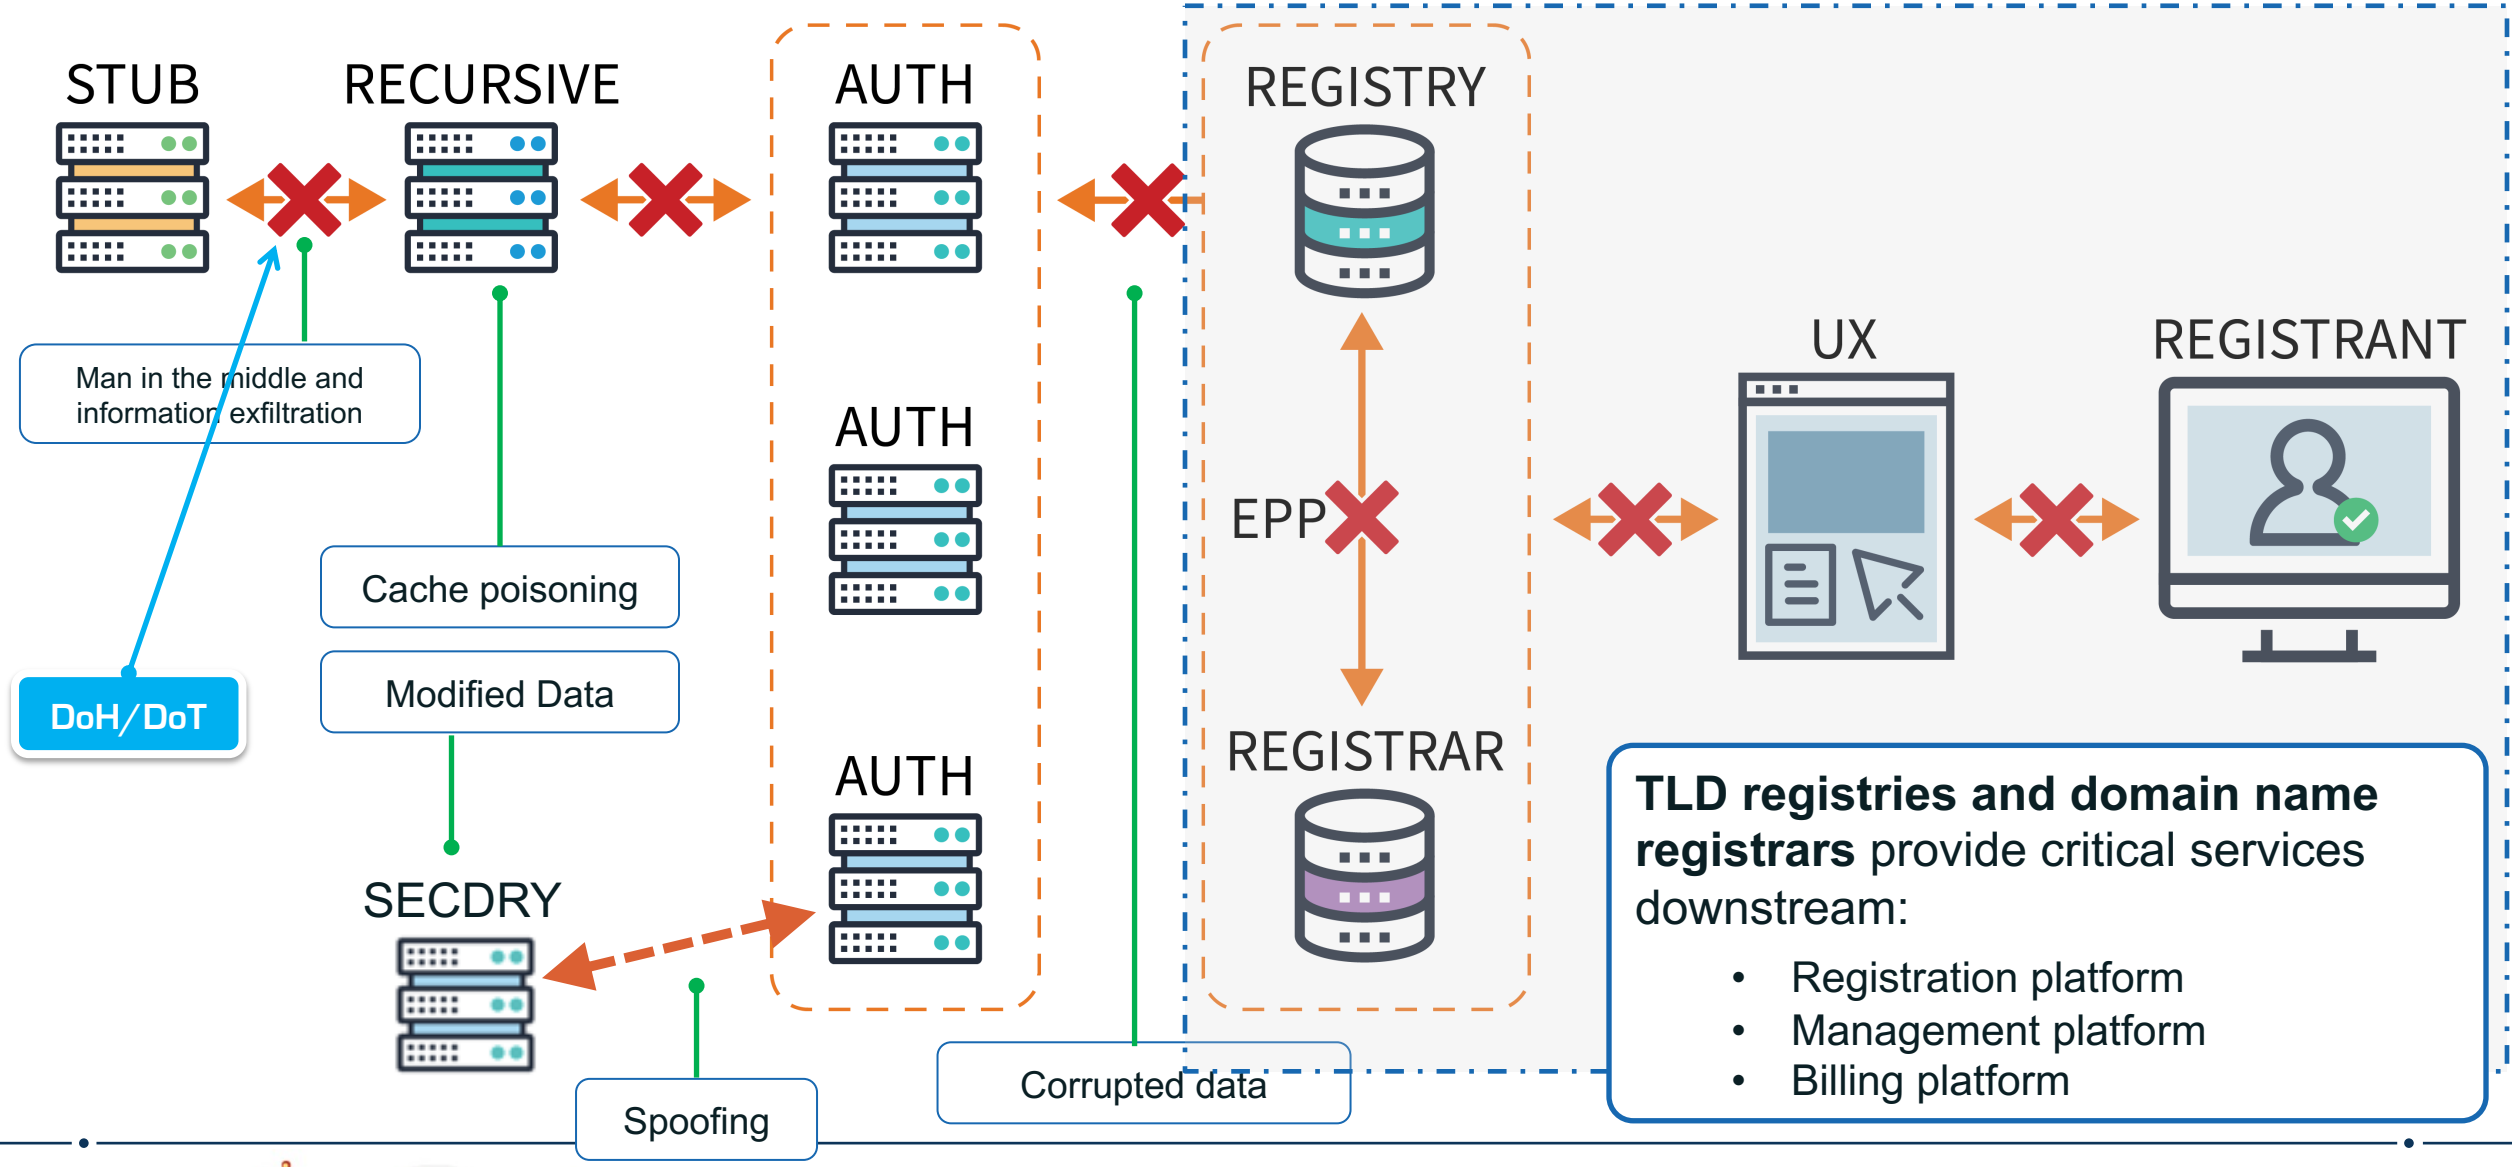

TLD registries and domain name registrars provide critical services downstream:

- Registration platform
- Management platform
- Billing platform

DoT y DoH… one slide ☺

The main idea behind DoT and DoH is to provide privacy by encrypting DNS queries and responses between the terminal equipment and the chosen recursive DNS server.

In this way, it increases the resilience against interception, blocking, interference and/or manipulation of that traffic (mainly the same as any point-to-point encryption method seeks).

- DoT stands for DNS over TLS.
- DoH stands for DNS over HTTPS.

Like almost all methods involving privacy issues, both DoT and DoH (and especially DoH) have sparked some discussion at both the political and technical levels...

…an important idea worth considering is the separation between standards and implementations thereof, which often leads to some debate. … from a technical perspective, it might be considered more convenient, from the point of view of security and resiliency of the global DNS system, to enable these mechanisms on their own recursive instead of forwarding all queries to a public one (thereby encouraging decentralization of DNS resolution).

Privacy: QNAME minimization …

DNS Ecosystem

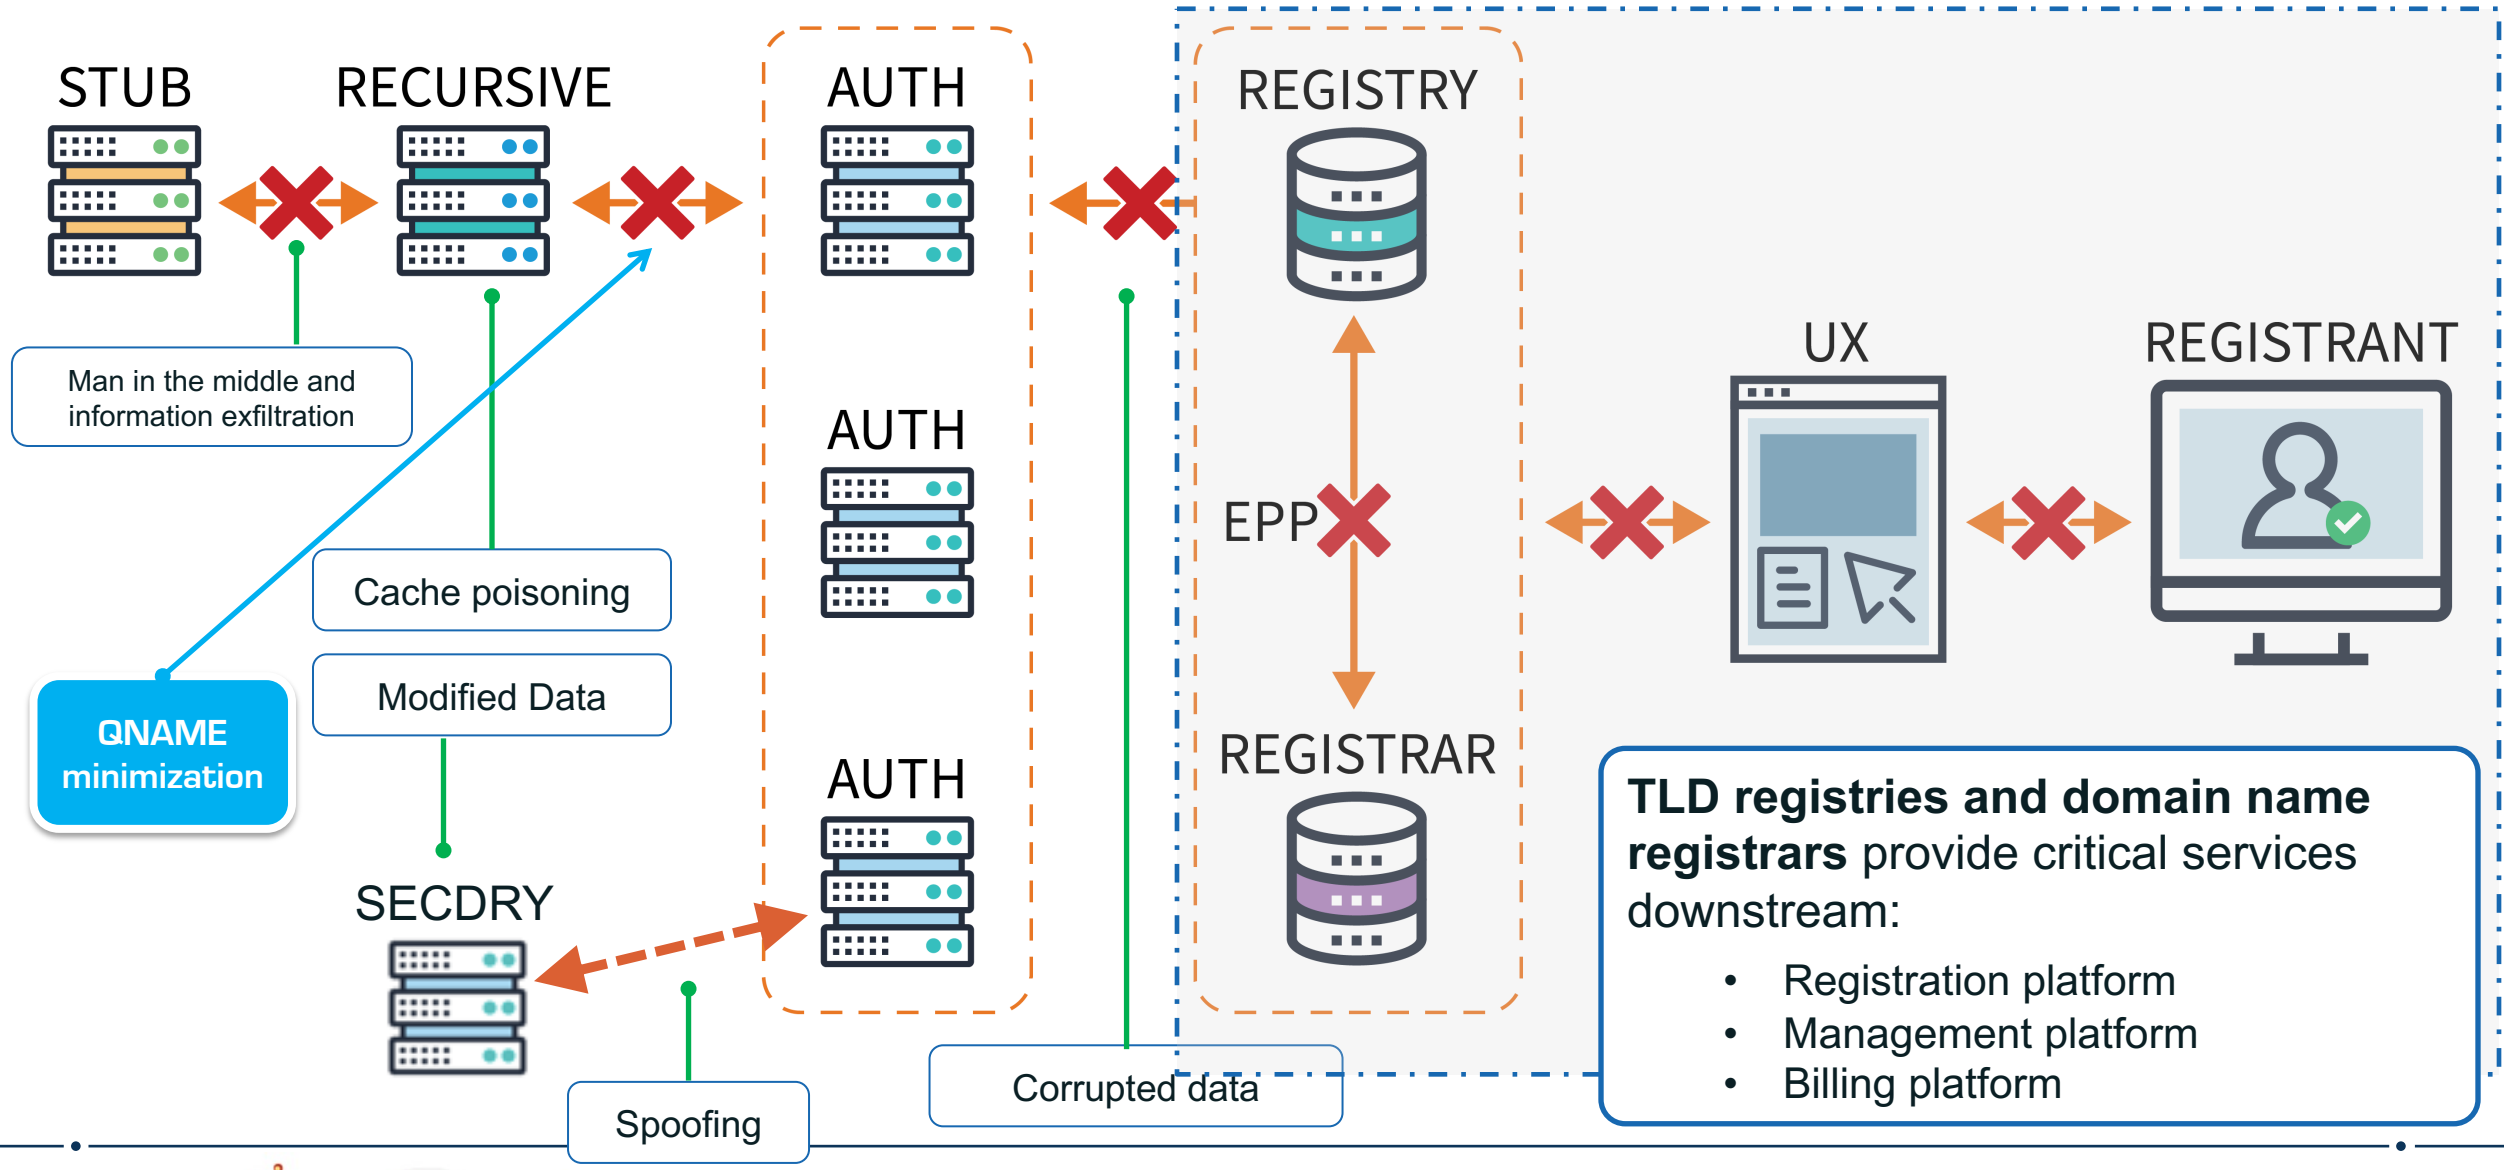

QNAME minimization… one slide ☺

QNAME minification follows the principle explained in Section 6.1 of [RFC6973]: the less data you send, the less privacy issues you have.

DNS Query Name (QNAME) minification is defined in RFC 7816 to improve end-user privacy in the DNS resolution process.

Change the recursive server's “standard” DNS queries to include only as much detail in each query as is necessary for that step in the resolution process. IETF RFC 7816 describes it as a technique "in which the DNS resolver no longer sends the full original QNAME to the authoritative name server."

… or two ☺

Remember there are two ways to configure

- Strict mode: **qnamemin** is used in the search for the response. In case of error, nxdomain, nodata, etc., the error is returned.

- Relaxed mode: **qnamemin** is used, but in case of error, the classic query (full name) is retried.

Speeding up resolution & improving privacy: Hyperlocal…

DNS Ecosystem

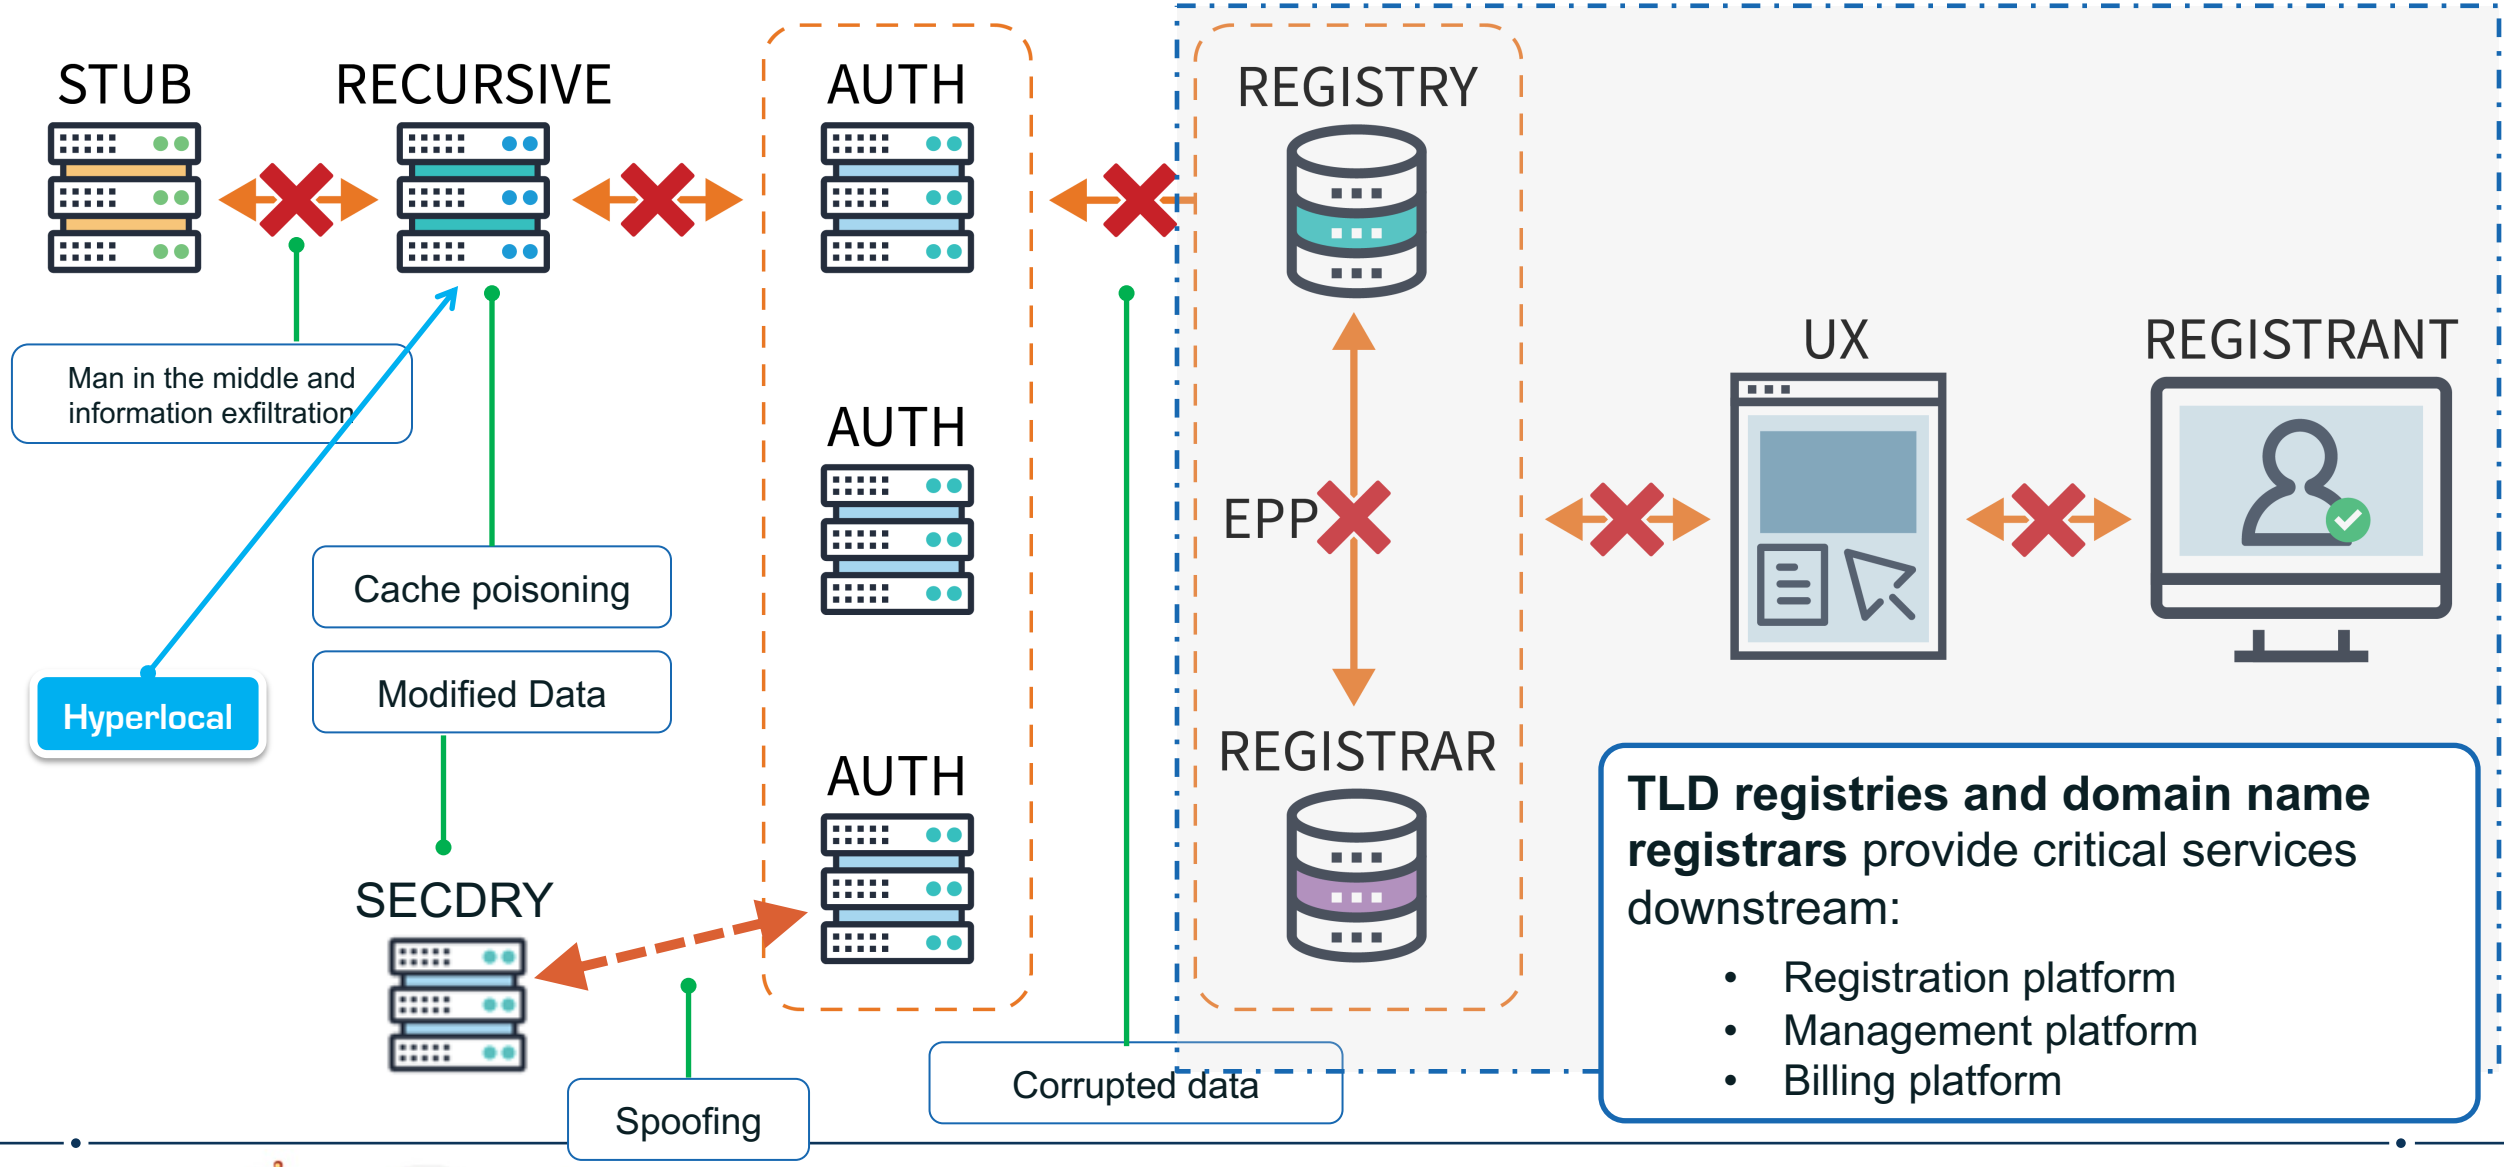

Hyperlocal

What is this?

- It's about keeping a local copy of the DNS Root on the same machine that runs recursive resolvers (recursive server).
- Included in that goal is ensuring that root zone data is always accessible.
- Steve Crocker named this technique Hyperlocal.

Standardized in RFC 8806, "Running a Root Server Locally in a Recursive"

- The root server must be running on the same machine as the recursive server.
- You can only answer queries from the local machine and no other machines.
- It is recommended to maintain and apply the standard resolution mechanism in case of failure or when the local copy is not available or out of date.

DNS & Hyperlocal

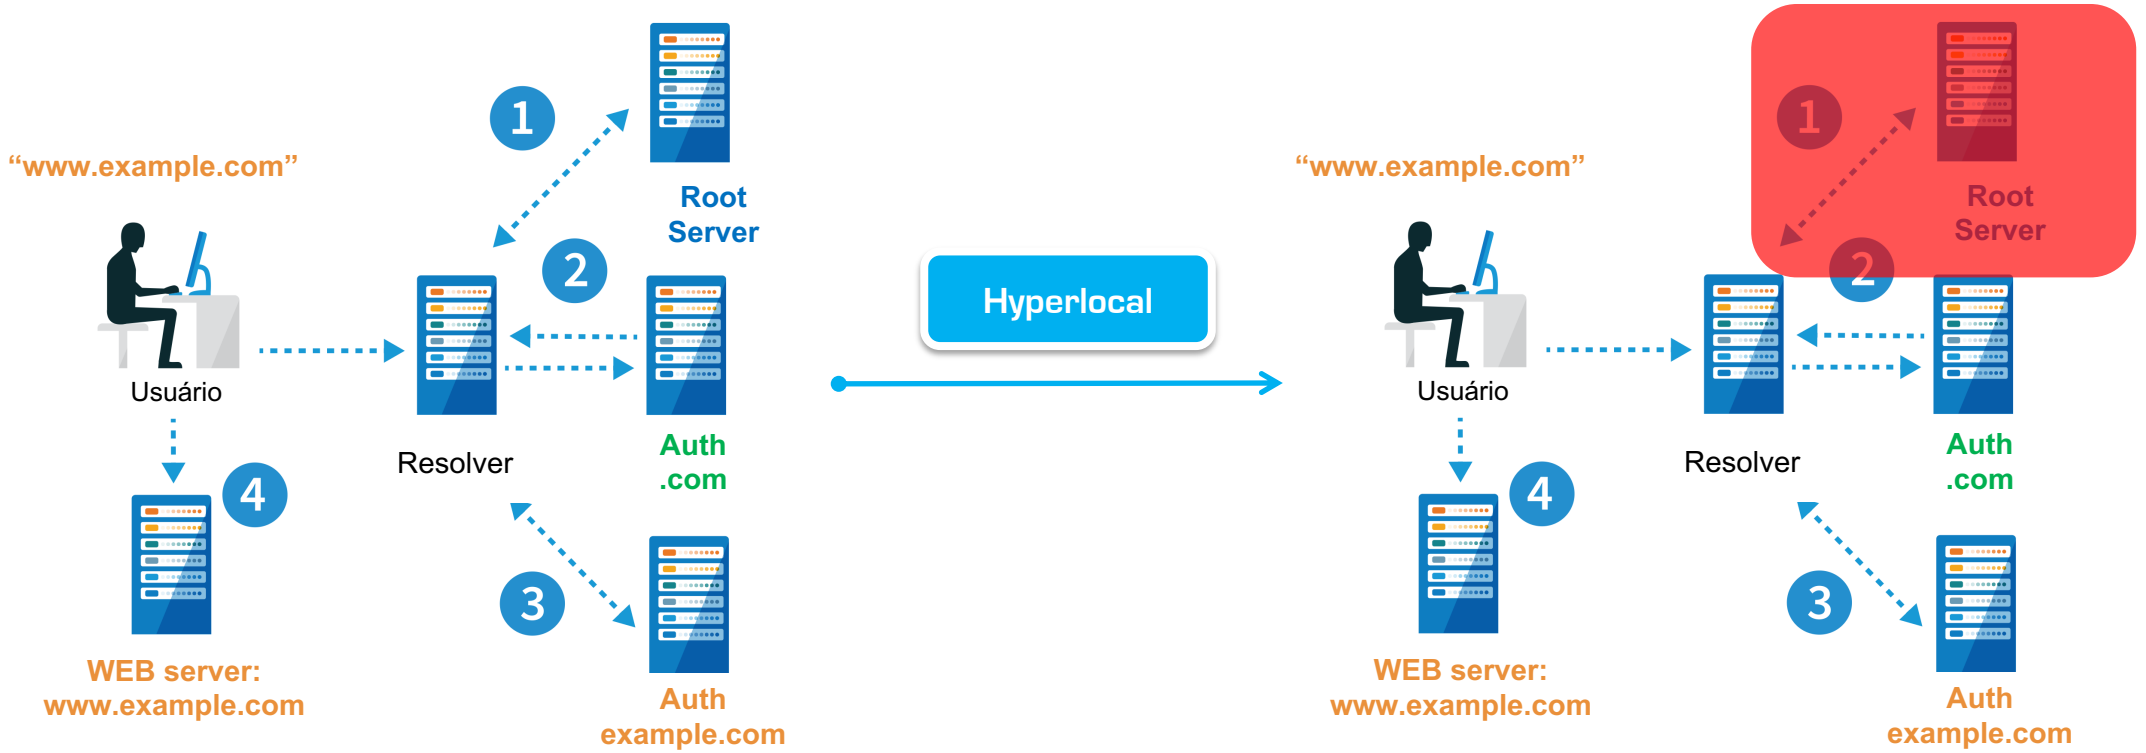

Some other good practices to consider…

KINDNS

Knowledge-sharing and
Instantiating
Norms for
DNS (Domain Name System) and
Naming
Security

What's this?

Knowledge-Sharing and Instantiating Norms for Domain Name System and Naming Security (KINDNS) is an initiative launched by the Internet Corporation for Assigned Names and Numbers (ICANN) to produce a simple and clear framework of operational best practices for DNS operators.

KINDNS

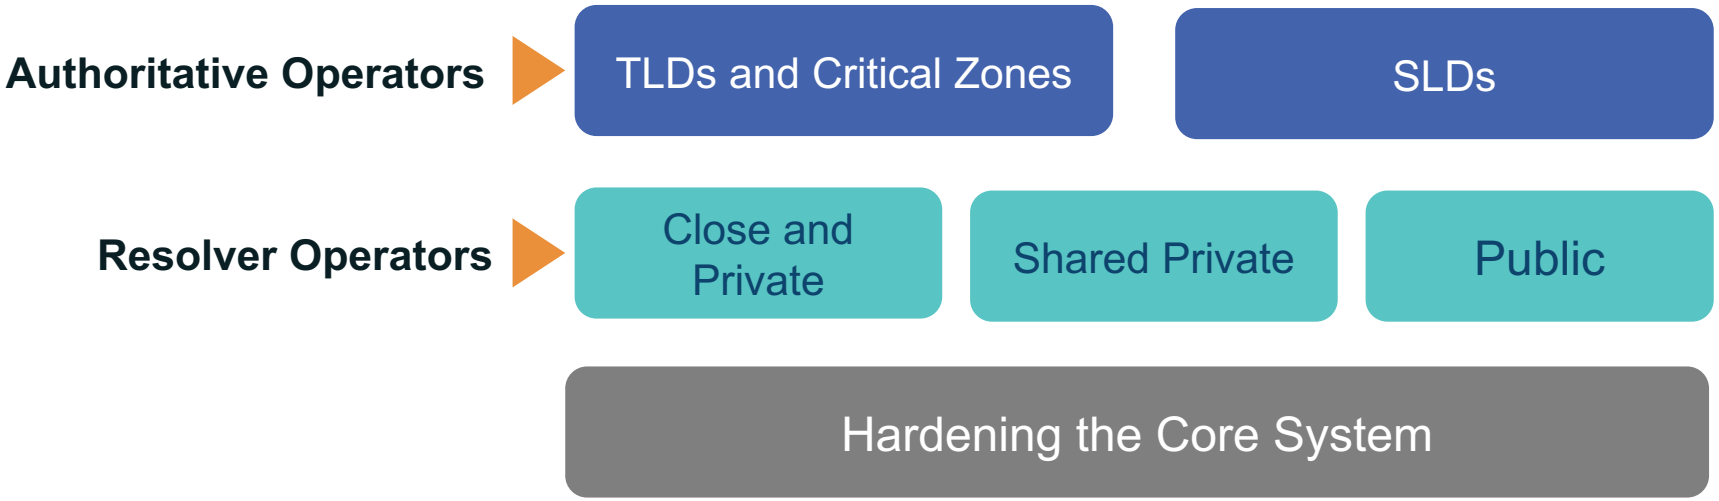

By joining the KINDNS initiative, DNS operators are voluntarily committing to adhere to the identified practices and act as “goodwill ambassadors” within the community.

KINDNS

1. Operators in each category can conduct a self-assessment of their operational practices against KINDNS and use the report to correct/adjust unaligned practices.

    - Self-Assessments will be anonymous, and reports will be directly downloaded from the website.

2. Operators can enroll to participate in one or many categories covered by KINDNS.

    - Participation in the KINDNS initiative means voluntarily committing to implement/adhere to agreed practices.
    - Participants becomes goodwill ambassadors and promote best practices.

KINDNS

- **The KINDNS discussion mailing list:**

kindns-discuss@icann.org

- **Wiki page** where we will share preliminary documents until the formal website is developed and launched

https://community.icann.org/display/KINDNS

Thank You and Questions

One World, One Internet

Visit us at **icann.org**

@icann

facebook.com/icannorg

youtube.com/icannnews

flickr.com/icann

linkedin/company/icann

slideshare/icannpresentations

soundcloud/icann