Turning the TabLeS – and how we got there

Ralph Holz, Thomas Riedmaier

Network Architectures and Services
Technische Universität München

Berlinsides 2011

We'll talk about SSL/TLS and X.509

SSL/TLS

- The backbone protocols for securing the WWW and e-mail
- Authentication, confidentiality, integrity
- Public-key cryptography

X.509: Public Key Infrastructure standard

- Certification Authorities (CAs) certify Web sites
- Non-forgeable signature:
  $Cert(X) = Sig_{CA}(id_X, pubkey_X)$

Our work is empirical

Part 1 of talk: the SSL landscape

- Background
- The state of the PKI for the WWW

Part 2 of talk: Man-in-the middle attacks on HTTPs

- Our tool: Crossbear
- We want hard data

What we already have done

Let us tell you a story: the SSL Landscape

Browser panic (Berlinsides)

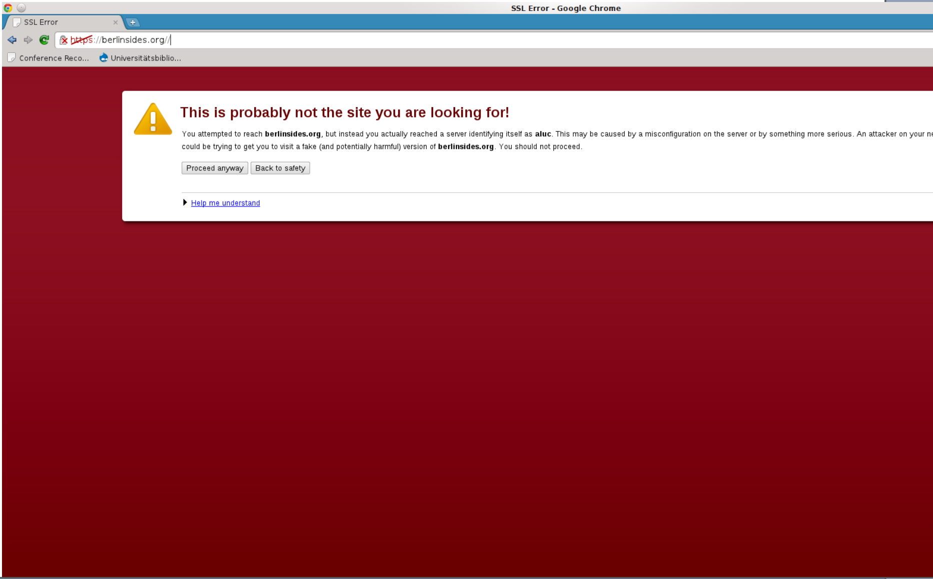

This is probably not the site you are looking for!

You attempted to reach **berlinsides.org**, but instead you actually reached a server identifying itself as **aluc**. This may be caused by a misconfiguration on the server or by something more serious. An attacker on your ne could be trying to get you to visit a fake (and potentially harmful) version of **berlinsides.org**. You should not proceed.

Basic idea of X.509 PKI: hierarchy

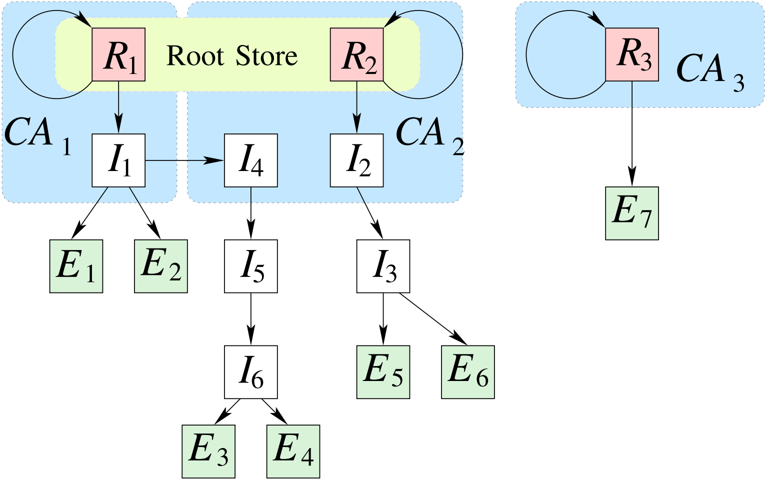

Basic idea of X.509

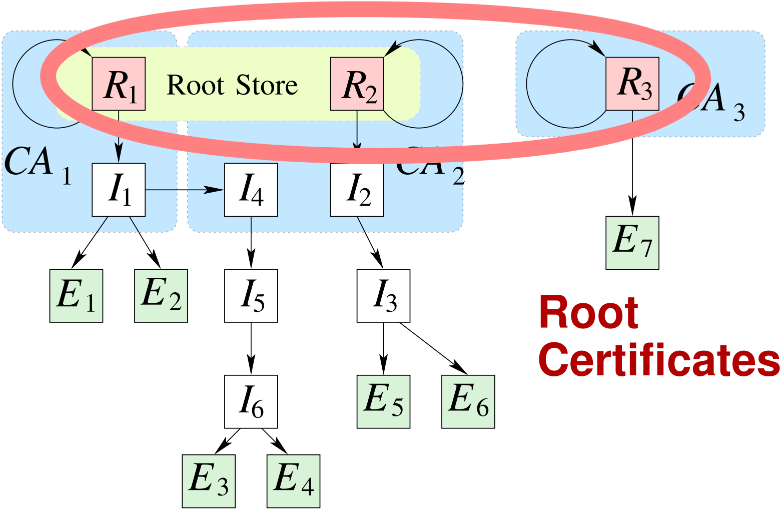

Root Certificates

Basic idea of X.509

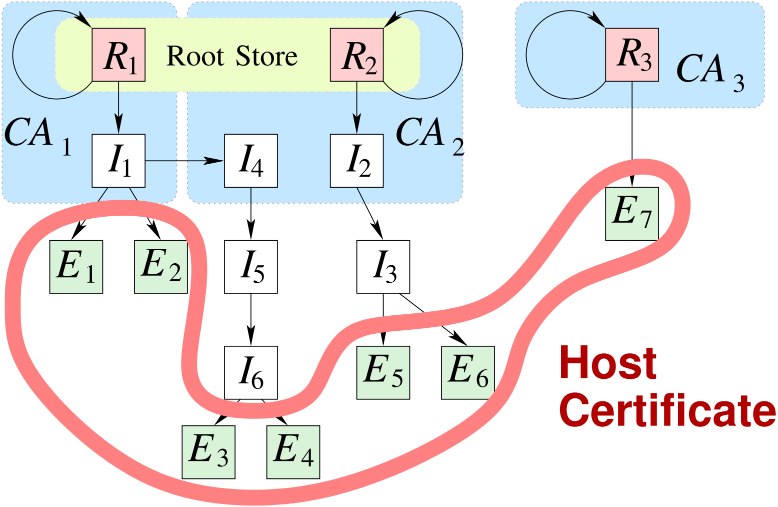

Host Certificates

Basic idea of X.509

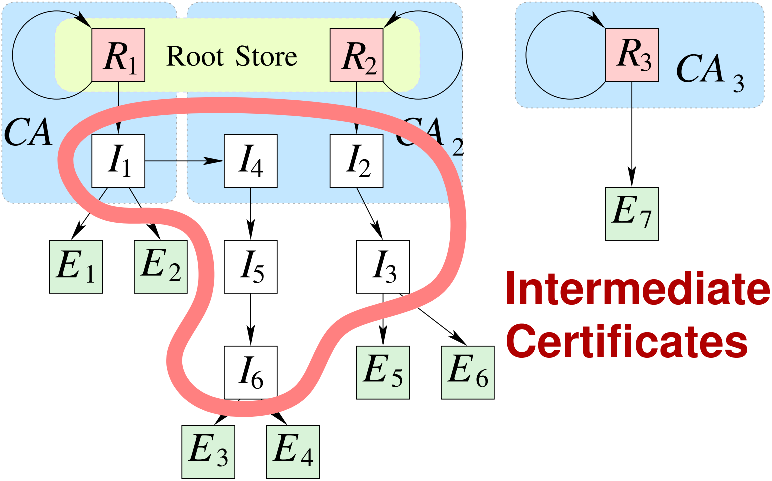

Intermediate Certificates

Basic idea of X.509

CAs in Root Store

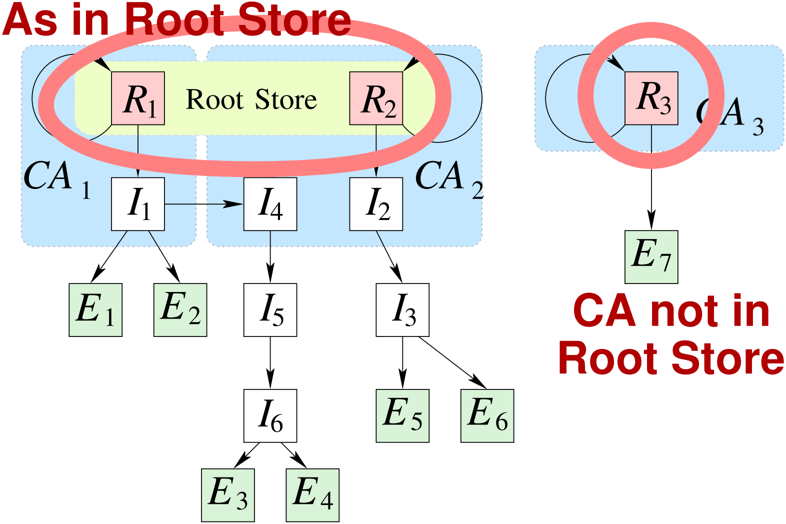

CA not in Root Store

Basic idea of X.509

Root certificate not in Root Store

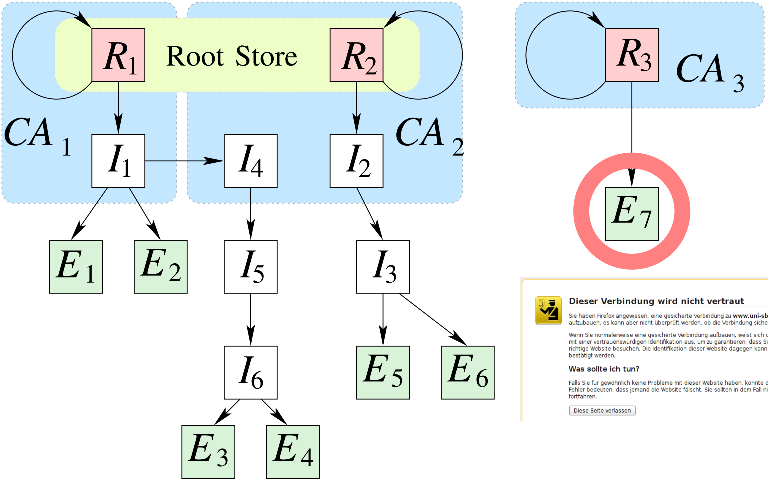

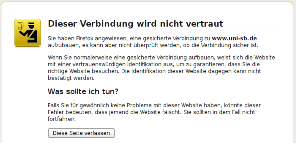

Dieser Verbindung wird nicht vertraut

Sie haben Firefox angewiesen, eine gesicherte Verbindung zu **www.uni-sb.de** aufzubauen, es kann aber nicht überprüft werden, ob die Verbindung sicher ist.

Wenn Sie normalerweise eine gesicherte Verbindung aufbauen, weist sich die Website mit einer vertrauenswürdigen Identifikation aus, um zu garantieren, dass Sie die richtige Website besuchen. Die Identifikation dieser Website dagegen kann nicht bestätigt werden.

Was sollte ich tun?

Falls Sie für gewöhnlich keine Probleme mit dieser Website haben, könnte dieser Fehler bedeuten, dass jemand die Website fälscht. Sie sollten in dem Fall nicht fortfahren.

An X.509 certificate

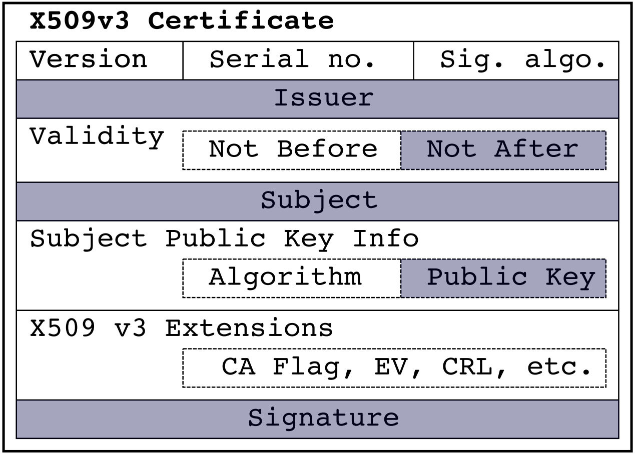

X509v3 Certificate

Version | Serial no. | Sig. algo.

Issuer

Validity: Not Before | Not After

Subject

Subject Public Key Info: Algorithm | Public Key

X509 v3 Extensions: CA Flag, EV, CRL, etc.

Signature

Root Stores contain CA certificates

CAs in Root Store

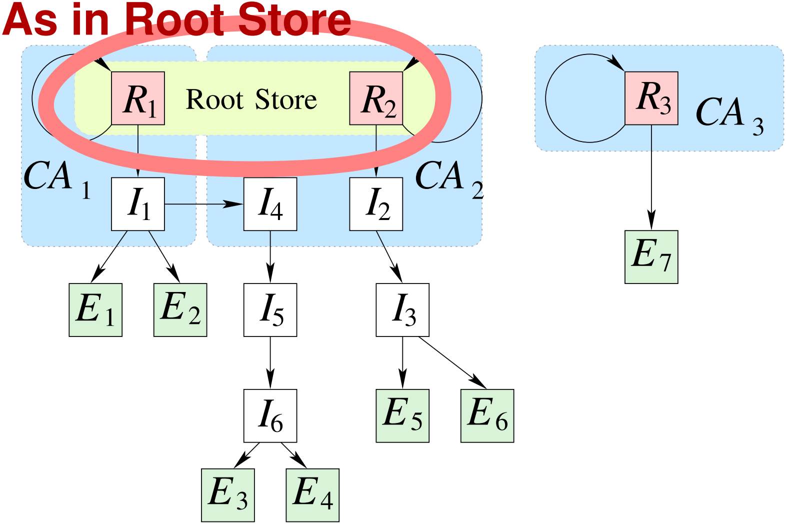

Browser (client) Root Stores

Your browser chooses the ‘trusted CAs’. Not you.

Any CA may issue a certificate for any domain.

This means the weakest CA determines the strength of the whole PKI.

Development of Mozilla Root Store

More than 150 trustworthy Root Certificates

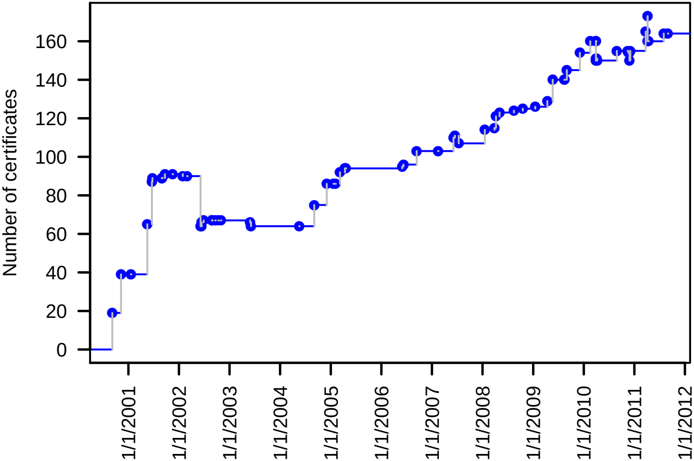

Certificate issuance

How is a certificate issued in practice?

- Domain Validation:
  - Send e-mail to (CA-chosen) mail address with code
  - Confirmed ownership of mail address = ownership of domain
- Organisational Validation (OV, rare)
- Extended Validation (later, rare)

Race to the bottom

- CAs have incentive to lower prices
- Translates into incentive to control less, not more

How this got our interest

PKI weaknesses in 2008

- Early December 2008:
  - 'Error' in Comodo CA: no identity check
  - Reported by Eddy Nigg of StartSSL (a CA)
  - A regional sub-seller just took the credit card number and gave you a certificate
  - No real reaction by Mozilla
- Late December 2008: whitehat hacks StartSSL CA
  - Technical report: simple flaw in Web front-end
  - Certificate for `mozilla.com` issued
  - Caught by 2nd line of defence:
    human checks for high-value domains

How this got our interest

PKI weaknesses in 2008

- Early December 2008:
  - ‘Error’ in Comodo CA: no identity check
  - Reported by Eddy Nigg of StartSSL (a CA)
  - A regional sub-seller just took the credit card number and gave you a certificate
  - No real reaction by Mozilla
- Late December 2008: whitehat hacks StartSSL CA
  - Technical report: simple flaw in Web front-end
  - Certificate for `mozilla.com` issued
  - Caught by 2nd line of defence:
    human checks for high-value domains

How this got our interest

PKI weaknesses in 2008

- Early December 2008:
  - ‘Error’ in Comodo CA: no identity check
  - Reported by Eddy Nigg of StartSSL (a CA)
  - A regional sub-seller just took the credit card number and gave you a certificate
  - No real reaction by Mozilla
- Late December 2008: whitehat hacks StartSSL CA
  - Technical report: simple flaw in Web front-end
  - Certificate for `mozilla.com` issued
  - Caught by 2nd line of defence:
human checks for high-value domains

How this got our interest

In 2011, the foundations of X.509 were rocked.

- March 2011: Comodo CA hacked (a sub-seller, again)
  - Attacker claims to come from Iran
  - $\approx$ 10 certificates for high-value domains issued
  - Browser reaction: blacklisting of those certificates *in code*
  - Neither CRLs nor OCSP trusted enough to work for victims
- July 2011: DigiNotar CA hacked
  - Attacker claims to be the same one as in March
  - 531 fake certificates, high-value domains
  - E.g., Google, Facebook, Mozilla, CIA, Mossad, Skype
  - Some hints pointed at Man-in-the-middle attack in Iran
  - For the first time, a Root CA is removed from a browser for being compromised

How this got our interest

In 2011, the foundations of X.509 were rocked.

- March 2011: Comodo CA hacked (a sub-seller, again)
  - Attacker claims to come from Iran
  - $\approx$ 10 certificates for high-value domains issued
  - Browser reaction: blacklisting of those certificates *in code*
  - Neither CRLs nor OCSP trusted enough to work for victims
- July 2011: DigiNotar CA hacked
  - Attacker claims to be the same one as in March
  - 531 fake certificates, high-value domains
  - E.g., Google, Facebook, Mozilla, CIA, Mossad, Skype
  - Some hints pointed at Man-in-the-middle attack in Iran
  - For the first time, a Root CA is removed from a browser for being compromised

How this got our interest

In 2011, the foundations of X.509 were rocked.

- March 2011: Comodo CA hacked (a sub-seller, again)
  - Attacker claims to come from Iran
  - $\approx$ 10 certificates for high-value domains issued
  - Browser reaction: blacklisting of those certificates *in code*
  - Neither CRLs nor OCSP trusted enough to work for victims
- July 2011: DigiNotar CA hacked
  - Attacker claims to be the same one as in March
  - 531 fake certificates, high-value domains
  - E.g., Google, Facebook, Mozilla, CIA, Mossad, Skype
  - Some hints pointed at Man-in-the-middle attack in Iran
  - For the first time, a Root CA is removed from a browser for being compromised

Can we assess the quality of this PKI?

A good PKI should

- ... allow HTTPs on all WWW hosts
- ... contain only valid certificates
- ... offer good cryptographic security
  - Long keys, only strong hash algorithms, ...
- ... have a sensible setup
  - Short validity periods (1 year)
  - Short certificate chains (but use intermediate certificates)
  - Number of issuers should be reasonable (weakest link!)

Acquiring our data sets

Active scans to measure *deployed* PKI

- Scan hosts on Alexa Top 1 million Web sites
- Nov 2009 – Apr 2011: scanned 8 times from Germany
- March 2011: scans from 8 hosts around the globe

Passive monitoring to measure *user-encountered* PKI

- Munich Research Network, monitored all SSL/TLS traffic
- Two 2-week runs in Sep 2010 and Apr 2011

EFF scan of IPv4 space in 2010

- Scan of 2-3 months, no *domain* information

Our data sets

Active Scans — Passive Monitoring — EFF IPv4 scan

Location	*Time (run)*	*Type*	*Certificates*
Tuebingen, DE	November 2009	Active scan	833,661
Tuebingen, DE	December 2009	Active scan	819,488
Tuebingen, DE	January 2010	Active scan	816,517
Tuebingen, DE	April 2010	Active scan	816,605
Munich, DE	September 2010	Active scan	829,232
Munich, DE	November 2010	Active scan	827,366
Munich, DE	April 2011	Active scan	829,707
Munich, DE	April 2011	Active scan with SNI	826,098
Shanghai, CN	April 2011	Active scan	798,976
Beijing, CN	April 2011	Active scan	797,046
Melbourne, AU	April 2011	Active scan	833,571
İzmir, TR	April 2011	Active scan	825,555
São Paulo, BR	April 2011	Active scan	833,246
Moscow, RU	April 2011	Active scan	830,765
Santa Barbara, US	April 2011	Active scan	834,173
Boston, US	April 2011	Active scan	834,054
Munich, DE	September 2010	Passive monitoring	183,208
Munich, DE	April 2011	Passive monitoring	989,040
EFF servers	March–June 2010	Active IPv4 scan	11,349,678

25 million certificates to evaluate.

Our data sets

Active Scans — Passive Monitoring — EFF IPv4 scan

Location	*Time (run)*	*Type*	*Certificates*
Tuebingen, DE	November 2009	Active scan	833,661
Tuebingen, DE	December 2009	Active scan	819,488
Tuebingen, DE	January 2010	Active scan	816,517
Tuebingen, DE	April 2010	Active scan	816,605
Munich, DE	September 2010	Active scan	829,232
Munich, DE	November 2010	Active scan	827,366
Munich, DE	April 2011	Active scan	829,707
Munich, DE	April 2011	Active scan with SNI	826,098
Shanghai, CN	April 2011	Active scan	798,976
Beijing, CN	April 2011	Active scan	797,046
Melbourne, AU	April 2011	Active scan	833,571
İzmir, TR	April 2011	Active scan	825,555
São Paulo, BR	April 2011	Active scan	833,246
Moscow, RU	April 2011	Active scan	830,765
Santa Barbara, US	April 2011	Active scan	834,173
Boston, US	April 2011	Active scan	834,054
Munich, DE	September 2010	Passive monitoring	183,208
Munich, DE	April 2011	Passive monitoring	989,040
EFF servers	March–June 2010	Active IPv4 scan	11,349,678

25 million certificates to evaluate.

Our data sets

Active Scans — Passive Monitoring — EFF IPv4 scan

Location	*Time (run)*	*Type*	*Certificates*
Tuebingen, DE	November 2009	Active scan	833,661
Tuebingen, DE	December 2009	Active scan	819,488
Tuebingen, DE	January 2010	Active scan	816,517
Tuebingen, DE	April 2010	Active scan	816,605
Munich, DE	September 2010	Active scan	829,232
Munich, DE	November 2010	Active scan	827,366
Munich, DE	April 2011	Active scan	829,707
Munich, DE	April 2011	Active scan with SNI	826,098
Shanghai, CN	April 2011	Active scan	798,976
Beijing, CN	April 2011	Active scan	797,046
Melbourne, AU	April 2011	Active scan	833,571
İzmir, TR	April 2011	Active scan	825,555
São Paulo, BR	April 2011	Active scan	833,246
Moscow, RU	April 2011	Active scan	830,765
Santa Barbara, US	April 2011	Active scan	834,173
Boston, US	April 2011	Active scan	834,054
Munich, DE	September 2010	Passive monitoring	183,208
Munich, DE	April 2011	Passive monitoring	989,040
EFF servers	March–June 2010	Active IPv4 scan	11,349,678

25 million certificates to evaluate.

Our data sets

Active Scans — Passive Monitoring — EFF IPv4 scan

Location	*Time (run)*	*Type*	*Certificates*
Tuebingen, DE	November 2009	Active scan	833,661
Tuebingen, DE	December 2009	Active scan	819,488
Tuebingen, DE	January 2010	Active scan	816,517
Tuebingen, DE	April 2010	Active scan	816,605
Munich, DE	September 2010	Active scan	829,232
Munich, DE	November 2010	Active scan	827,366
Munich, DE	April 2011	Active scan	829,707
Munich, DE	April 2011	Active scan with SNI	826,098
Shanghai, CN	April 2011	Active scan	798,976
Beijing, CN	April 2011	Active scan	797,046
Melbourne, AU	April 2011	Active scan	833,571
İzmir, TR	April 2011	Active scan	825,555
São Paulo, BR	April 2011	Active scan	833,246
Moscow, RU	April 2011	Active scan	830,765
Santa Barbara, US	April 2011	Active scan	834,173
Boston, US	April 2011	Active scan	834,054
Munich, DE	September 2010	Passive monitoring	183,208
Munich, DE	April 2011	Passive monitoring	989,040
EFF servers	March–June 2010	Active IPv4 scan	11,349,678

25 million certificates to evaluate.

Our data sets

Active Scans — Passive Monitoring — EFF IPv4 scan

Location	*Time (run)*	*Type*	*Certificates*
Tuebingen, DE	November 2009	Active scan	833,661
Tuebingen, DE	December 2009	Active scan	819,488
Tuebingen, DE	January 2010	Active scan	816,517
Tuebingen, DE	April 2010	Active scan	816,605
Munich, DE	September 2010	Active scan	829,232
Munich, DE	November 2010	Active scan	827,366
Munich, DE	April 2011	Active scan	829,707
Munich, DE	April 2011	Active scan with SNI	826,098
Shanghai, CN	April 2011	Active scan	798,976
Beijing, CN	April 2011	Active scan	797,046
Melbourne, AU	April 2011	Active scan	833,571
İzmir, TR	April 2011	Active scan	825,555
São Paulo, BR	April 2011	Active scan	833,246
Moscow, RU	April 2011	Active scan	830,765
Santa Barbara, US	April 2011	Active scan	834,173
Boston, US	April 2011	Active scan	834,054
Munich, DE	September 2010	Passive monitoring	183,208
Munich, DE	April 2011	Passive monitoring	989,040
EFF servers	March–June 2010	Active IPv4 scan	11,349,678

25 million certificates to evaluate.

Results

Most results in our paper

- The SSL Landscape – A thorough analysis of the X.509 PKI using active and passive measurements
- Here: brief tour-de-force over the most interesting stuff

Validity of end-host certificates

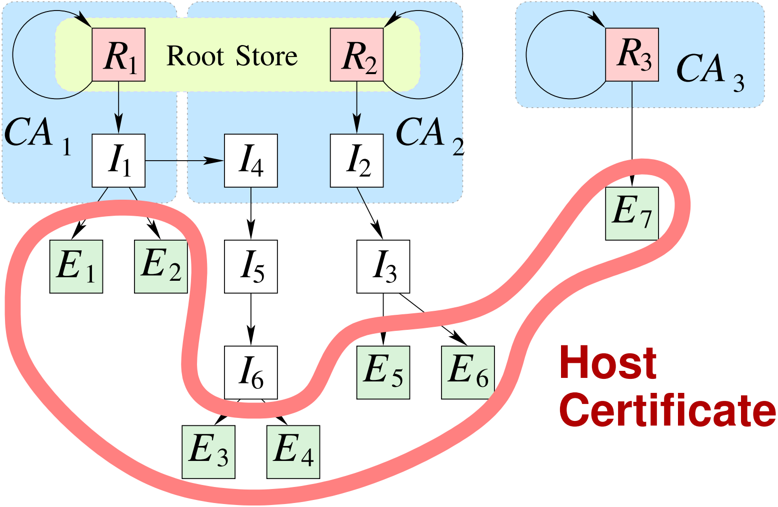

Host Certificates

Validation of certificate chains

Just check chains, not host names

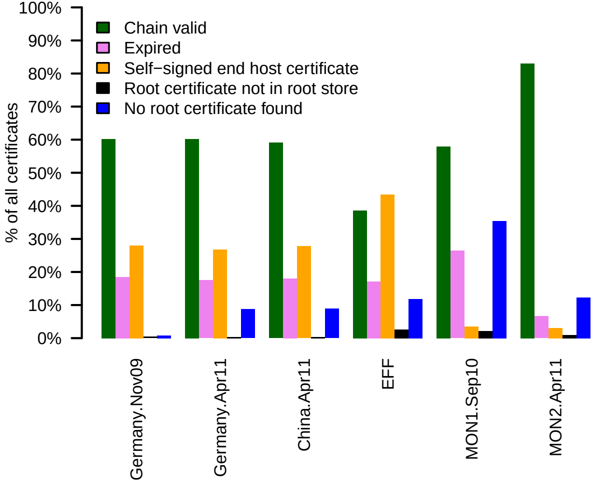

Correct domain name in certificate

Now also check host names

- Look in Common Name (CN) and Subject Alternative Name (SAN)
- Munich, April 2011, only valid chains:
  - 12.2% correct CN
  - 5.9% correct SAN

Only 18% of certificates are fully verifiable

- Positive ‘trend’: from 14.9% in 2009 to 18% in 2011

Host names in self-signed certificates

Self-signed means:

- Issuer the same as subject of certificate
- Requires out-of-band distribution of certificate

Active scan

- **2.2%** correct Common Name (CN)
- **0.5%** correct Subject Alternative Name

Certificate quality

We defined 3 categories

- ‘Good’:
  - Correct chains, correct host name
  - Chain ≤ 2
  - No MD5, strong key of $>$ 1024 bit
  - Validity $\leq$ 13 months
- ‘Acceptable’
  - Chain ≤ 3, validity $\leq$ 25 months
  - Rest as above
- ‘Poor’: the remainder

Certificate quality

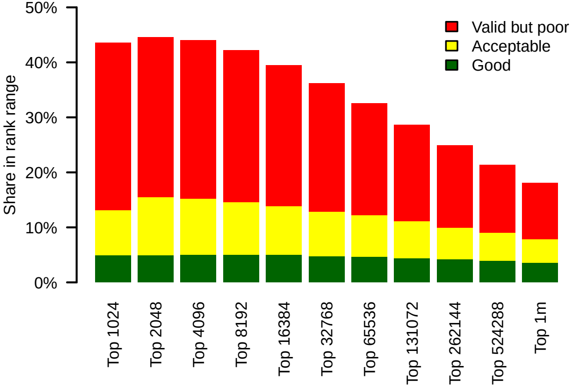

Validity correlates with rank

- Share of ‘poor’ certificates higher among high-ranking sites

X.509 for the WWW is a mess

Many more results in the paper.

In great part, the X.509 PKI is in a sorry state.

- 18% of certs in Top 1m fully valid
- Much carelessness

So far, so (not) good

Coming slowly to 2nd part of talk: Men-in-the-middle

- Question: what do users experience?
- Can we find attacks?
- Can we find proof for attacks?

THE CROSSBEAR SYSTEM

- Distributed data aquirement for detection and localization of TLS Men-In-The-Middle

Connecting to a TLS server guarded by Crossbear

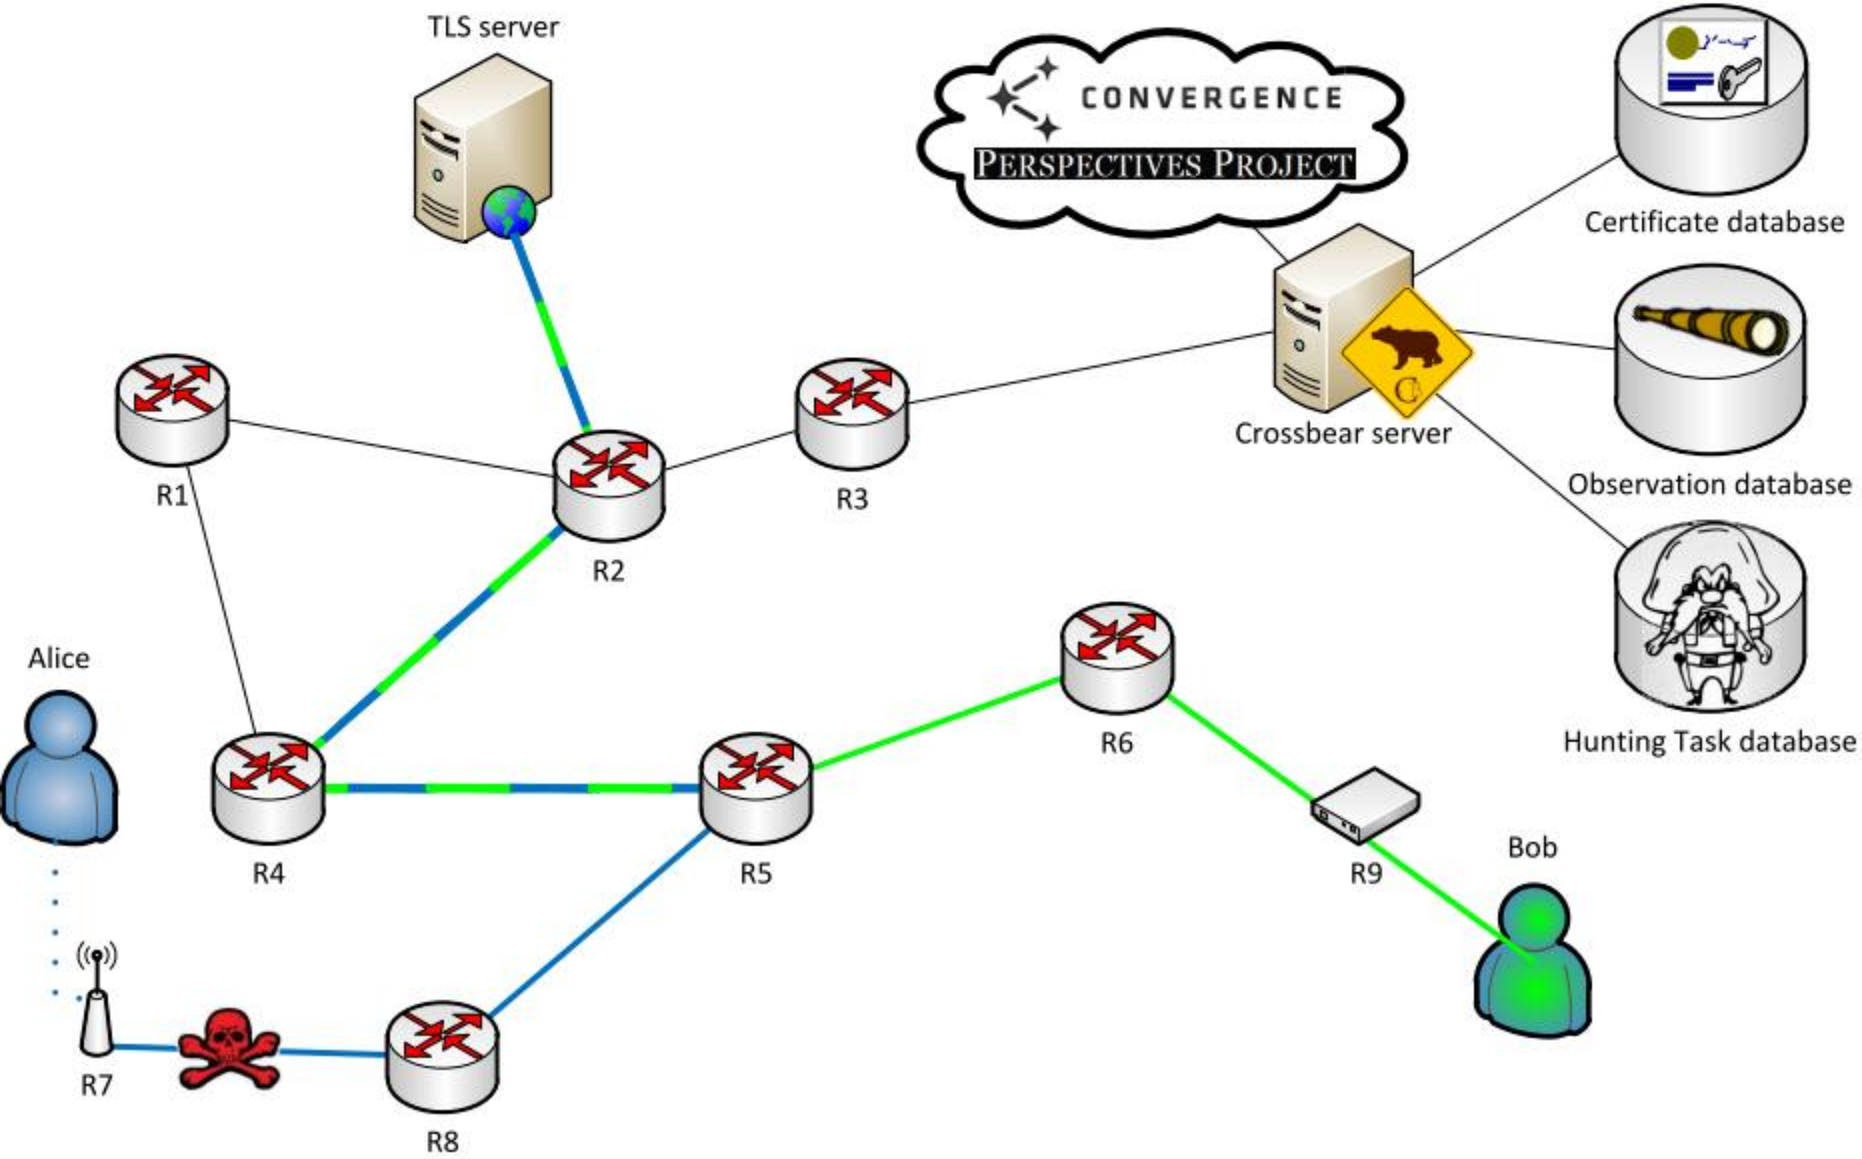

Connecting to a TLS server guarded by Crossbear

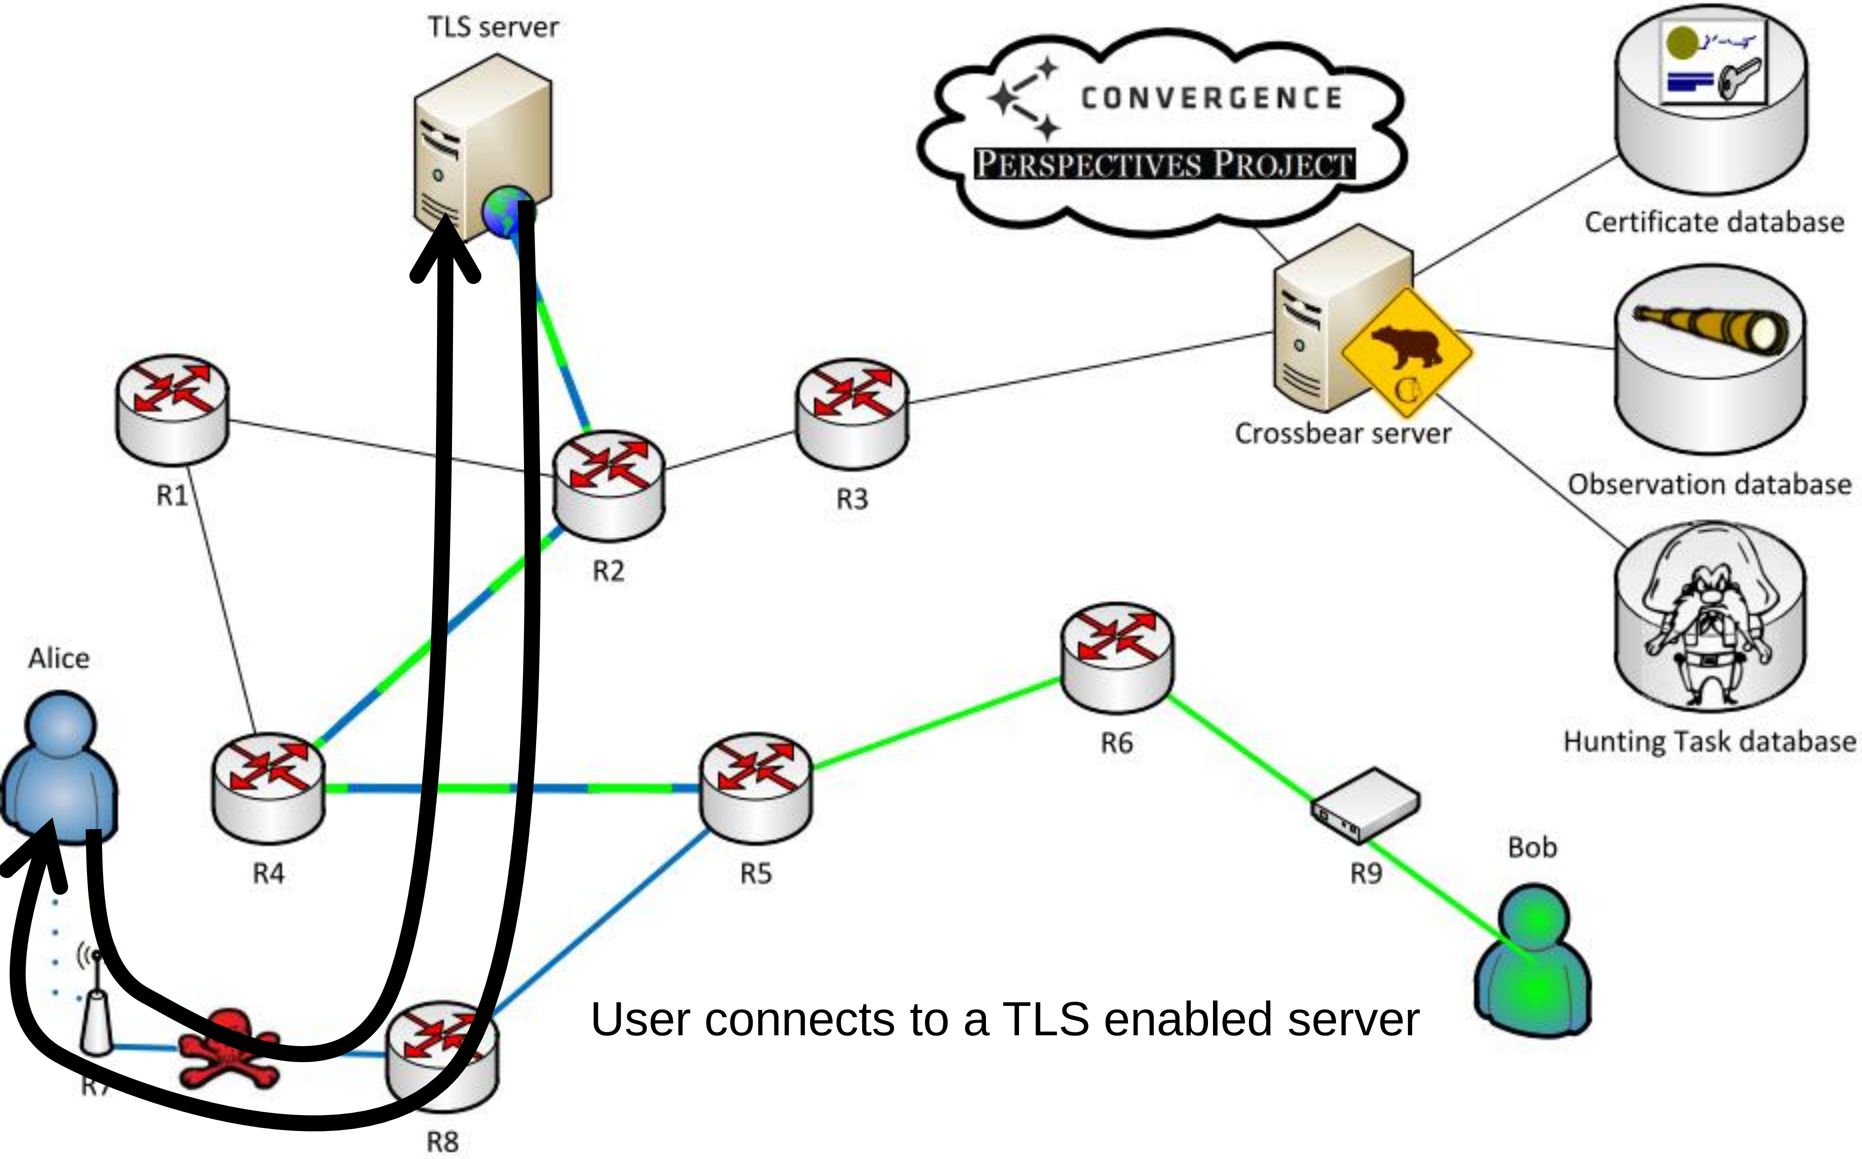

User connects to a TLS enabled server

Connecting to a TLS server guarded by Crossbear

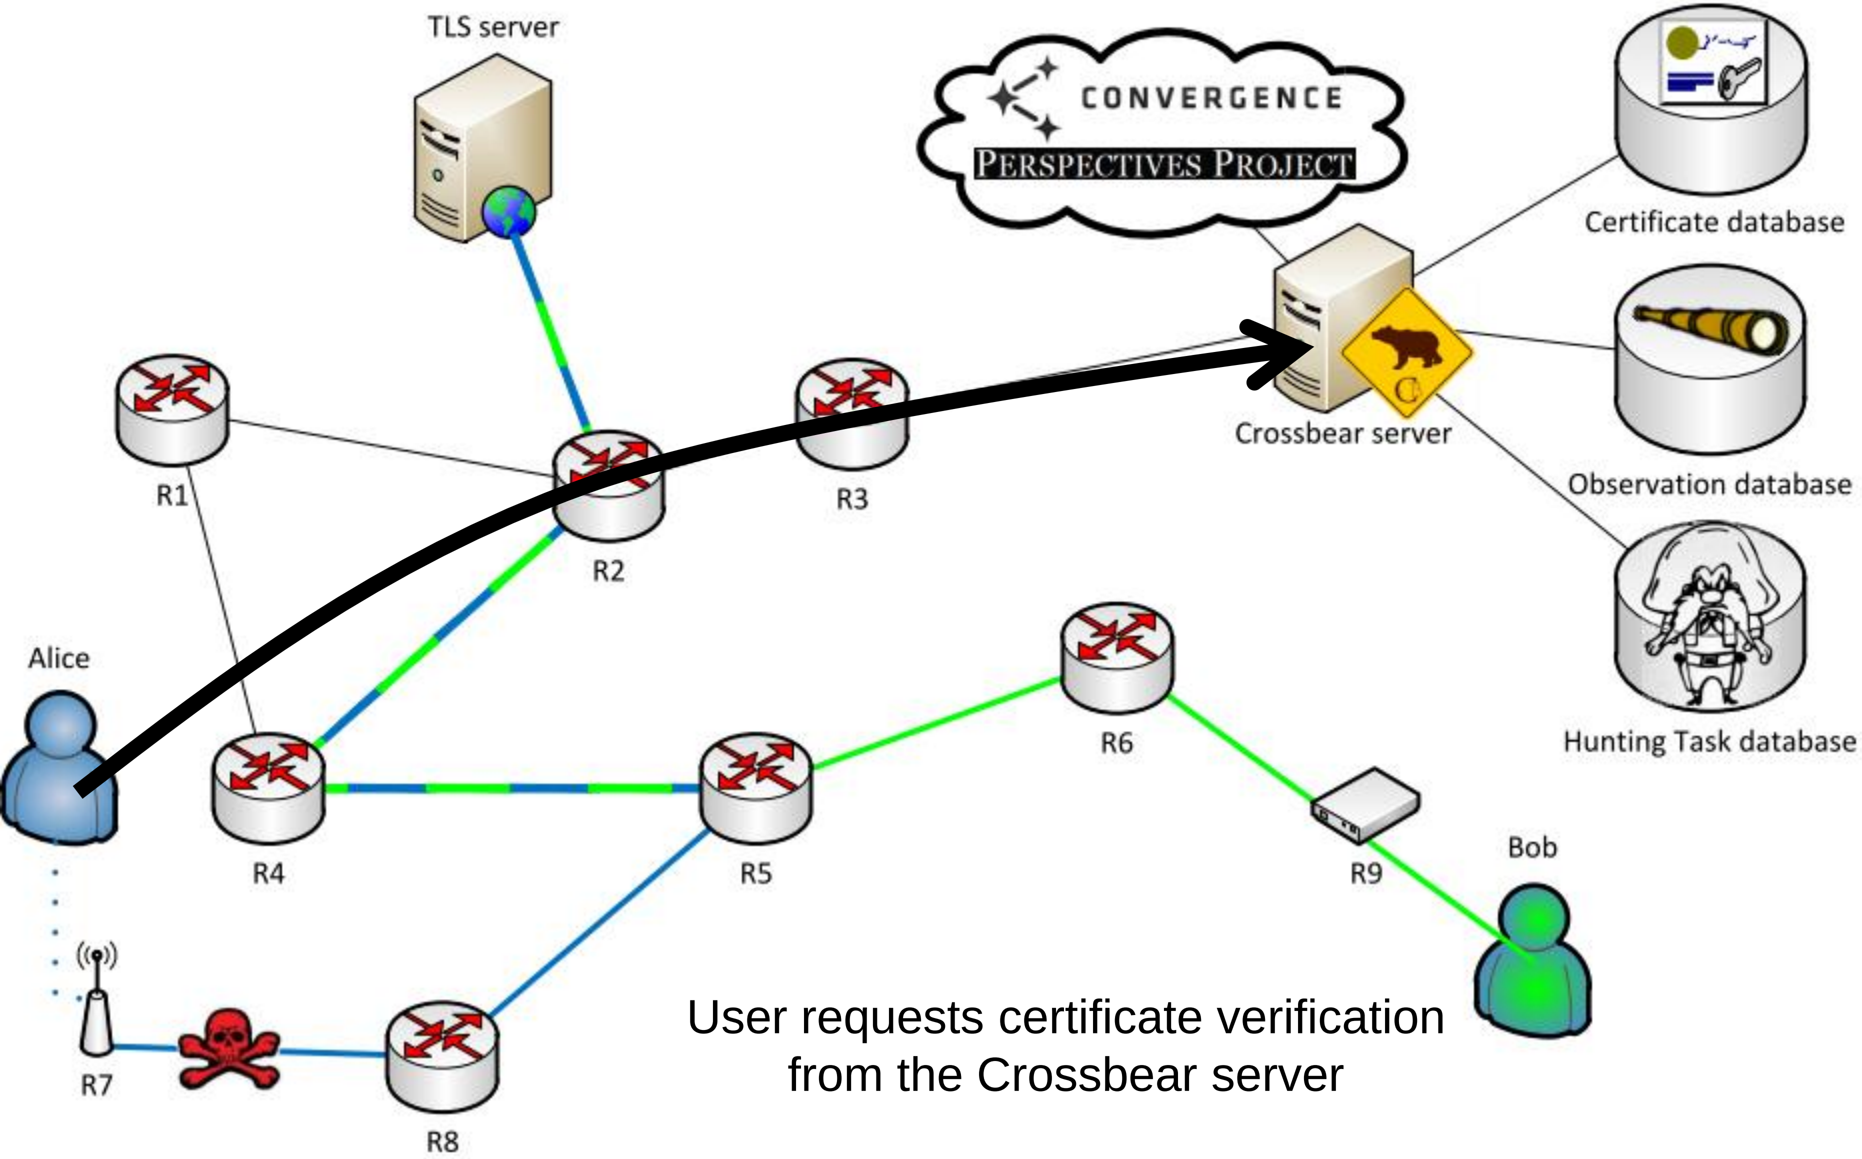

User requests certificate verification from the Crossbear server

Connecting to a TLS server guarded by Crossbear

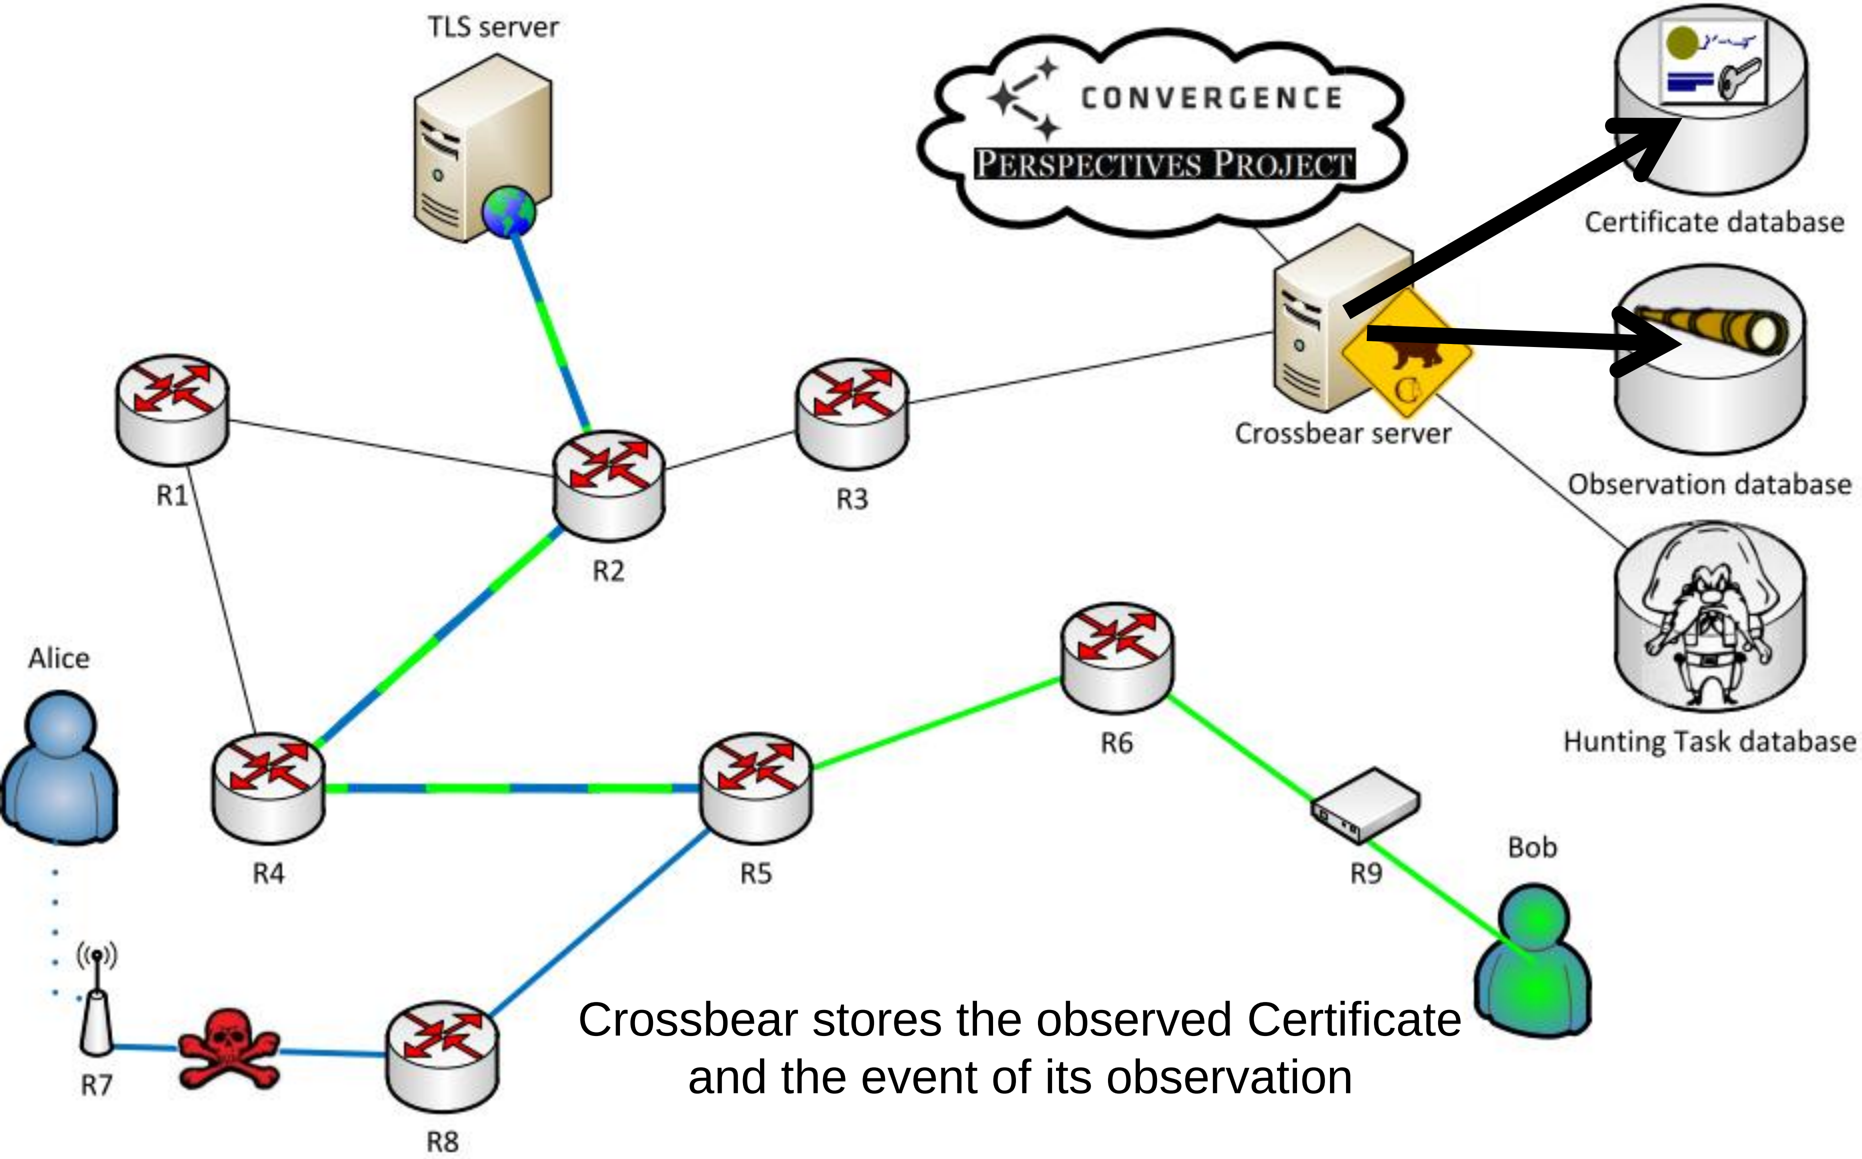

Crossbear stores the observed Certificate and the event of its observation

Connecting to a TLS server guarded by Crossbear

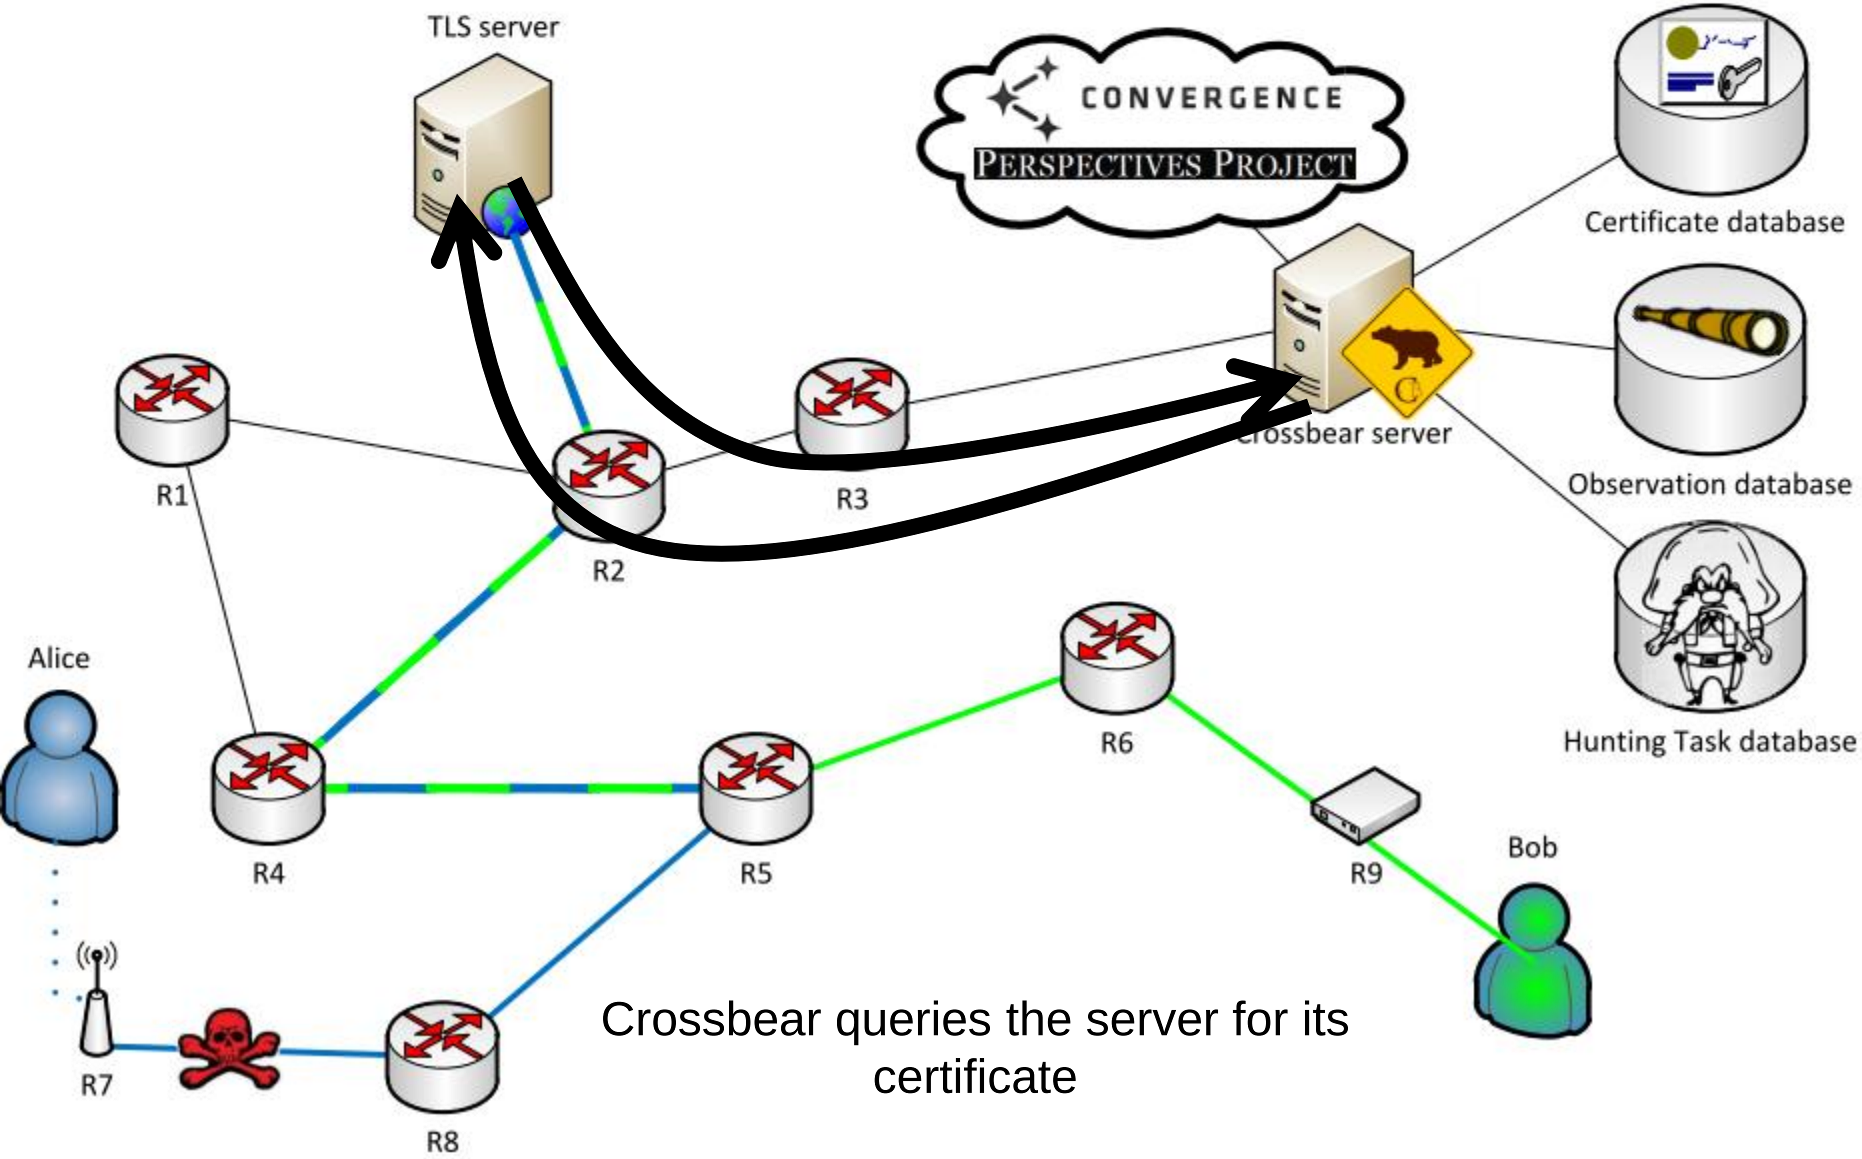

Crossbear queries the server for its certificate

Connecting to a TLS server guarded by Crossbear

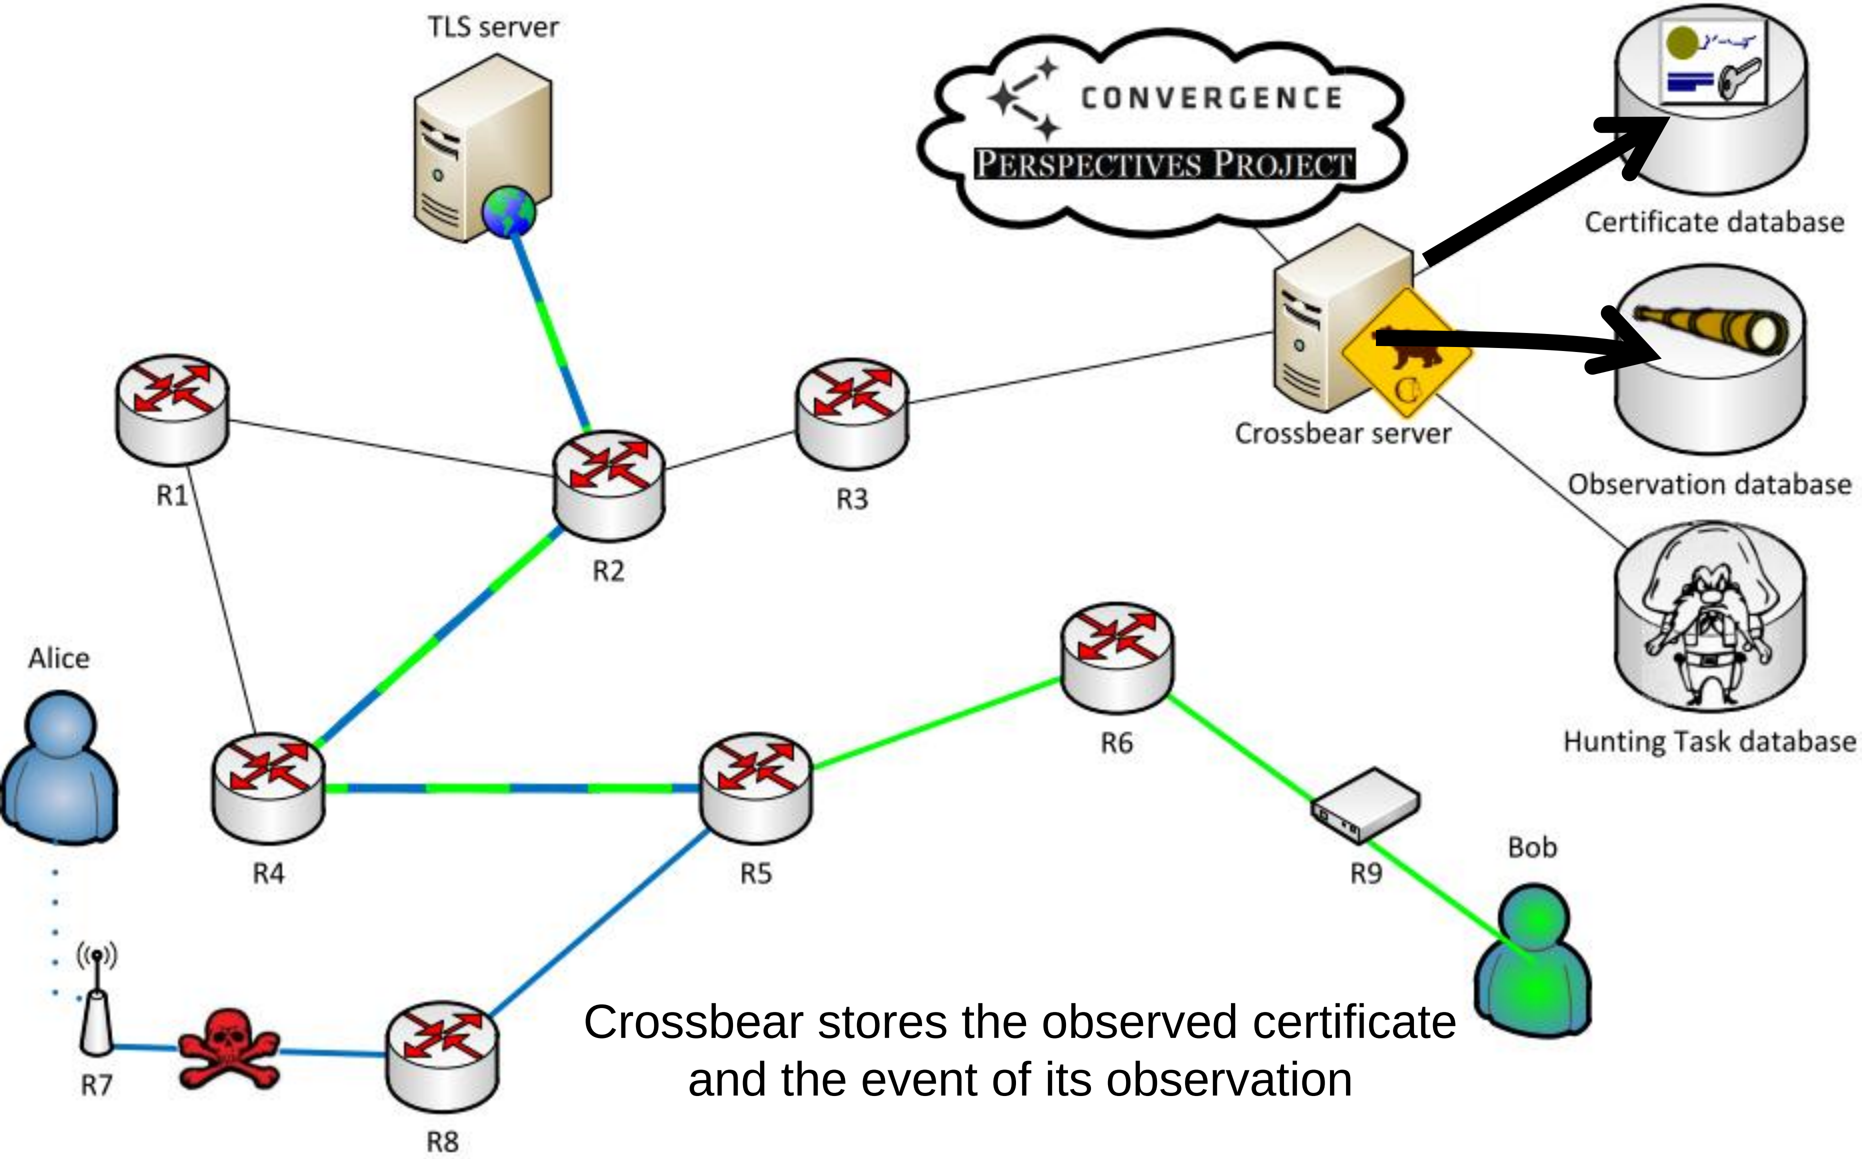

Crossbear stores the observed certificate and the event of its observation

Connecting to a TLS server guarded by Crossbear

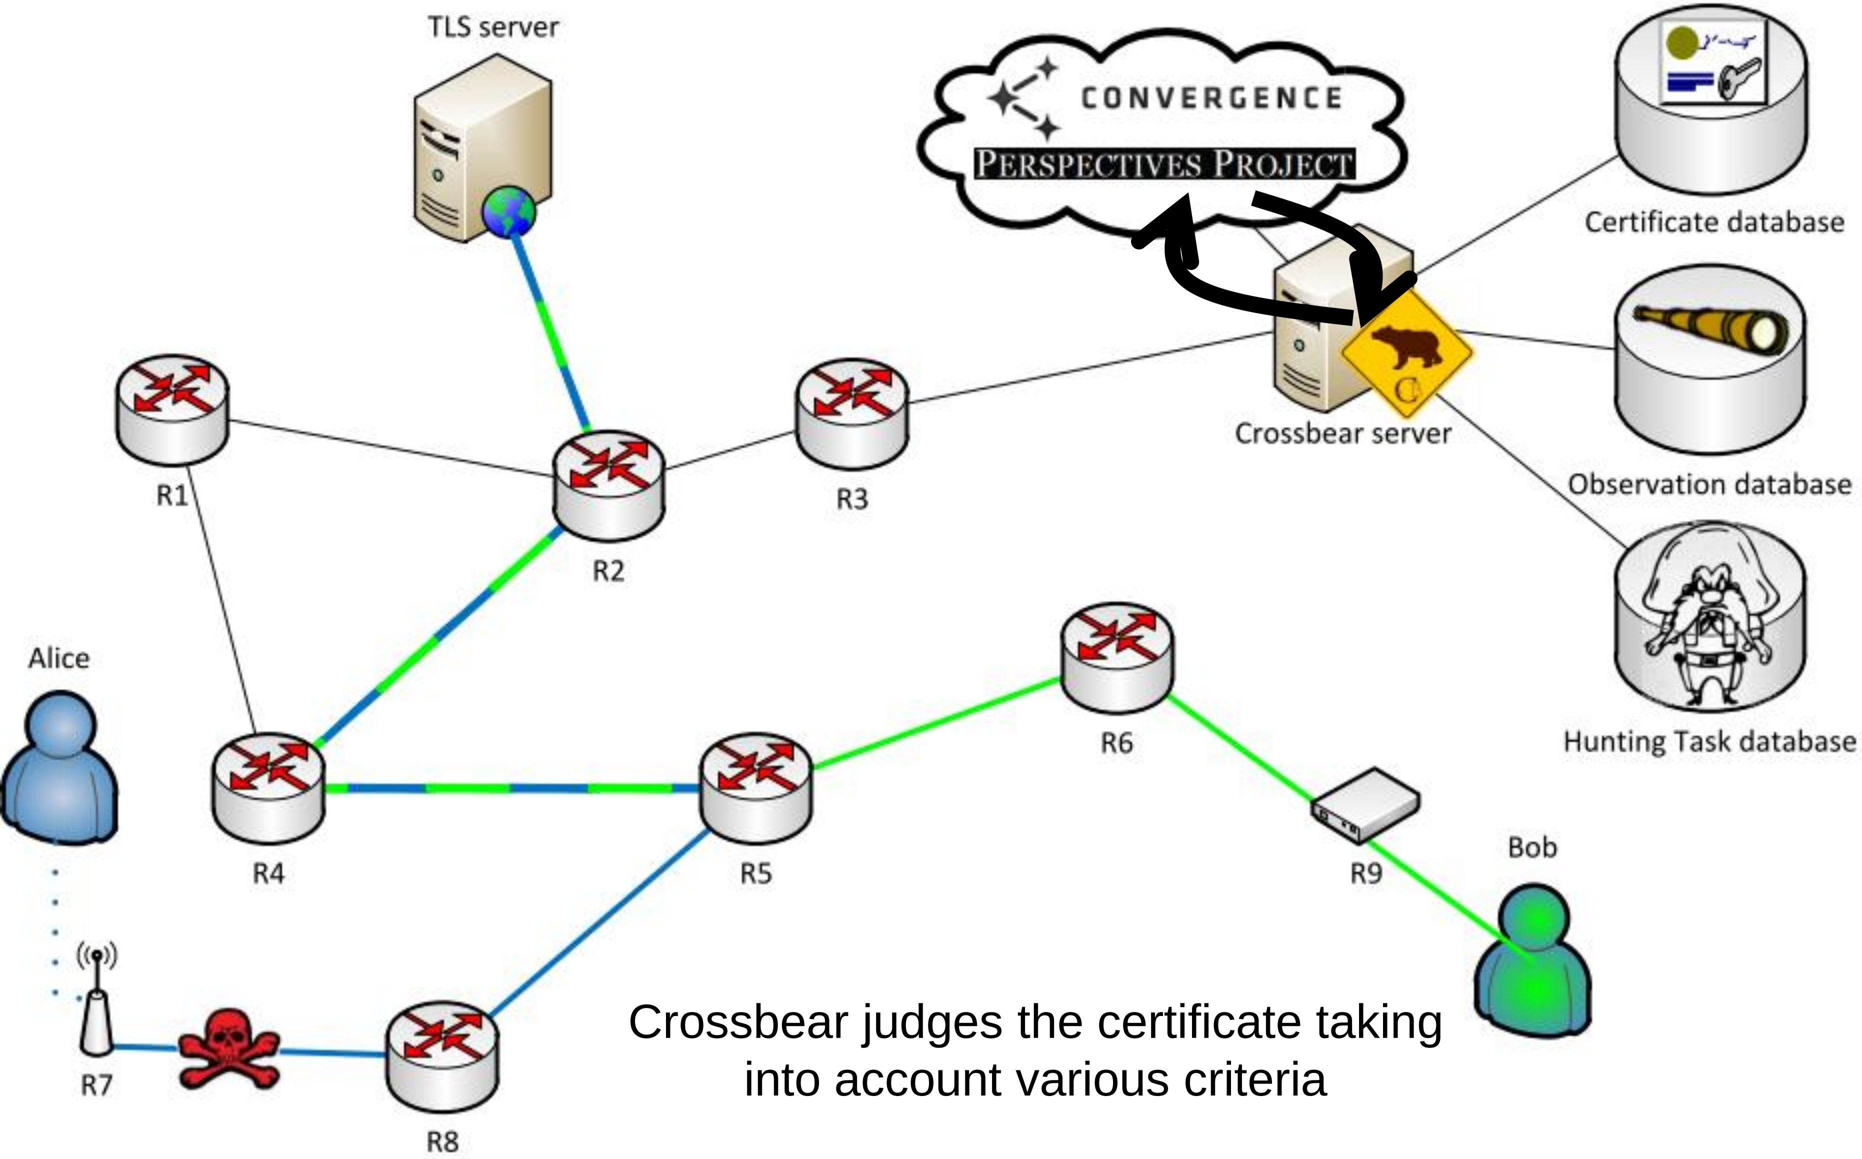

Crossbear judges the certificate taking into account various criteria

Connecting to a TLS server guarded by Crossbear

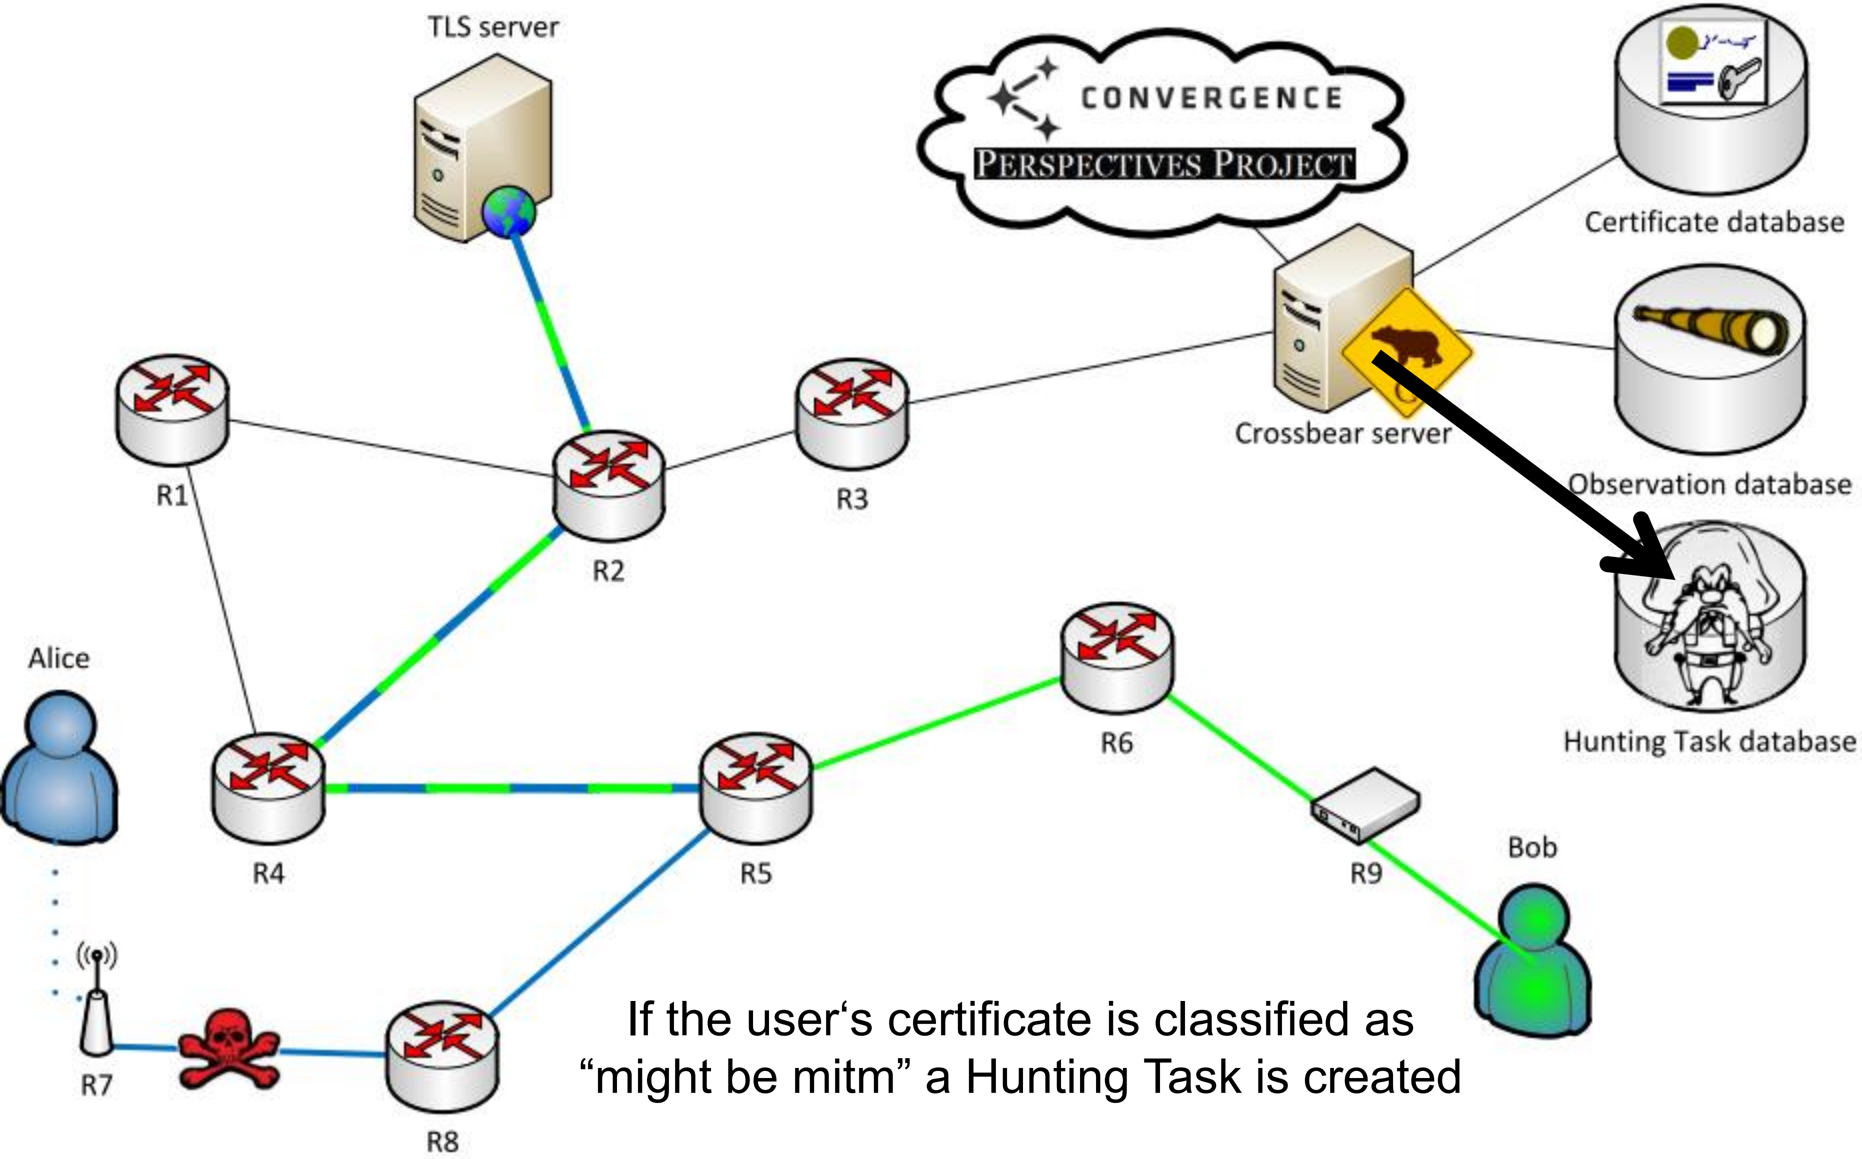

If the user‘s certificate is classified as “might be mitm” a Hunting Task is created

Connecting to a TLS server guarded by Crossbear

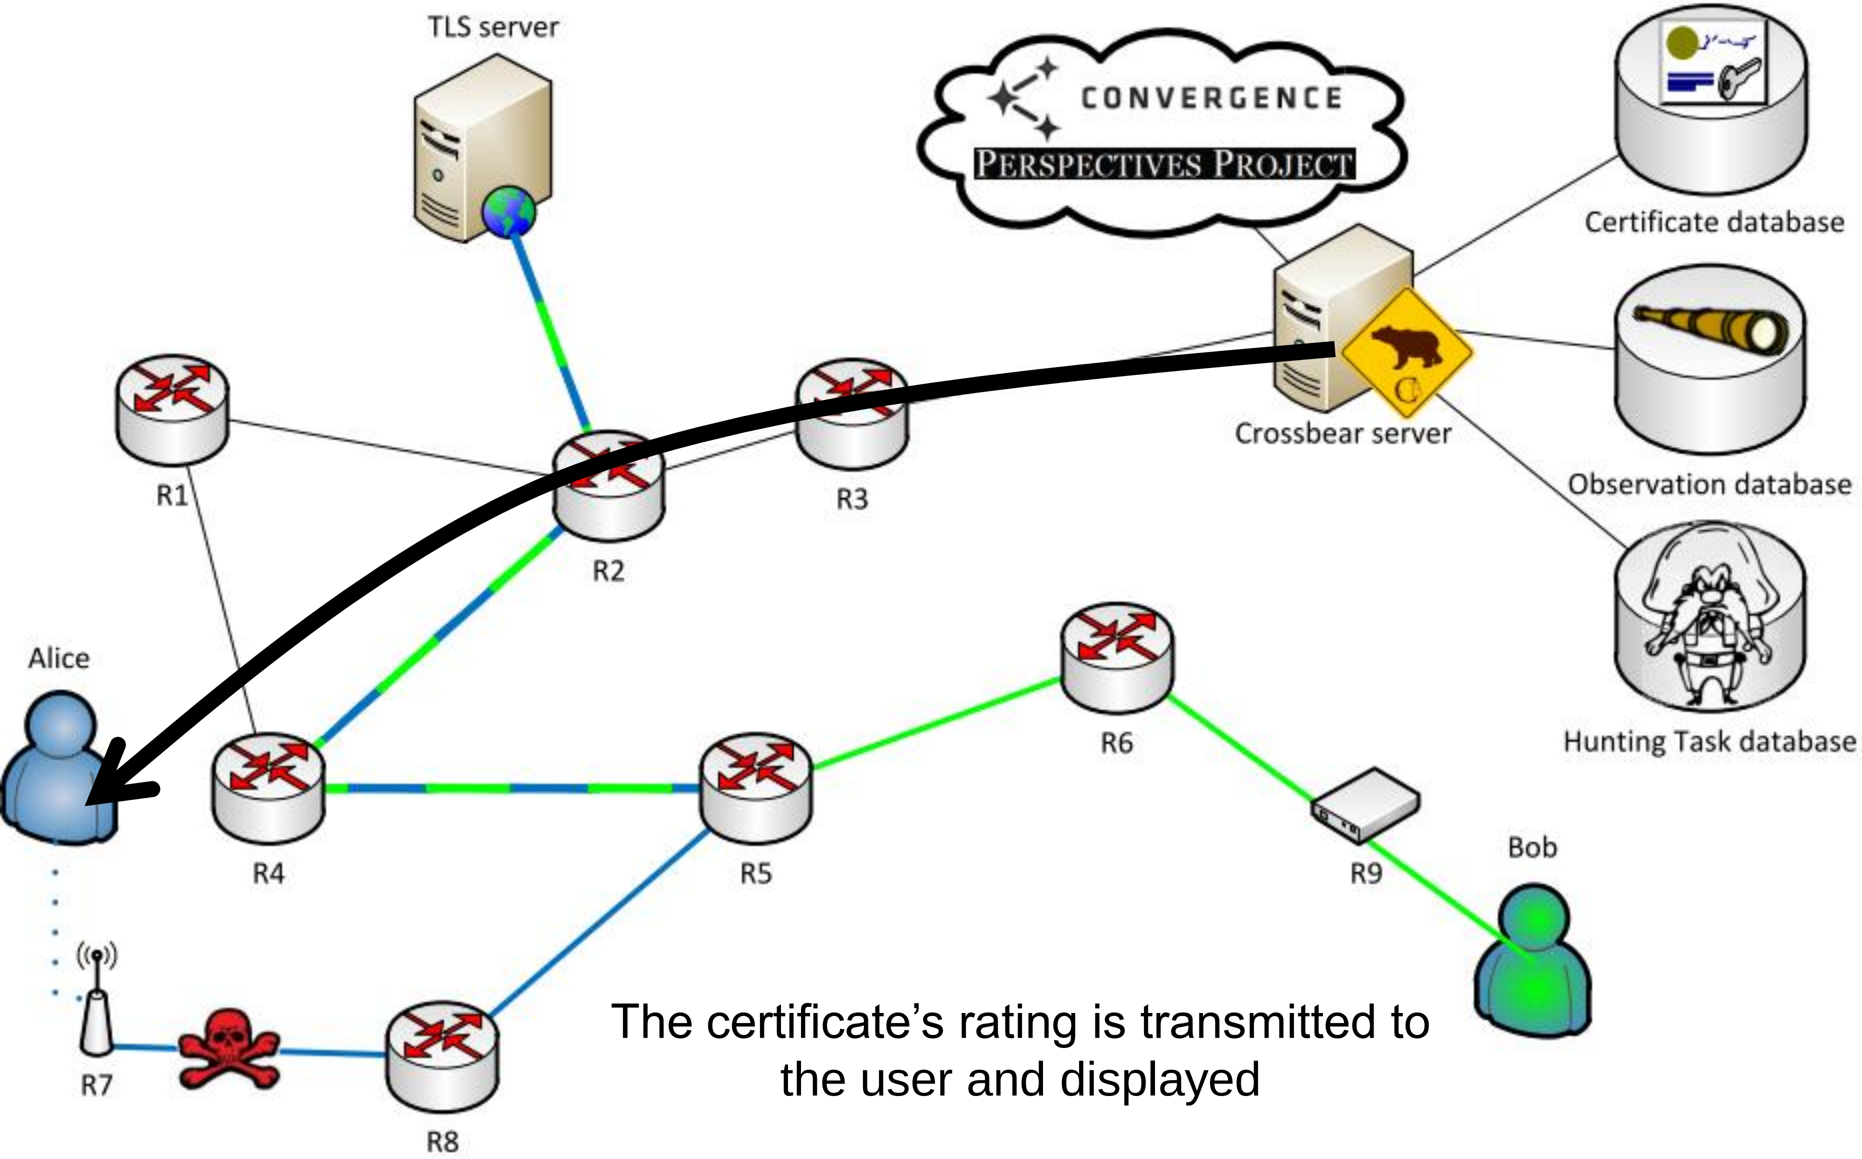

The certificate's rating is transmitted to the user and displayed

Rating scheme

- Black & White rating is not flexible enough
  - What about certificate that changed recently?
  - What about pages with several certificates?
  - What about certificates not issued for a page?

- Better: Grayscale rating (0-255)
  - Result of the certificate comparison
  - Last continuous observation period
  - Total number of observations
  - Is the certificate valid for the domain?
  - Is the certificate valid today?
  - Used Algorithms and keylength
  - What do Perspectives/Convergence think about it?

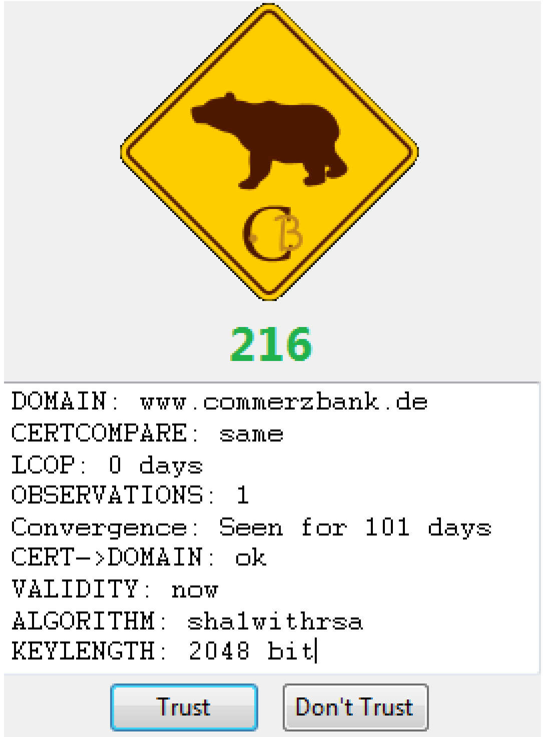

```
DOMAIN: www.commerzbank.de
CERTCOMPARE: same
LCOP: 0 days
OBSERVATIONS: 1
Convergence: Seen for 101 days
CERT->DOMAIN: ok
VALIDITY: now
ALGORITHM: sha1withrsa
KEYLENGTH: 2048 bit
```

Hunting with Crossbear

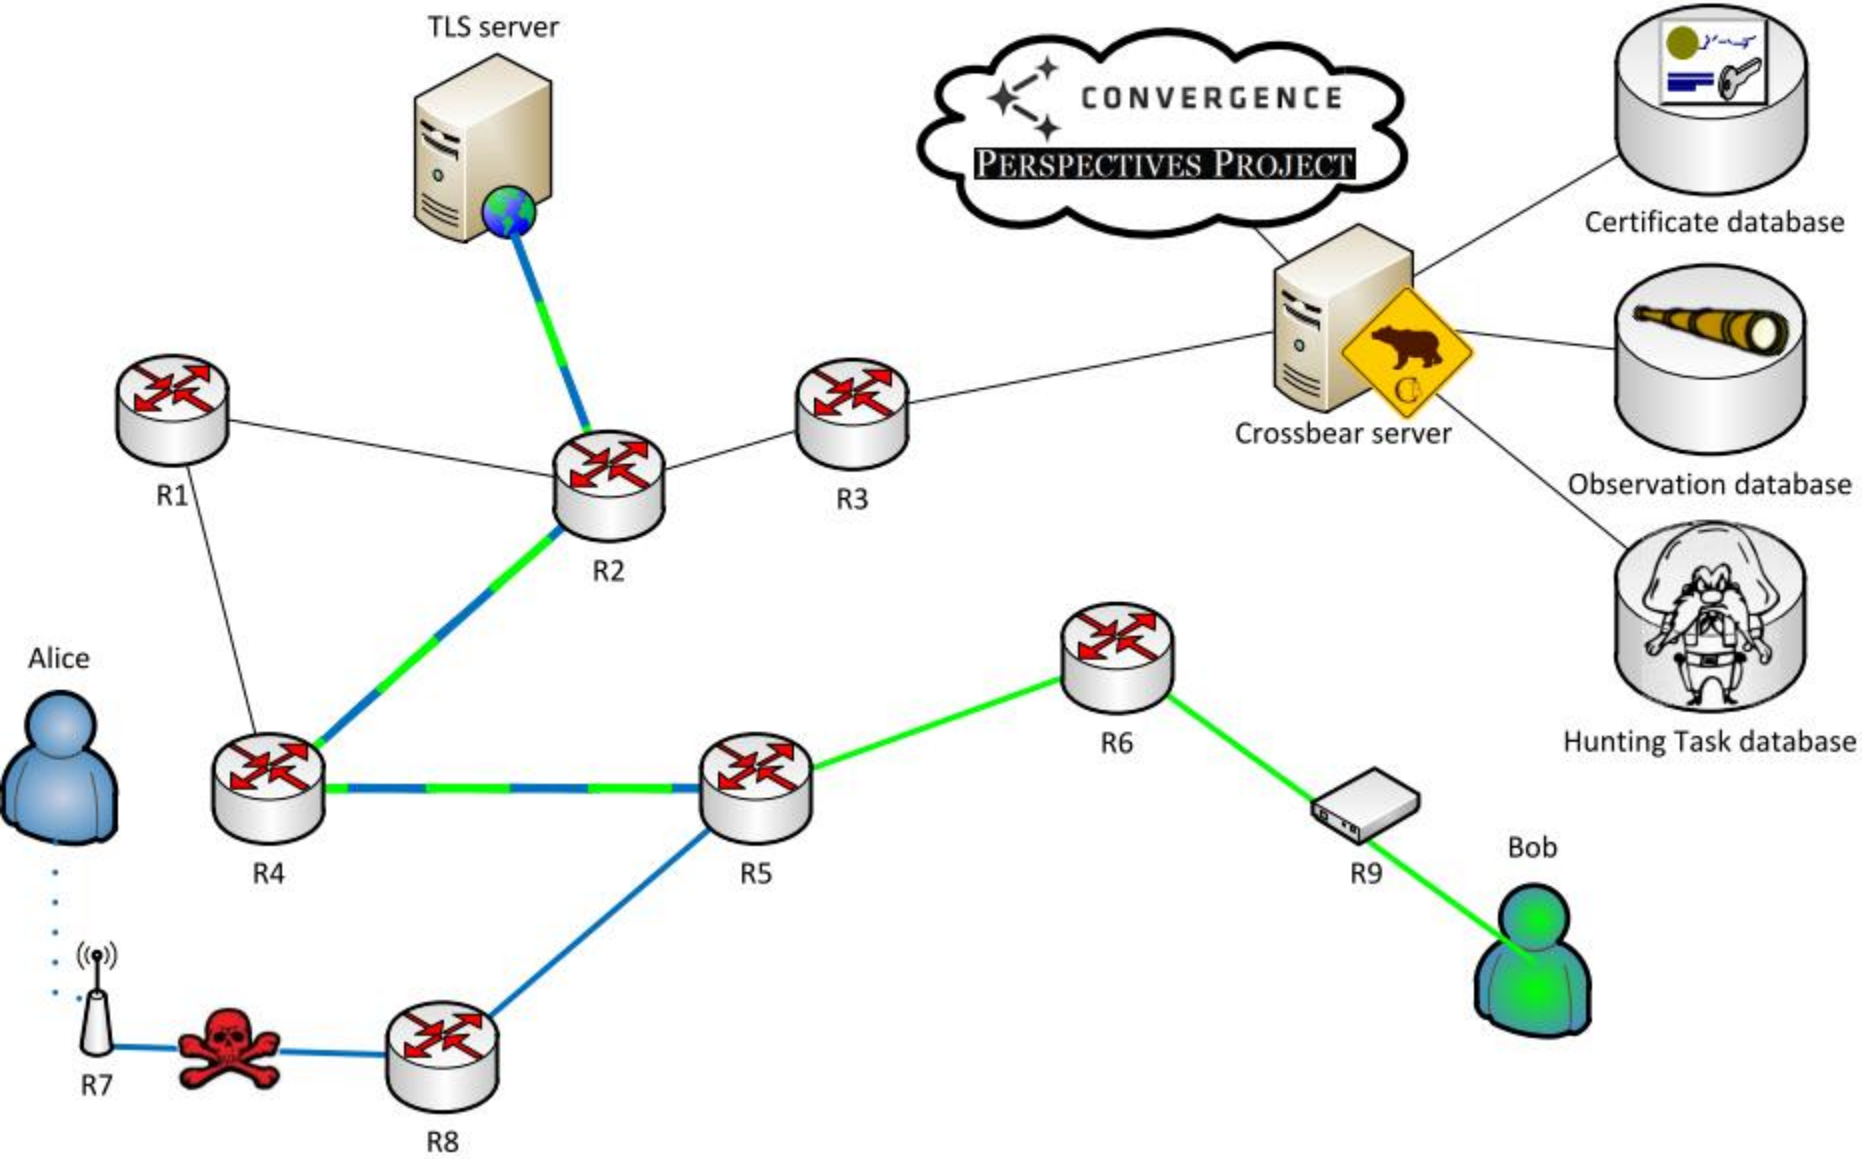

Hunting with Crossbear

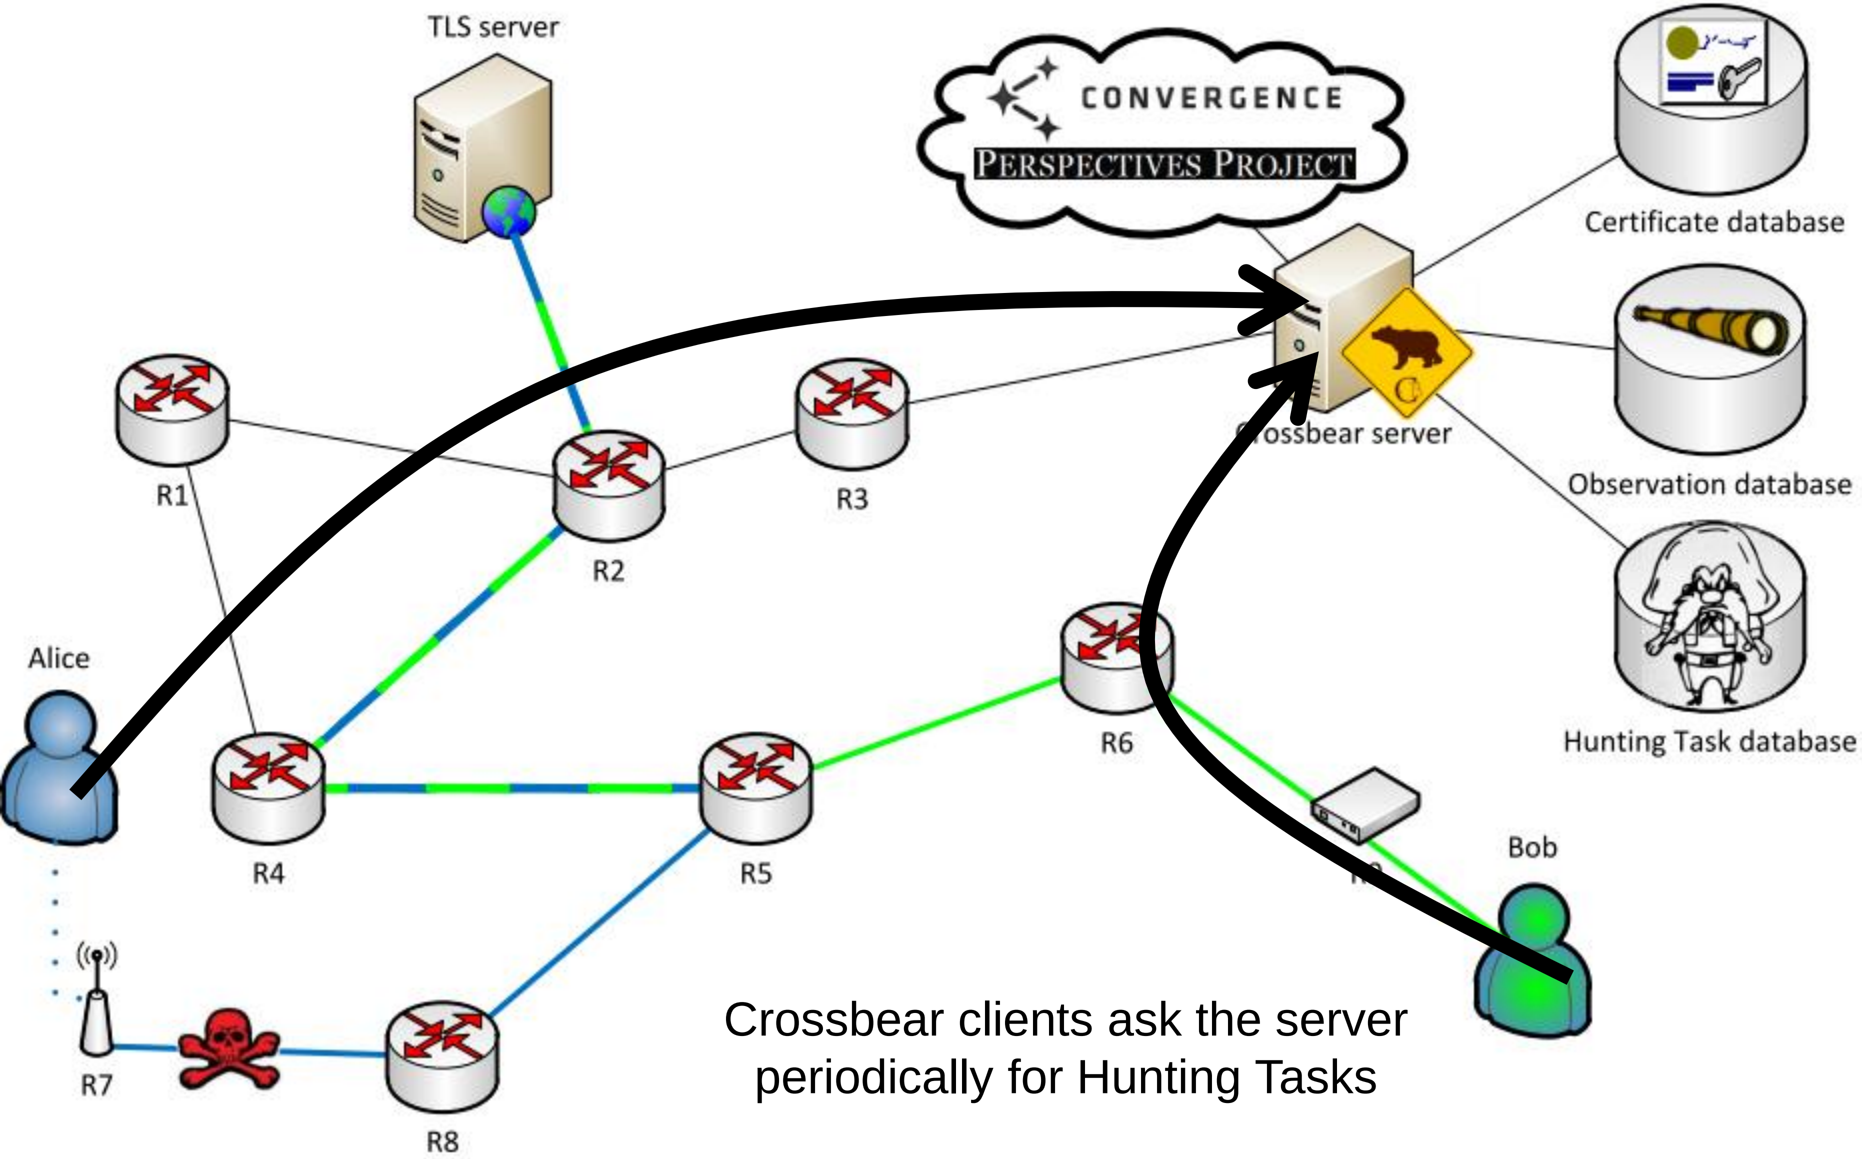

Crossbear clients ask the server periodically for Hunting Tasks

Hunting with Crossbear

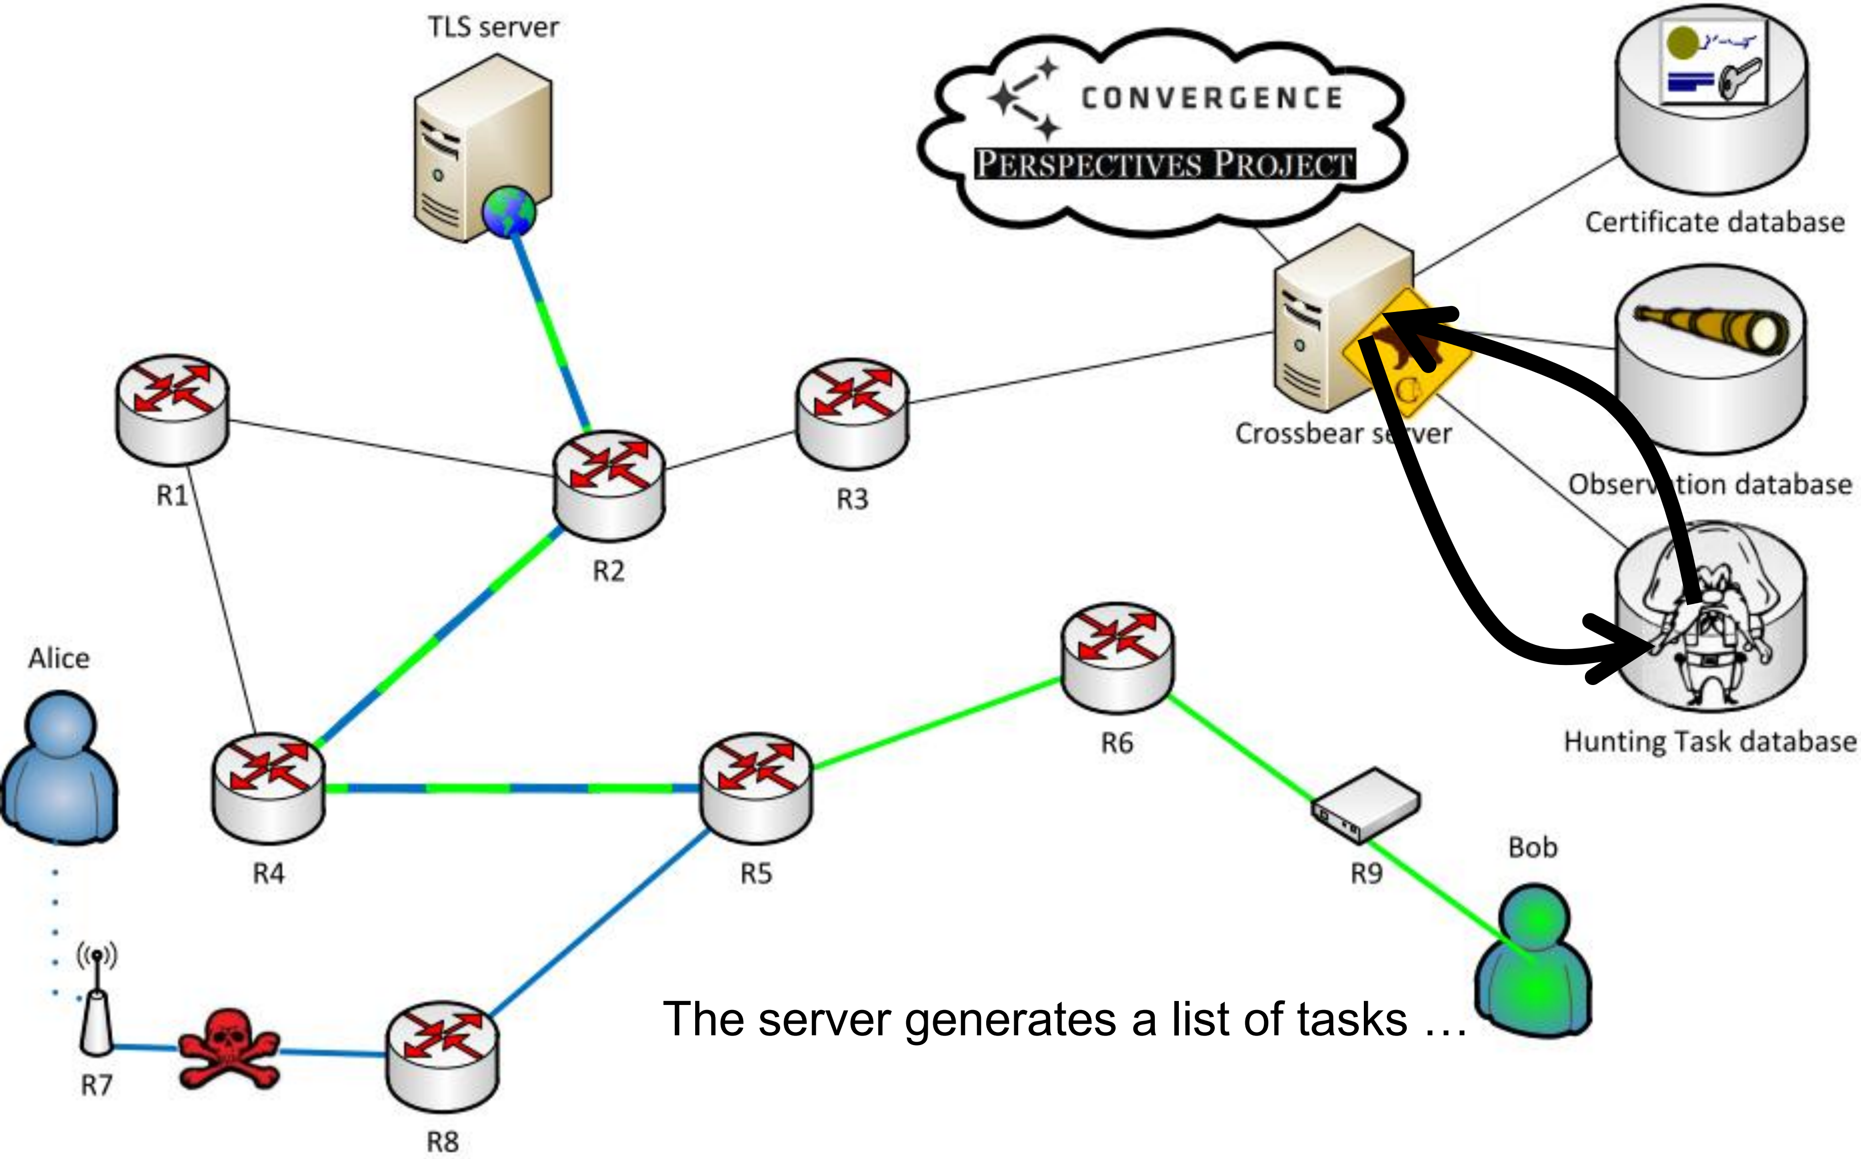

The server generates a list of tasks …

Hunting with Crossbear

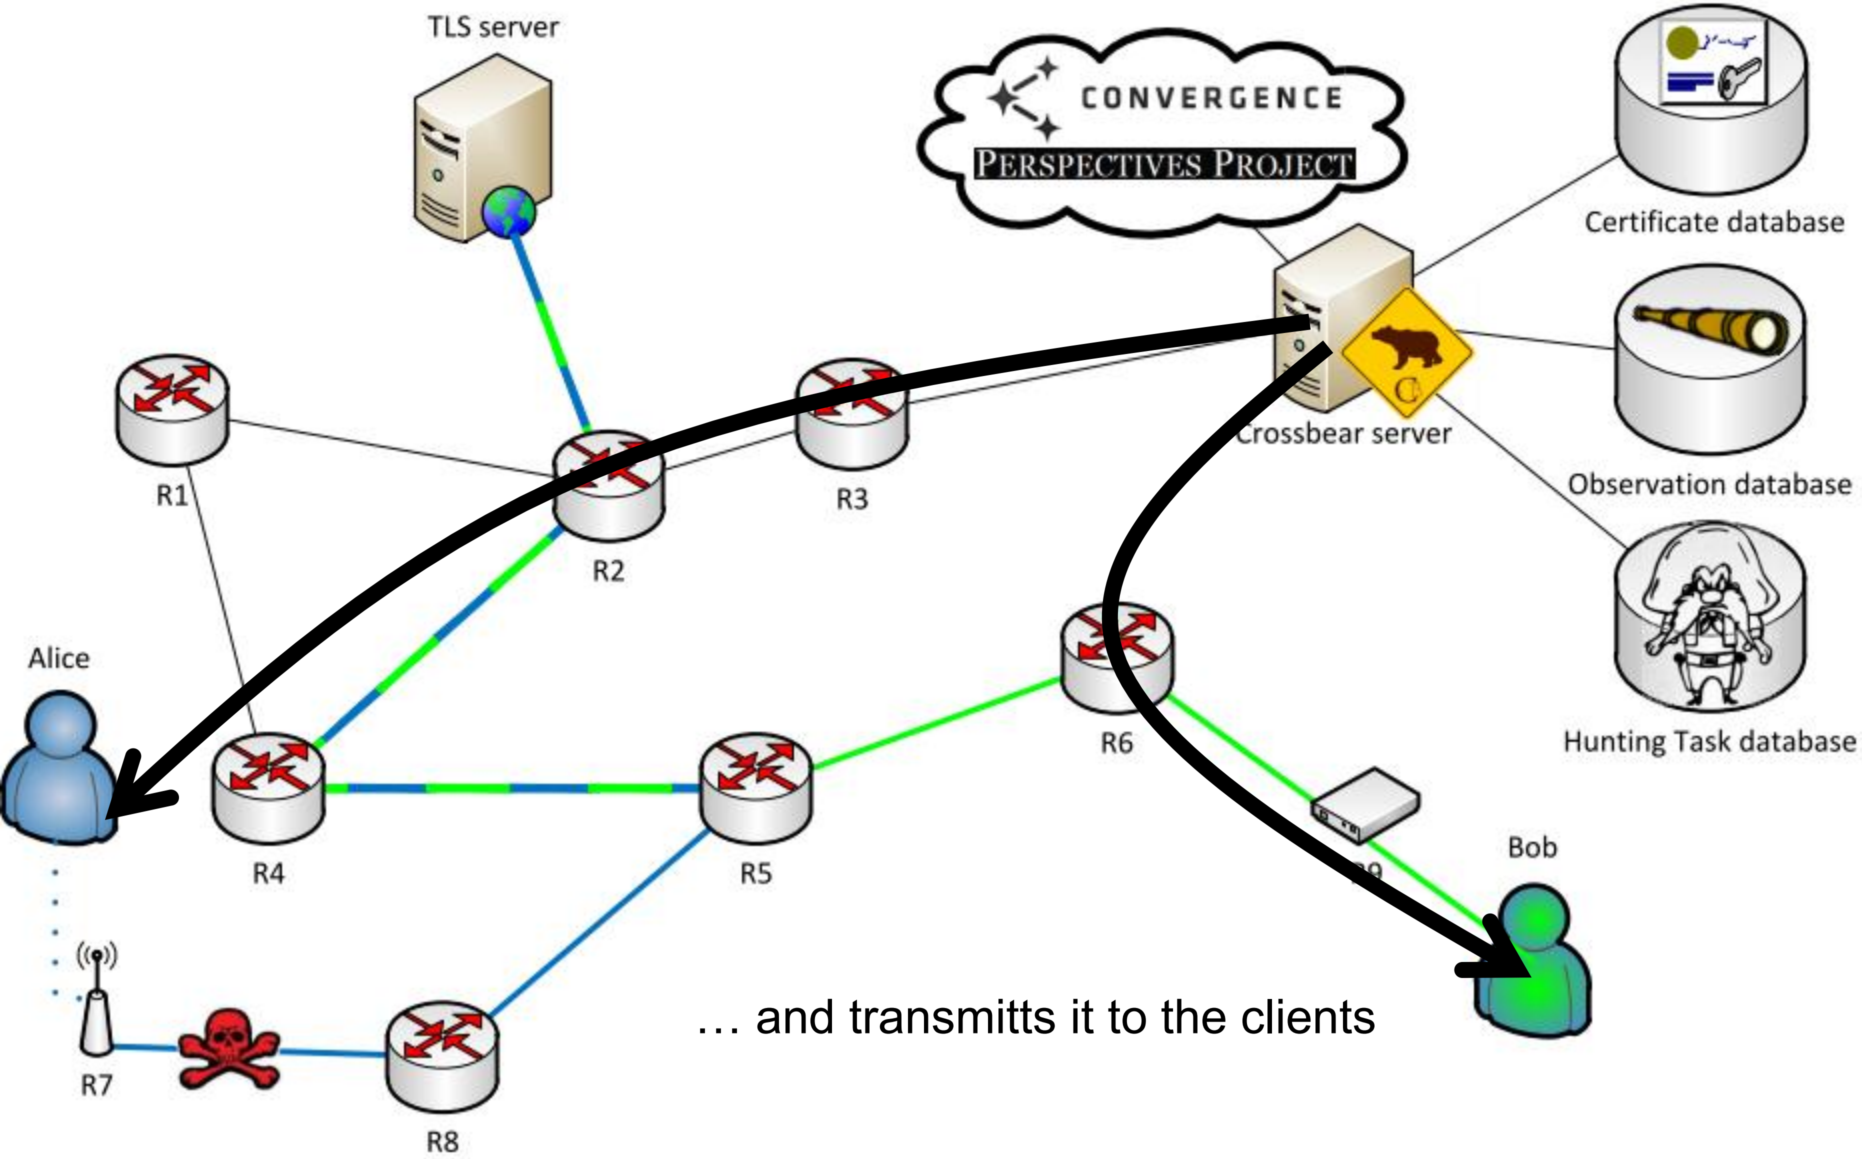

… and transmitts it to the clients

Hunting with Crossbear

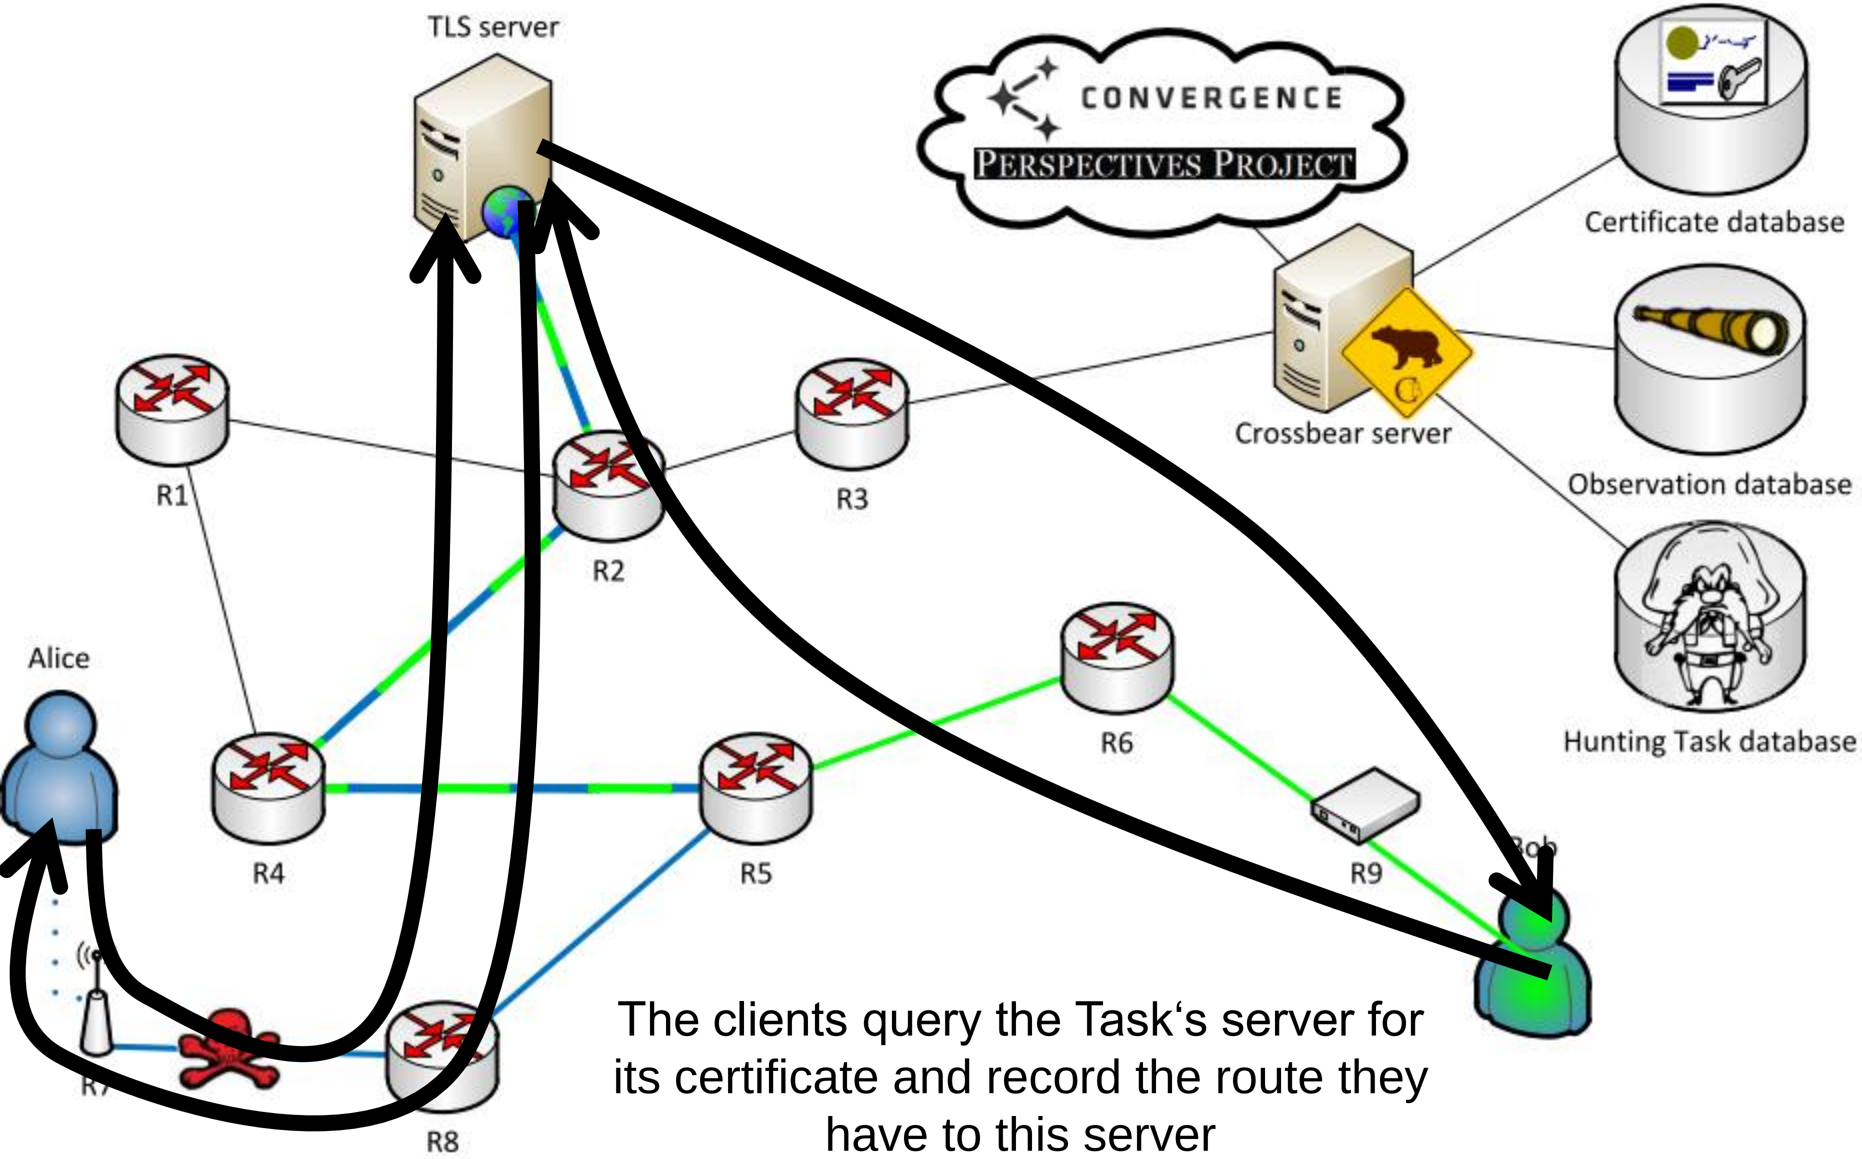

The clients query the Task's server for its certificate and record the route they have to this server

Hunting with Crossbear

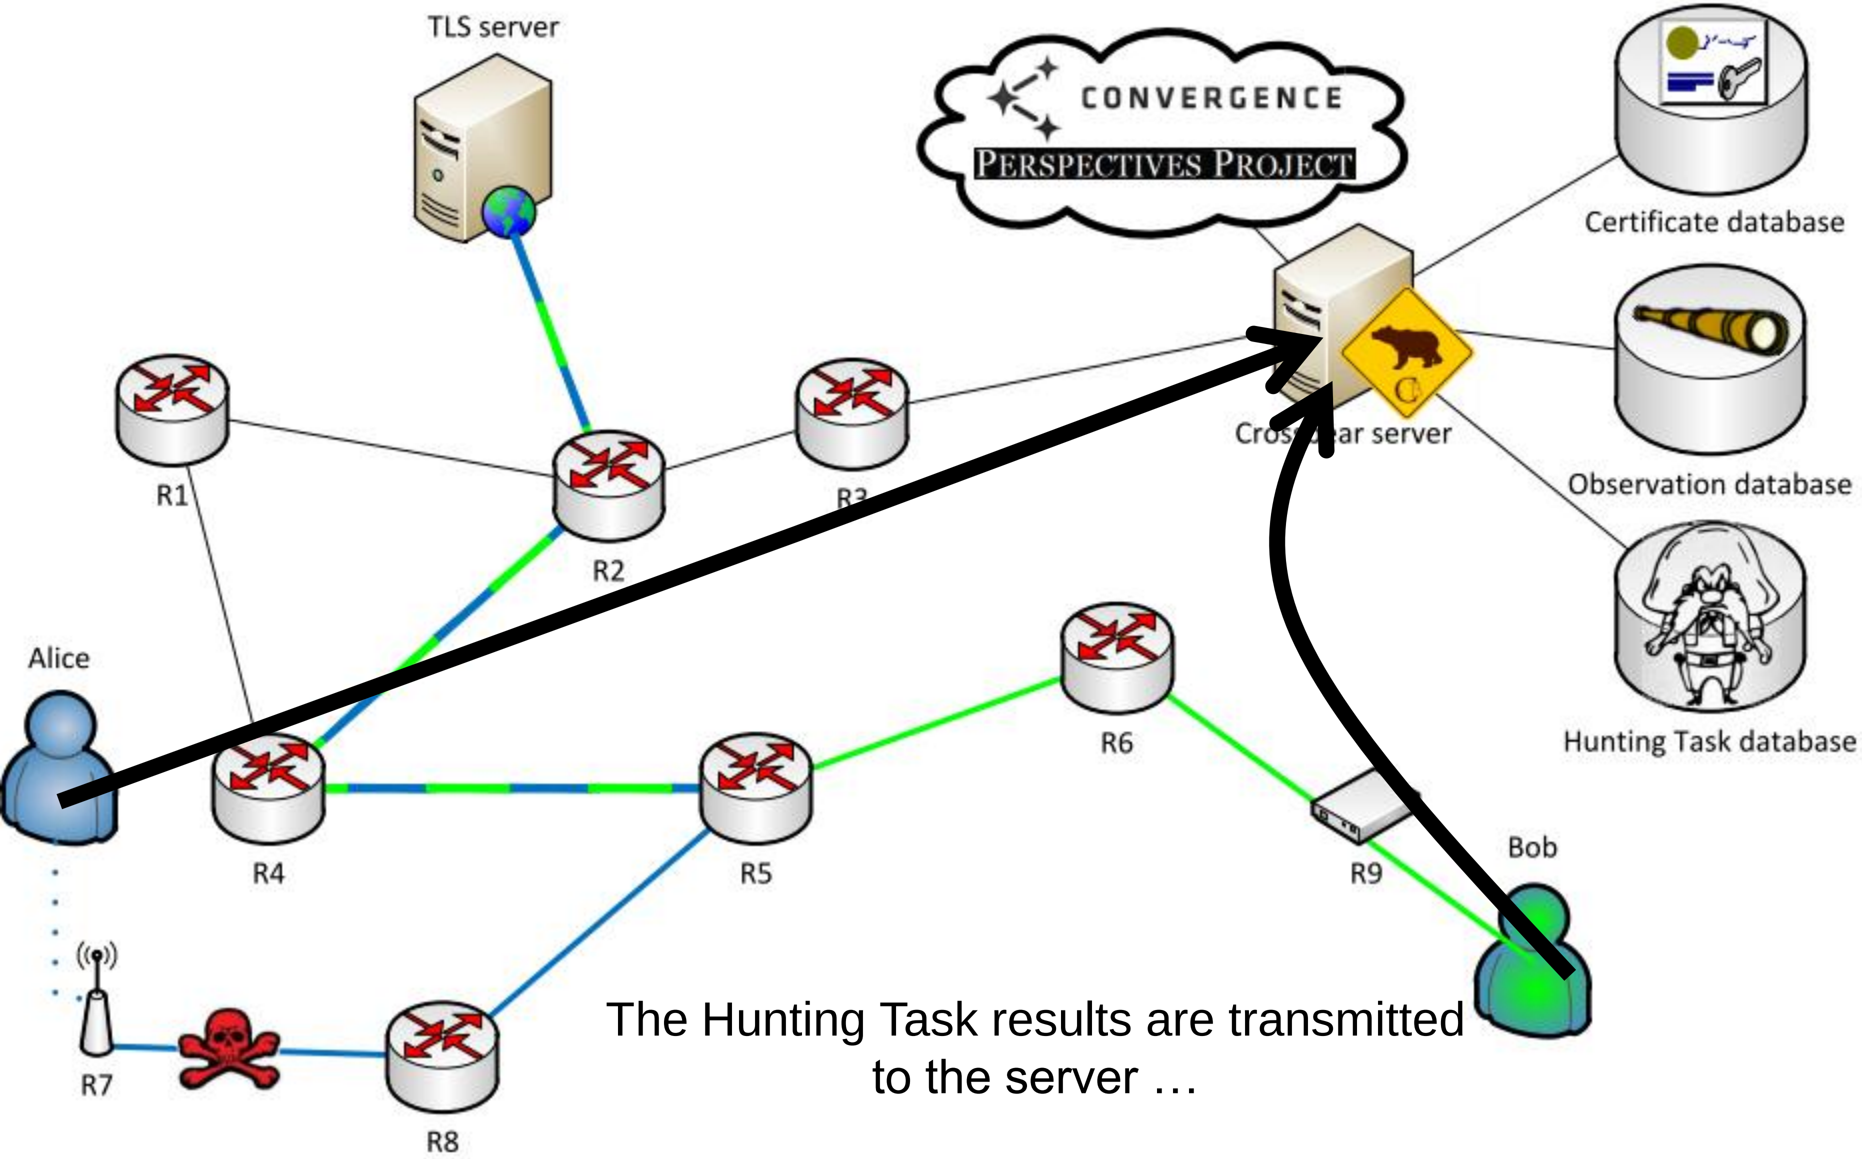

The Hunting Task results are transmitted to the server …

Hunting with Crossbear

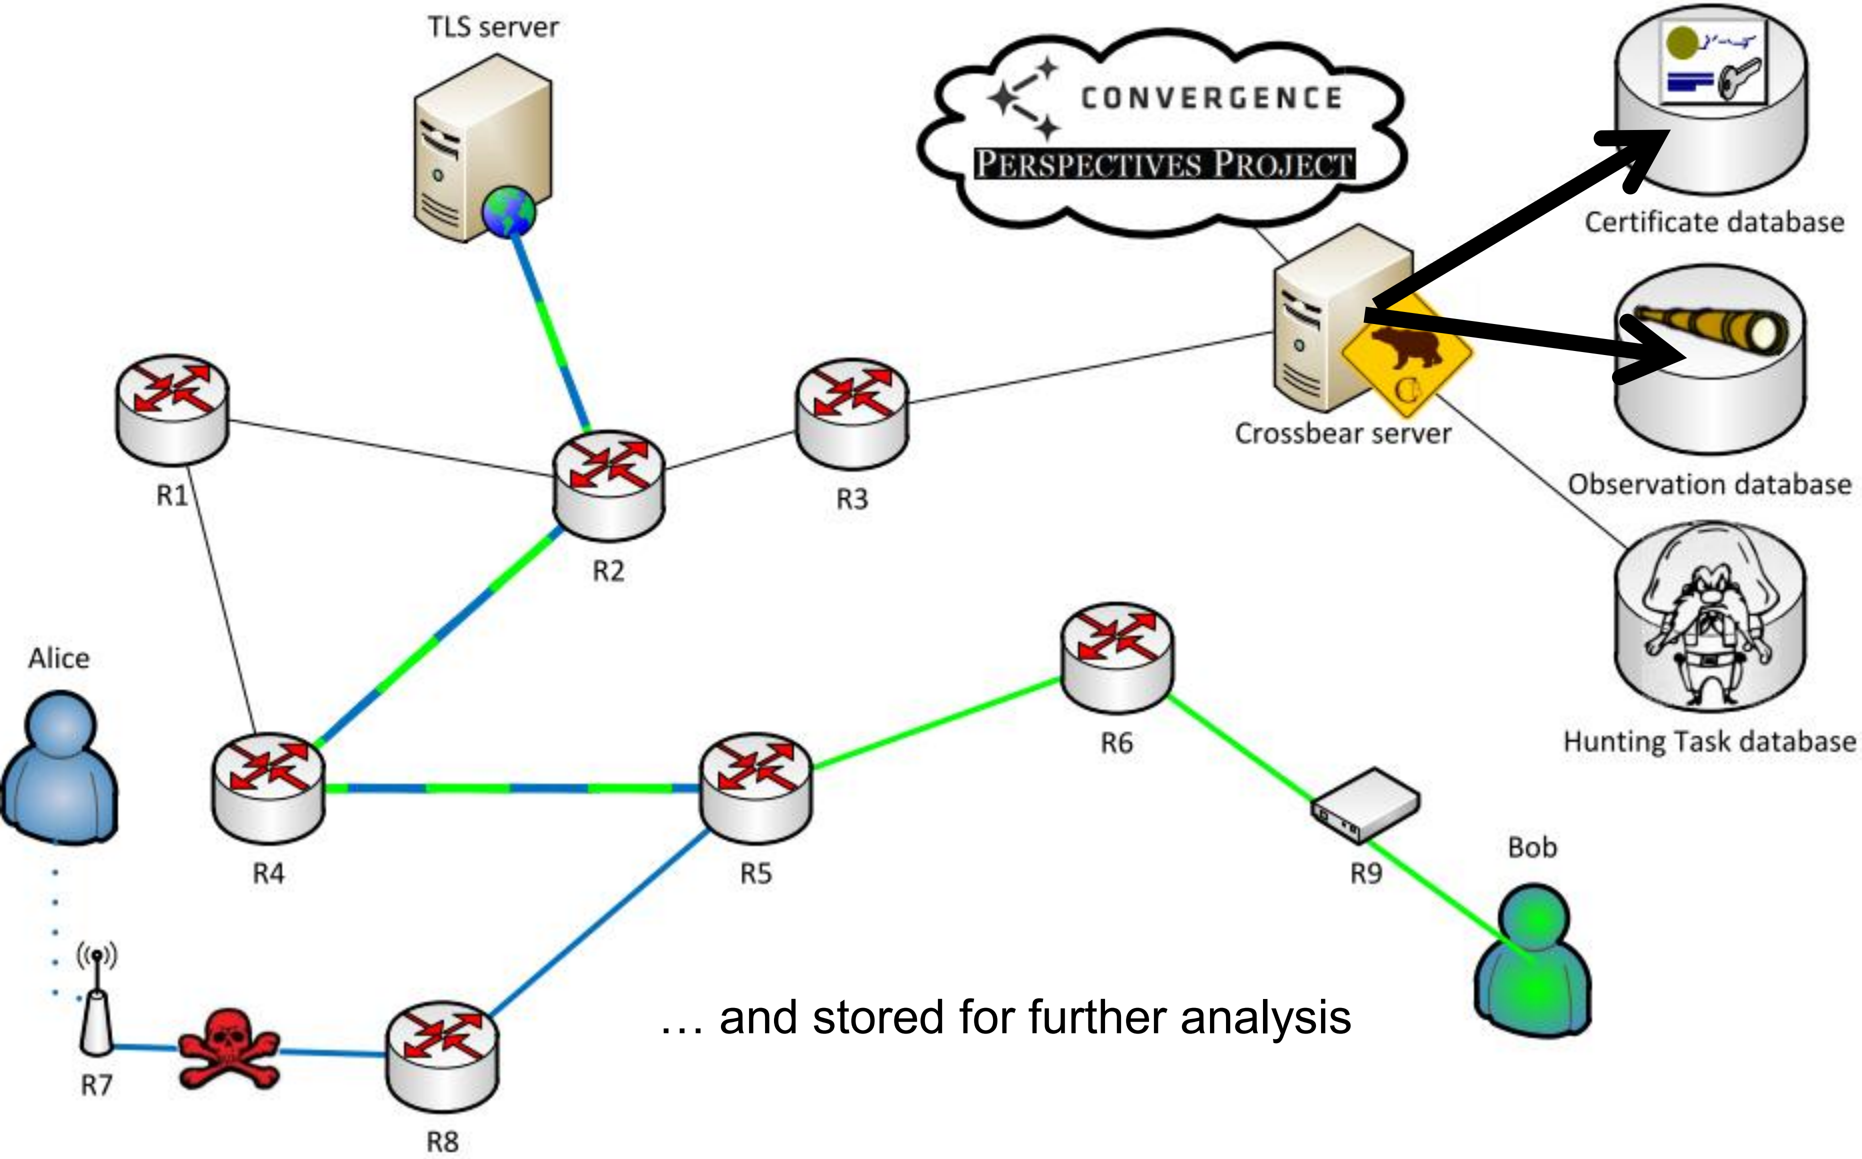

… and stored for further analysis

Constraints that shaped the system

- ❑ Usability
  - Client implemented as Firefox Plug-in
  - No external dependencies (out-of-the-box)
- ❑ Security
  - Data Confidentiality
  - Data Integrity
  - User Privacy

- ❑ Performance & Efficiency
- ❑ State-of-the-art protocols
  - Full support for IPv4 and IPv6
  - Full support for SNI
  - SHA-256 / RSA-OAEP-2048 / AES-256

How it is implemented

- Firefox Plug-in
  - Javascript extended by the Mozilla API
  - Native c-library calls through the c-types interface
    - Downloading certificate chains by the use of Firefox internal libraries
    - Traceroute by the use of Iphlpapi.dll on Windows
    - Traceroute by the use of ping and ping6 on Linux

- Server
  - Tomcat and JSP
    - JSP performs better than PHP
    - Java libraries like bouncy-castle available
    - Java code more readable than PHP
  - Located in the Faculty of Computer Science in the TU-München
    - Unlikely to be compromised by a local Mitm
    - 1GB/s uplink

Crossbear vs. Convergence/Perspectives

- Convergence/Perspectives
  - Problems with some pages
    - SNI-enabled pages
    - Non-TLS legacy systems
  - Focus on users’ privacy and protection
  - Guard functionality only

CONVERGENCE Beta

PERSPECTIVES PROJECT

- Crossbear
  - Works with all pages that Firefox can show
  - Focus on collection of Mitm-related data
    - IPs are stored (partly anonymized)
    - Observations will be published
  - Guard and hunting functionality

Threats to data integrity

- Crossbear Firefox plug-in is freely available
  - No user authentication (anybody can use it)
  - Source code is known to potential attackers (Open Source)

- Attackers could send invalid Hunting Task Replies
  - False positive: forged certificate instead of correct one
  - False negative: correct one instead of forged one
  - False routes

- Why one would do that
  - Accidentally (e.g. because of proxies)
  - To cover the position of ones Mitm
  - To make somebody else look like a Mitm
  - …

Verification of server traces

- Route verification using knowledge about Internet topology
- Assert first IP in trace equals client’s public IP
  - Reduce attacker’s options
  - Might not be the IP sending the Hunting Task Reply (IPv4 <-> IPv6)
  - Implementation:
    - PublicIP-Notification-Messages contain HMAC of public IP

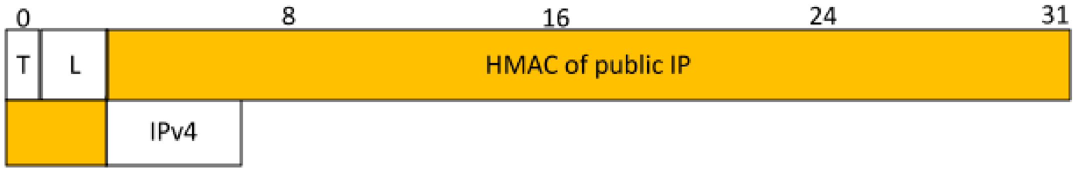

0 8 16 24 31

T | L | HMAC of public IP

IPv4

T: Message Type; L: Message Length

    - Hunting Task results contain that HMAC, too:

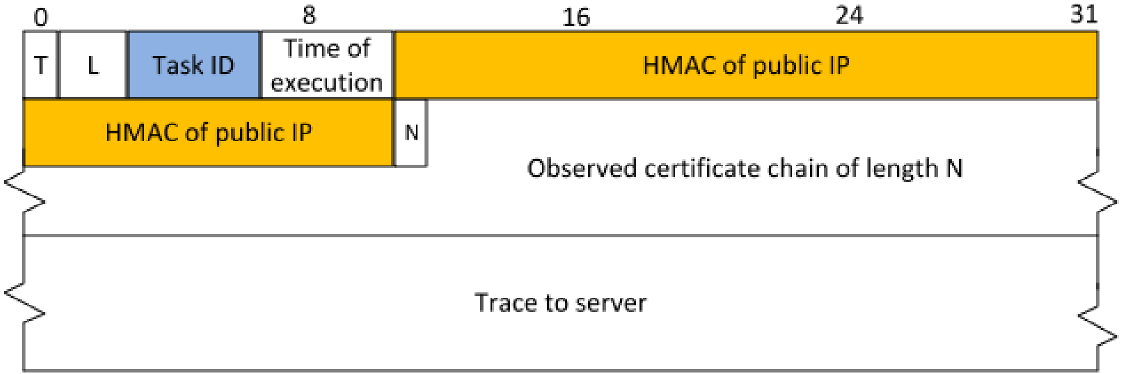

0 8 16 24 31

T | L | Task ID | Time of execution | HMAC of public IP

HMAC of public IP | N | Observed certificate chain of length N

Trace to server

T: Message Type; L: Message Length

Verification of certificate chains

- General Problem: It is unknown which certifcate should be observed
  - Client might or might not be behind a Mitm
  - Mitm might or might not attack a connection
  - Websites like Facebook use
    - Multiple certificates at the same time
    - Multiple Root CAs

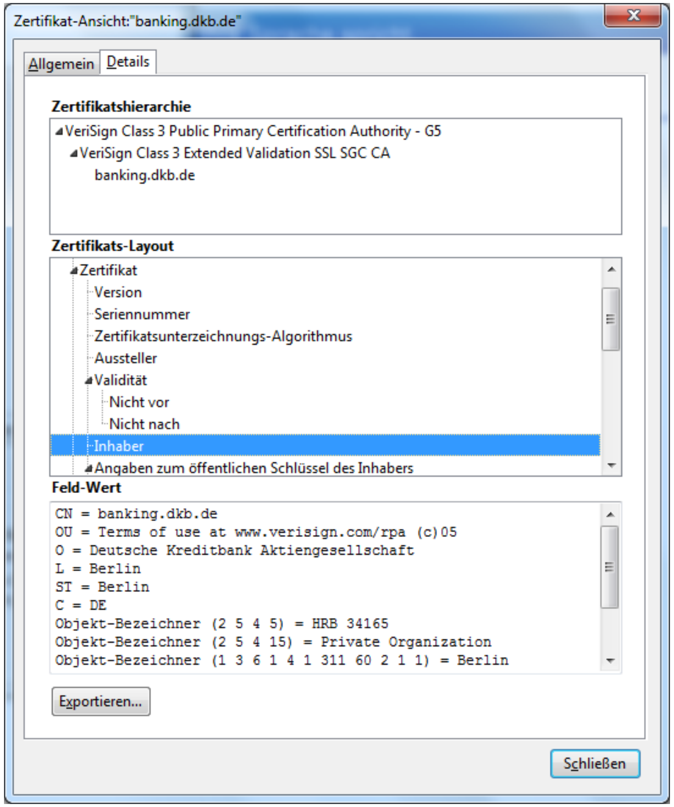

Zertifikatshierarchie

```
◢VeriSign Class 3 Public Primary Certification Authority - G5
  ◢VeriSign Class 3 Extended Validation SSL SGC CA
      banking.dkb.de
```

Zertifikats-Layout

```
◢Zertifikat
    Version
    Seriennummer
    Zertifikatsunterzeichnungs-Algorithmus
    Aussteller
  ◢Validität
      Nicht vor
      Nicht nach
    Inhaber
  ◢Angaben zum öffentlichen Schlüssel des Inhabers
```

Feld-Wert

```
CN = banking.dkb.de
OU = Terms of use at www.verisign.com/rpa (c)05
O = Deutsche Kreditbank Aktiengesellschaft
L = Berlin
ST = Berlin
C = DE
Objekt-Bezeichner (2 5 4 5) = HRB 34165
Objekt-Bezeichner (2 5 4 15) = Private Organization
Objekt-Bezeichner (1 3 6 1 4 1 311 60 2 1 1) = Berlin
```

- What can be done
  - Check if the sent chains are sane
  - Statistics: Identification of outliers
  - Manual Certificate chain inspection

Current level of implementation

- Hunter basic functionality: fully implemented
- Guard basic functionality: fully implemented
- Firefox-Plug-in GUI: fully implemented

- Dozens of little improvements: partially implemented
- Source code documentation: almost done
- Usage of Perspectives / Convergence: partially implemented

- Crossbear website: not yet created
- Evaluation of measured data: not yet done due to missing data

You Could Help!

Why you should use Crossbear

- Crossbear protects you against Mitm attacks
  - Setting up a Mitm is very easy (and attractive)!
  - Frequent travelers are likely to run into one of them (hotels, cafés, …)

- Crossbear contributes to a safer internet
  - Detection and location of Mitm
    - Warn users
    - Notify authorities
  - Possible discovery of new threats on X.509 PKI
  - Collection of data which will be publically available for security research

- Crossbear is a young project and needs users to improve

Thank you for your attention!

URL: pki.net.in.tum.de

Mail: crossbear@pki.net.in.tum.de

Twitter: @crossbearteam

Indication of Sources

- [1]: Performance Comparison of PHP and JSP as Server-Side Scripting Languages by Scott Trent et al.
- [2]: Fortinet FortiGate®: http://www.scribd.com/doc/49908929/31/Table-3-SSL-content-scanning-and-inspection-settings
- [3]: Packet Forensics 5-series : https://www.packetforensics.com/pfli5b.safe
- [4]: sslsniff: http://www.thoughtcrime.org/software/sslsniff/

Indication of image sources

- Cross-Bearing: http://de.wikipedia.org/w/index.php?title=Datei:Rueckwaertseinschneiden_a6.png&filetimestamp=20110602141415
- Firefox: http://de.wikipedia.org/w/index.php?title=Datei:FirefoxLogo3.5.png&filetimestamp=20090630200721
- IPv6: http://www.futurenews.at/wp-content/uploads/2011/02/ipv6.jpg
- Tomcat: http://tomcat.apache.org/images/tomcat.gif
- Perspectives: http://perspectives-project.org/about-us/
- Convergence: http://convergence.io/imgs/logo.png
- Hacker: http://denis-l.de/wp-content/uploads/hacker.gif
- Good/Average/Excellent: http://ipwatchdog.com/images/excellent-good-average.jpg
- All images that are not listed explicitely are created by myself using non-copyrighted material.

Distributed Position Estimation

The Cross-Bearing Method

- **Output**: Position of a ship
- **Given**: Observers along the coast, with
  - » Known Position
  - » Direction towards the target
- **Method**: Intersect the observations

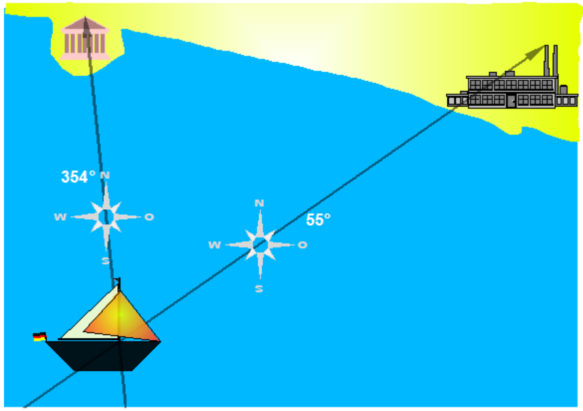

The CrossBear-System

- **Output**: Position of a Man-in-the-middle
- **Given**: Observers around the world, with
  - » Known IP-Address
  - » Route to an attacked TLS-server
  - » Knowledge if that route is poisoned
- **Method**: Compare & Intersect the routes