The SSL Landscape

Ralph Holz, Thomas Riedmaier,
Lothar Braun, Nils Kammenhuber

Network Architectures and Services
Technische Universität München

Hagenberg, 15 Mar 2012

Significance of 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:
  $Cert(X) = Sig_{CA}(id_X, pubkey_X)$

Browser Panic (but AT is pretty good)

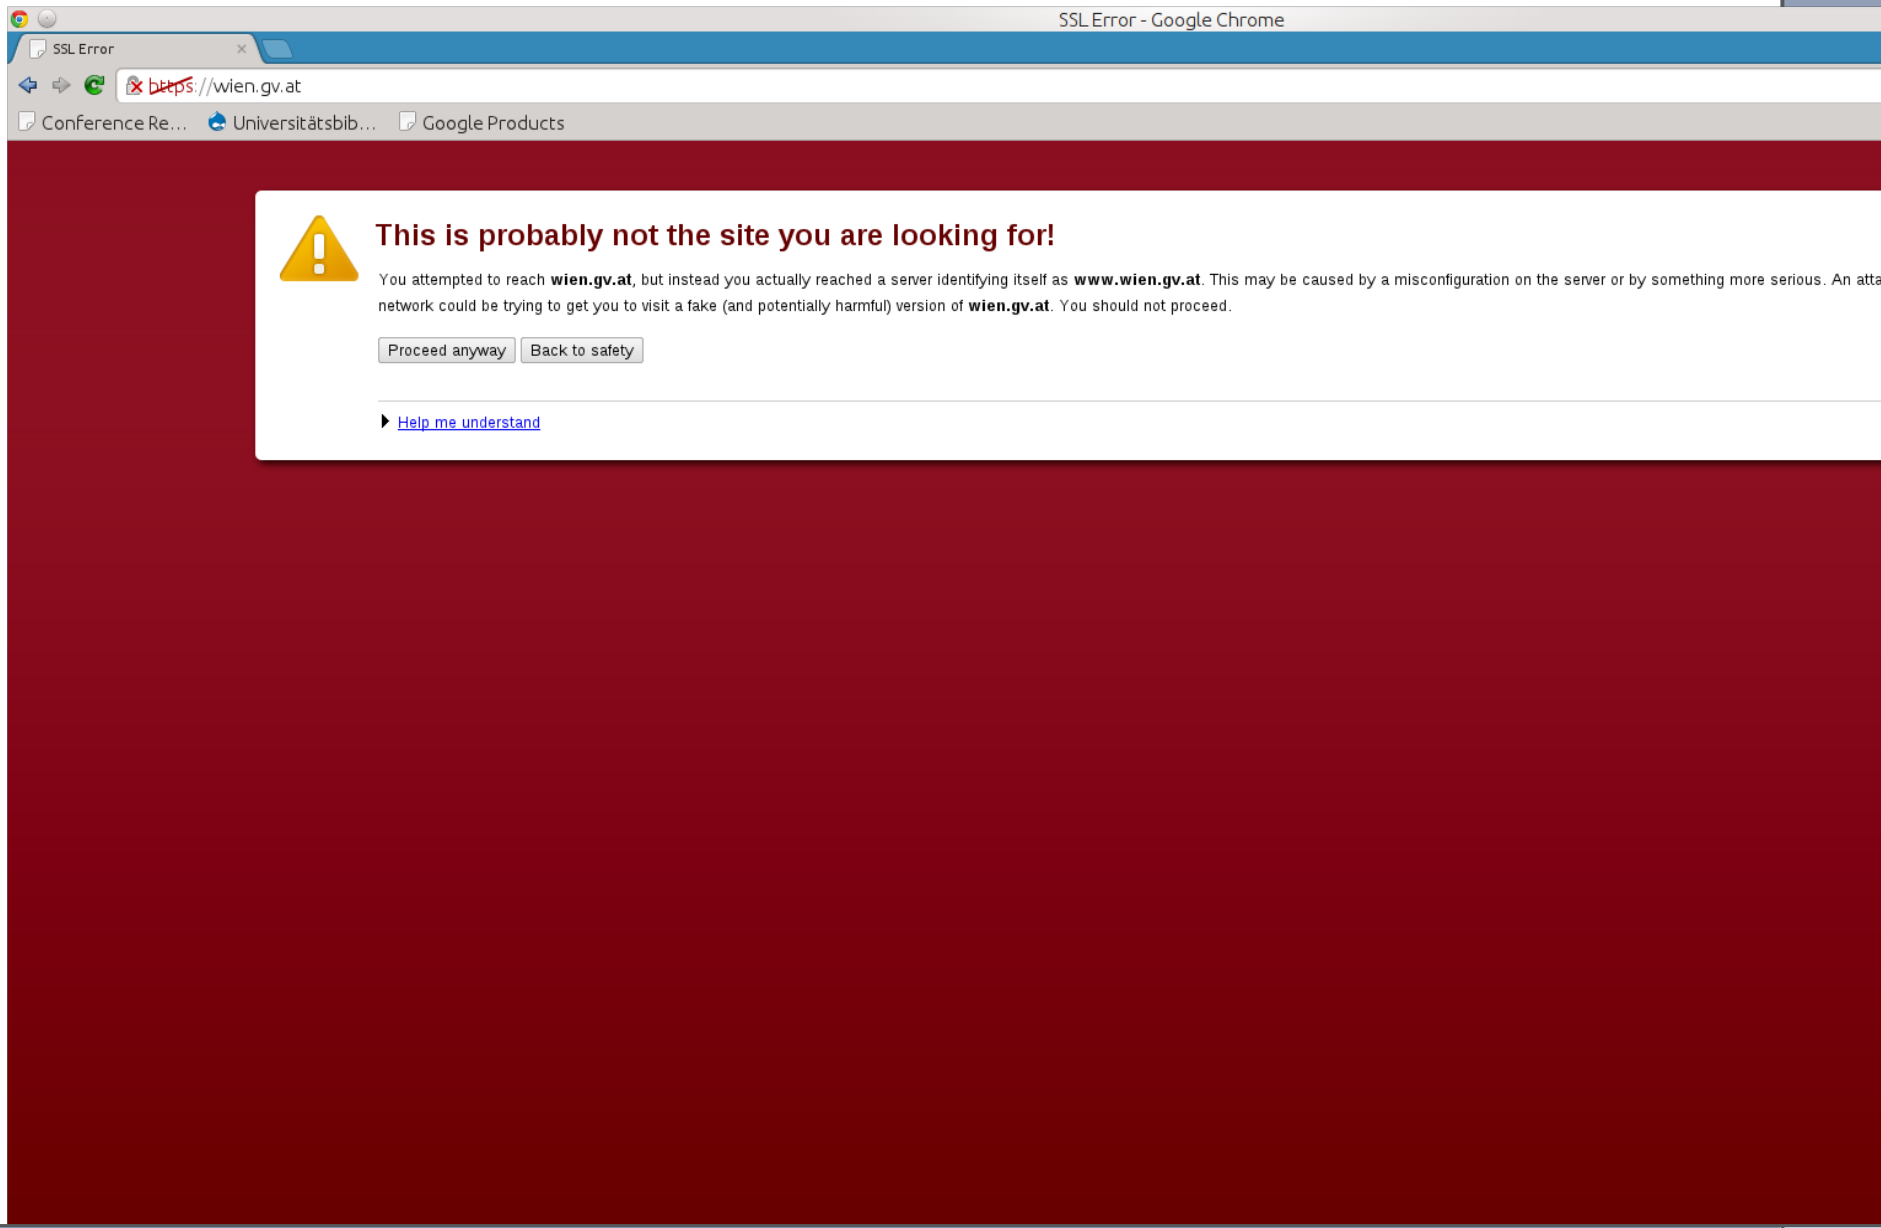

Basic Idea of X.509 PKI: Hierarchy

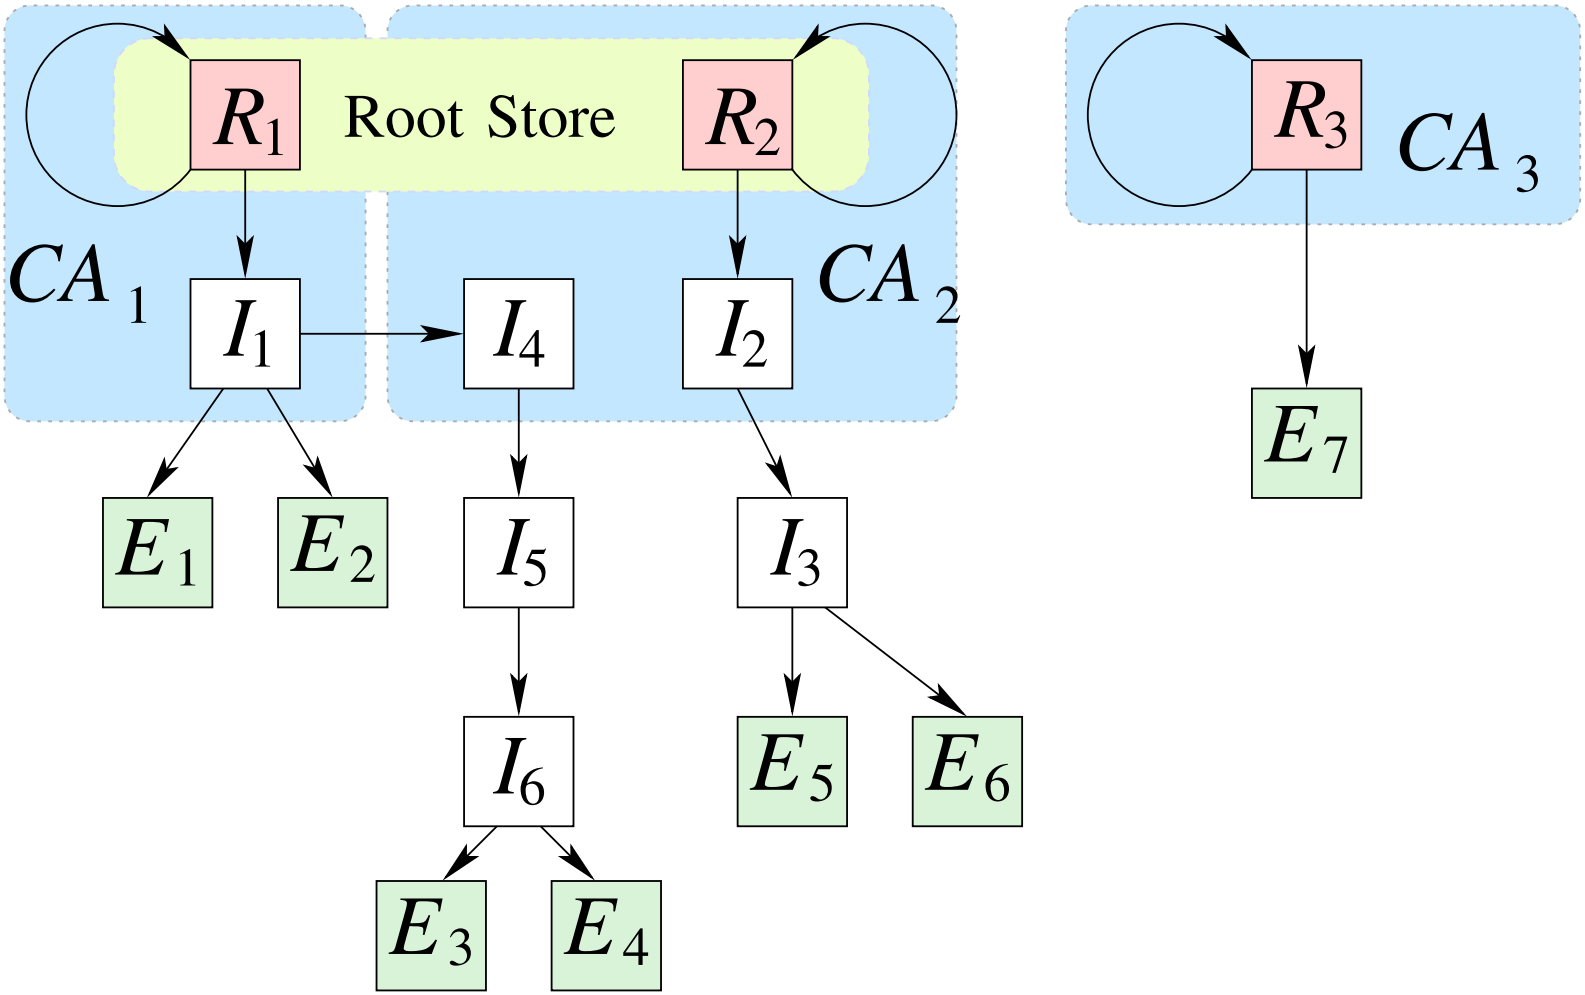

Basic Idea of X.509

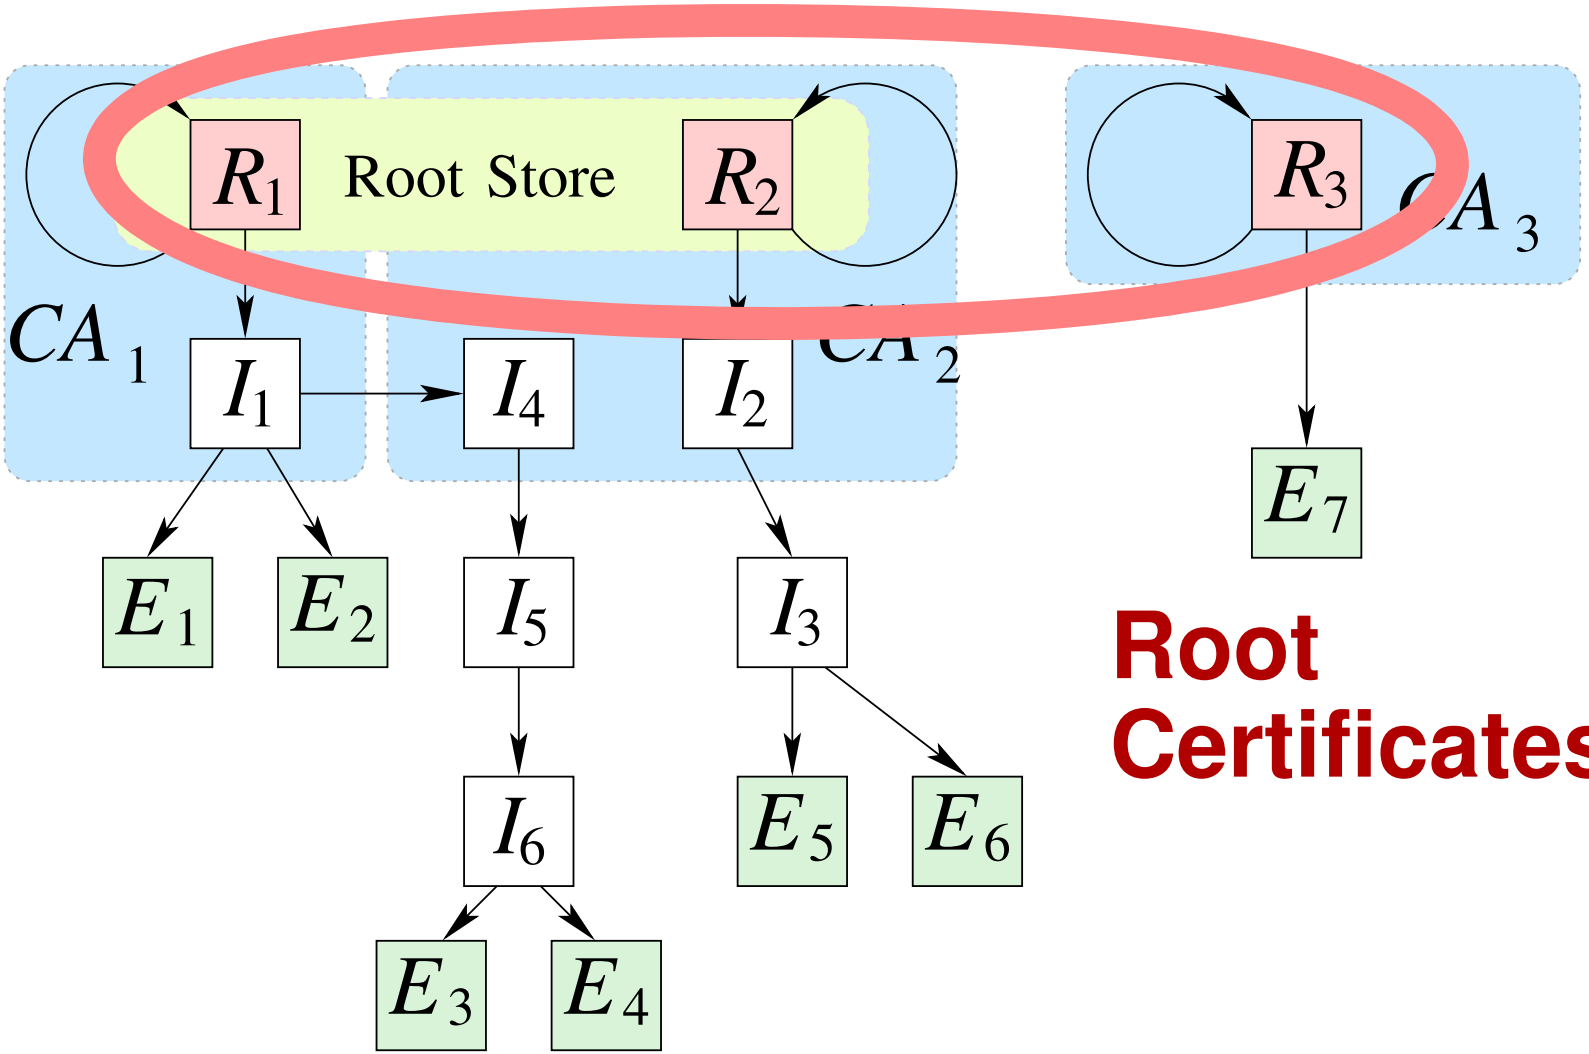

Root Certificates

Basic Idea of X.509

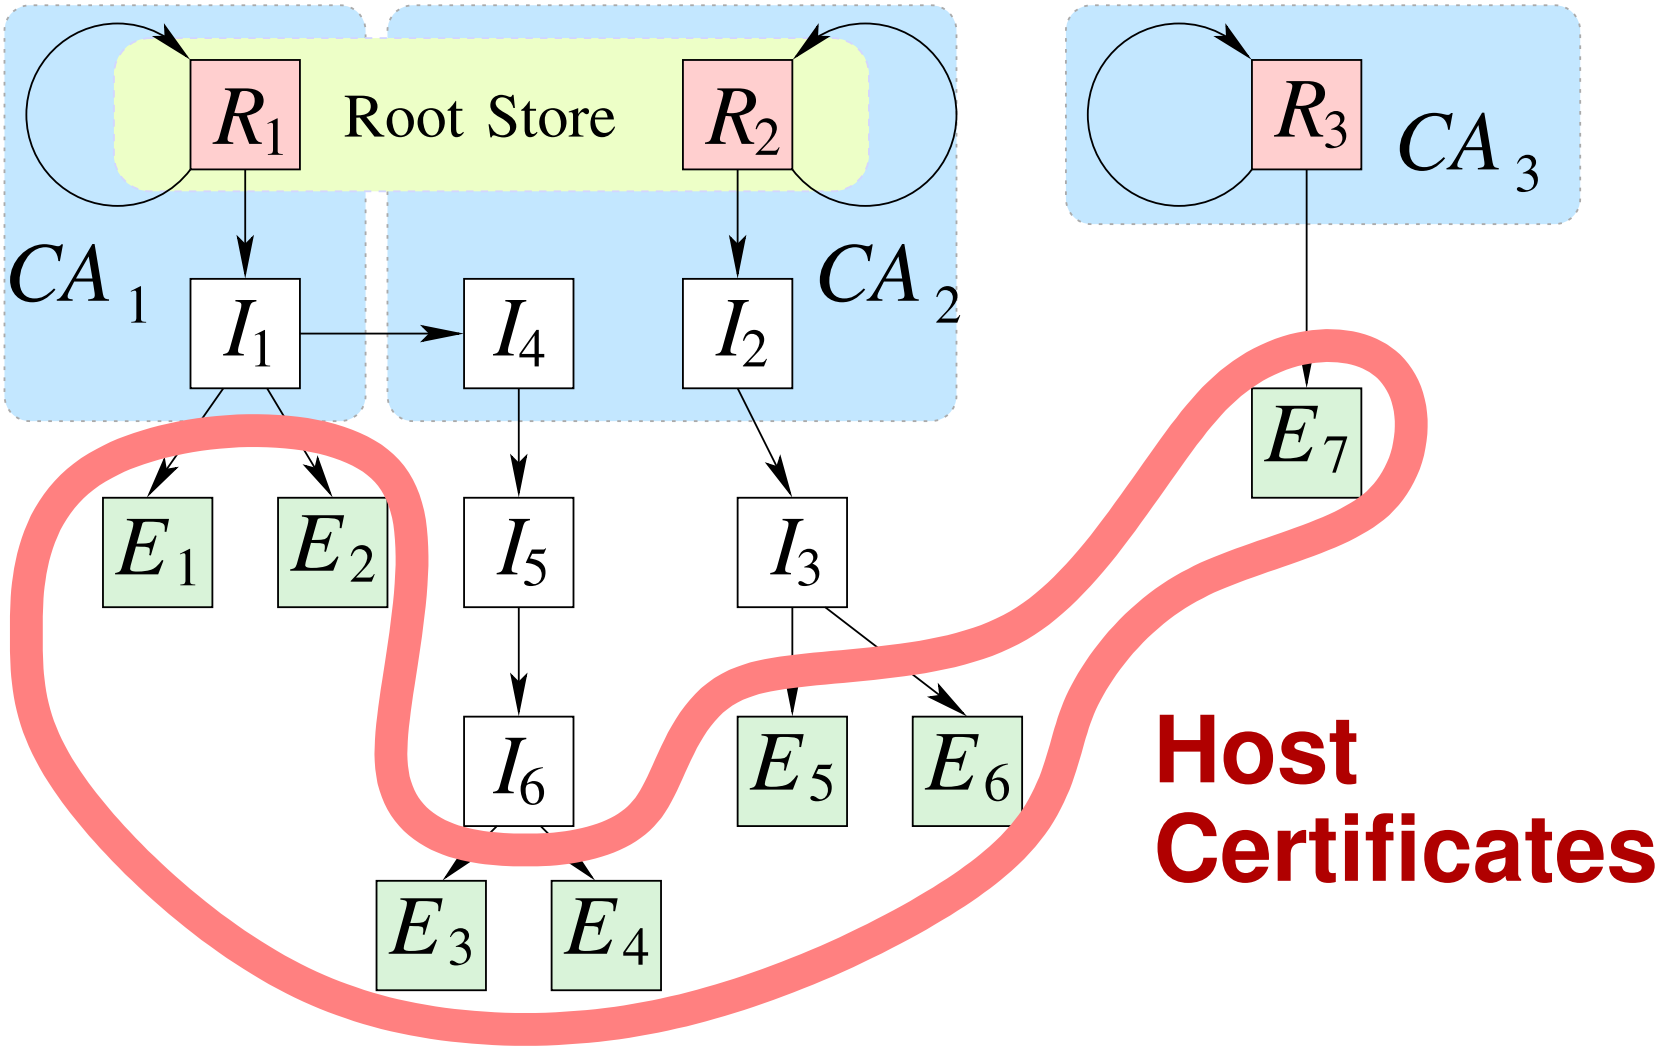

Basic Idea of X.509

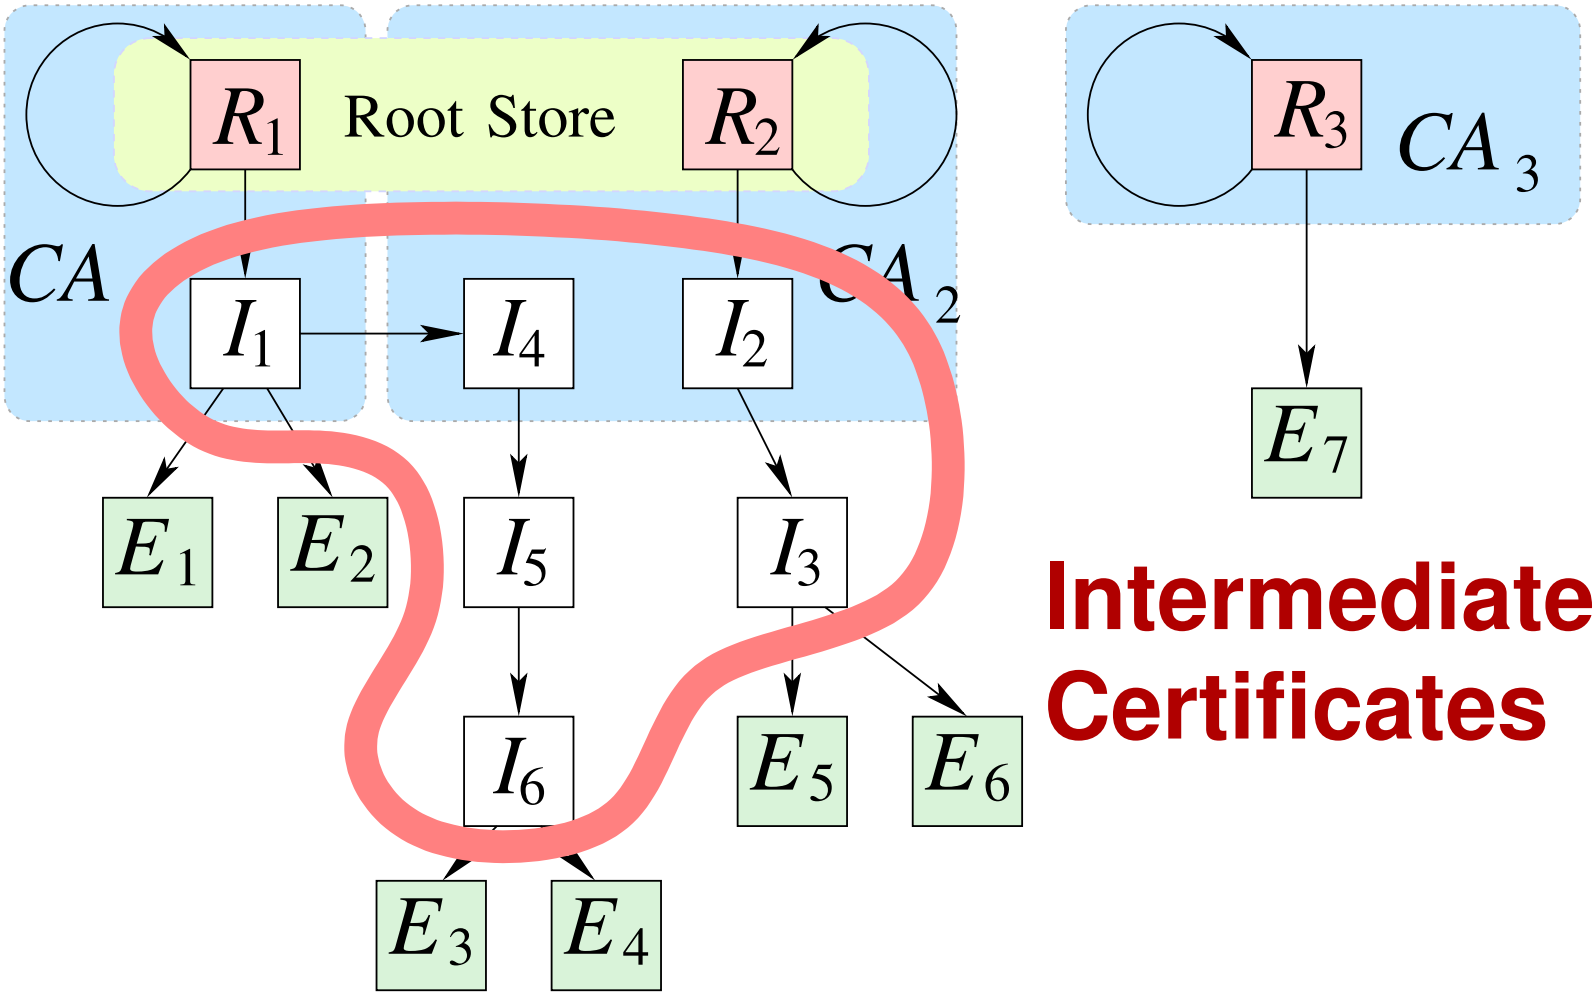

Intermediate Certificates

Basic Idea of X.509

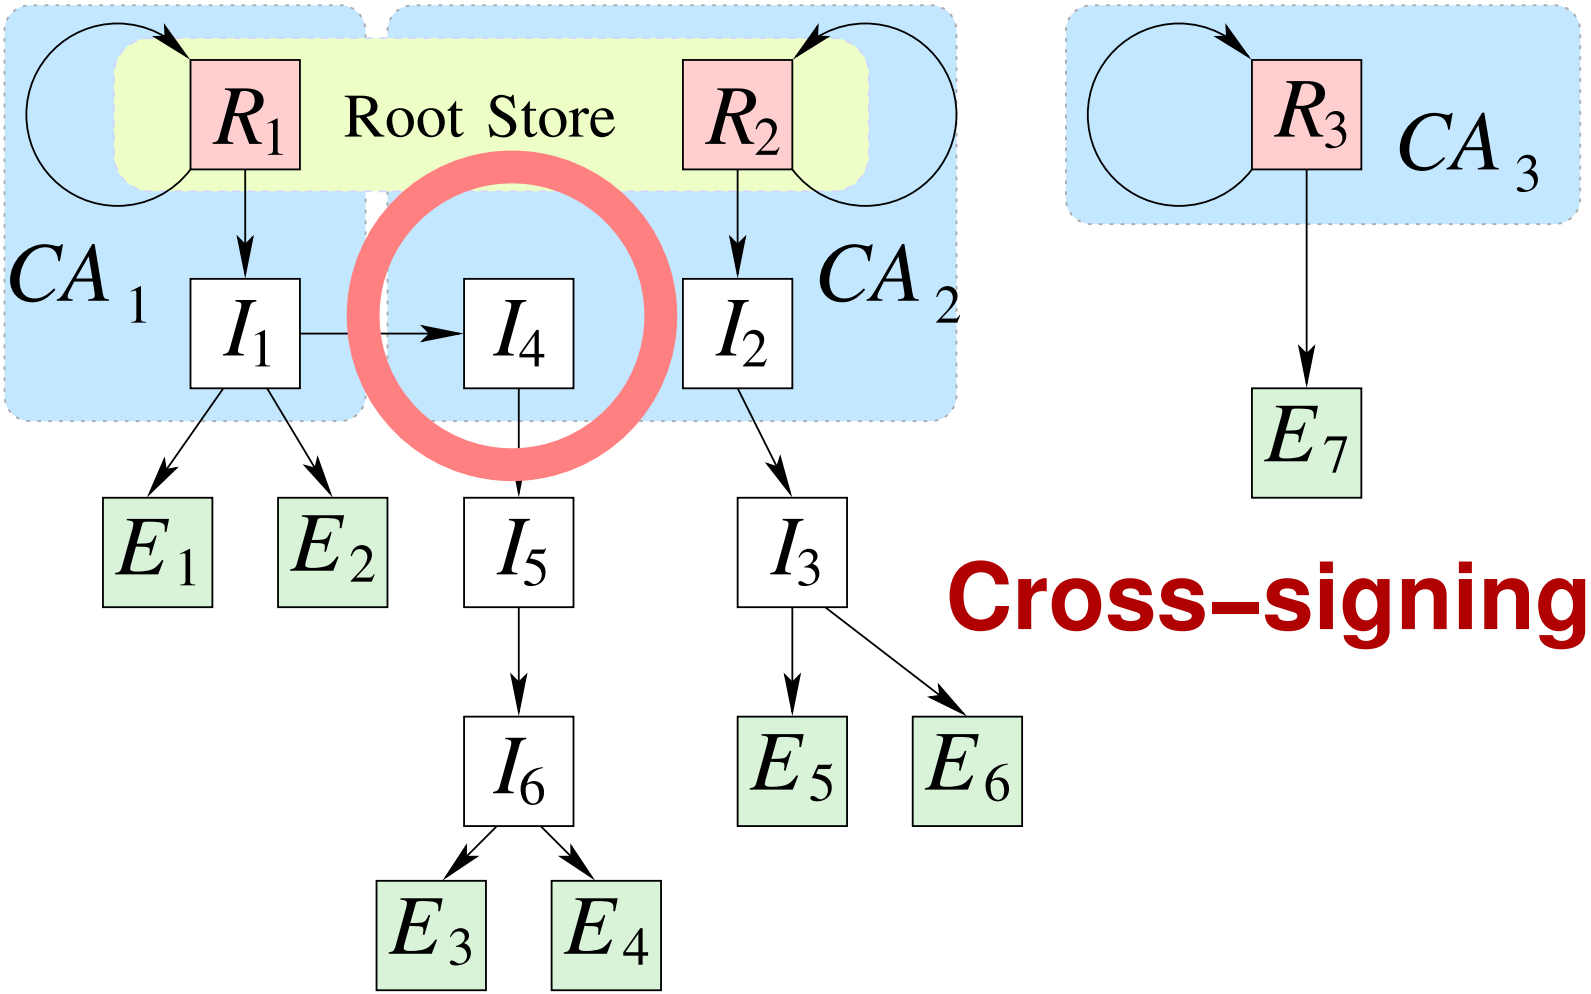

Cross–signing

Basic Idea of X.509

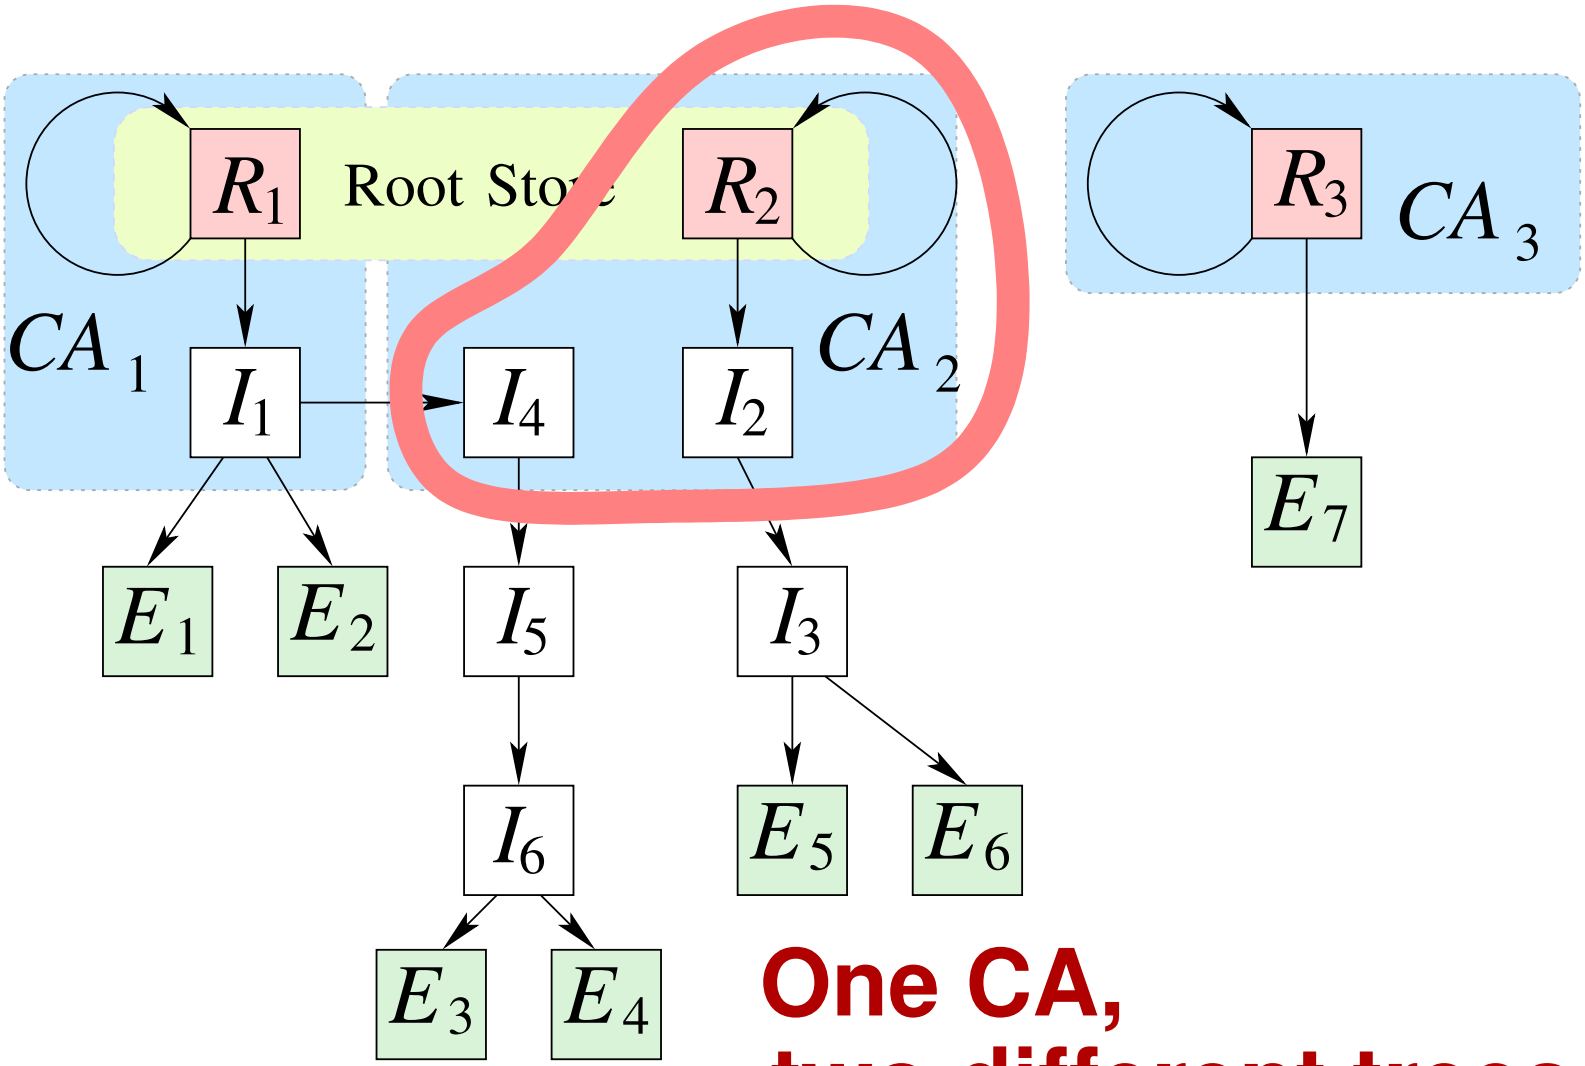

One CA,
two different trees

Basic Idea of X.509

CAs in Root Store

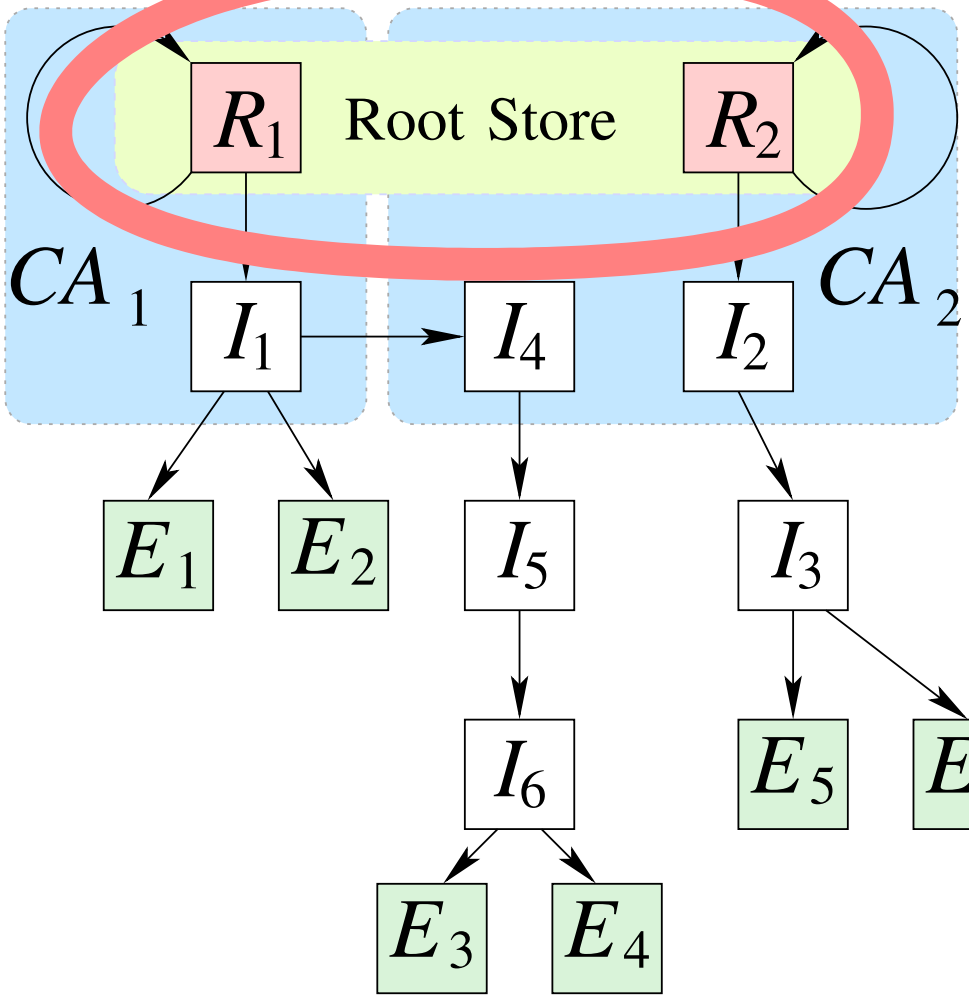

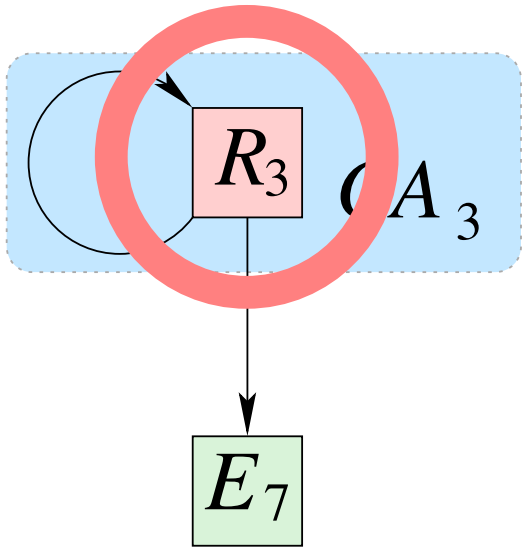

CA not in Root Store

Basic Idea of X.509

Root certificate not in Root Store

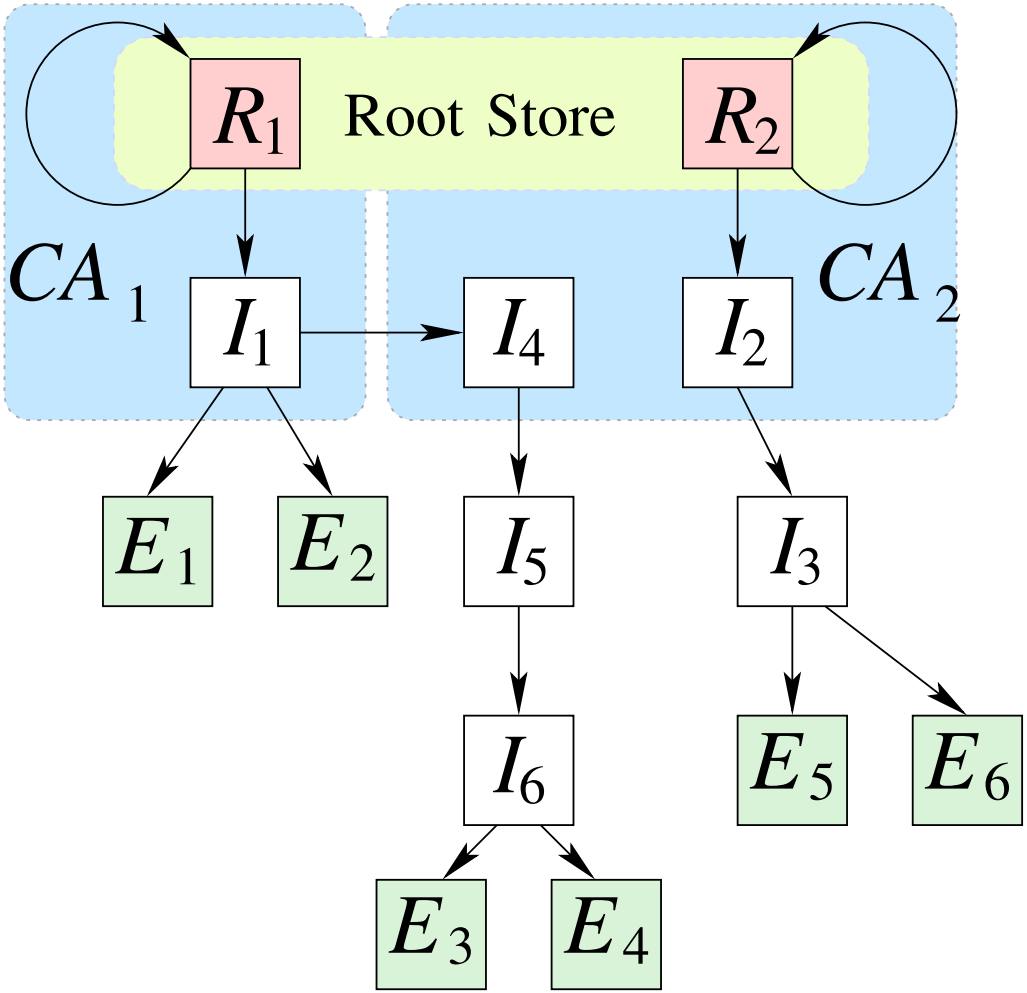

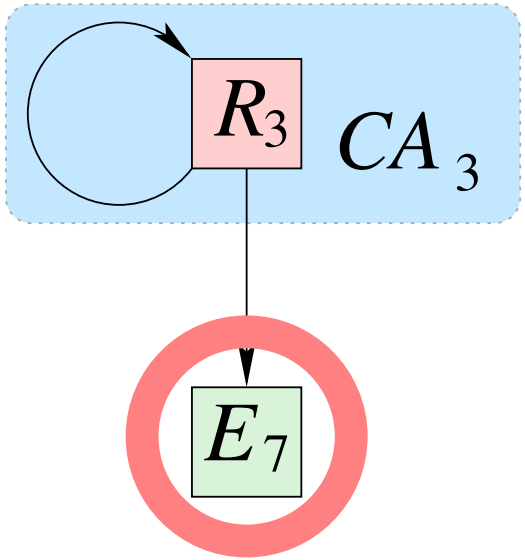

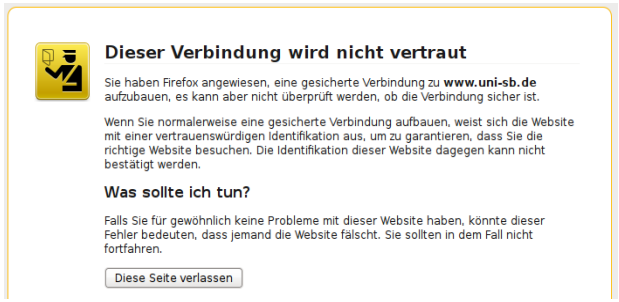

An X.509 Certificate

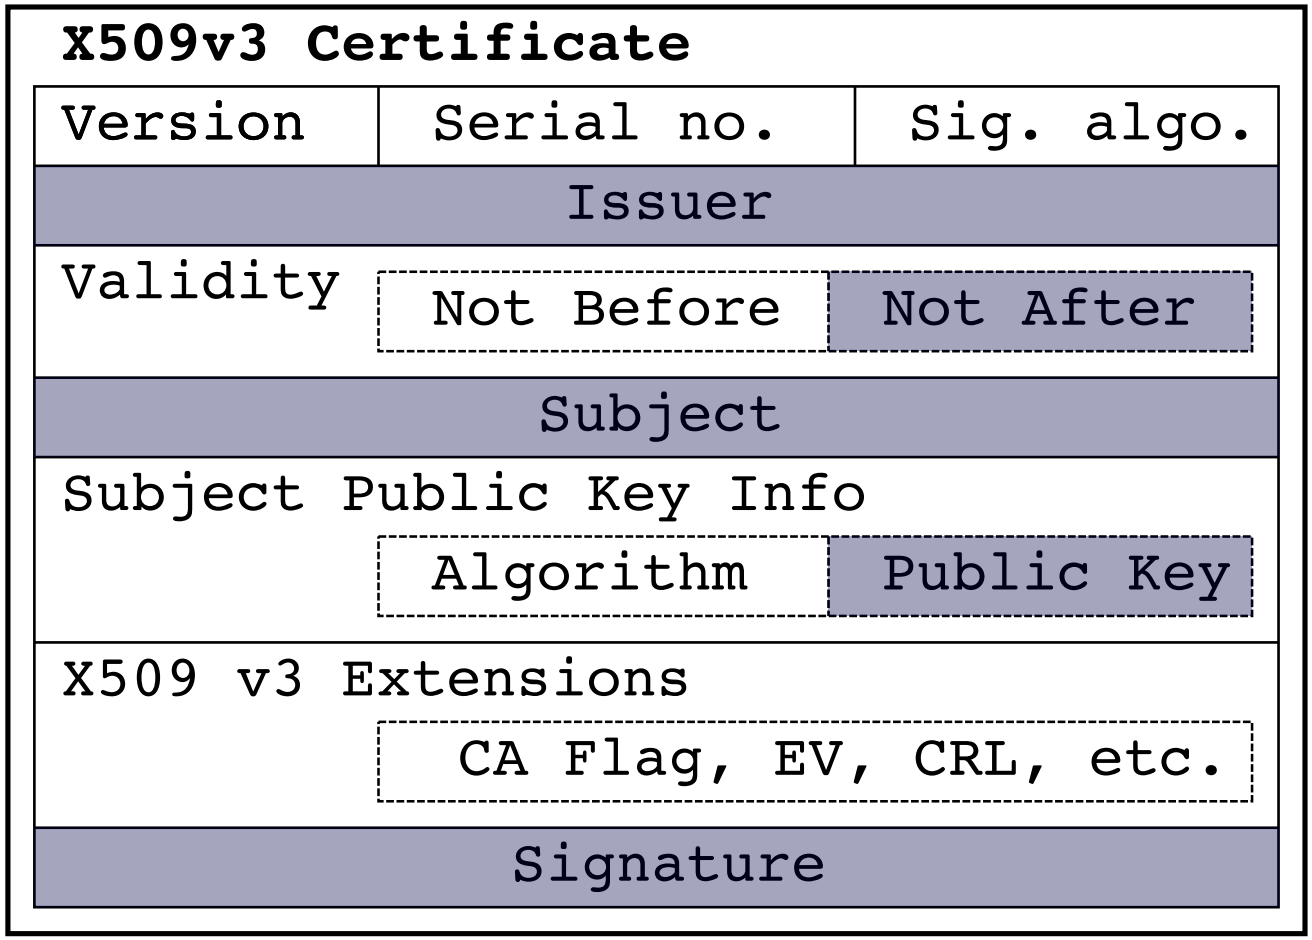

Root Stores Contain CA Certificates

CAs in Root Store

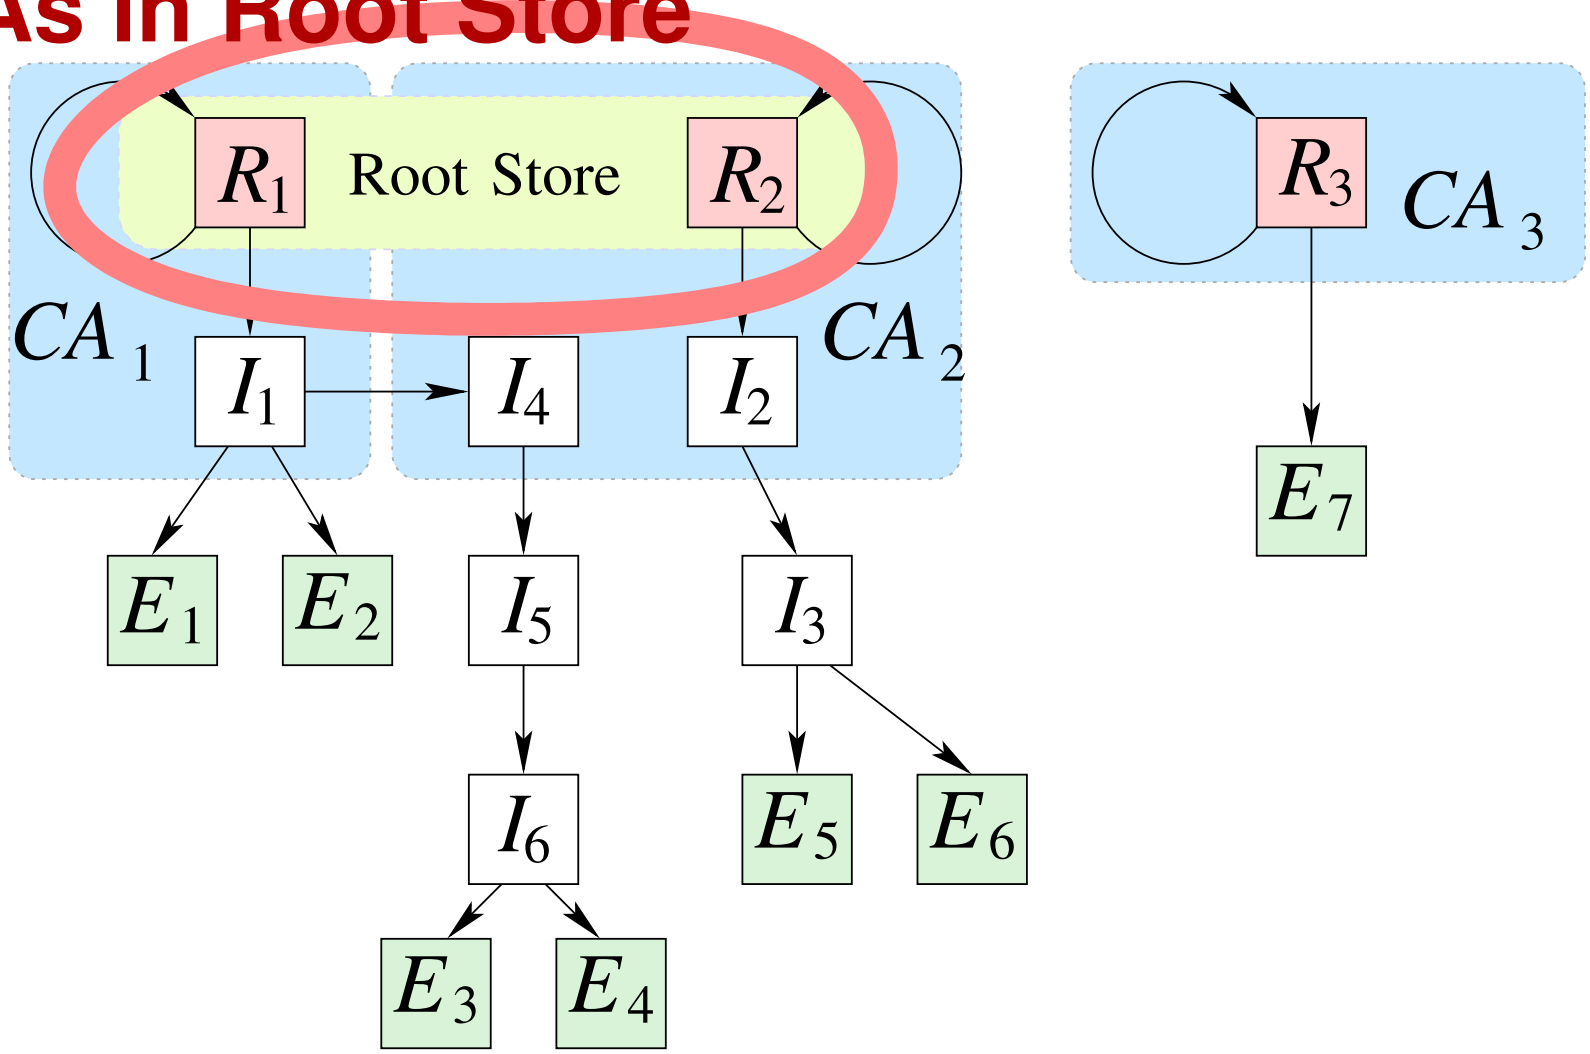

Browser (Client) Root Stores

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

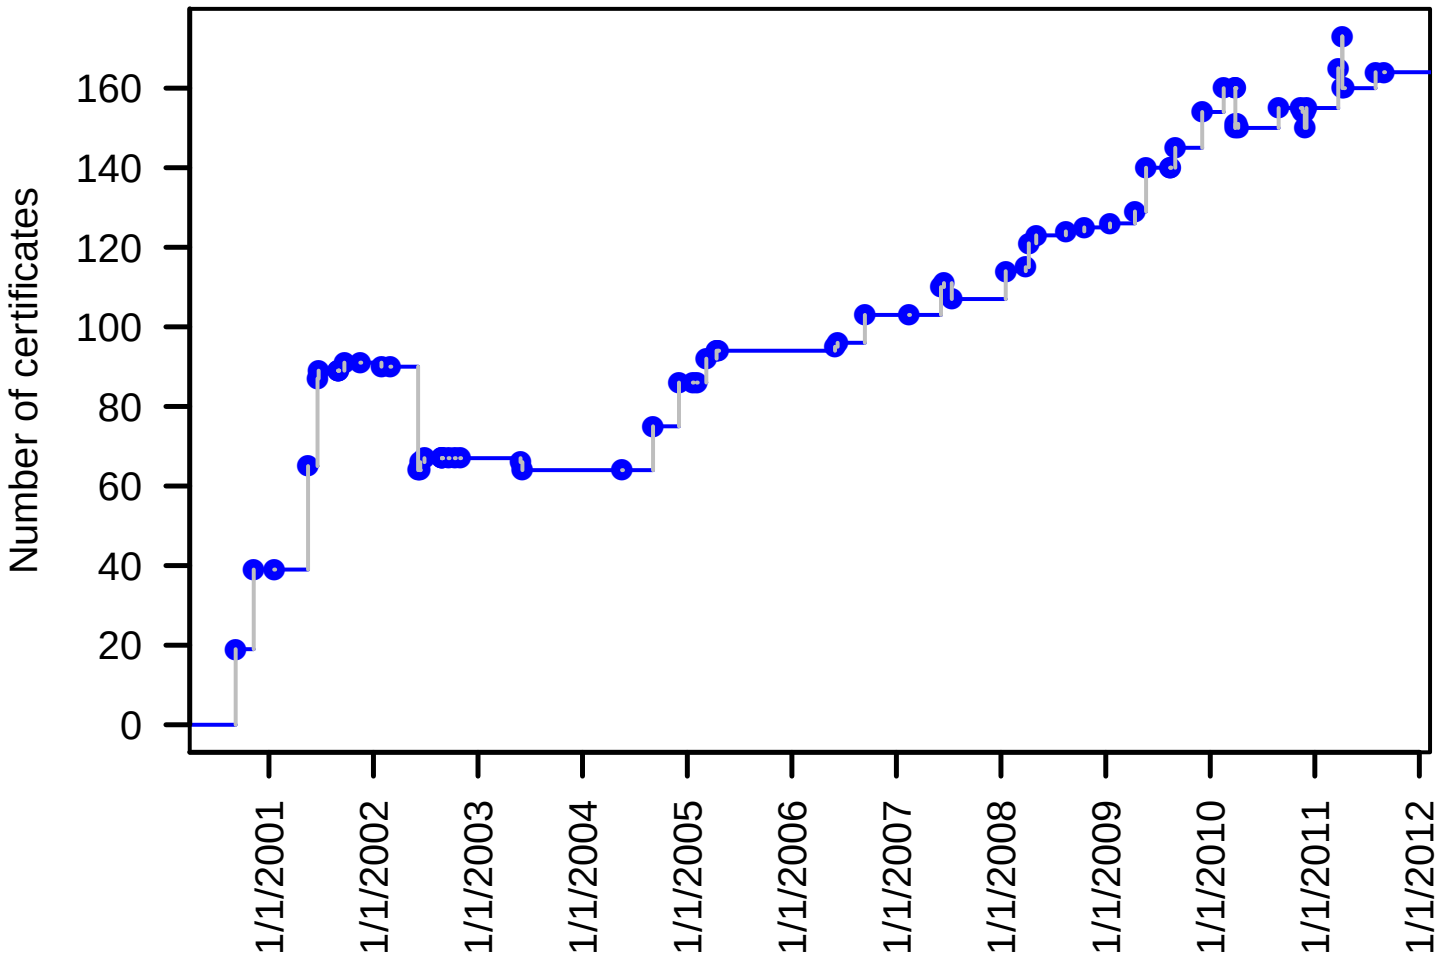

Certificate Issuance

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
  - The Netherlands PKI was operated by DigiNotar...
  - 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
  - The Netherlands PKI was operated by DigiNotar...
  - 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
  - The Netherlands PKI was operated by DigiNotar...
  - For the first time, a Root CA is removed from a browser for being compromised

And the latest entrant to the contest...

Trustwave issued MitM certs for companies

- Goal: transparent SSL proxying
  - Defence against industry espionage
  - Defeats whole end-to-end authentication
  - OK for companies? What about nation states? Iran? China?
- Early 2012: Mozilla asks all CAs to report and revoke MitM certs
  - Quite a few CAs admit to having issued them, but cite contractual obligation
  - Are you happy to trust those contracts?

And the latest entrant to the contest...

Trustwave issued MitM certs for companies

- Goal: transparent SSL proxying
  - Defence against industry espionage
  - Defeats whole end-to-end authentication
  - OK for companies? What about nation states? Iran? China?
- Early 2012: Mozilla asks all CAs to report and revoke MitM certs
  - Quite a few CAs admit to having issued them, but cite contractual obligation
  - Are you happy to trust those contracts?

And the latest entrant to the contest...

Trustwave issued MitM certs for companies

- Goal: transparent SSL proxying
  - Defence against industry espionage
  - Defeats whole end-to-end authentication
  - OK for companies? What about nation states? Iran? China?
- Early 2012: Mozilla asks all CAs to report and revoke MitM certs
  - Quite a few CAs admit to having issued them, but cite contractual obligation
  - Are you happy to trust those contracts?

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.

Errors in TLS Connection Setup

Scans from Germany, Nov 2009 and Apr 2011

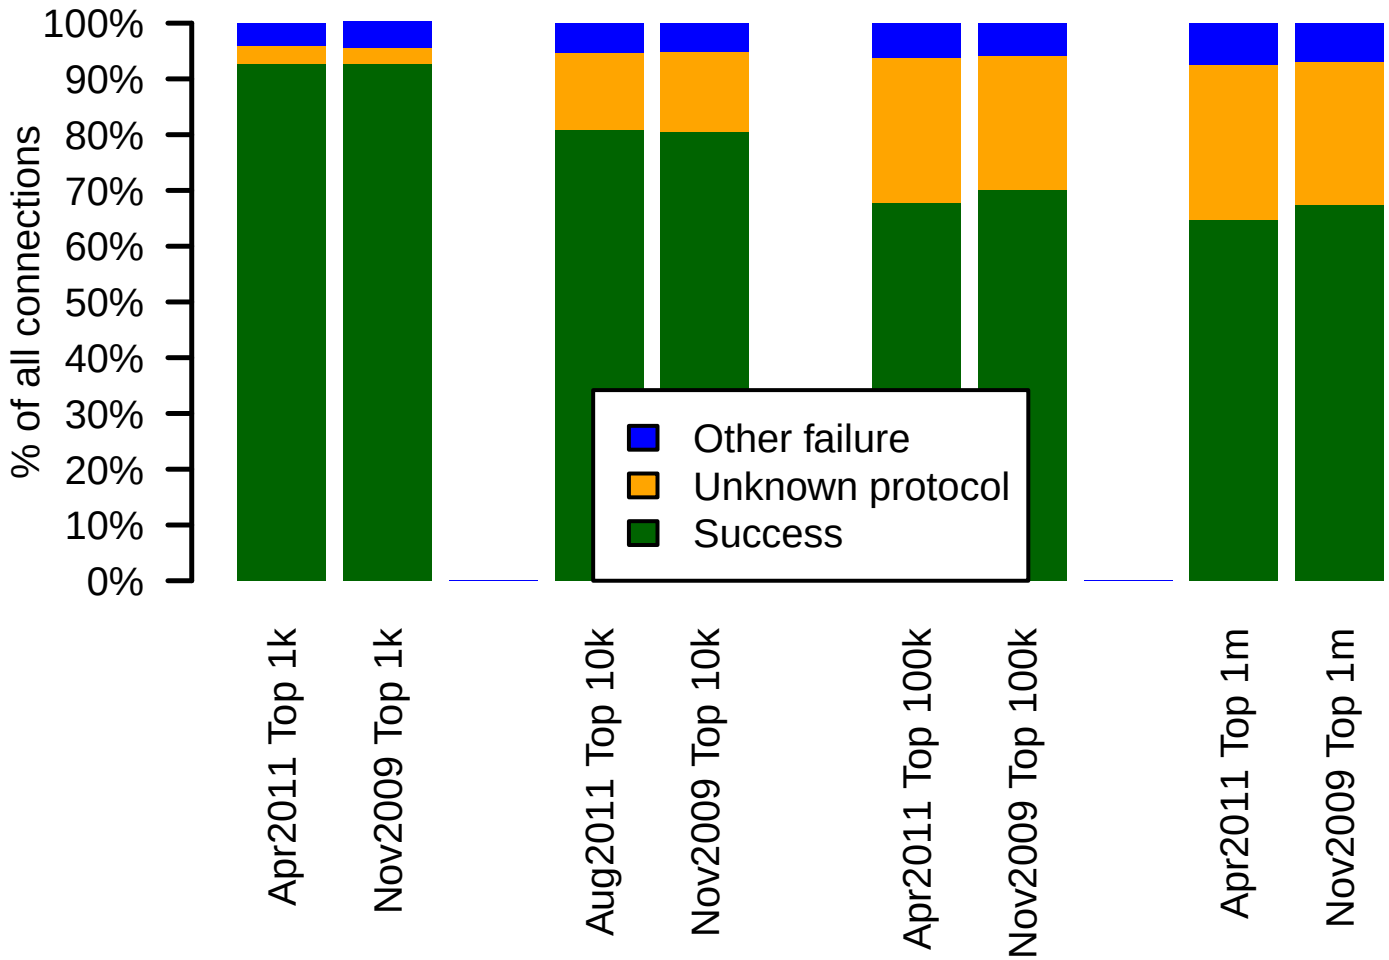

Validity of End-Hosts Certificates

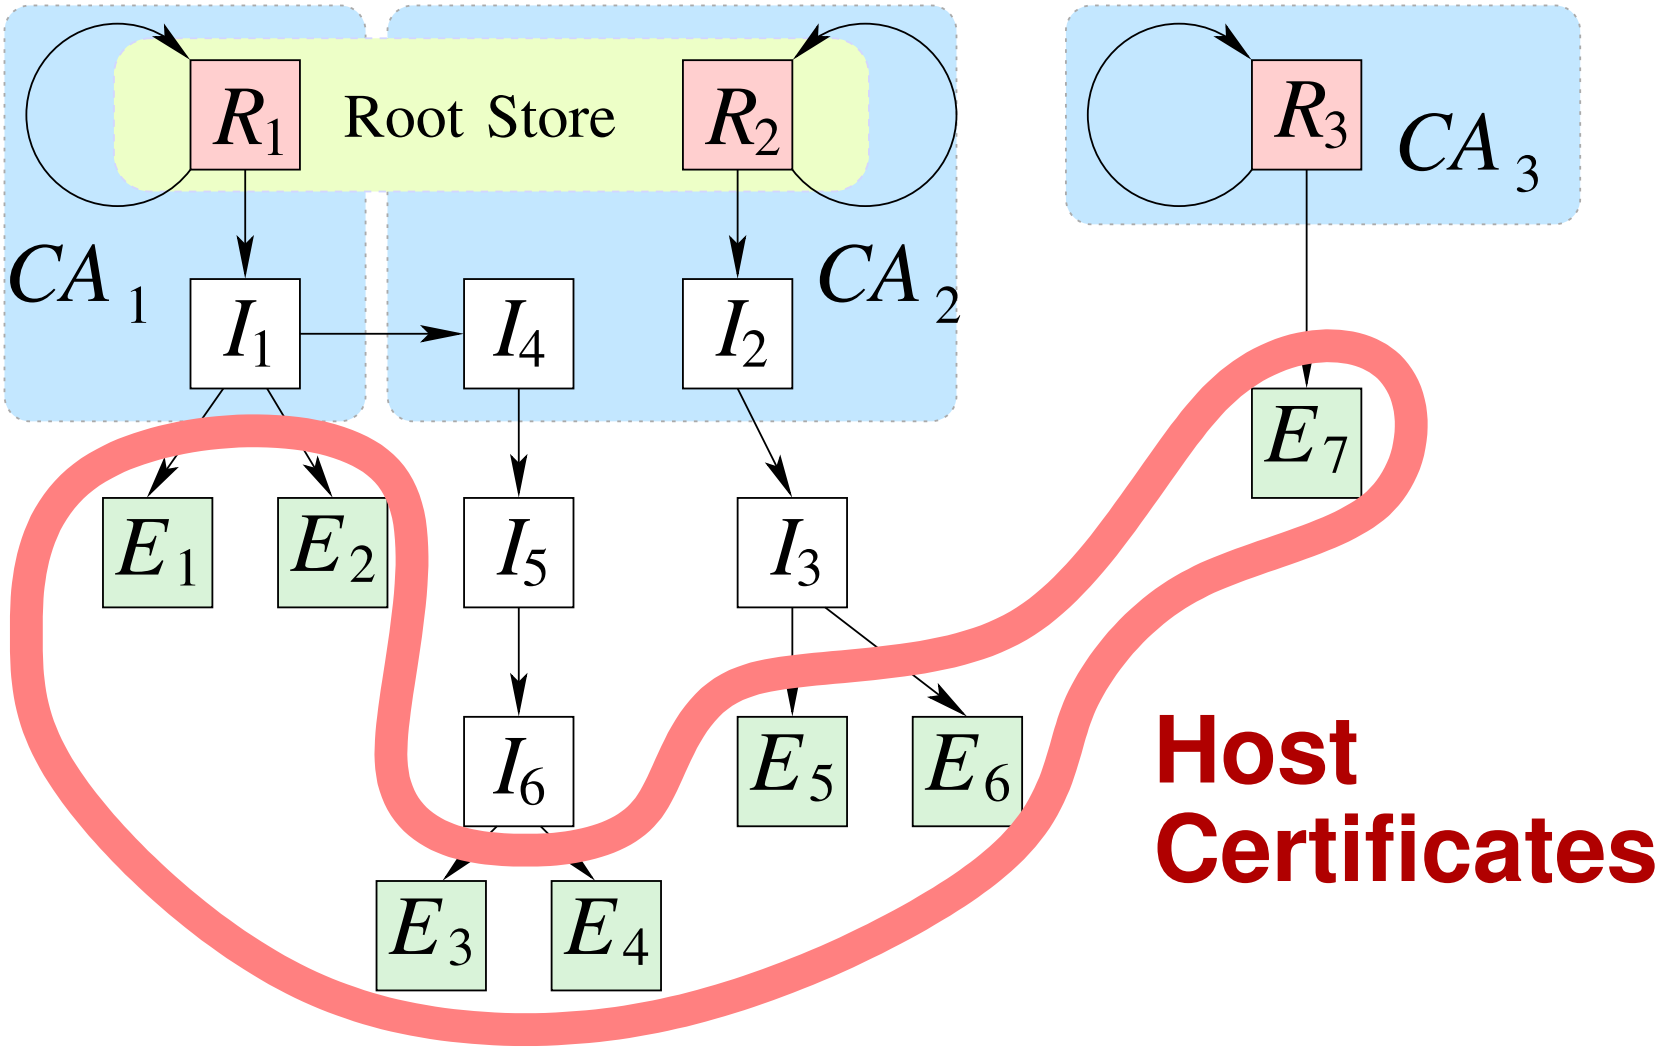

Validation of Certificate Chains

Just check chains, not host names

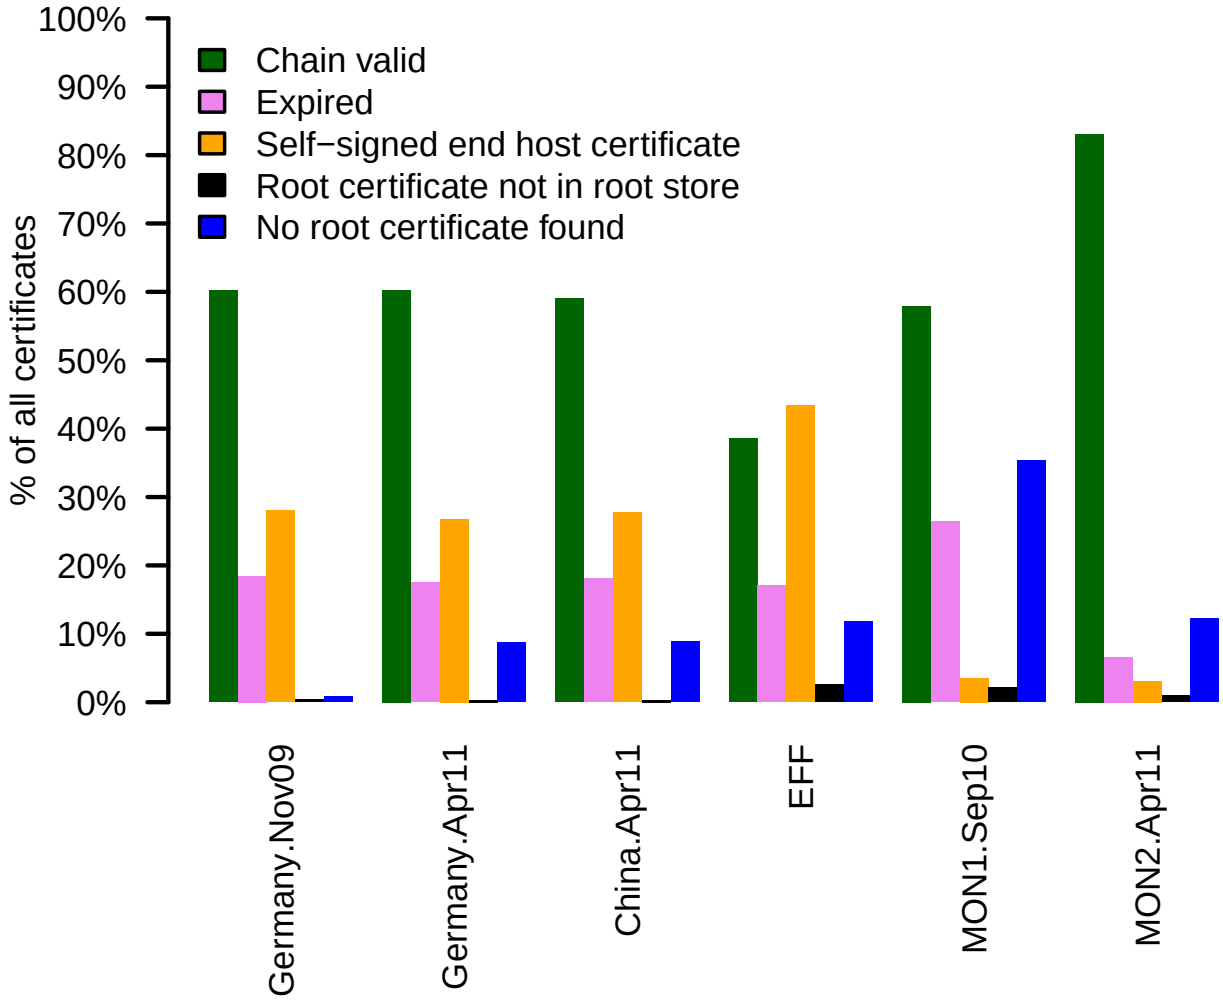

Correct Domain Name in Certificate

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

Unusual Host Names

CN=plesk or similar

- Found in 7.3% of certificates
- Verified: Plesk/Parallels panels

CN=localhost

- 4.7% of certificates
- Very common: redirection to HTTP after HTTPs

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

Top 3 most frequent CNs account for $>$ 50%

- `plesk` or similar in 27.3%
- `localhost` or similar in 25.4% – standard installations?

Certificate Occurrences

Many certificates valid for more than one domain

- Domains served by same IP
- Some certificates issued for dozens of domains
- Certificate reuse on multiple machines increases options for attacker

Often found on hosters

- E. g. `*.blogger.com`, `*.wordpress.com`

Certificate Occurrences

How often does a certificate occur on X hosts?

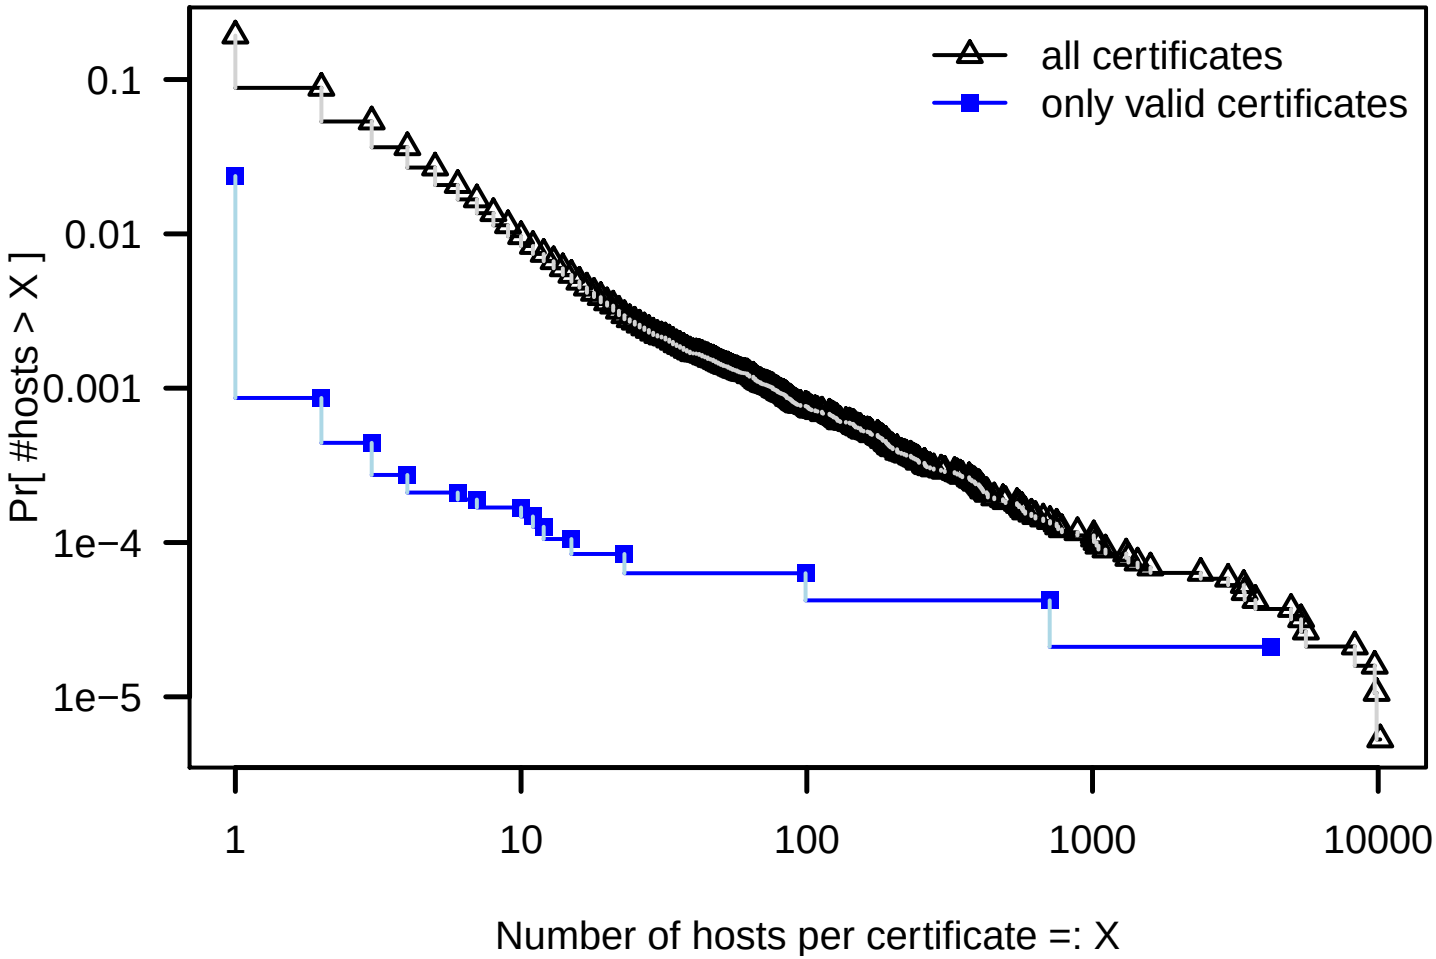

Certificate Chains

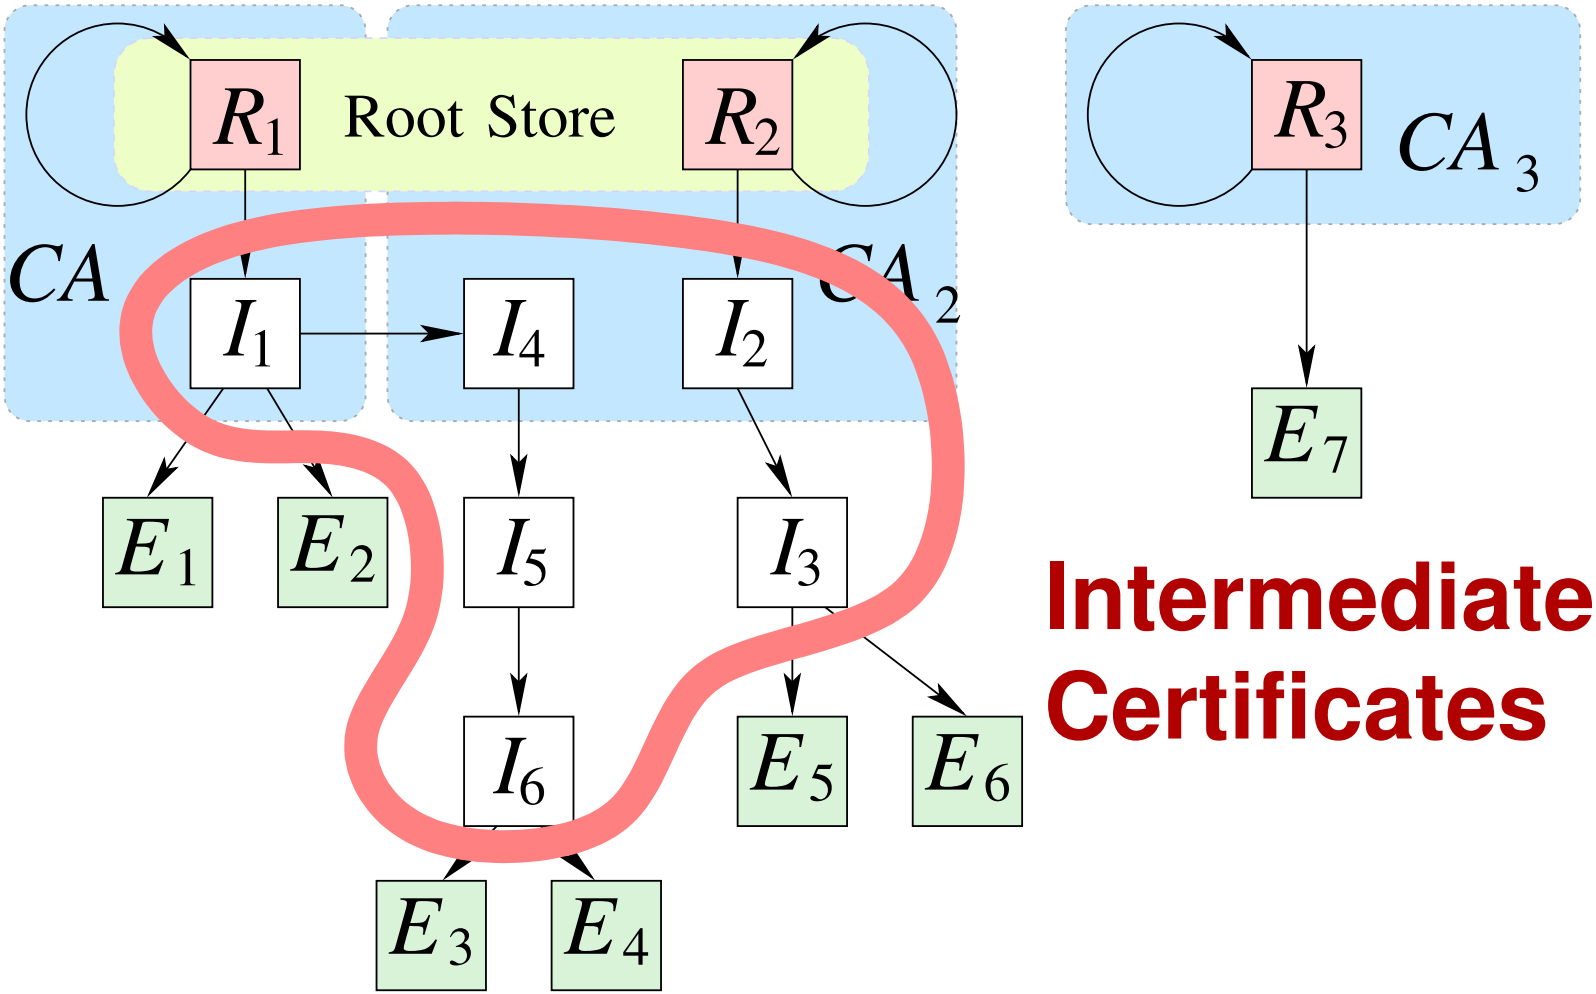

Intermediate Certificates

Certificate Chain Lengths

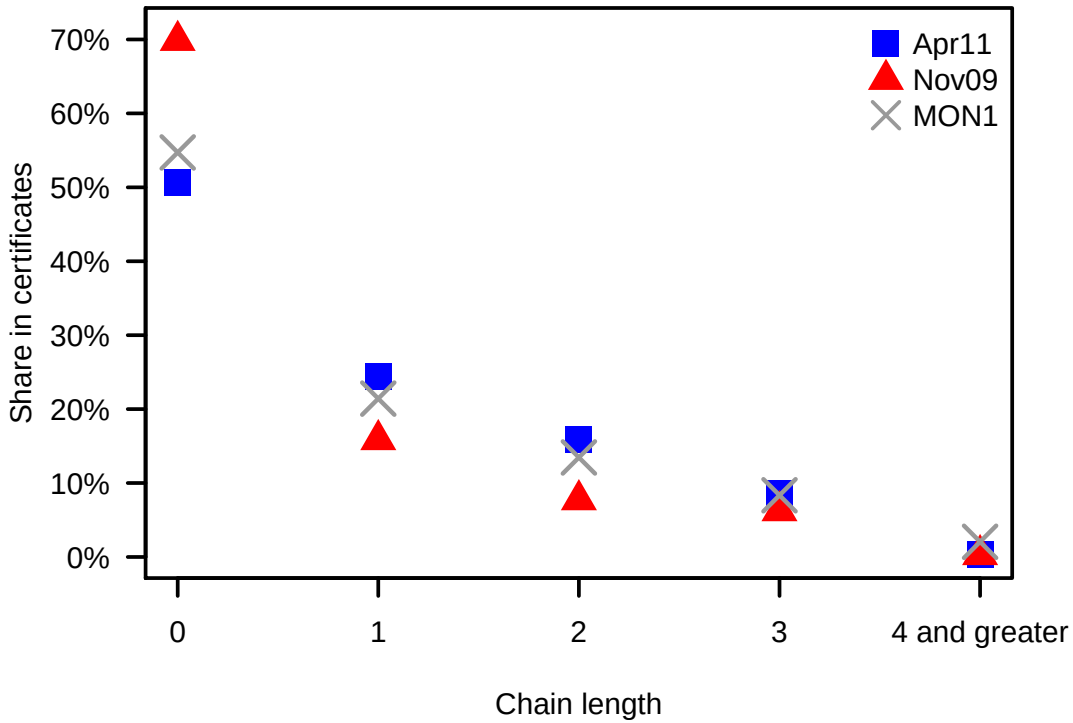

Finding more positive than negative:

- Trend to use intermediate certificates more often
- Allows to keep Root Certificates offline
- But chains still reasonably short

Validity Periods

CDF of validity periods, scans and monitoring

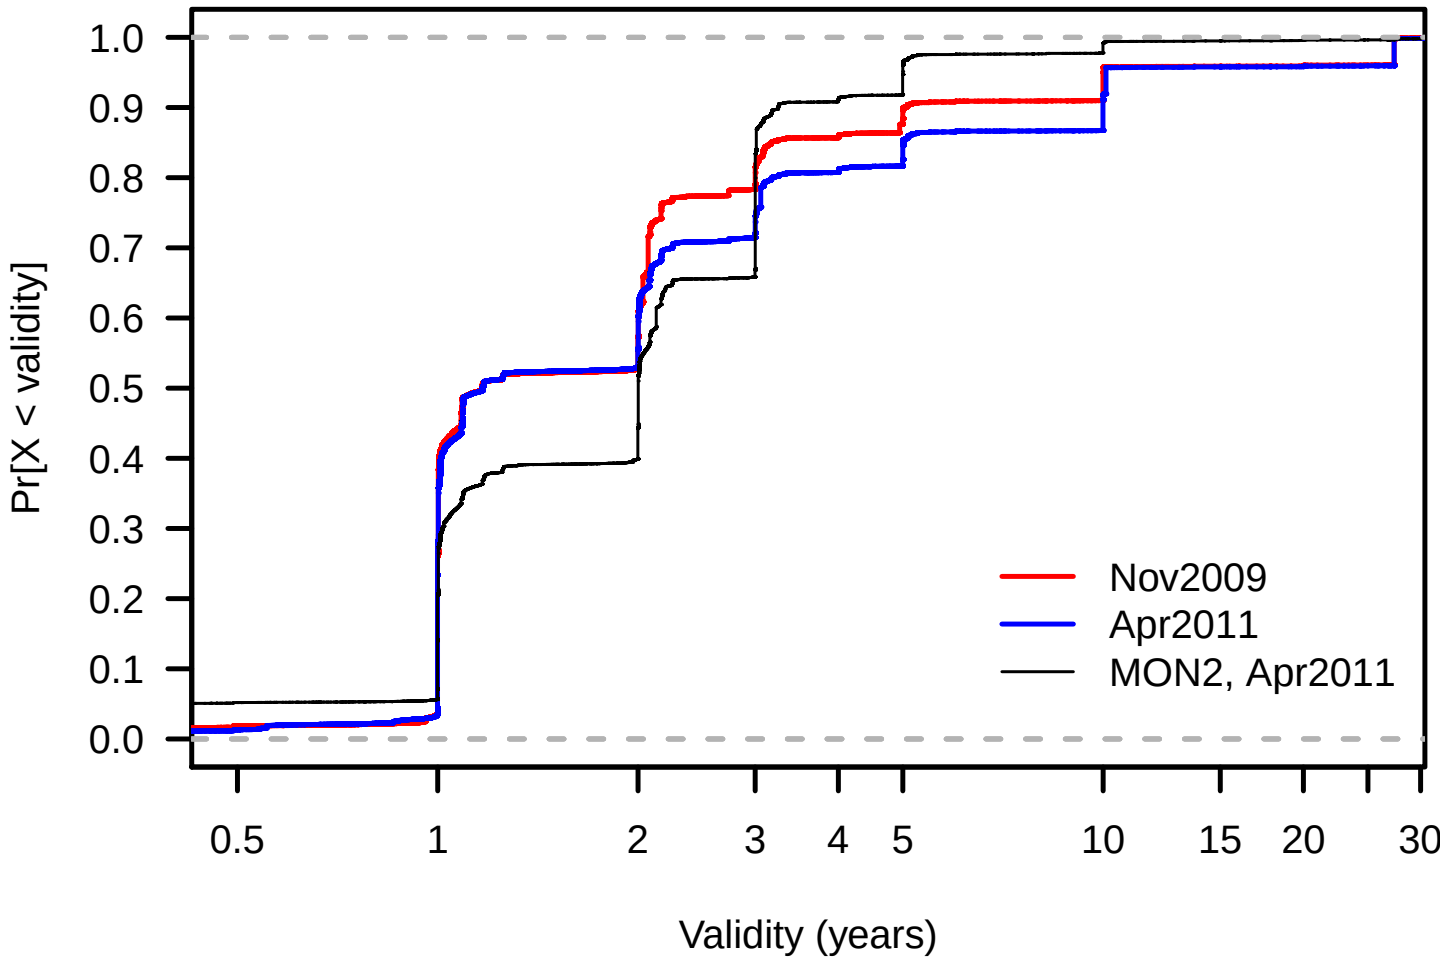

Public Key Properties

Key types

- RSA: 99.98% (rest is DSA)
- About 50% have length 1,024 bit
- About 45% have length 2,048 bit
- Clear trend from 1,024 to 2,048 bit (2009–2011)

Weird encounters

- 1,504 distinct certificates that share another certificate's key
- Could be traced to a handful of hosting companies

Signature Algorithms

MD5 is being phased out

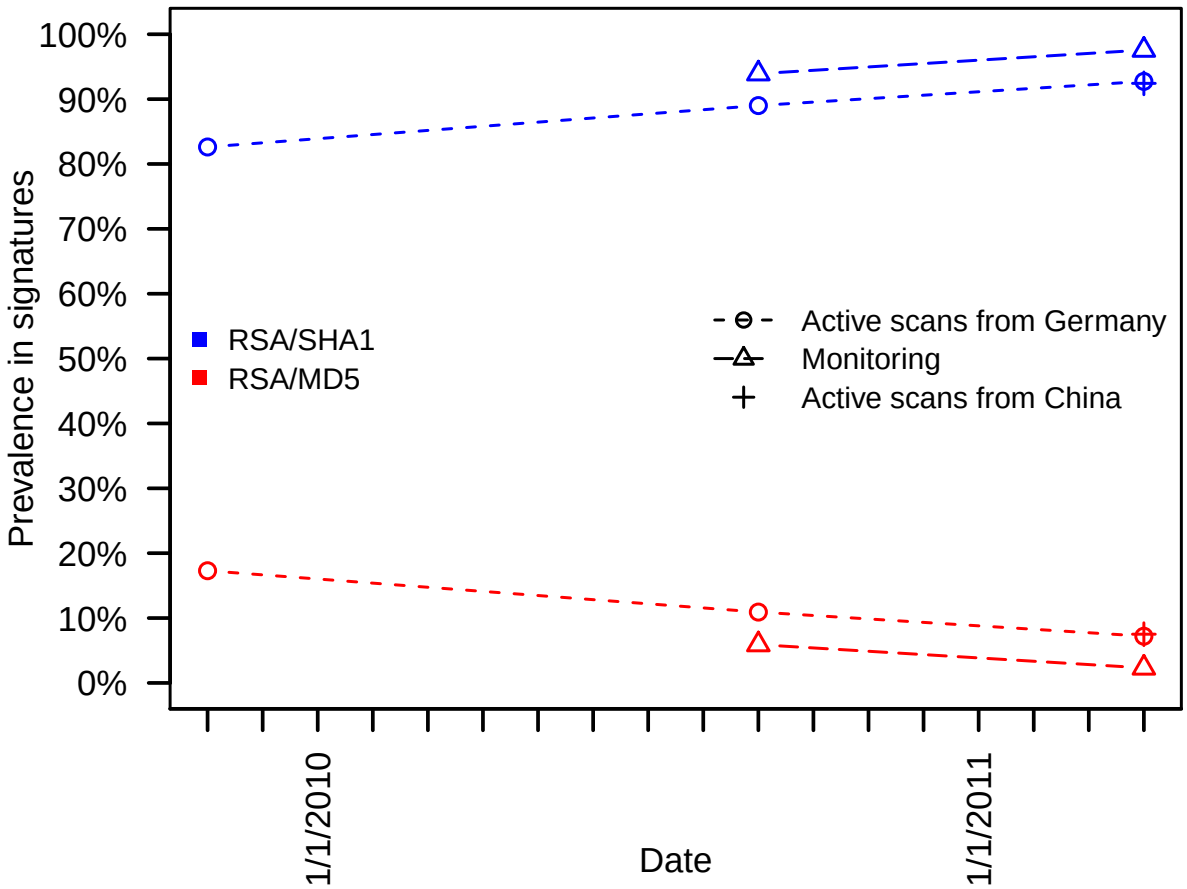

Certificate Issuers

Very few CAs account for $>$ 50% of certificates

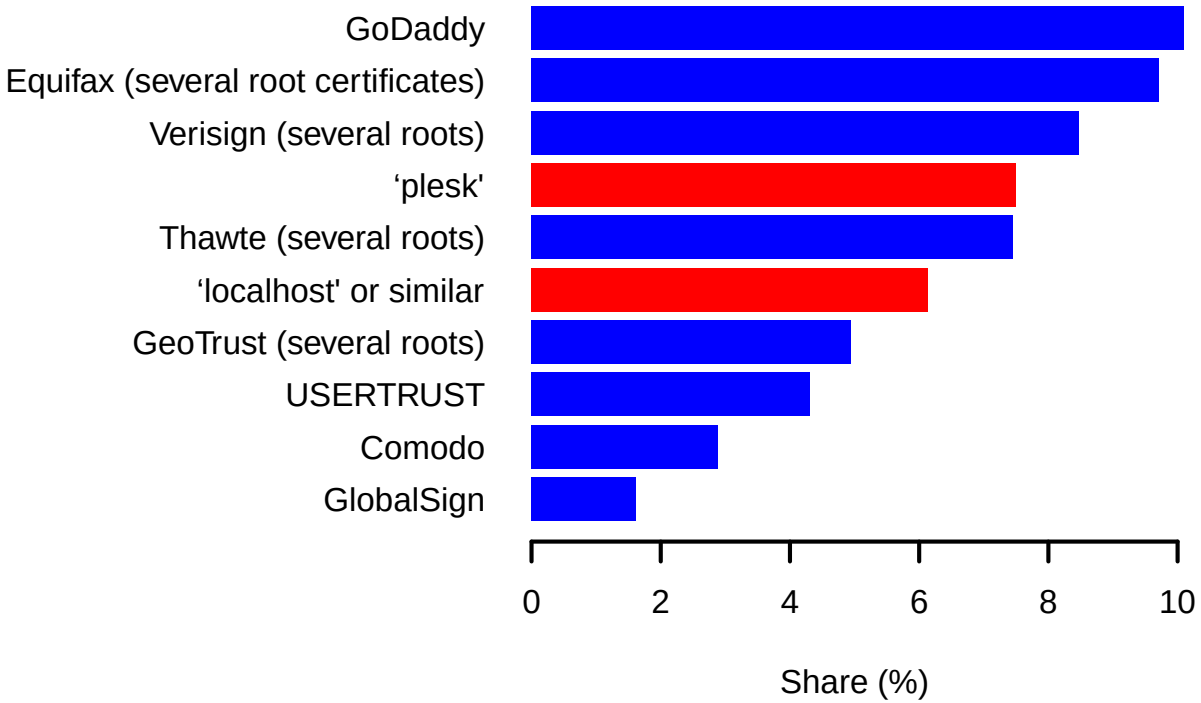

But there are 150+ Root Certificates in Mozilla.

Certificate Quality

We defined 3 categories

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

Certificate Quality

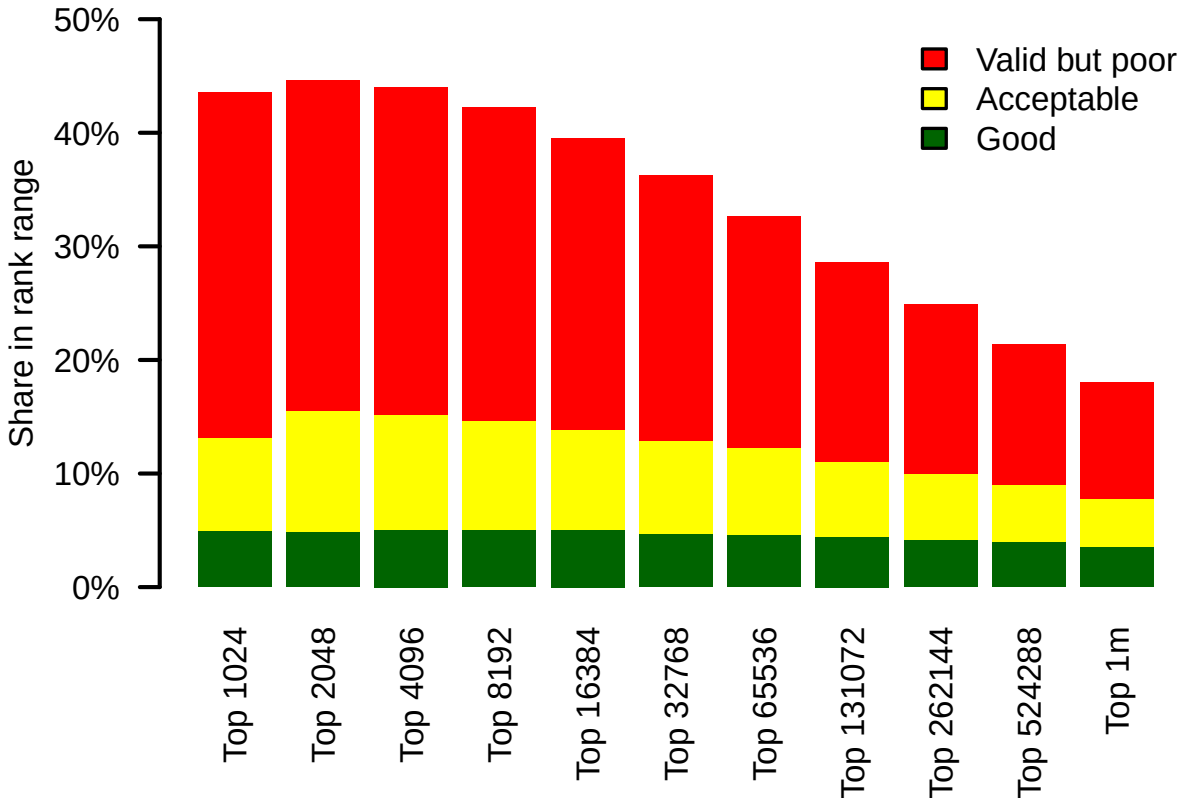

Validity correlates with rank

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

What to do about these problems?

No silver bullet known

- Part of the problem: SSL meant to protect stuff like credit card numbers
- But state-scale attacks were not in scope back in the 1990s

Several proposals:

- Extended Validation, Base Line Requirements
- Pinning Information
- Keys in DNSSEC (DANE)
- Perspectives/Convergence
- Public Log schemes: Sovereign Keys, Auditable CAs

PKI-me-harder

Extended Validation

- CAs to require state-issued documents before certification
- More expensive
- Rarely bought by customers

Base Line Requirements

- CA/Browser forum standard
- Absolute minimum requirements for validation
- Audit-based, rules for audits

Pinning Information

Idea

- Browser stores last-seen public key of a site
- Alternatively: store issuing CA
- Recognise again upon next visit

Discussion

- Does not help against attack on first contact
- False alarms when certificates change (not rare!)
- How many certs to store?

Keys in DNSSEC: DANE

Idea

- DNSSEC already a hierarchical **state-level** PKI
- Verification from Root Server down to end-host
- Resolve DNS name to IP address
- New Resource Record in DNSSEC: public key of site

Discussion

- Straight-forward and strong
- Performance? Caching?
- Countries control their own TLDs. Think `bit.ly`!
- Defence against country-level attack?

Idea: Notaries

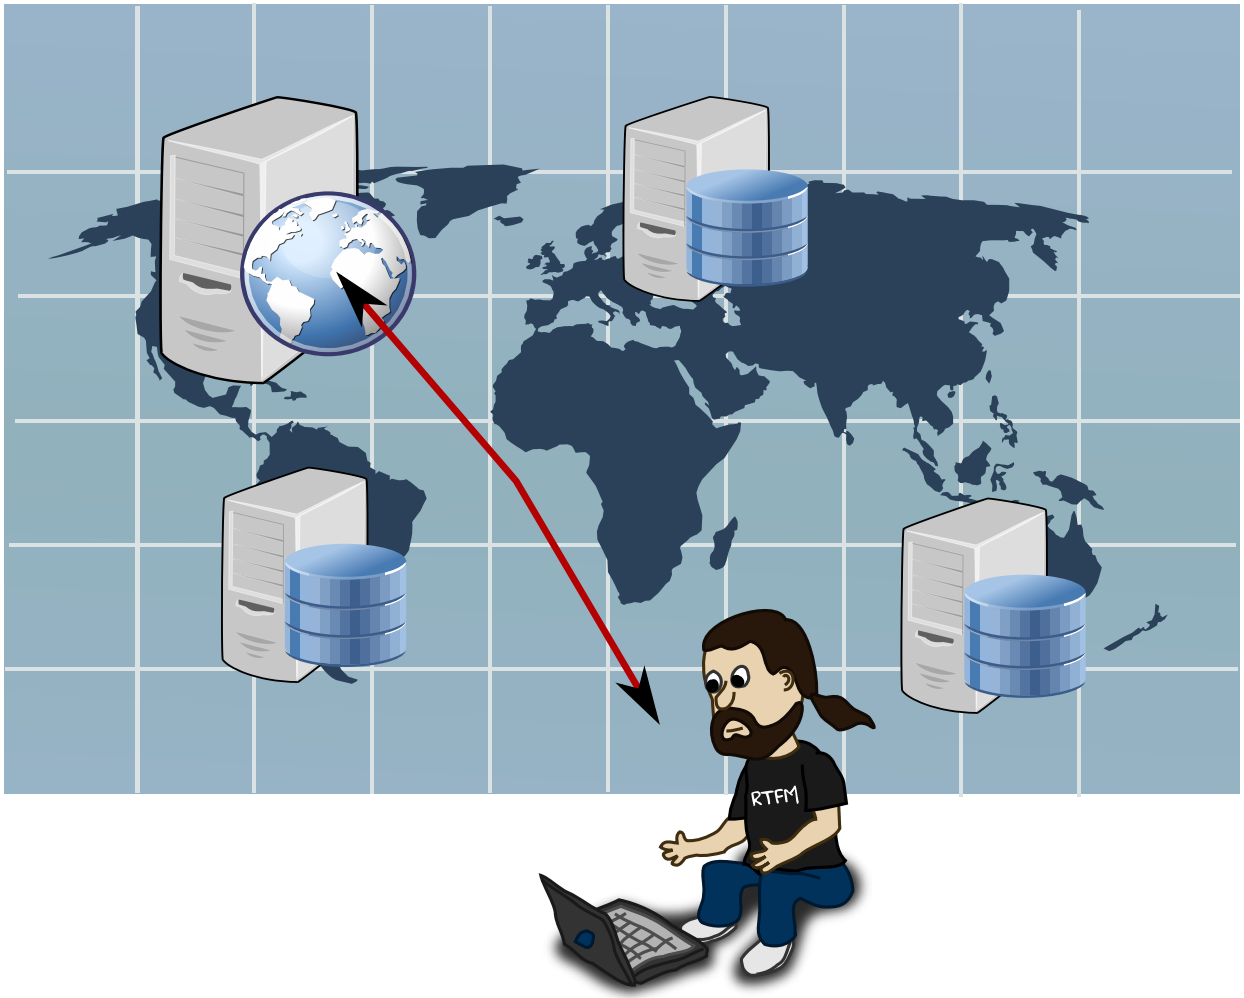

Reconfirm with notaries

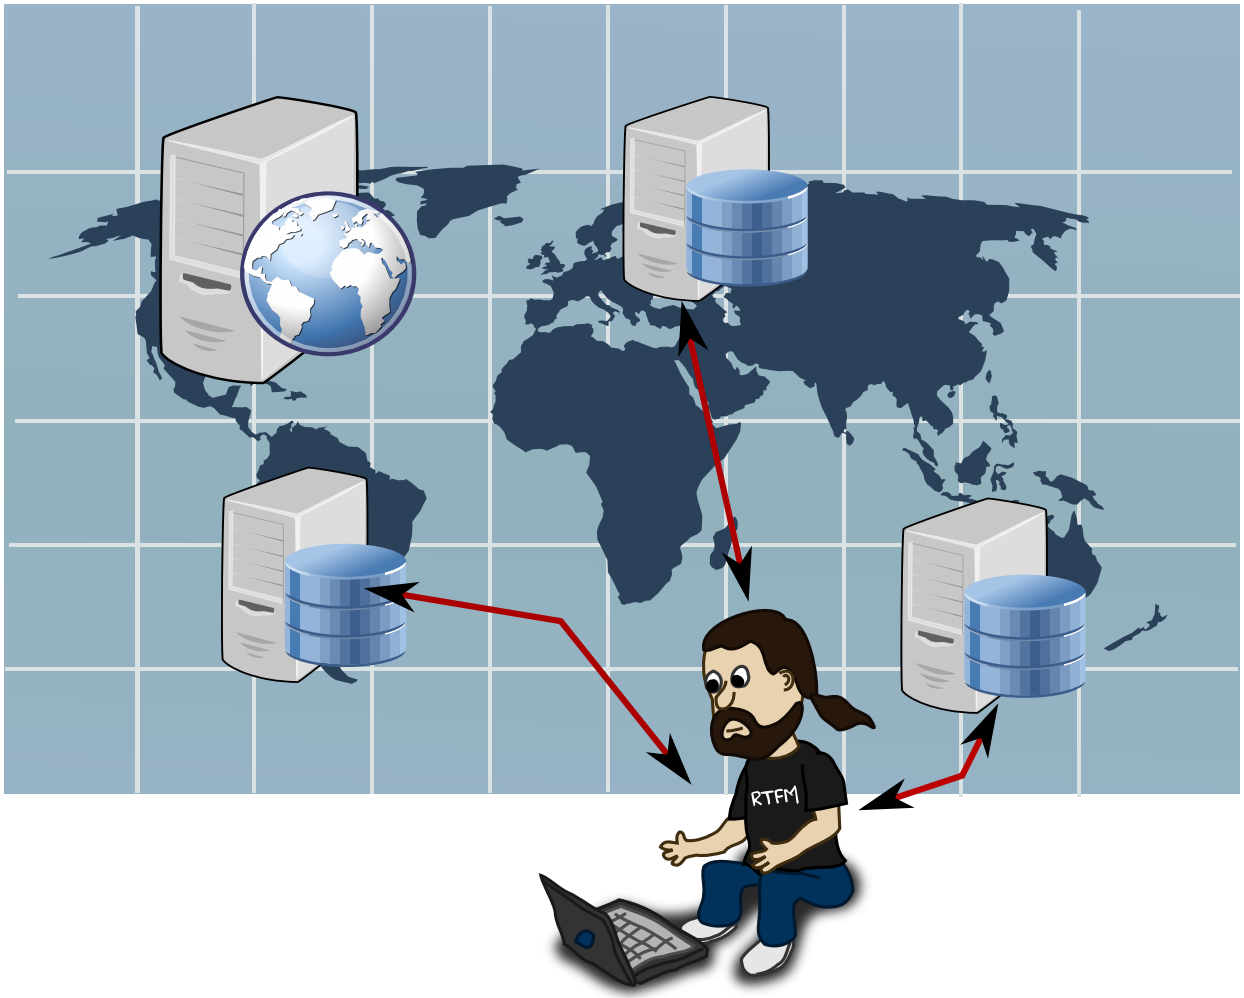

Discussion of Notary System

Advantages

- Works well against MitM in the network
- Reinforcement of CAs system

Possible problems

- Privacy!
- False positives: some sites change certificates frequently
- Content Distribution Networks?
- Indication of site identity still necessary

Public Log

Store information

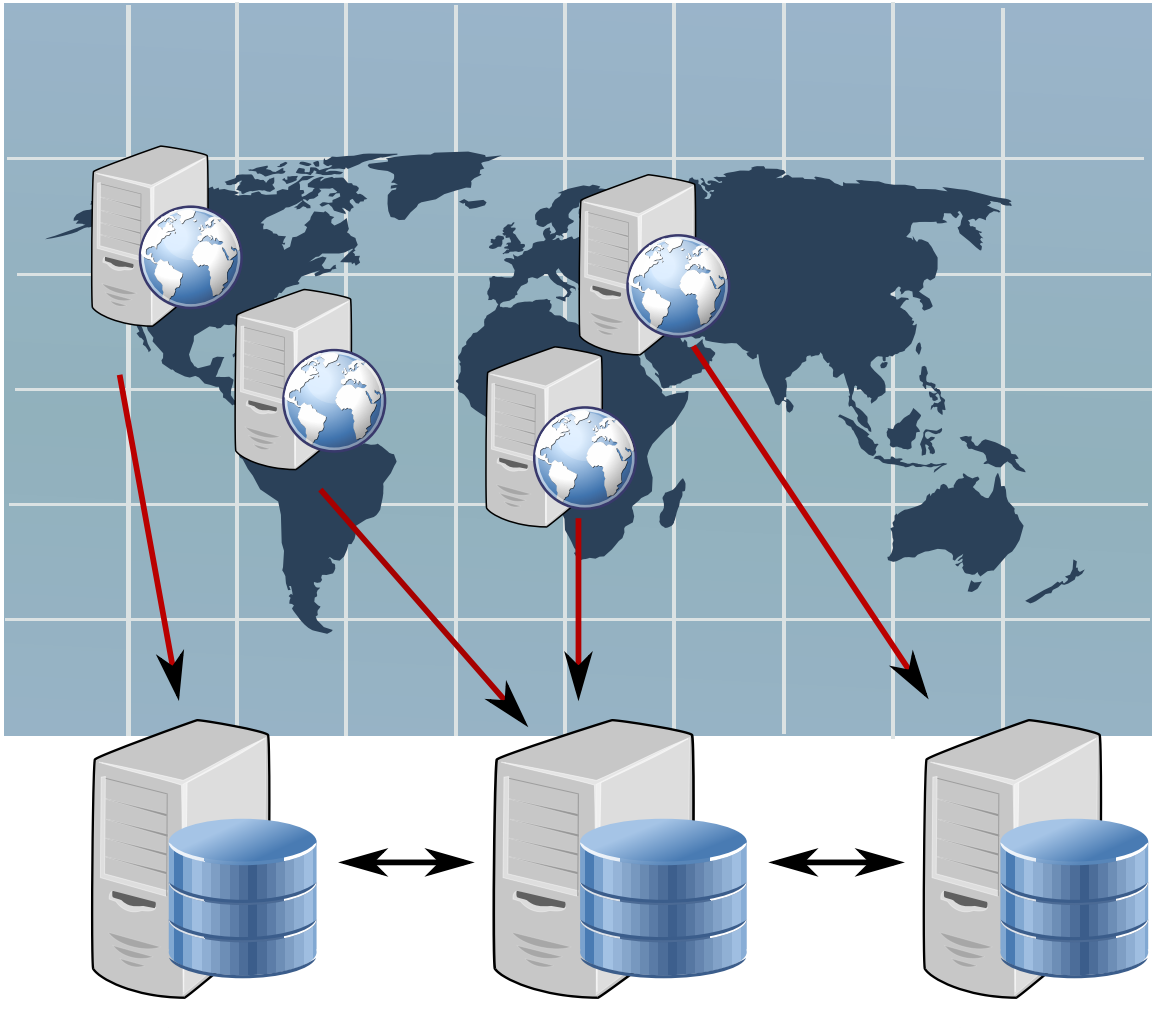

Public Log

Retrieve information

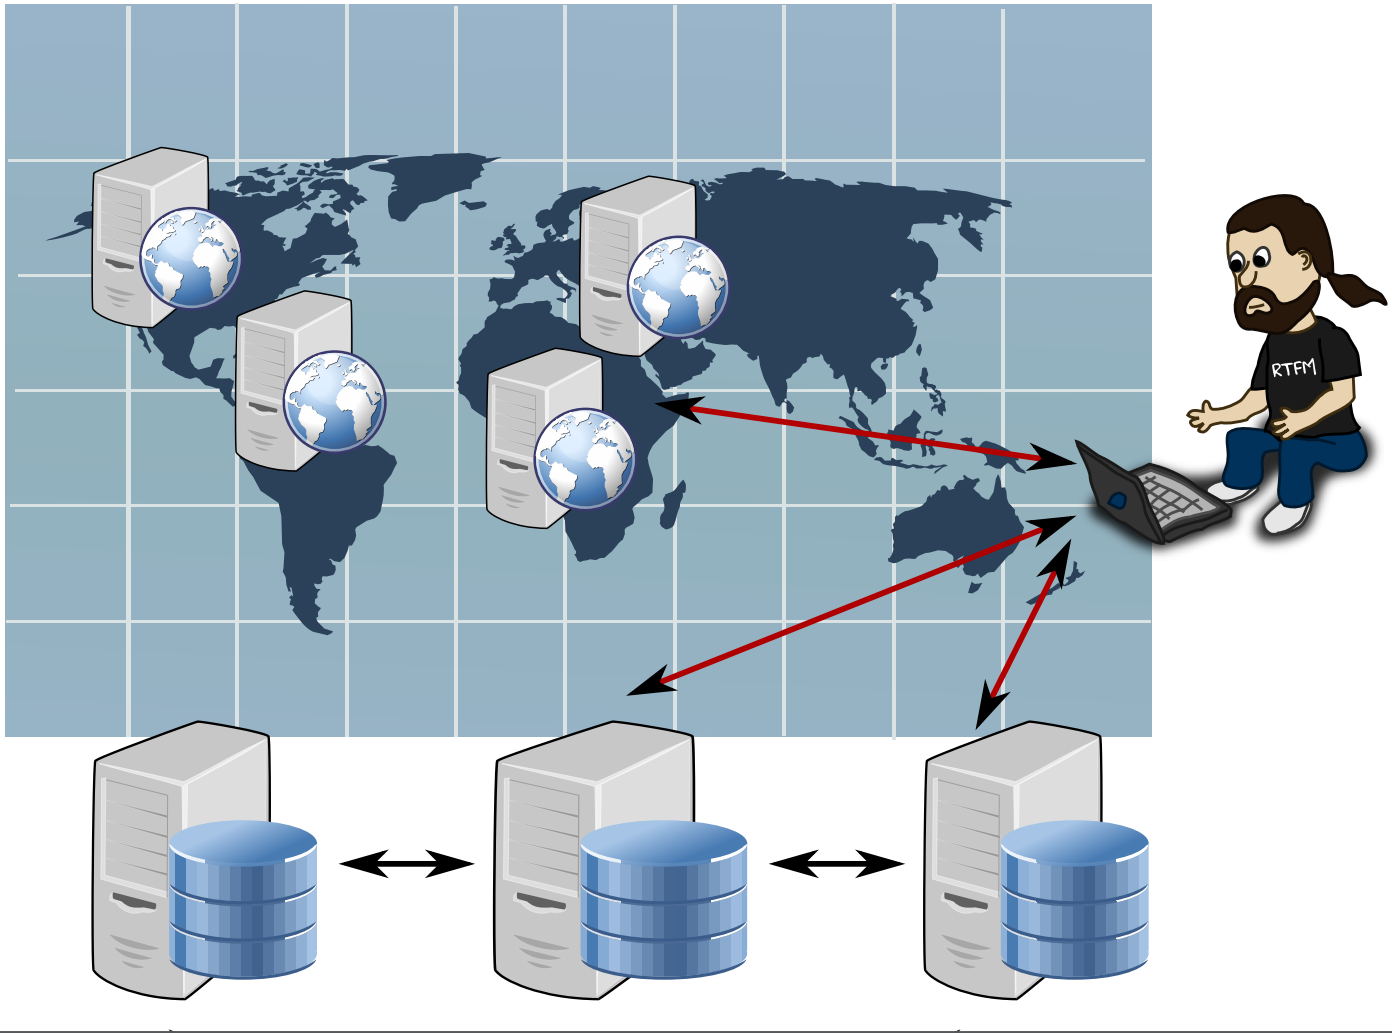

Public Log Schemes

Idea: store information publicly and append-only

- Sovereign Keys
  - Sites stores authoritative key to cross-sign its certificates
  - Goal: cross-certification and cross-validation of keys
- Auditable CAs
  - CAs store info about whom they certify
  - Goal: detect rogue CA issuing key for a site

Schemes are very new

- Both from November 2011
- Sovereign Keys to be presented at 28C3 in 2 weeks

Sovereign Keys (EFF)

Sites store information on $<$ 30 timeline servers

timestamp	name	key	protocols	evidence
1322736203	A	0x427E8A	https, smtps	$Sig_{CA}(A, \cdots)$
1323254603	B	0x7389FB	https:8080	$Sig_{B}(B, \cdots)$
1323657143	C	0x49212A	imaps	$Sig_{C}(C, \cdots)$
1413787143	A	0x427E8A	https, smtps	$Sig_{CA}(A, \cdots)$
...	...	...	...	...

Work-in-progress

- Timeline is auditable by clients
- Mirrors proposed
- https://www.eff.org/sovereign-keys

Sovereign Keys: Discussion

Pros

- Does not need CA support
- Evidence can be based on DANE DNSSEC, CAs, . . .
- Performance and bandwidth?

Cons

- Continous monitoring of timeline server needed
- Maintain list of timeline servers
- Entries are not space-efficient
- Privacy (suggested remedy: TOR-like proxying)
- Key loss

Auditable CAs (Ben Laurie et al.)

CAs store certification proof on public servers

timestamp	name	cert	evidence
1322736203	A	Cert by Verisign	$MSig(hashes)$
1323254603	B	Cert by GoDaddy	$MSig(hashes)$
1323657143	C	Cert by CACert	$MSig(hashes)$
...	...	...	$MSig(hashes)$

Very-much-work-in-progress (IANAC!)

- Log = certificates, periodically Merkle signatures
- Induces small latency
- `http://www.links.org/files/CertificateAuthorityTransparencyandAuditability.pdf`

Discussion of Auditable CAs

Pros

- Protects against rogue/hacked CAs
- Efficient data structure

Cons

- CAs do not like to disclose their customers
- Requires continous monitoring of logs
- Revocation unclear

Crossbear: hunting the MitM

Our own contribution: find the MitM

- Idea: combine Notary concept with tracerouting
- Goal: gain reliable data about MitM
  - How often do they occur? Are they common?
  - Where are the attackers located?
  - Maybe: who are the attackers?

Alice is surfing...

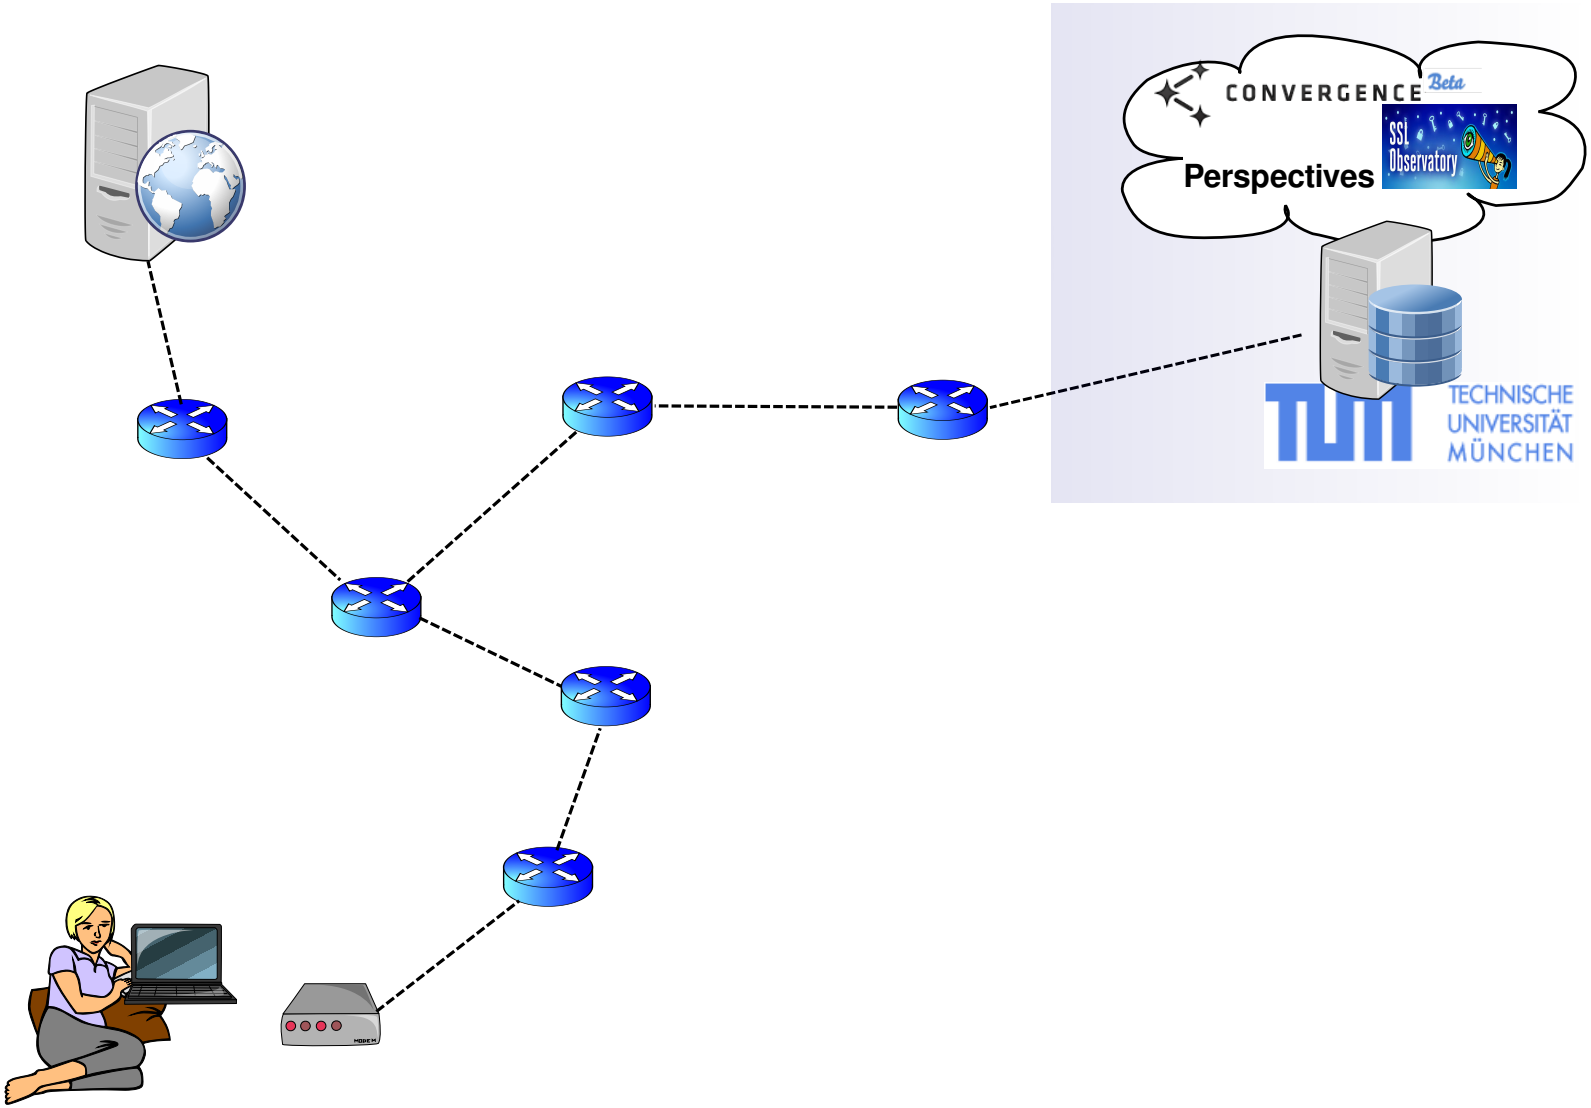

Man-in-the-middle

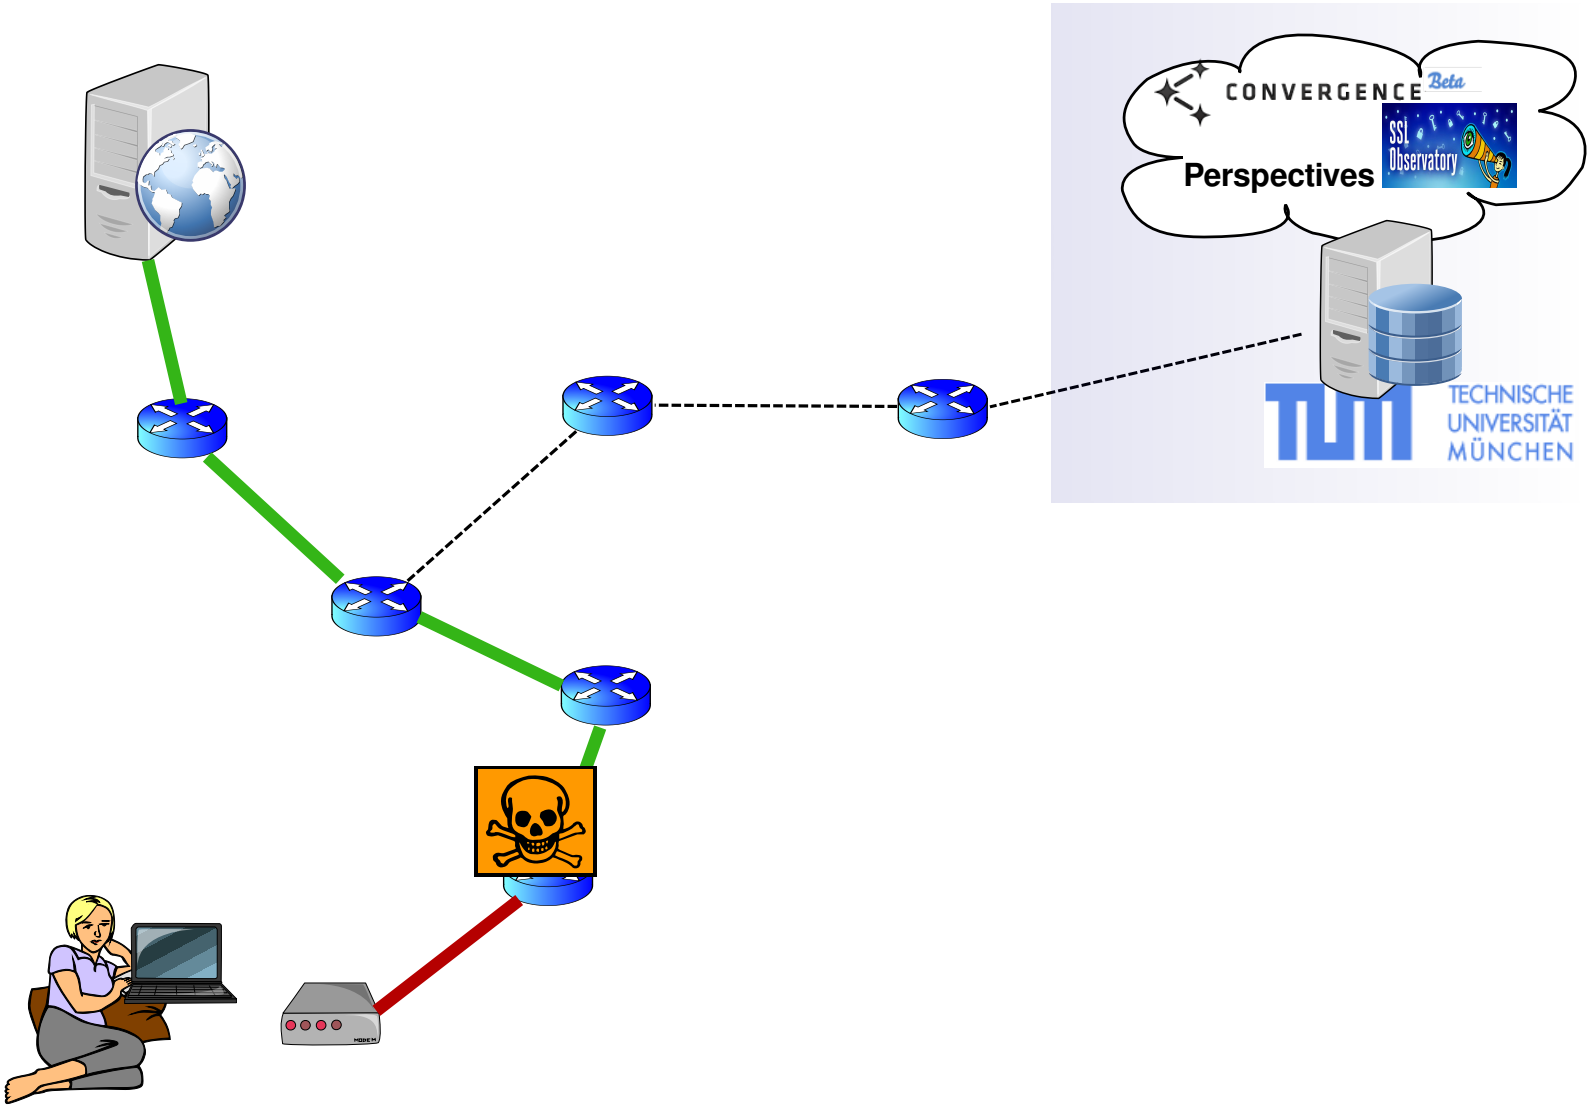

Alice queries CrossBear

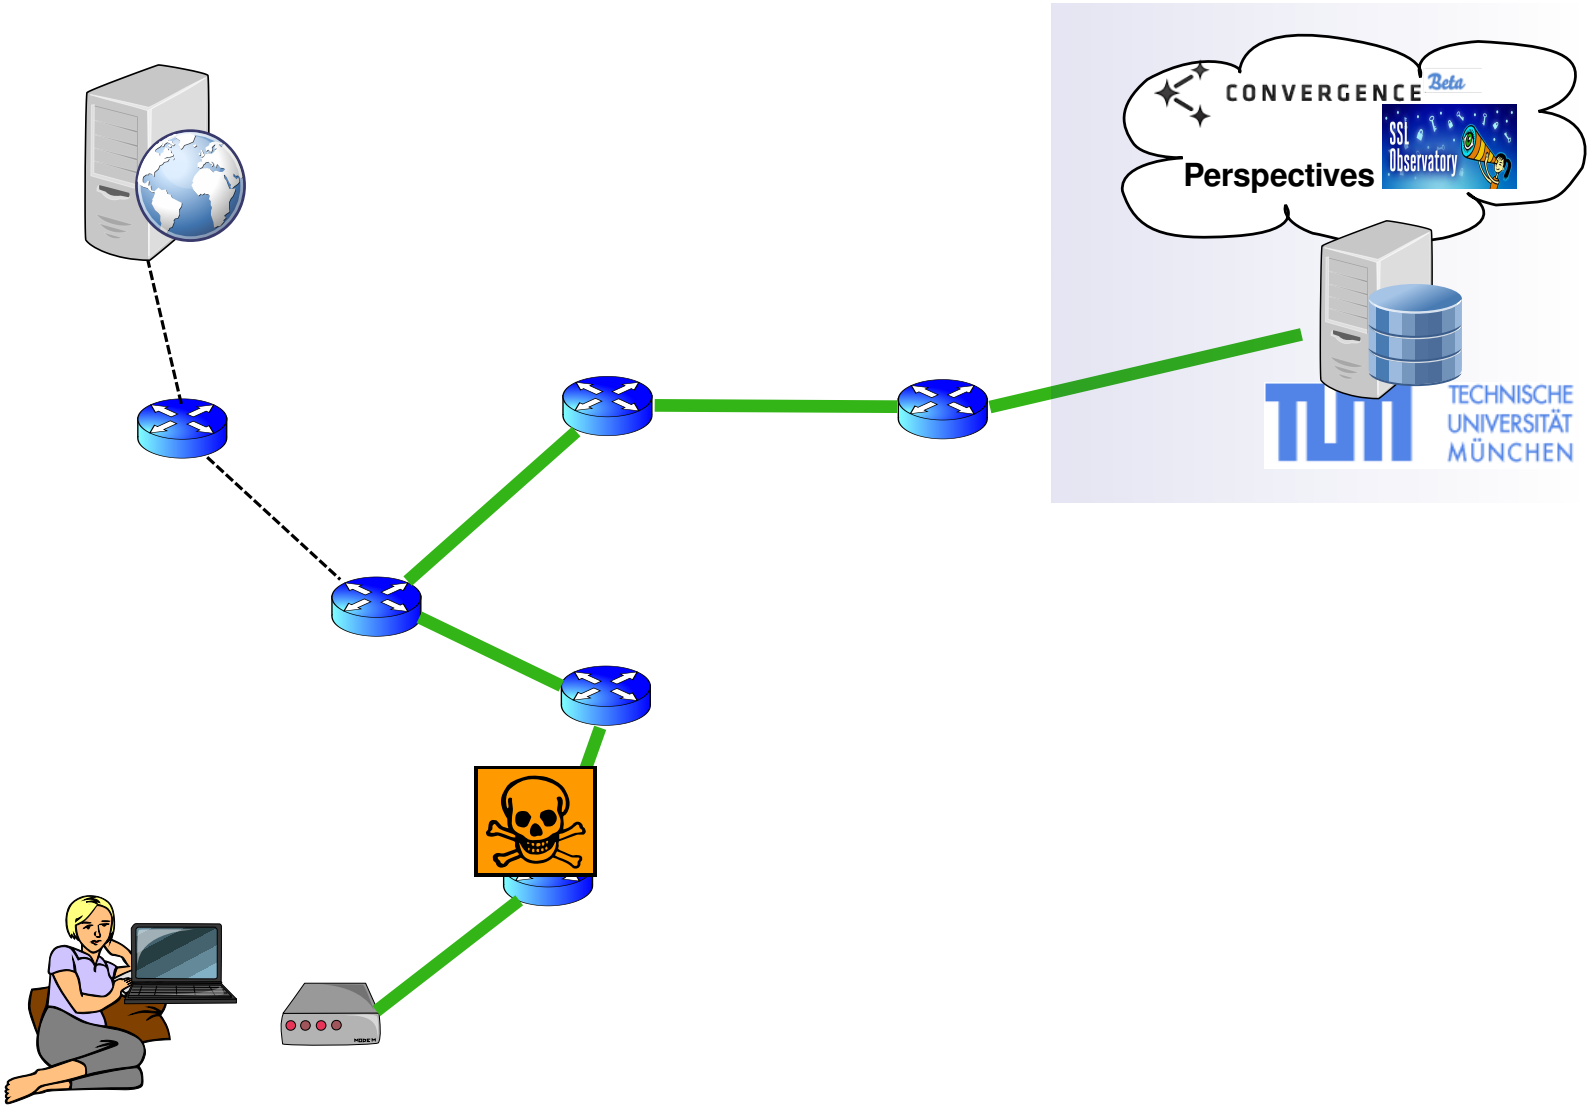

CrossBear checks the server

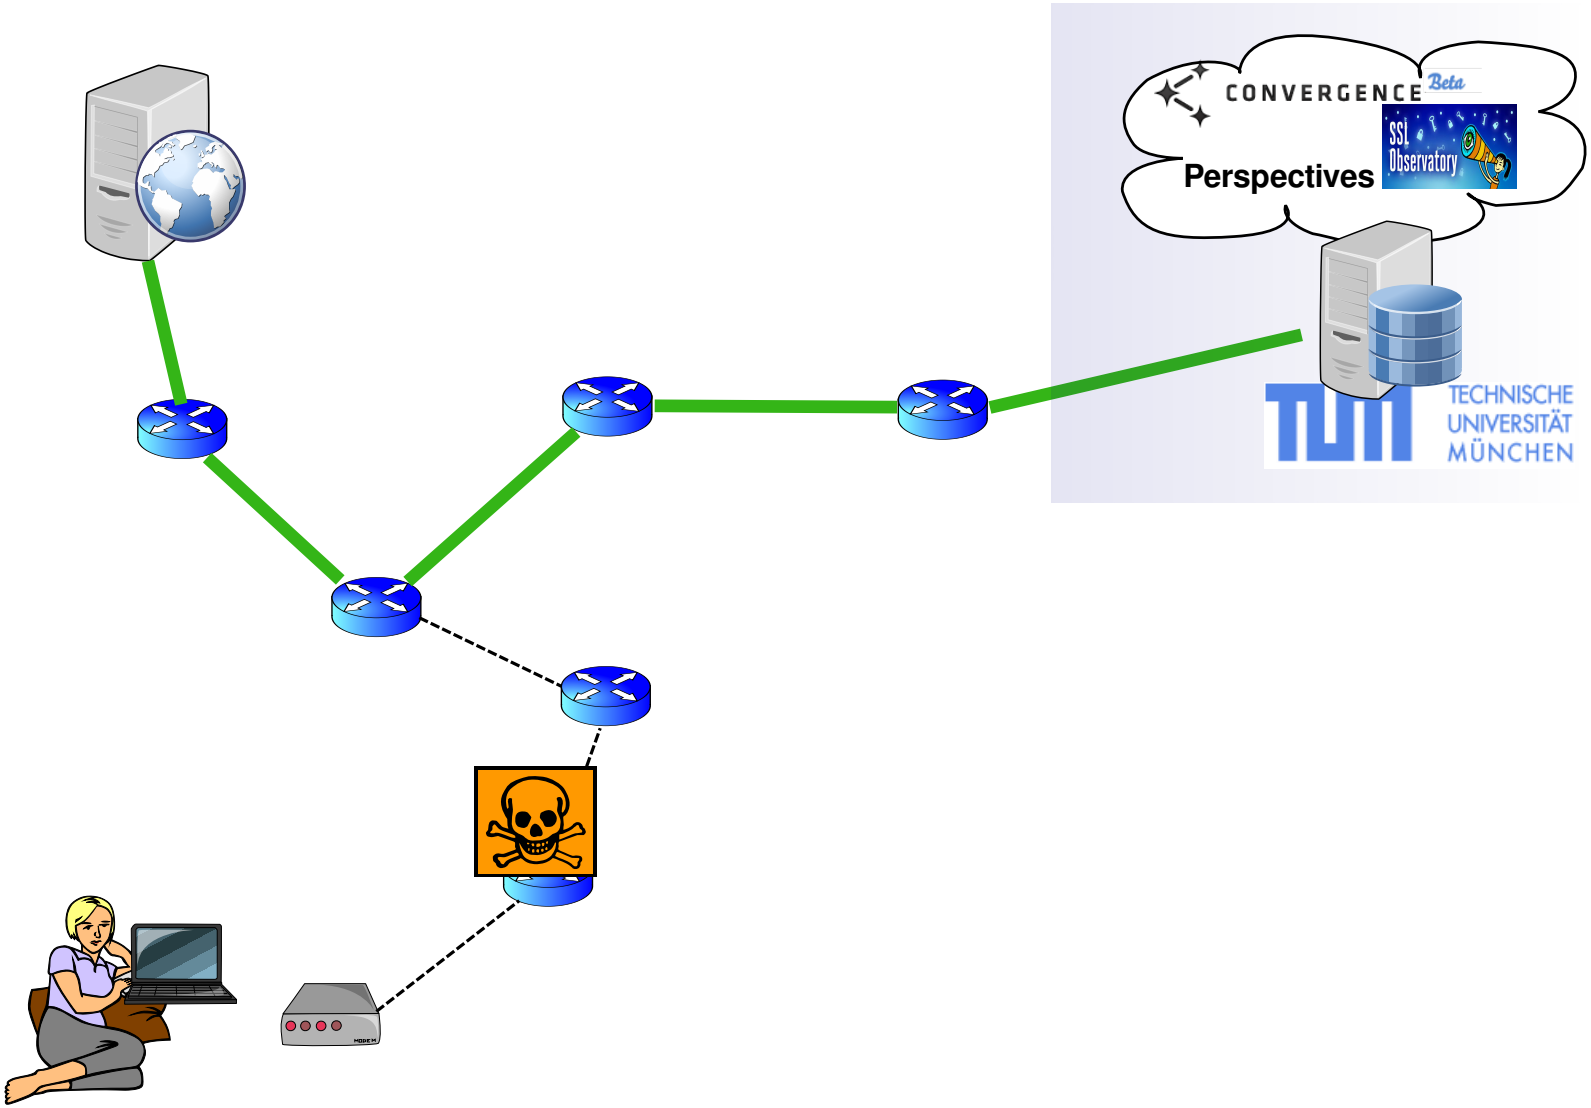

CrossBear reports result

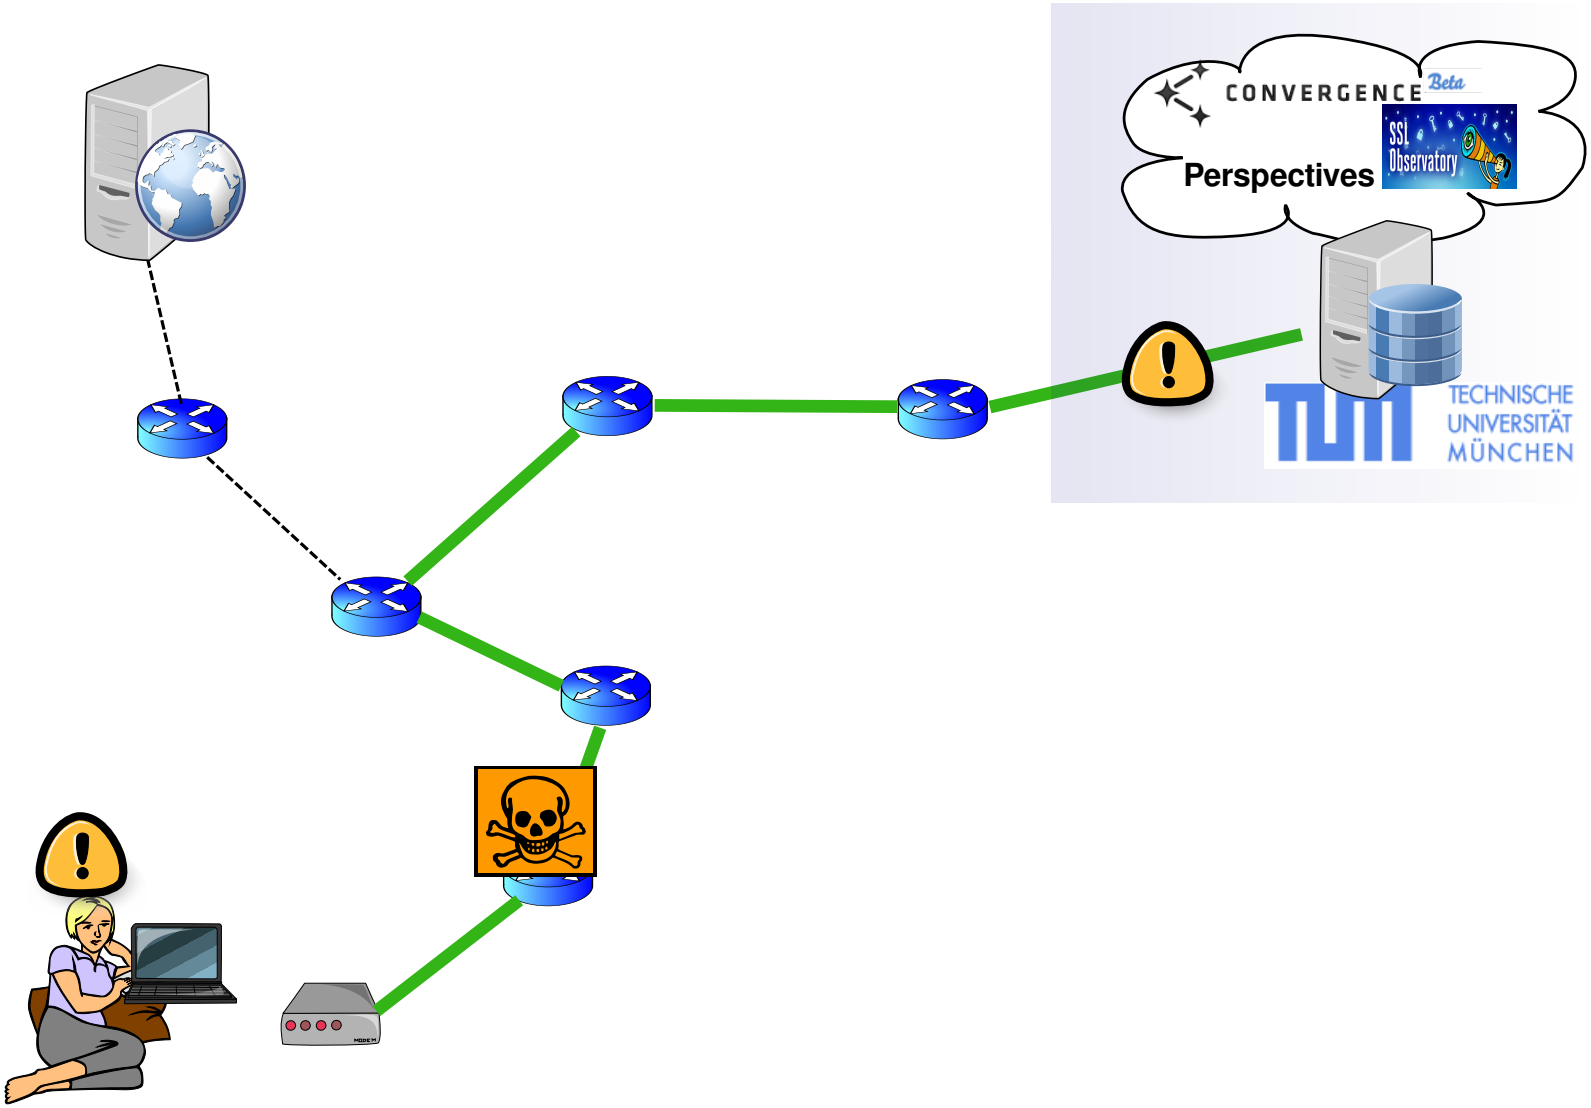

Alice traceroutes to server

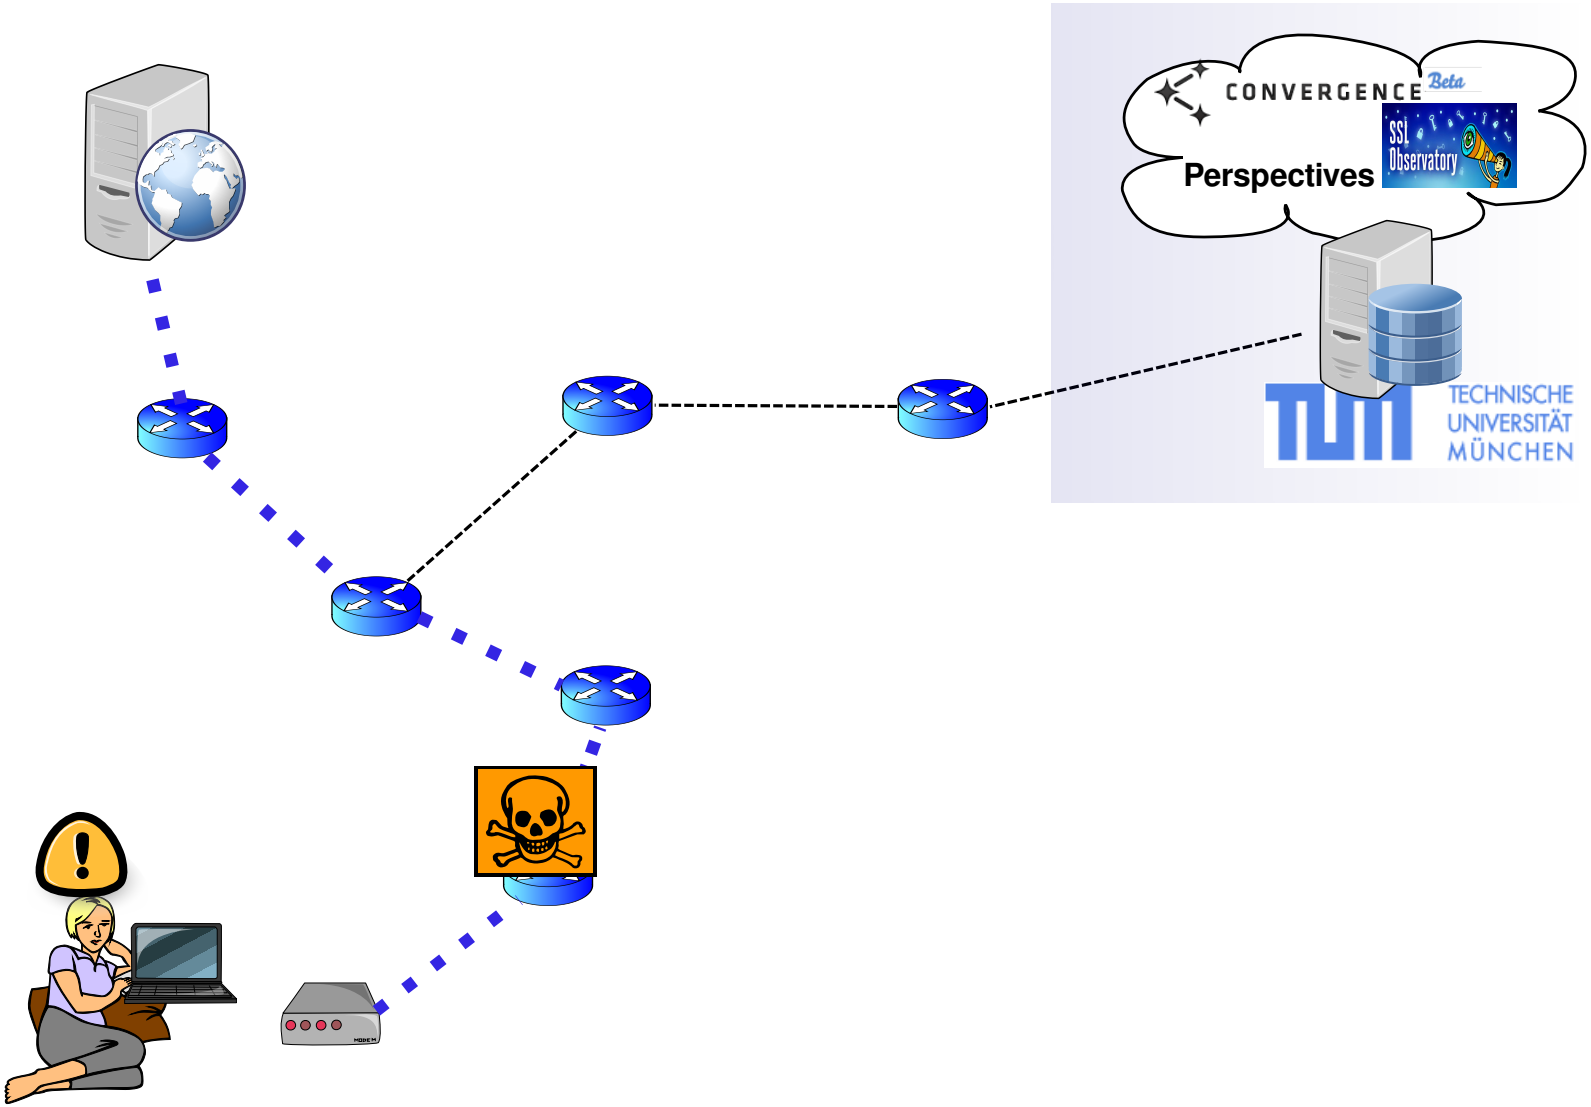

Alice reports to CrossBear

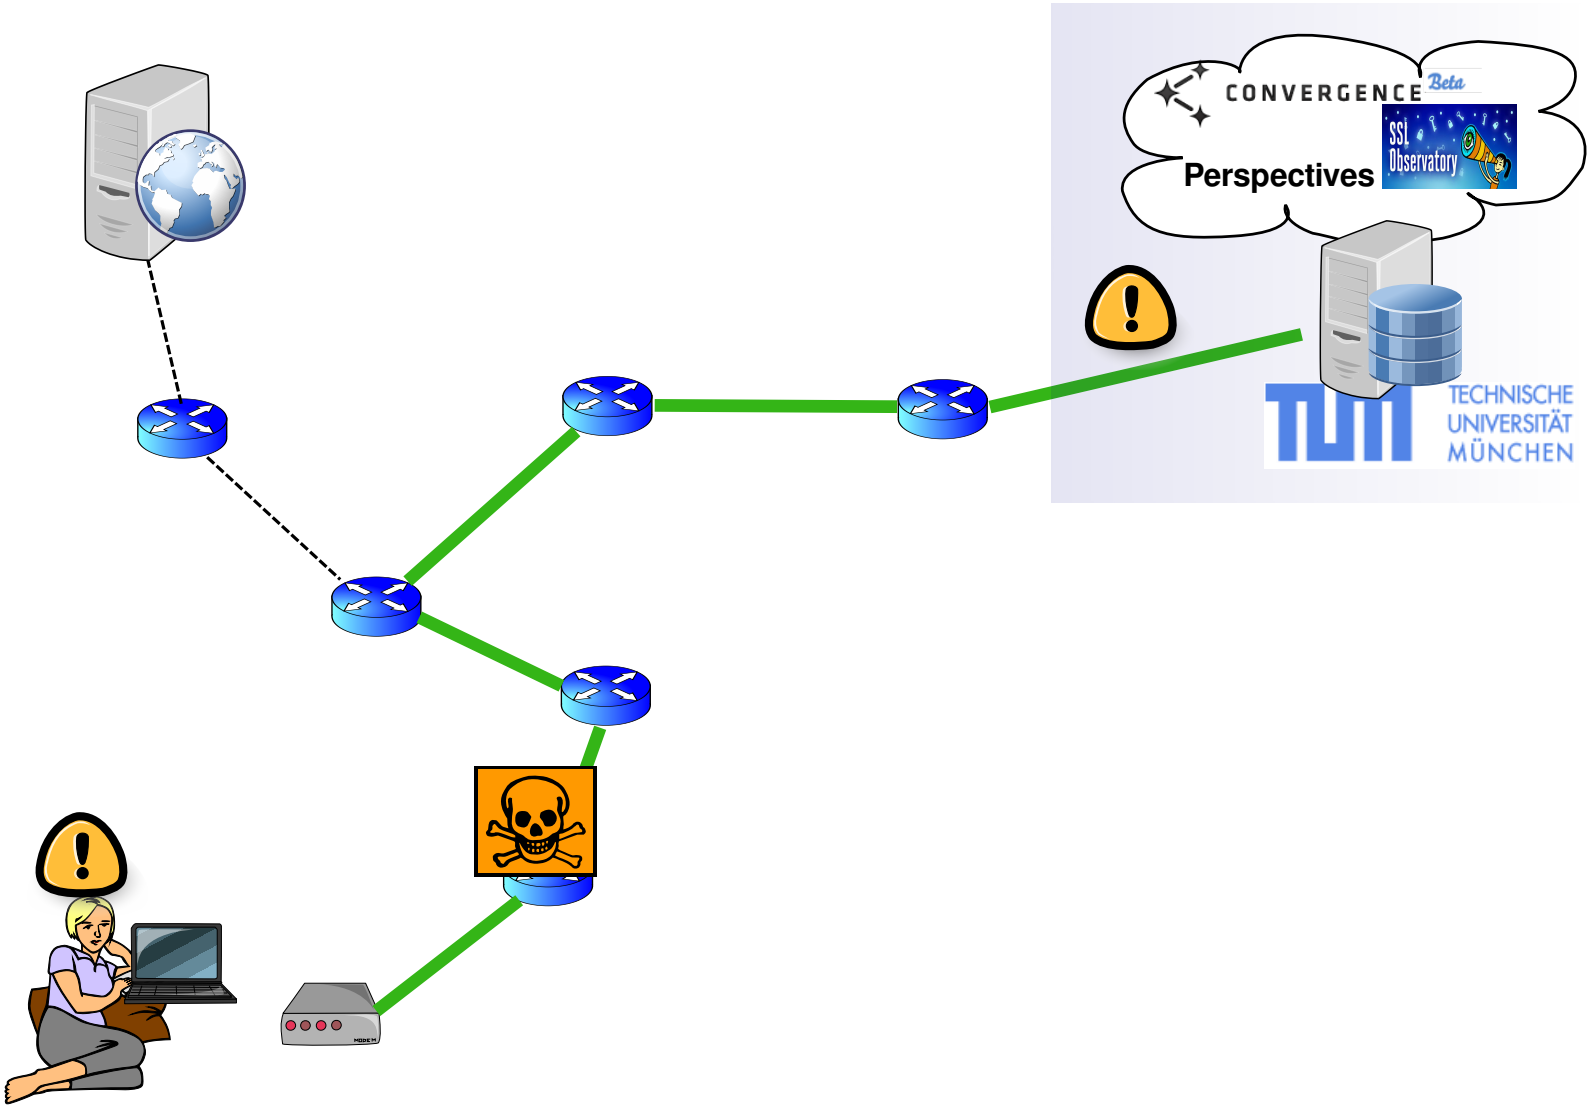

Distribute hunting tasks

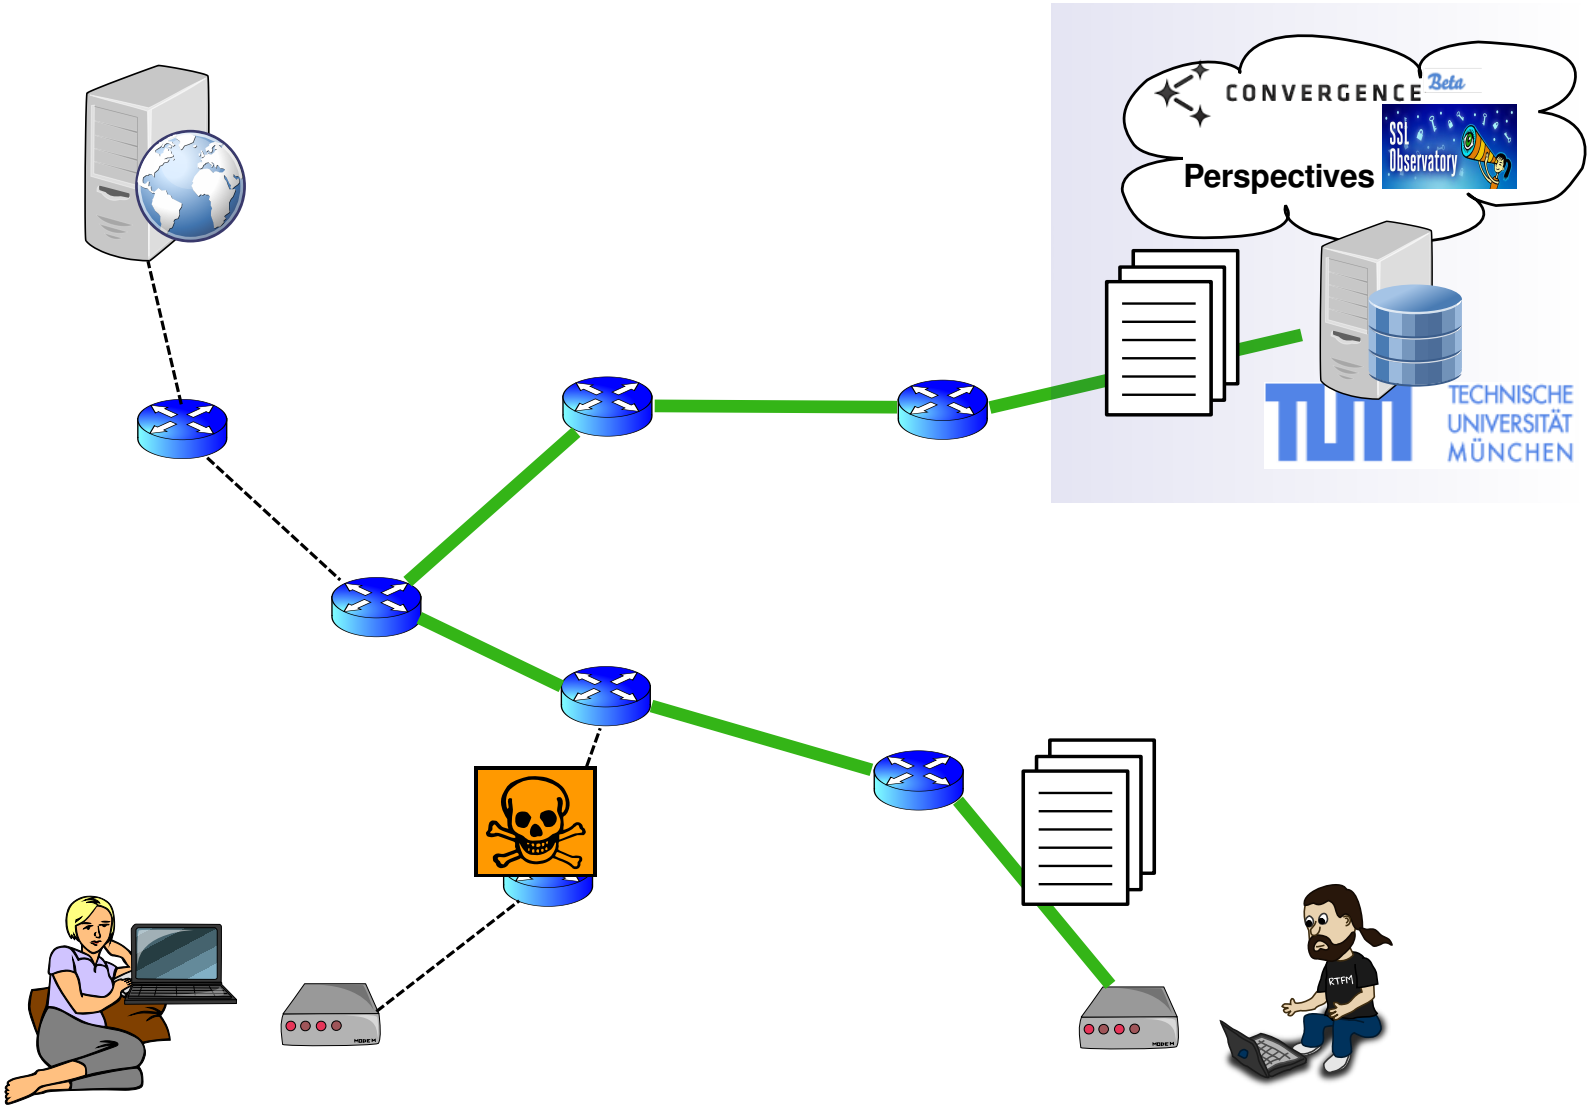

Bob goes hunting

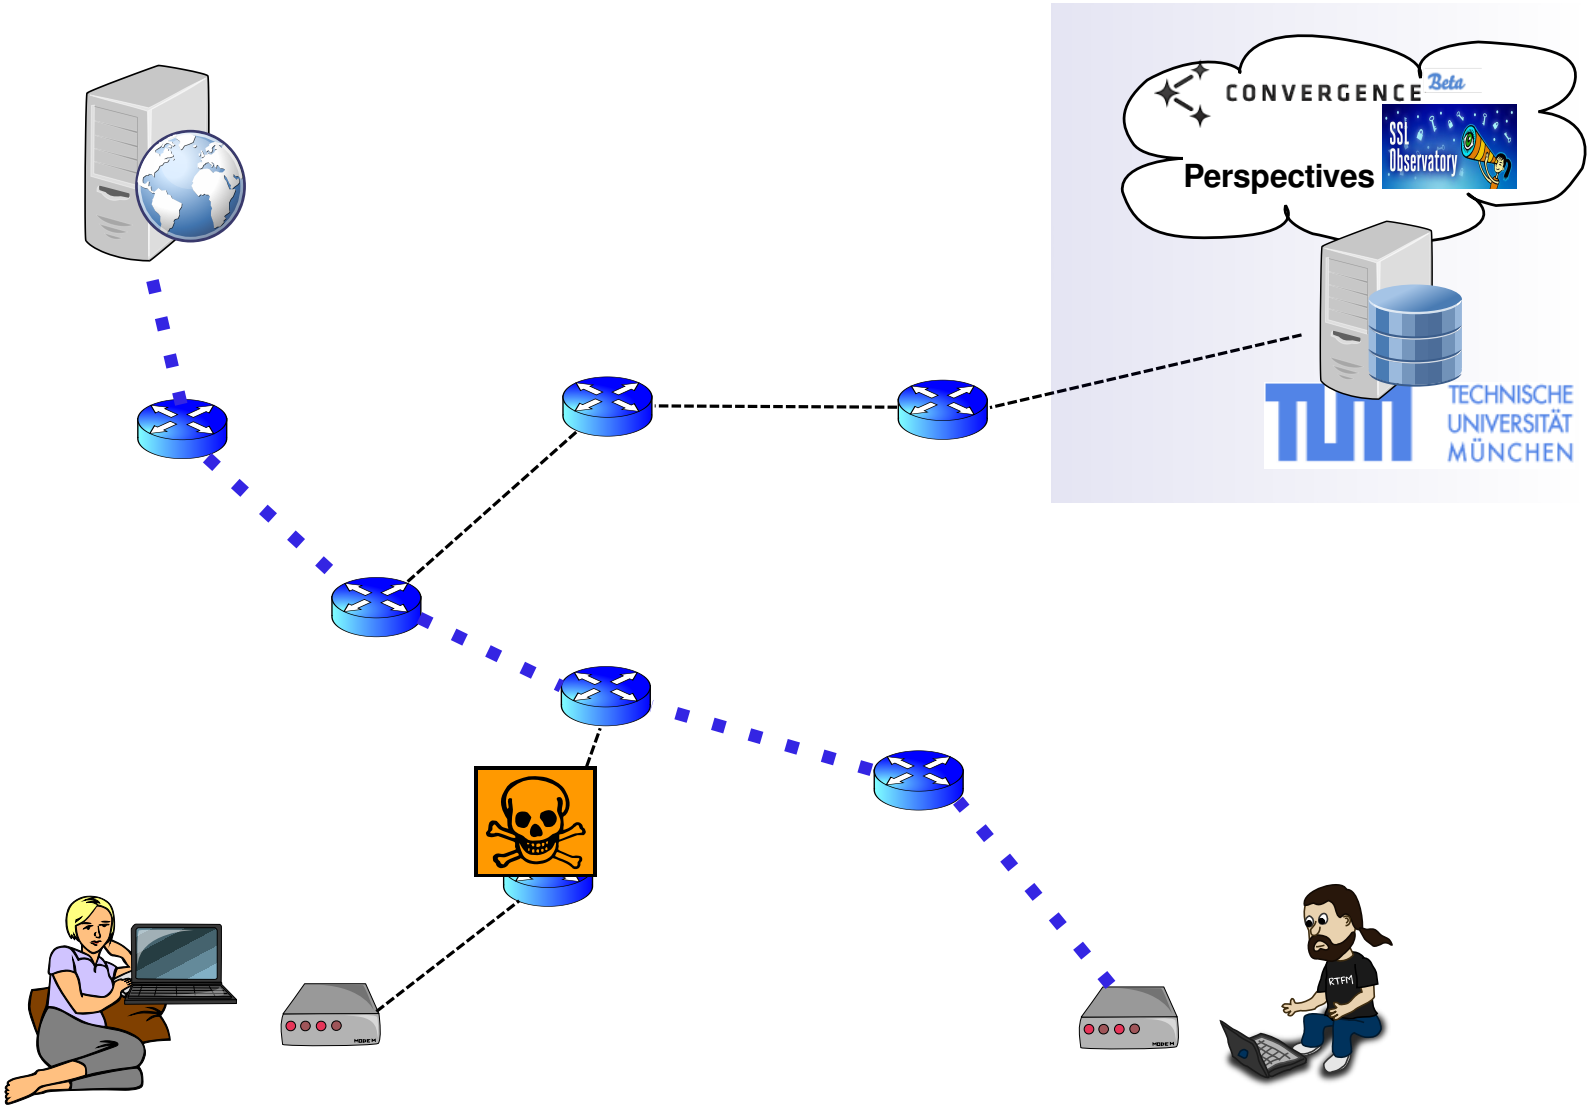

Bob reports

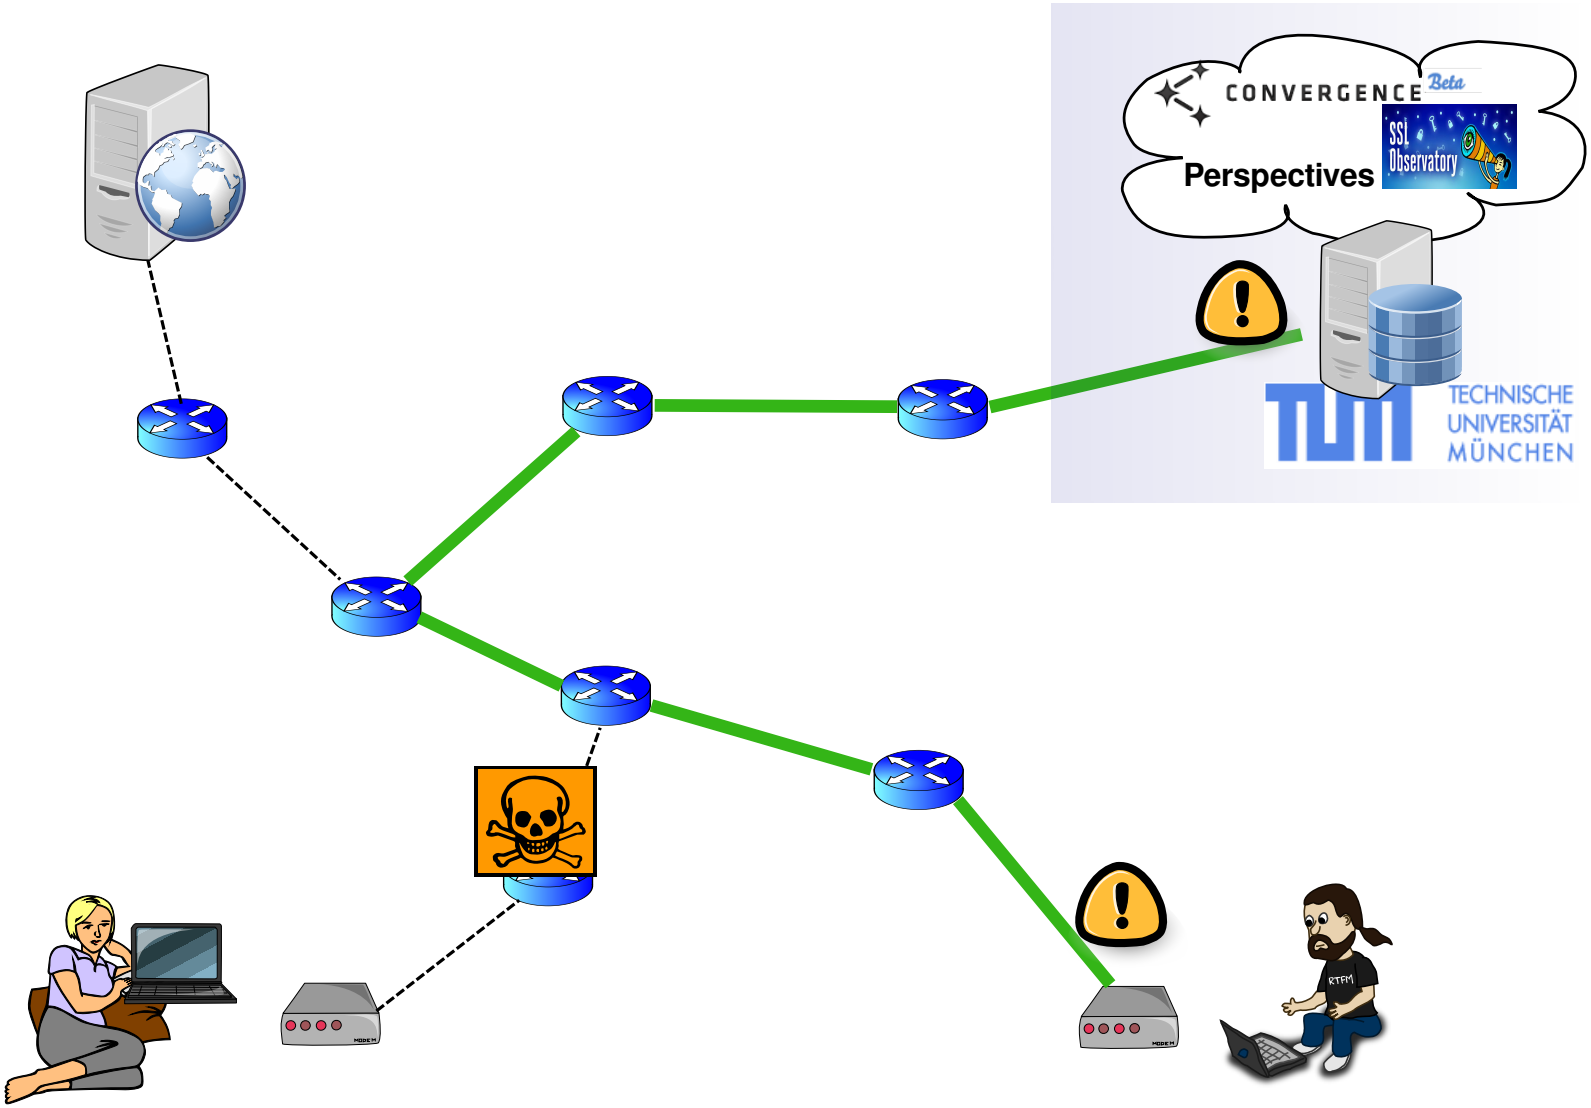
CONVERGENCE
Beta
SSL Observatory
Perspectives
TECHNISCHE UNIVERSITÄT MÜNCHEN
RTFM

There are many Bobs

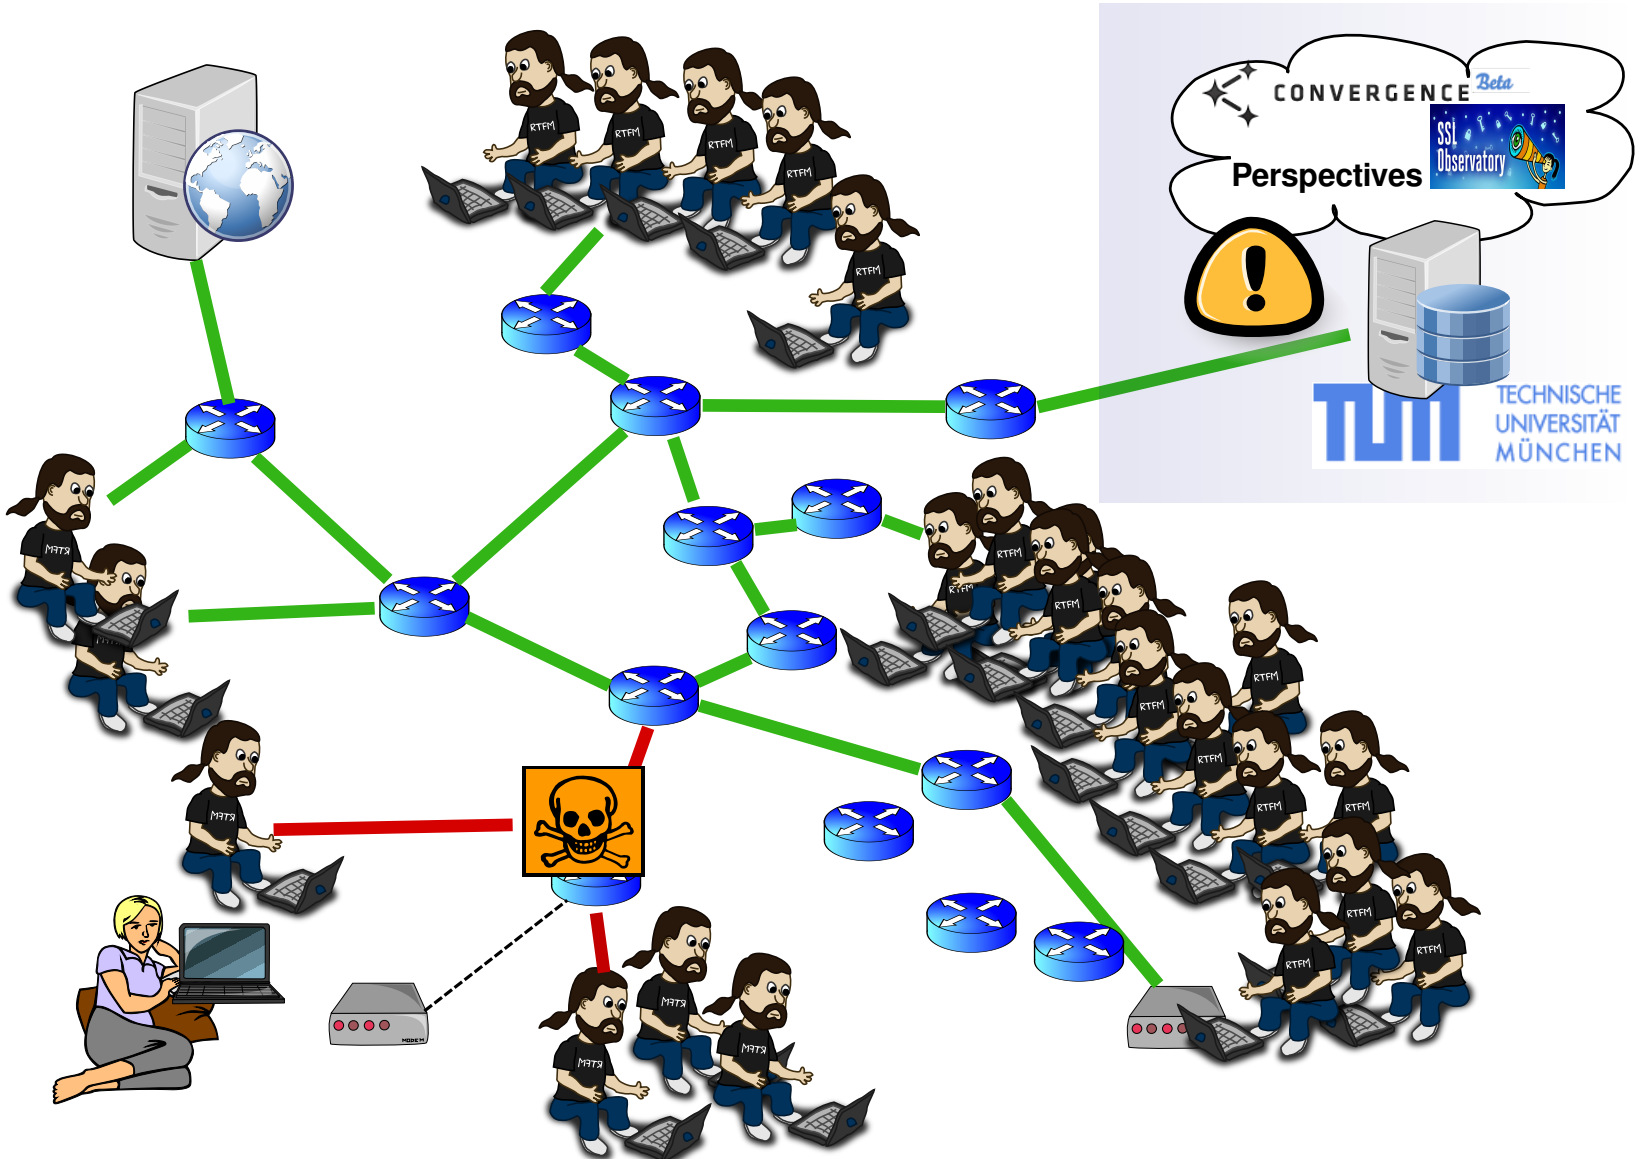

Crossbear summary

Status

- Running
- 150 hunters deployed on PlanetLab, waiting for action
- Project needs Bobs (and Bobinas) now
- Get it from `https://github.com/crossbear/Crossbear`

Contact us

- Twitter: @crossbearteam
- WWW: `https://pki.net.in.tum.de`

Summary

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

- Only 18% of the Top 1 Million Web sites show fully valid certificates
- Much carelessness

Some positive developments

- Slight trend towards fully valid certificates
- Crypto OK

Remedies? Remains to be seen.

Thank you!

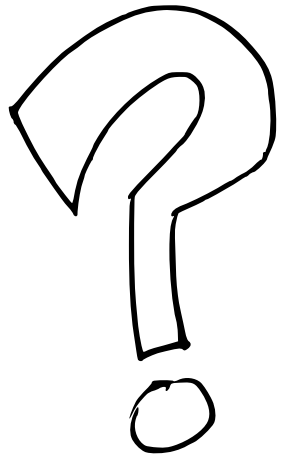

- Crossbear: `https://github.com/crossbear/Crossbear`
- Web presence (download datasets!):
  `https://pki.net.in.tum.de`
- Twitter: `@crossbearteam`

Errors in TLS Connection Setup

`UNKNOWN PROTOCOL`

- Rescanned those hosts and manual sampling
- Always plain HTTP...
- ... and always an `index.html` with HTML 2 ...
- Hypothesis: old servers, old configurations
- More likely to happen in the lower ranks

Symmetric Ciphers

Results from monitoring

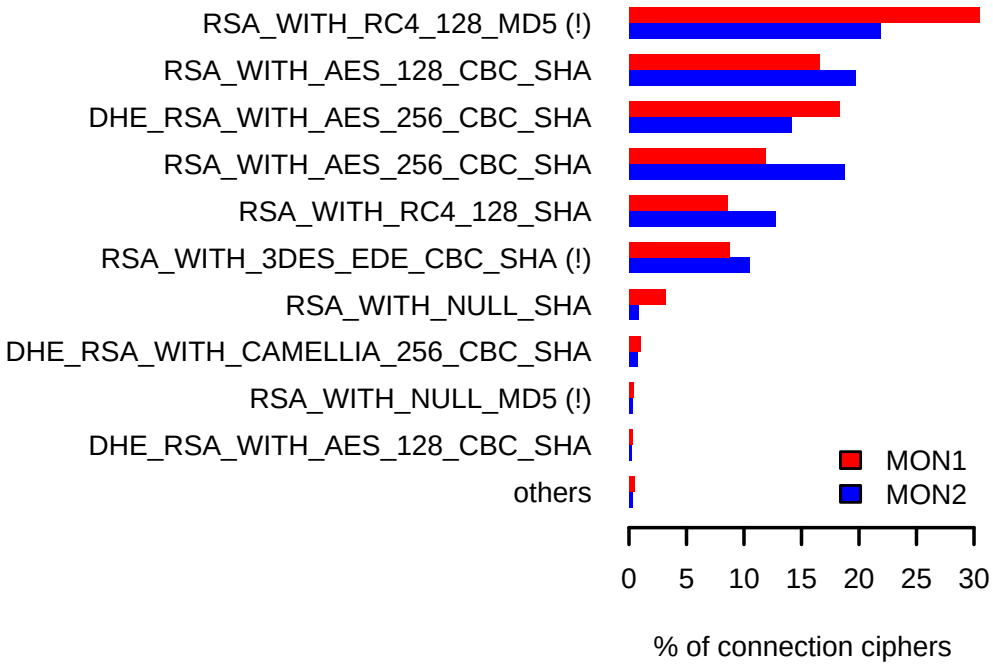

(Mostly) in line with results from 2007 by Lee et al.

- Order of AES and RC4 has shifted, RC4-128 most popular

Public Key Lengths

CDF for RSA key lengths – double-log Y axis

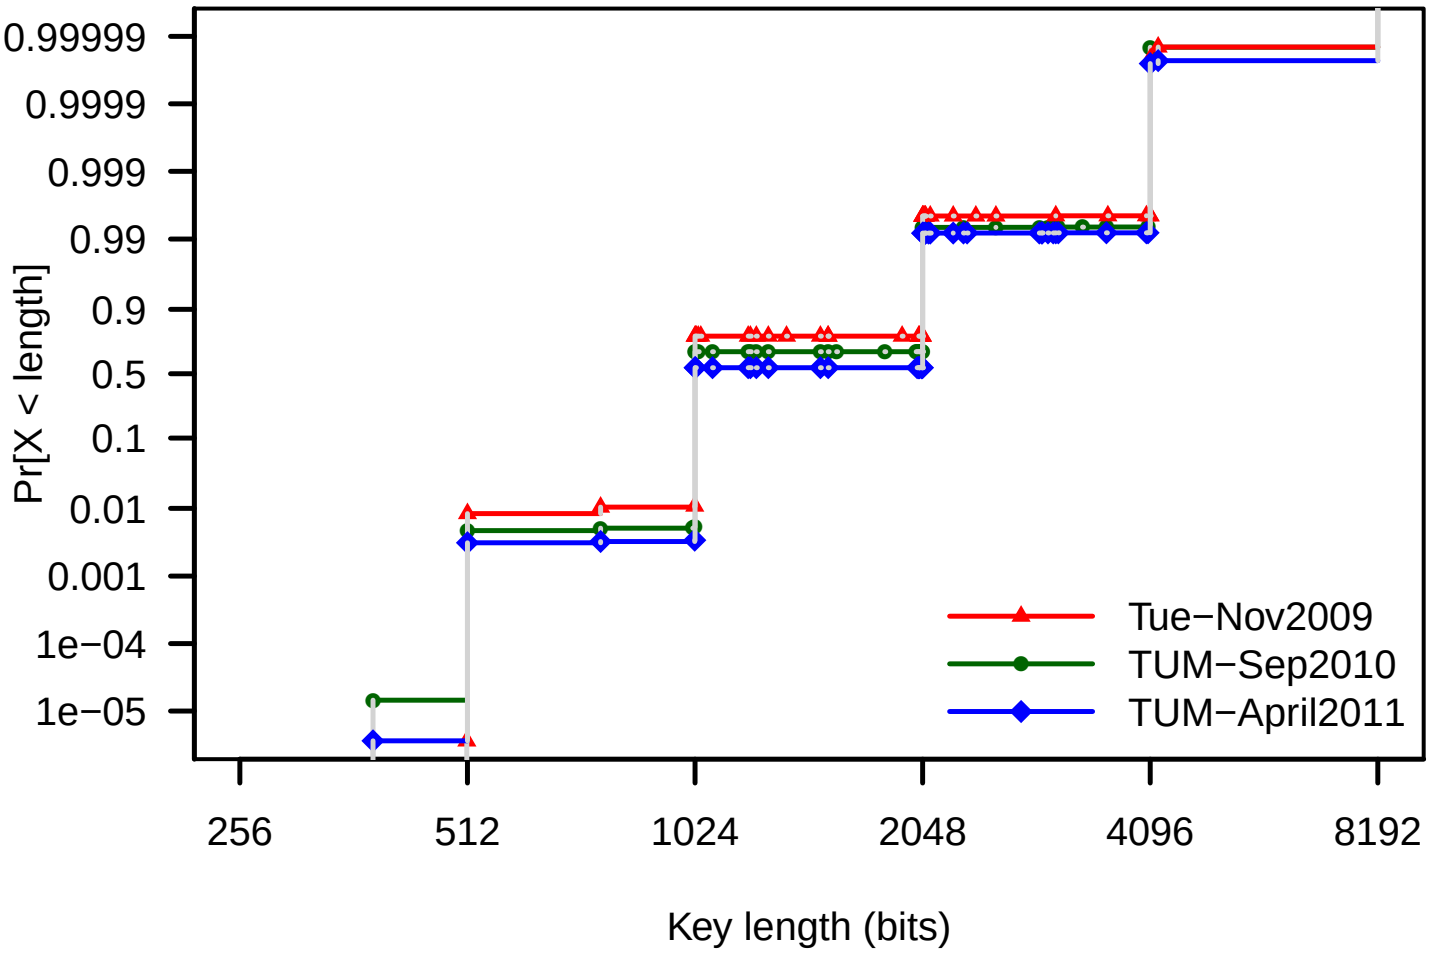

Debian Weak Keys (1)

Bug of 2008

- Generation of random numbers weak (bad initialisation)
- Only 2^{16} public/private key-pairs generated
- Allows pre-computation of private keys
- Debian ships blacklist of keys

Debian Weak Keys (2)

Weak randomness in key generation – serious bug of 2008

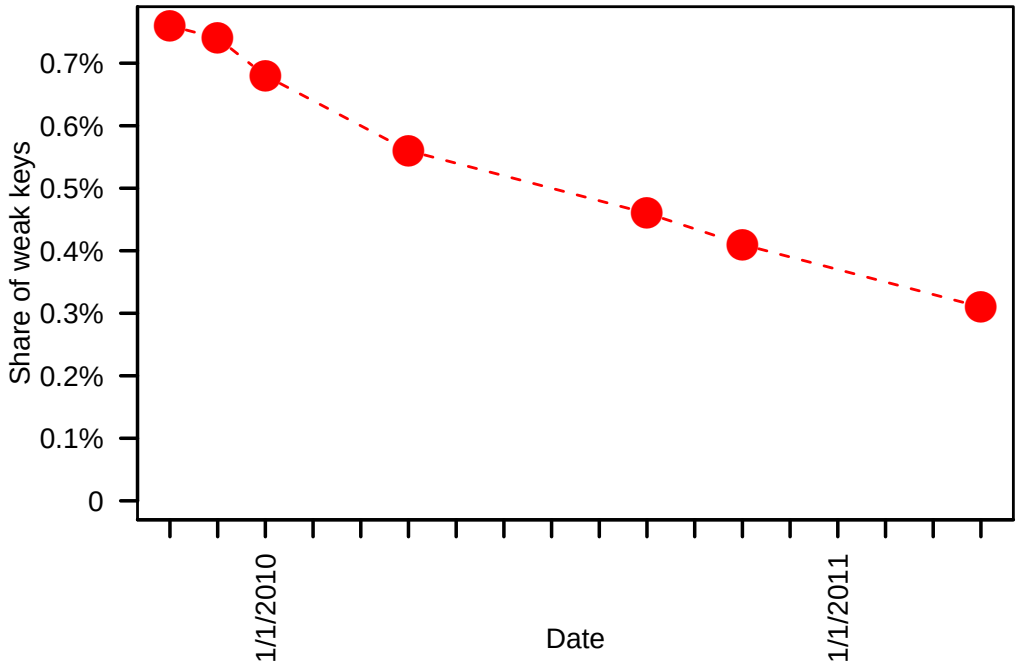

Work by colleagues

EFF scan presented at 27C3

- Focuses on CA certification structure
- Scan of IP addresses:
  does not allow to check match of host names
- No temporal distribution
- EFF project: SSL Observatory

Ivan Ristic of Qualys presents similar scan

- Smaller data basis
- Data set not published as raw data
- No temporal distribution
- Could not include it in our analysis