A First Look at the CT Landscape: Certificate Transparency Logs in Practice

Josef Gustafsson, *Linköping University*
Gustaf Overier, *Linköping University*
Martin Arlitt, University of Calgary, Canada
Niklas Carlsson, *Linköping University*

Proc. PAM, Sydney, Australia, Mar. 2017

LINKÖPING UNIVERSITY

Motivation and high-level problem

- Private and confidential communication important
  -
  -

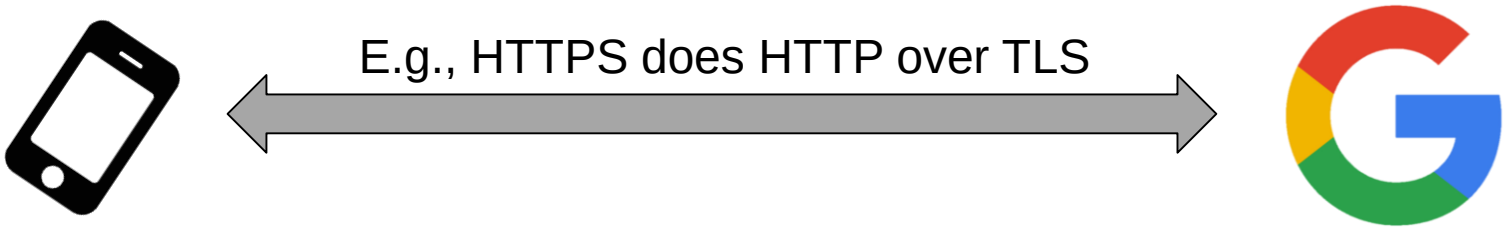

Motivation and high-level problem

- Private and confidential communication important
  -
  -

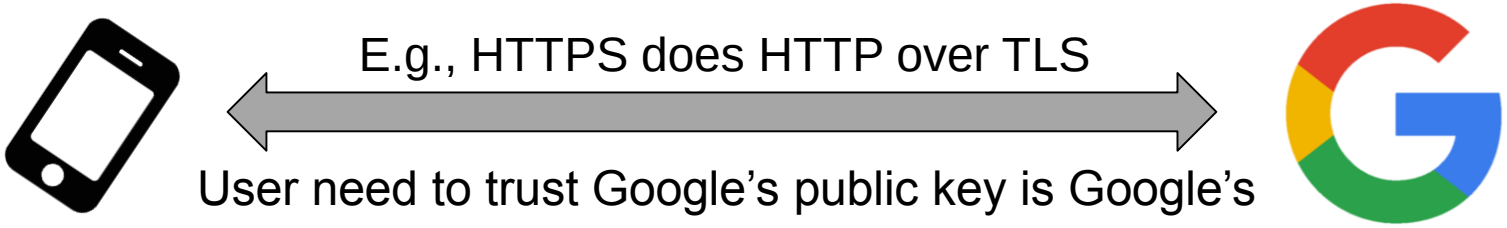

Motivation and high-level problem

- Private and confidential communication important
  - Billions of devices
  - Millions of services
-
  -

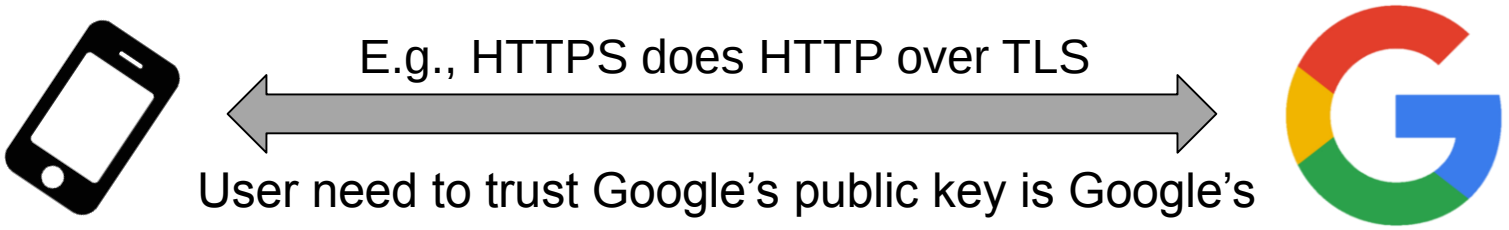

Motivation and high-level problem

- Private and confidential communication important
  - Billions of devices
  - Millions of services
-
  -

User need to trust FB's public key is FBs

E.g., HTTPS does HTTP over TLS

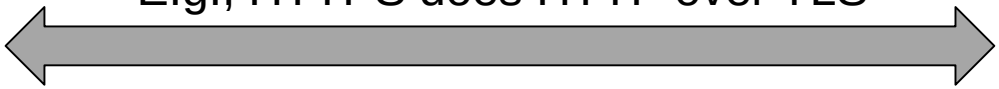

User need to trust Google's public key is Google's

Motivation and high-level problem

- Private and confidential communication important
  - Billions of devices
  - Millions of services
-
  -

User need to **trust** FB's public key is FB's

E.g., HTTPS does HTTP over TLS

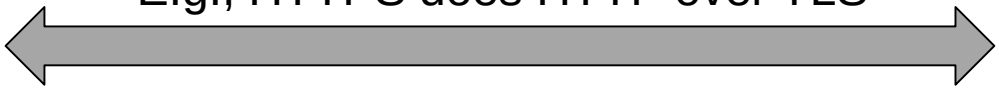

User need to **trust** Google's public key is Google's

Motivation and high-level problem

- Private and confidential communication important
  - Billions of devices
  - Millions of services
- Certification Authorities (CAs) issue certificates
  - Proof of identity (signed with their private key)

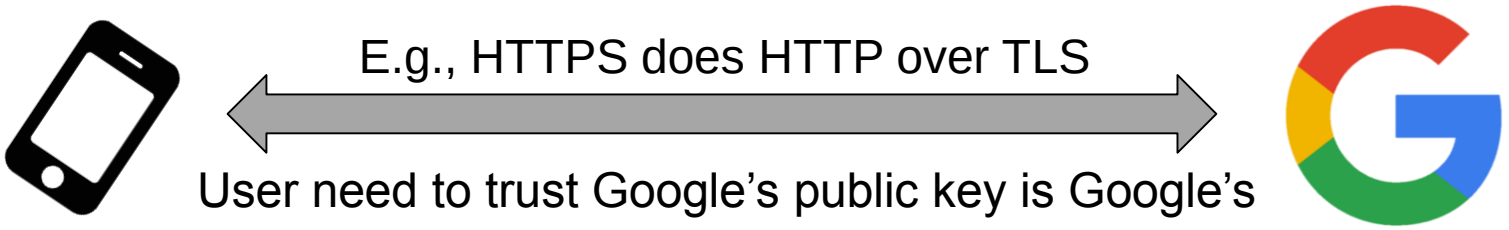

Motivation and high-level problem

- If CAs in our trust (root) store (e.g., Symantec/ Verisign) tells us that a public key belongs to Google, our browsers (and us) trust that this is the case

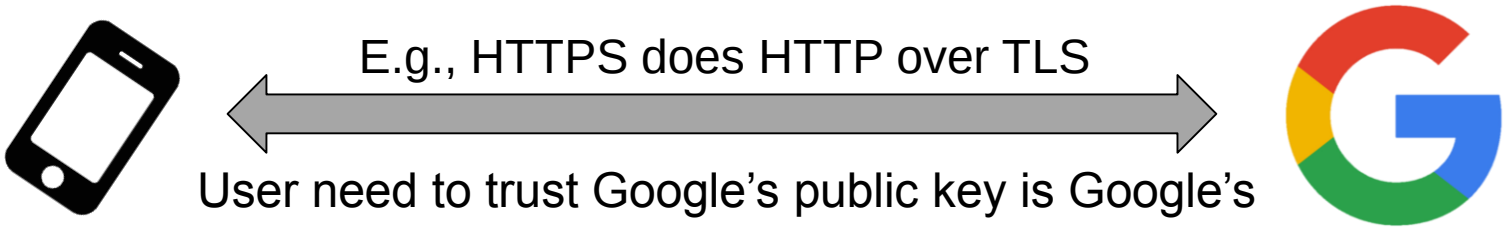

Motivation and high-level problem

- If CAs in our trust (root) store (e.g., Symantec/ Verisign) tells us that a public key belongs to Google, our browsers (and us) trust that this is the case

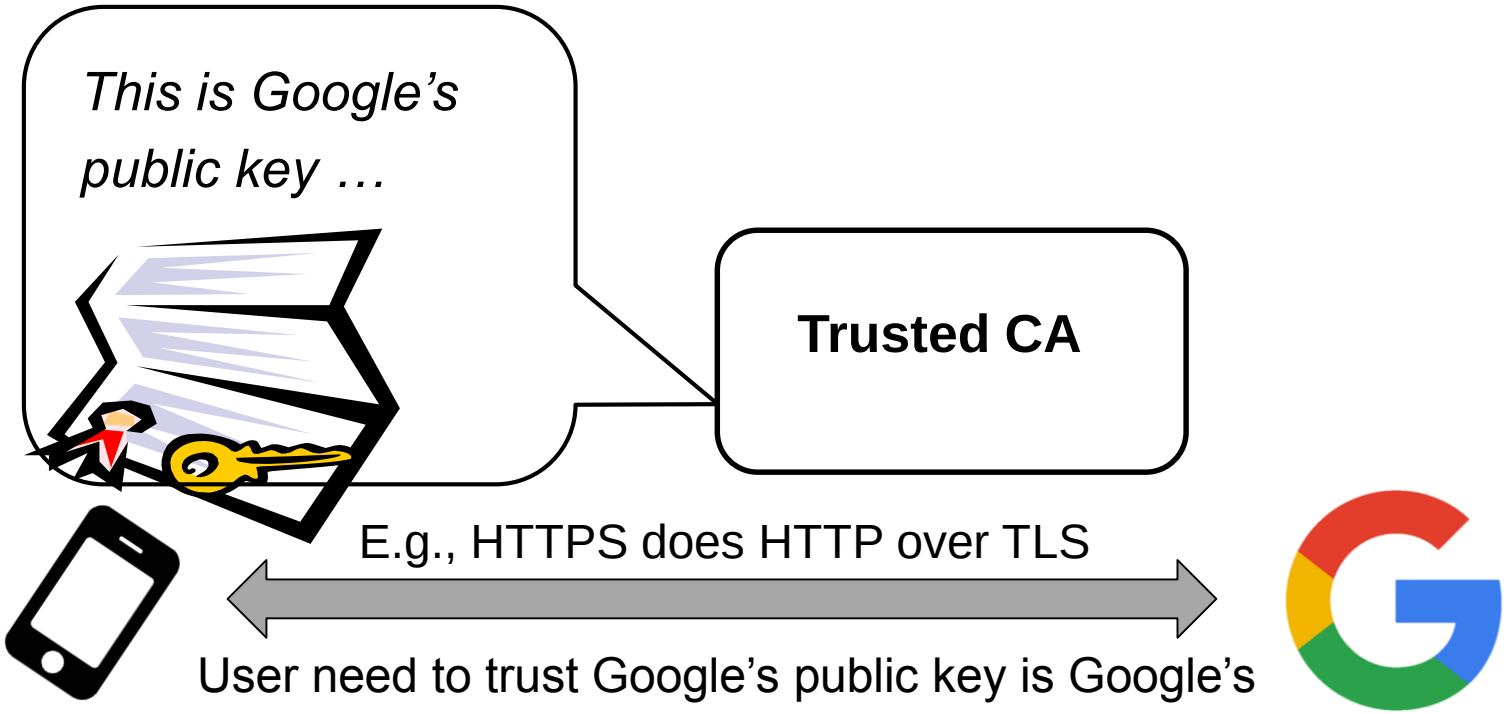

Motivation and high-level problem

- However, mistakes happen ...
  - E.g., in Oct. 2015, Google discovered (using CT) that Symantec had issued test certificates for 76 domains that they did not own (including Google domains) and another 2,458 unregistered domains …

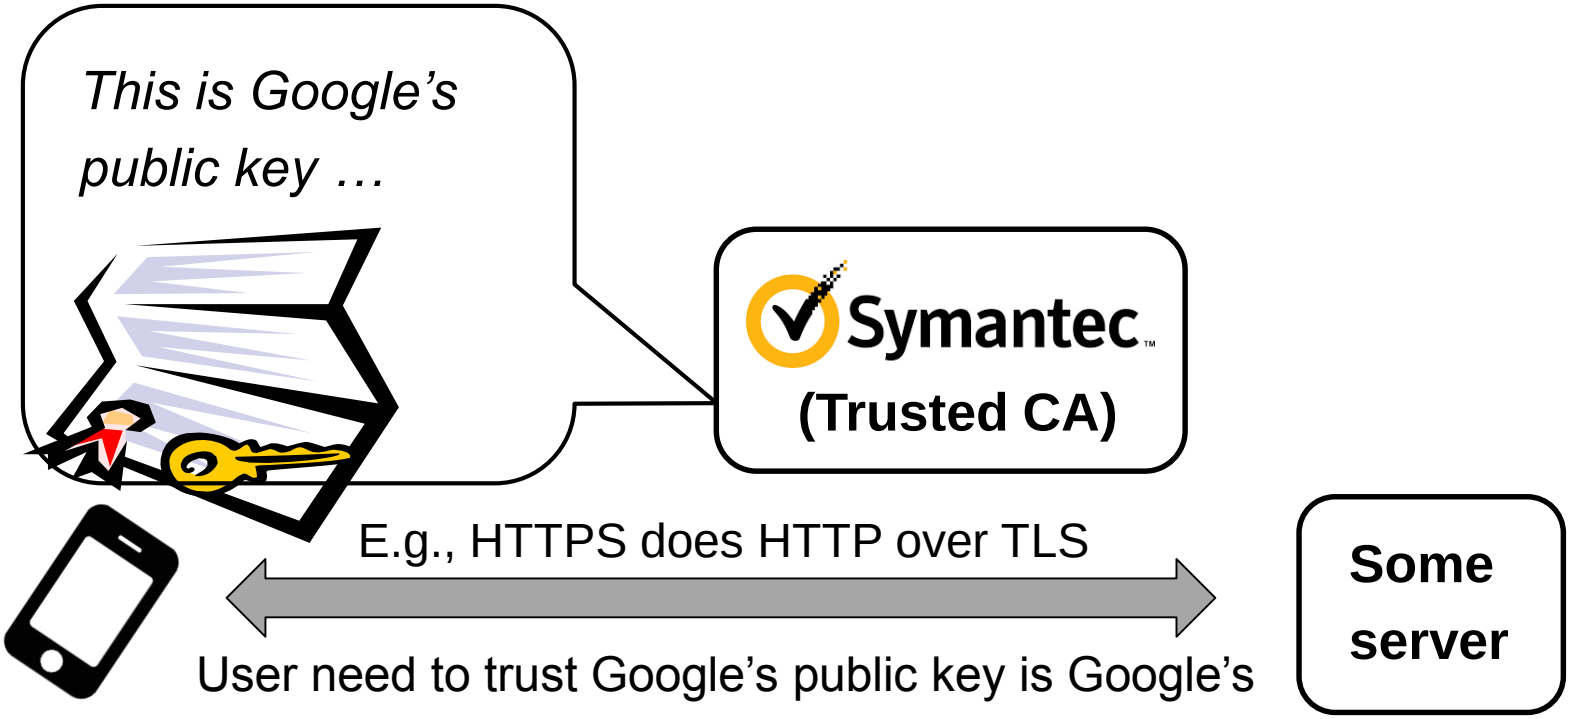

CT: Emerging trust-monitoring solution

- Since then, Google has demanded that Symantec logs all their certificates in public (append-only) CT logs
- Since Jan. 2015, the Chrome browser requires all EV certificates be logged in 1 Google log and 1 other log
  - Mozilla planning to make similar demands
  - Both Chrome and Mozilla expected policies to DV certificates too …

CT: Emerging trust-monitoring solution

- Since then, Google has demanded that Symantec logs all their certificates in public (append-only) CT logs
- Since Jan. 2015, the Chrome browser requires all EV certificates be logged in 1 Google log and 1 other log
  - Mozilla planning to make similar demands
  - Both Chrome and Mozilla expected policies to DV certificates too …
- In this paper, we present the first large-scale characterization of the CT landscape

Certification of public keys

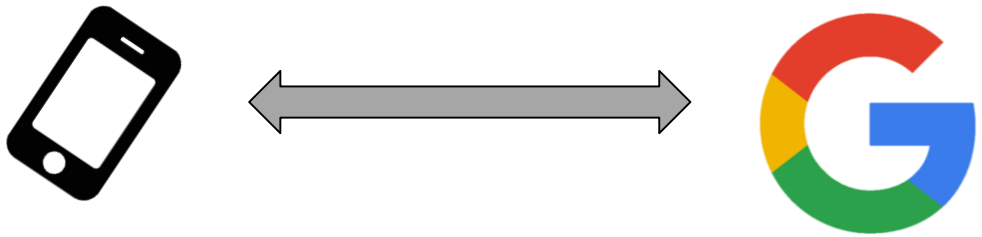

Certification of public keys

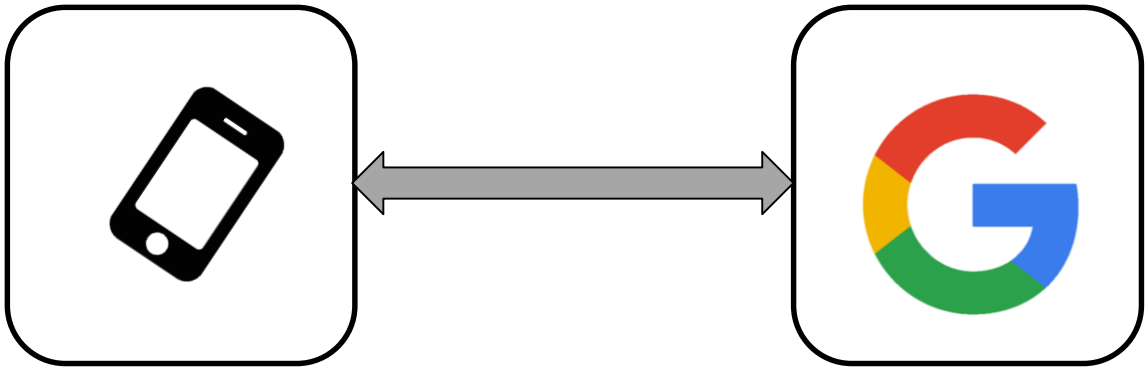

Certification of public keys

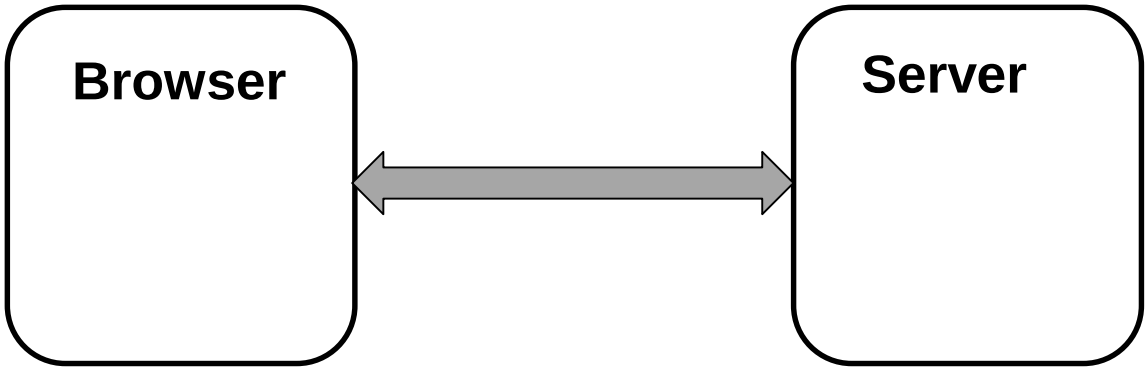

Certification of public keys

- Browsers have trust stores with root certs (of CAs)

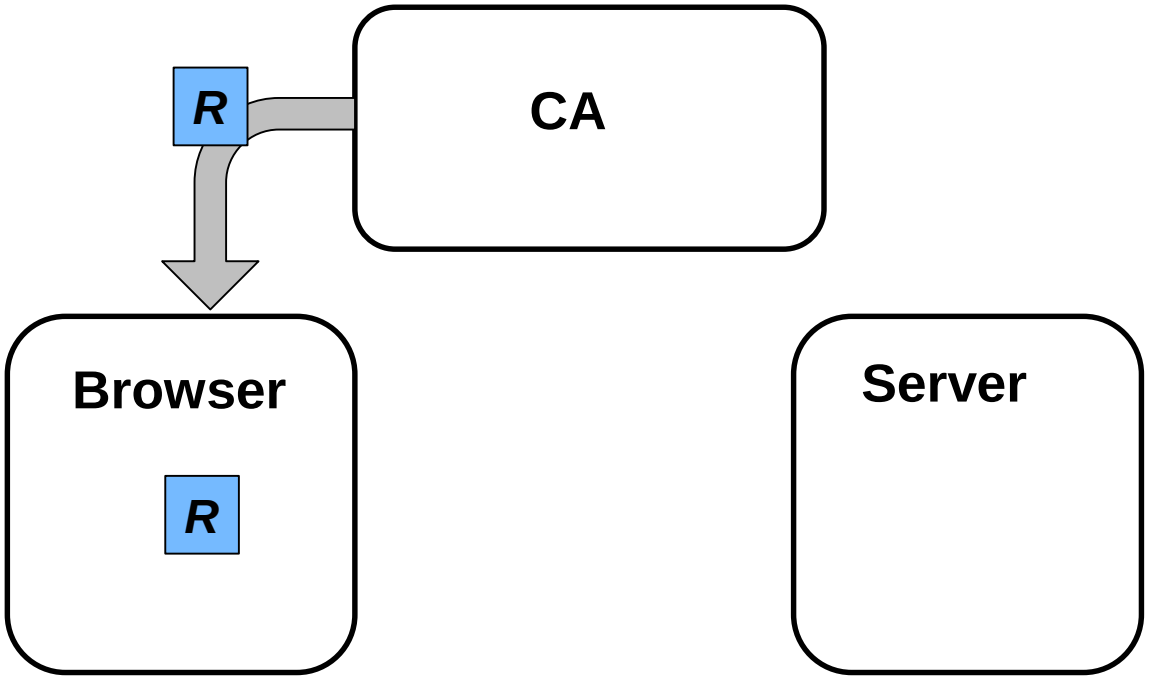

Certification of public keys

- Browsers have trust stores with root certs (of CAs)

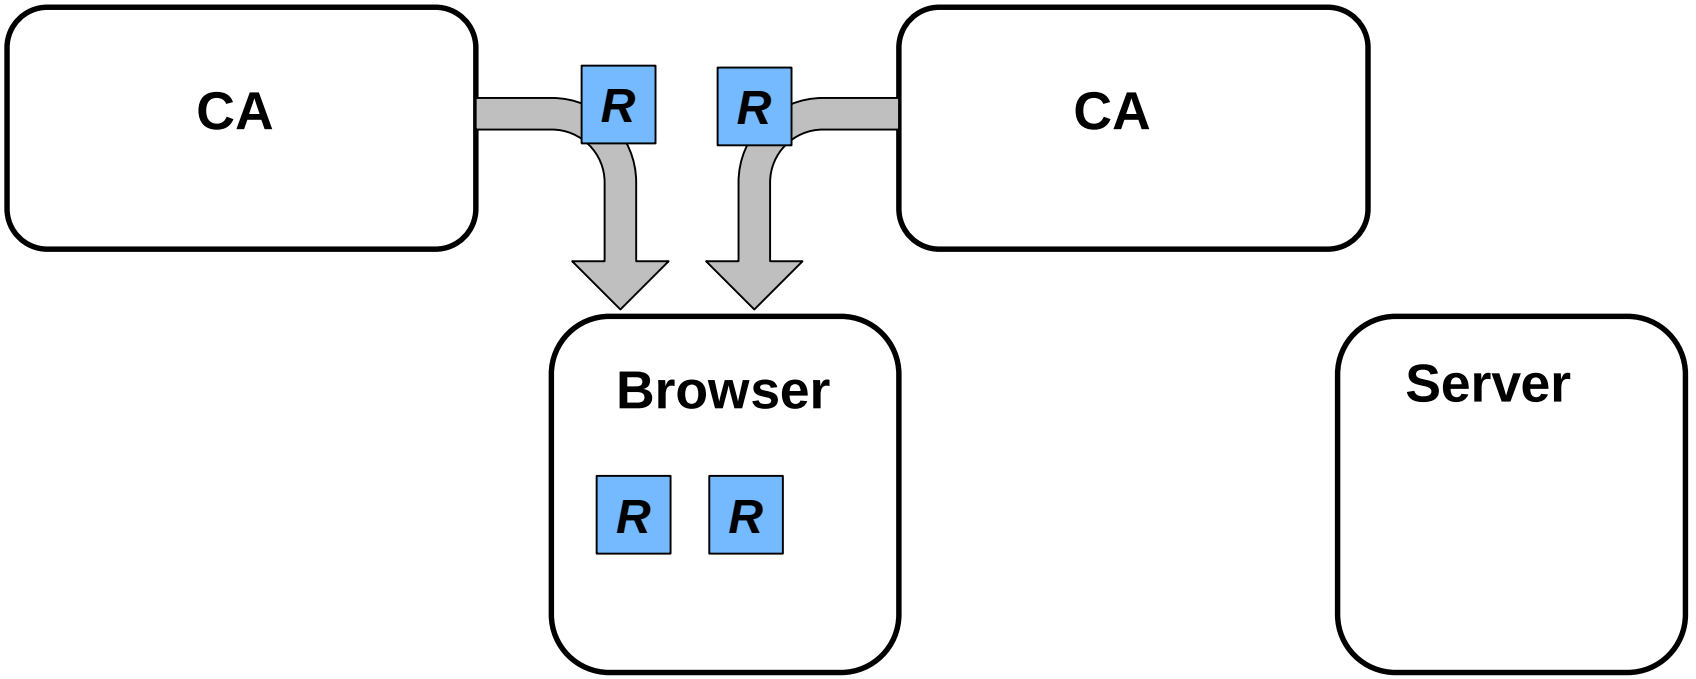

Certification of public keys

- Browsers have trust stores with root certs (of CAs)

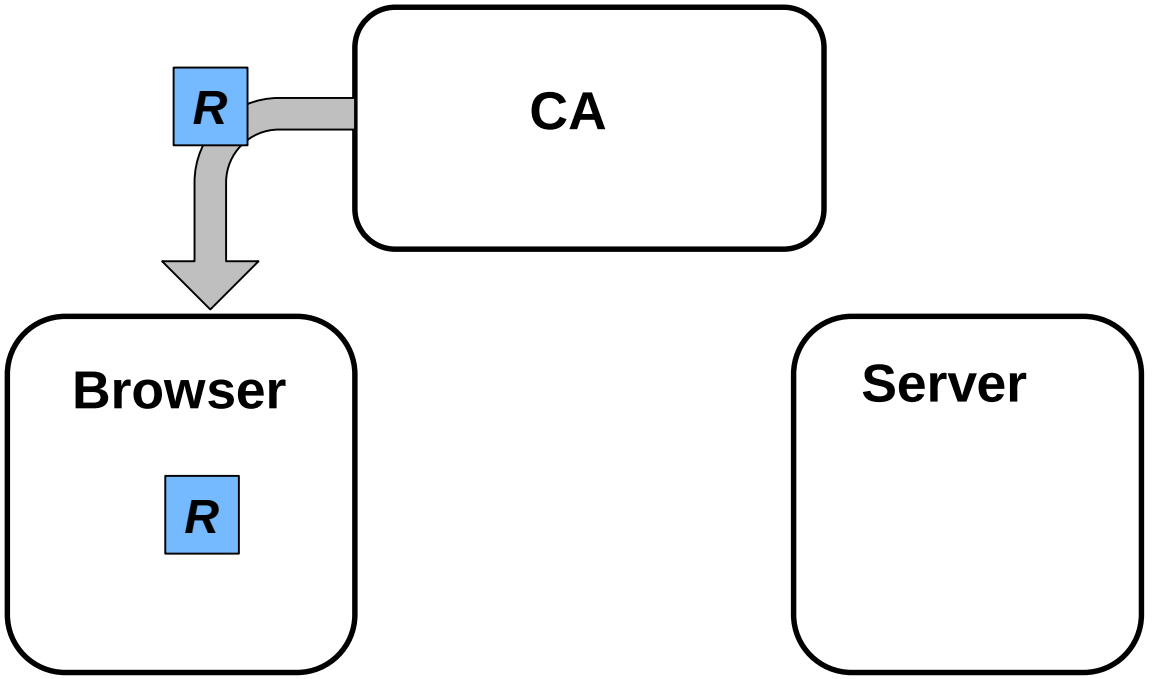

Certification of public keys

- Browsers have trust stores with root certs (of CAs)
- CAs use private key to sign certs for servers/domains
  - Certs are proof that public key belongs to server/domain

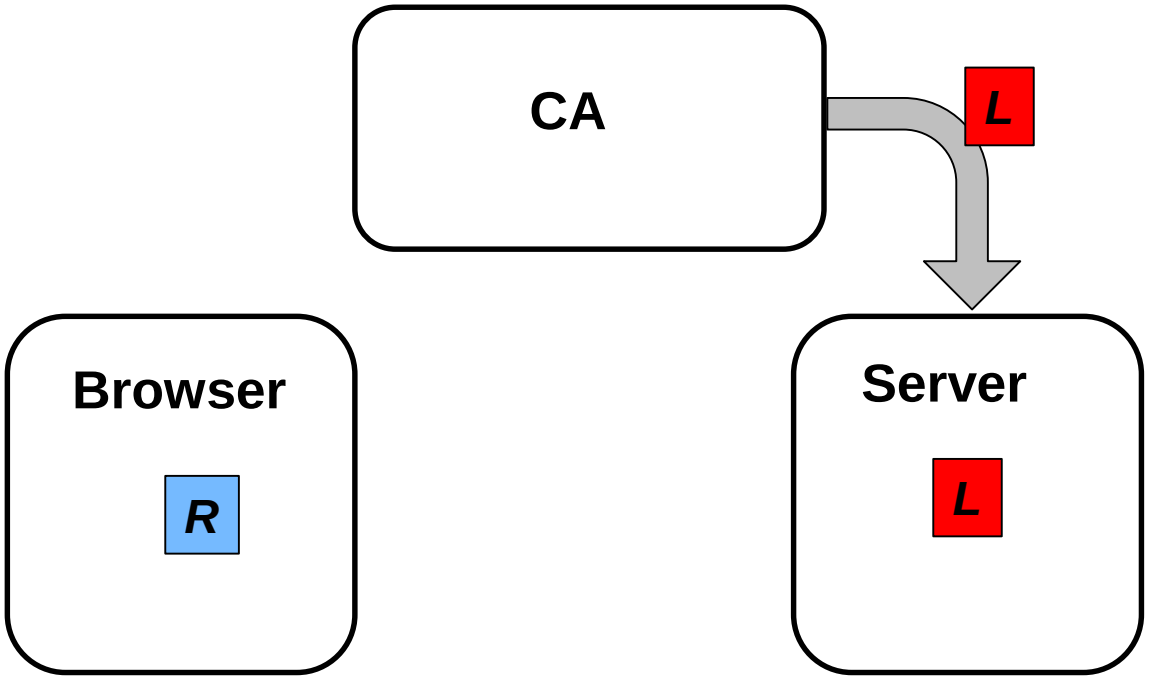

Certification of public keys

- Browsers have trust stores with root certs (of CAs)
- CAs use private key to sign certs for servers/domains
  - Certs are proof that public key belongs to server/domain
  - Signature of certs can be validated using keys in root store

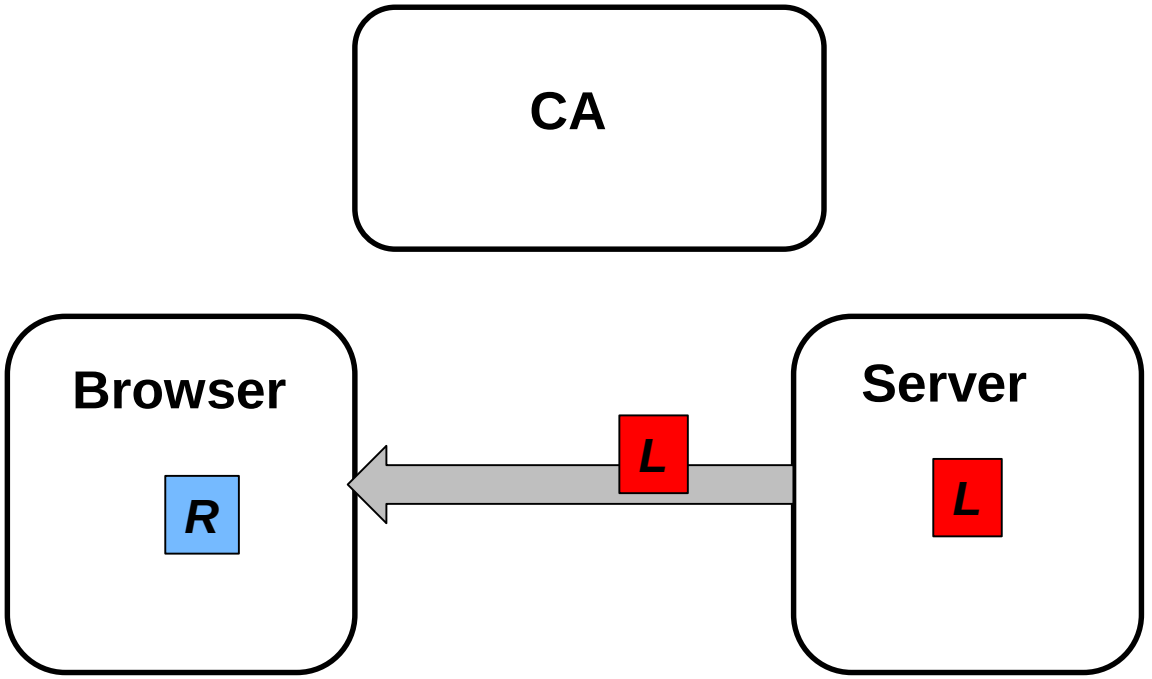

Certification of public keys

- Browsers have trust stores with root certs (of CAs)
- CAs use private key to sign certs for servers/domains
  - Certs are proof that public key belongs to server/domain
  - Signature of certs can be validated using keys in root store

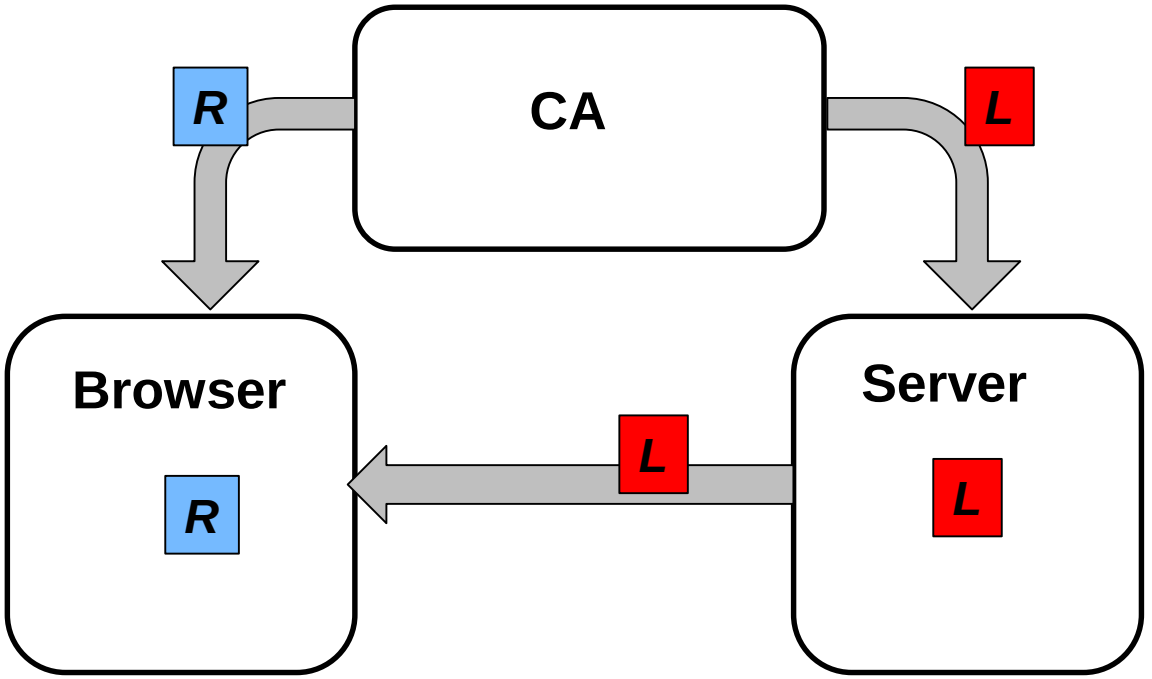

Certification of public keys

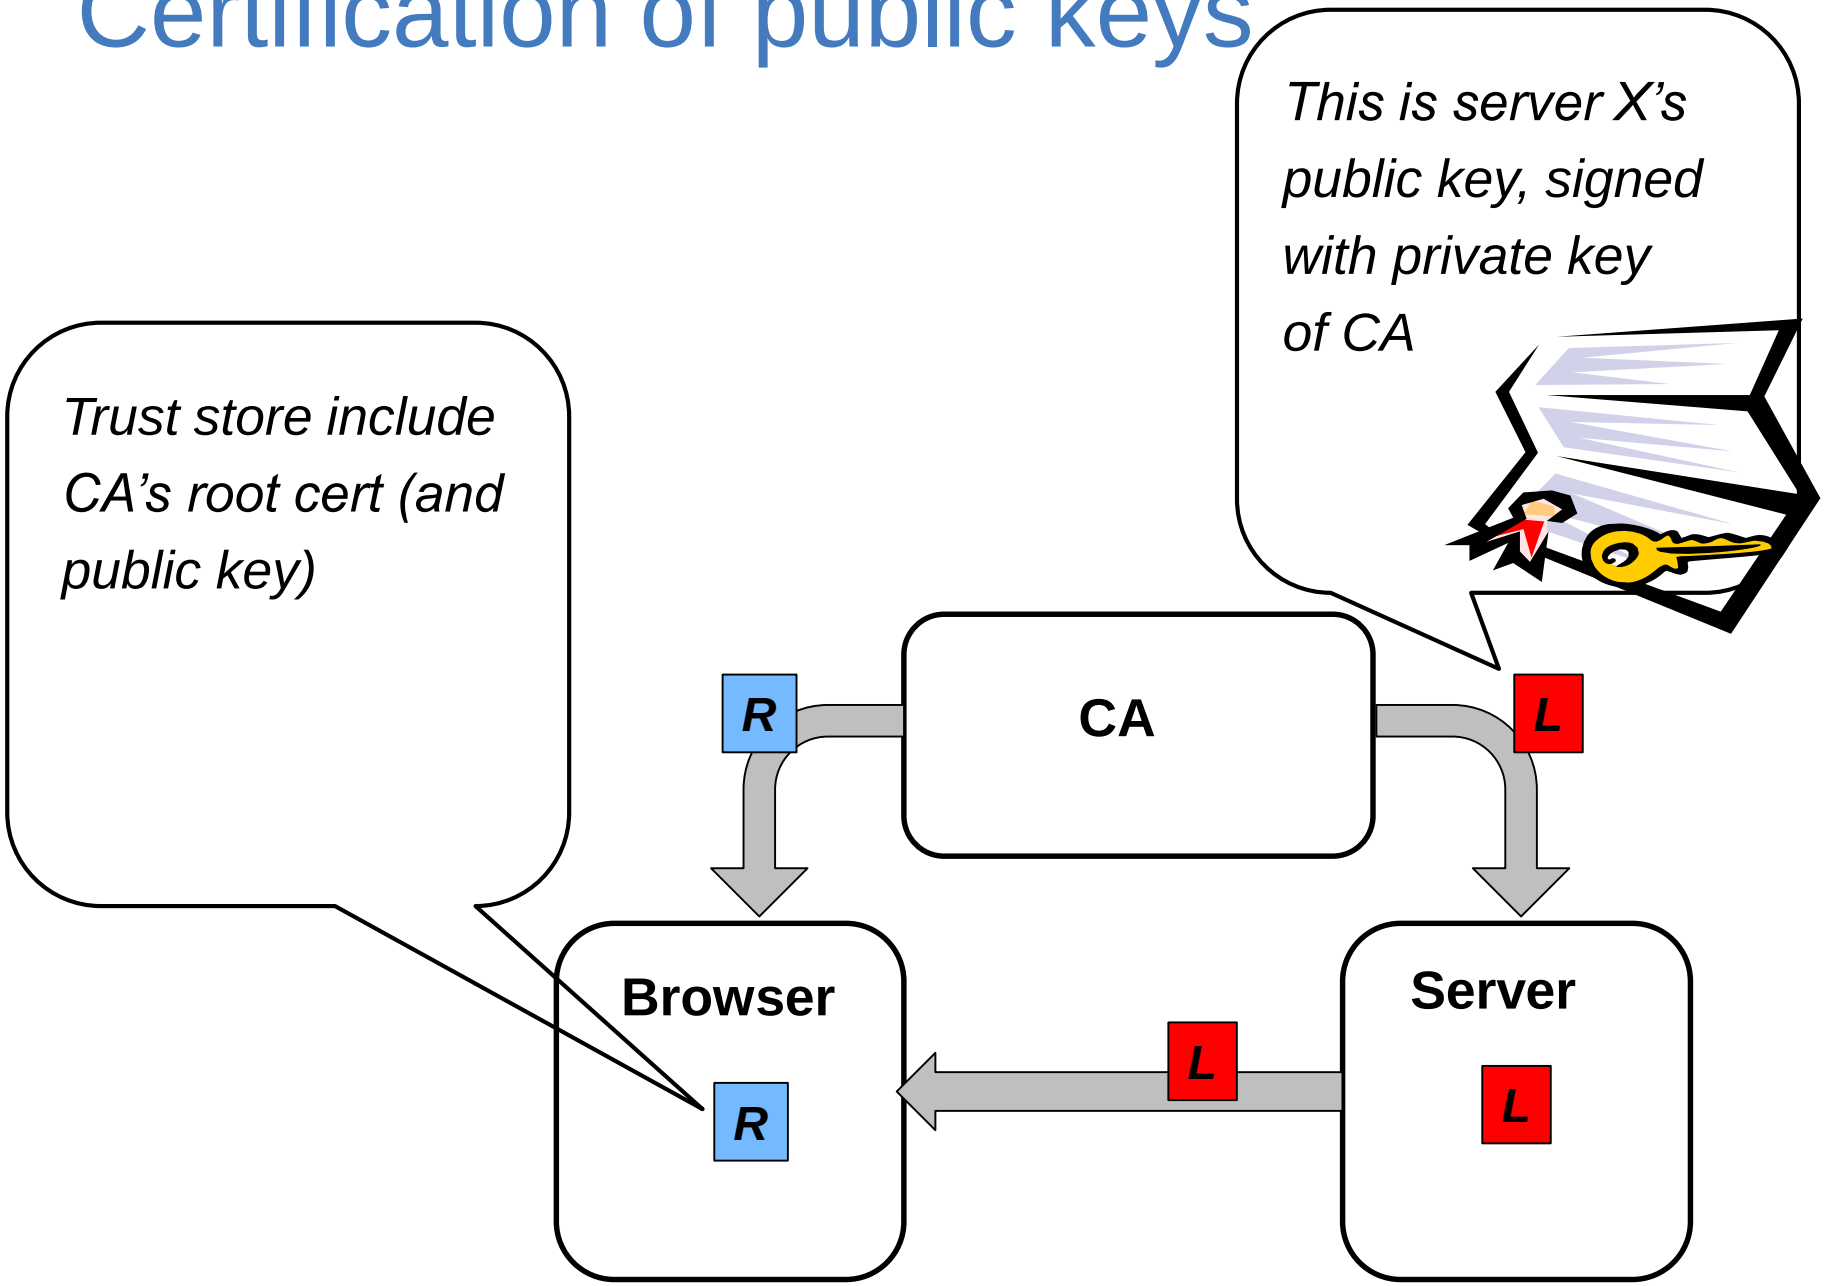

Certification of public keys

- Browsers have trust stores with root certs (of CAs)
- CAs use private key to sign certs for servers/domains
  - Certs are proof that public key belongs to server/domain
  - Signature of certs can be validated using keys in root store
- In practice, many
  - Many CAs, servers
  - Varying trust+security

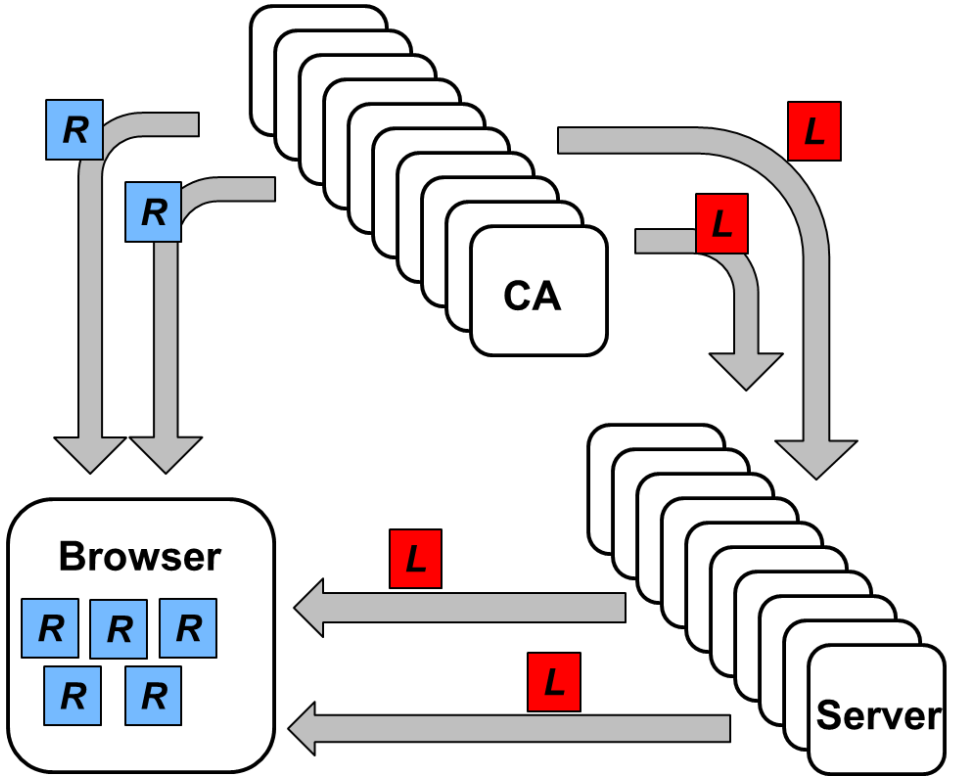

Certification of public keys

- Browsers have trust stores with root certs (of CAs)
- CAs use private key to sign certs for servers/domains
  - Certs are proof that public key belongs to server/domain
  - Signature of certs can be validated using keys in root store
- In practice, many
  - Many CAs, servers
  - Varying trust+security
- Trust can be undermined
  - Human error
  - Intentional fraud
  - Compromised CAs
  - …

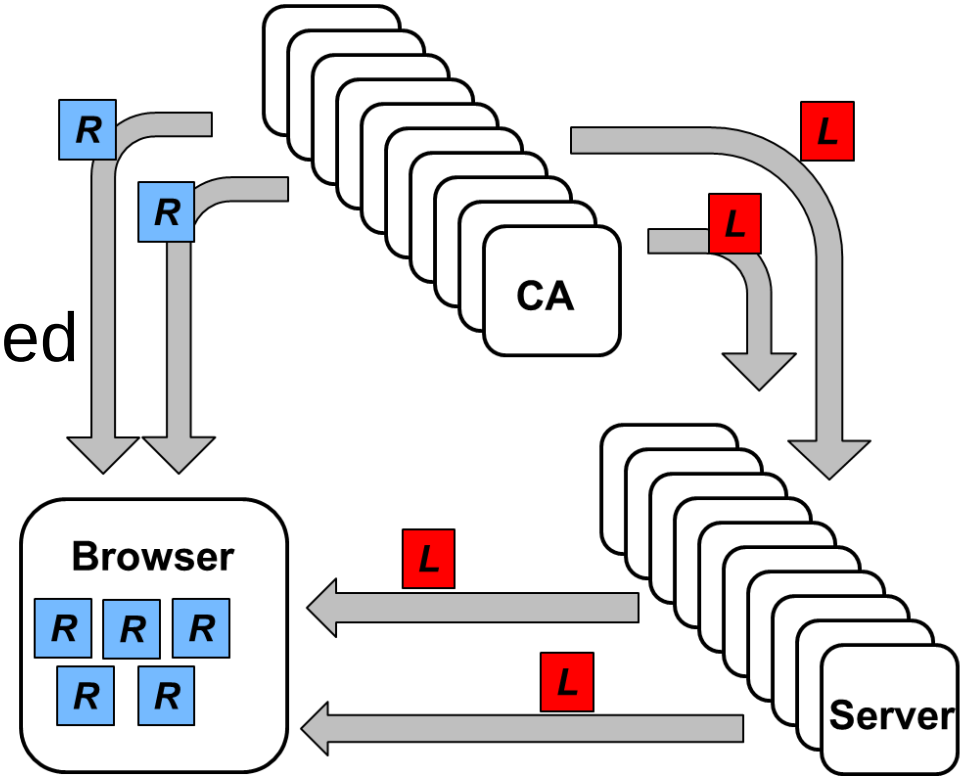

Trust landscape

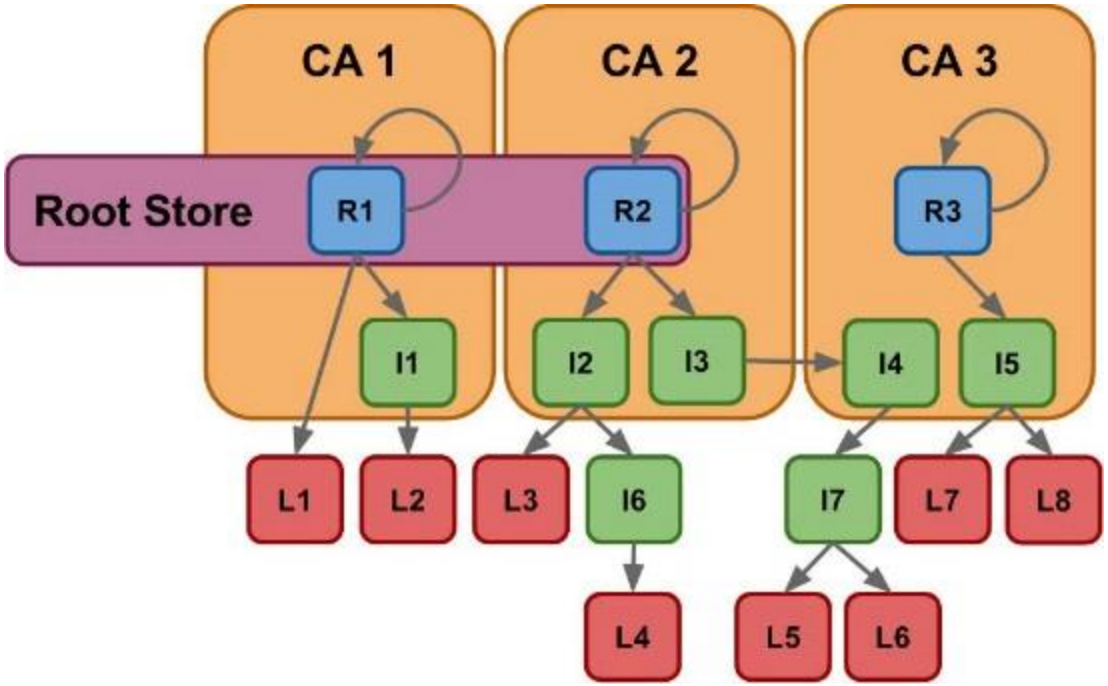

- Delegation of trust to intermediates (Ii)
- Browsers trust that the servers that can present certs (Li) that map to (trusted) root certs are who they claim to be
- Impersonation
  - Any trusted CA (Ri) or intermediate (Ii) can issue rogue certs
  - Very difficult to know all certs issued in ones name

Certification Transparency (CT)

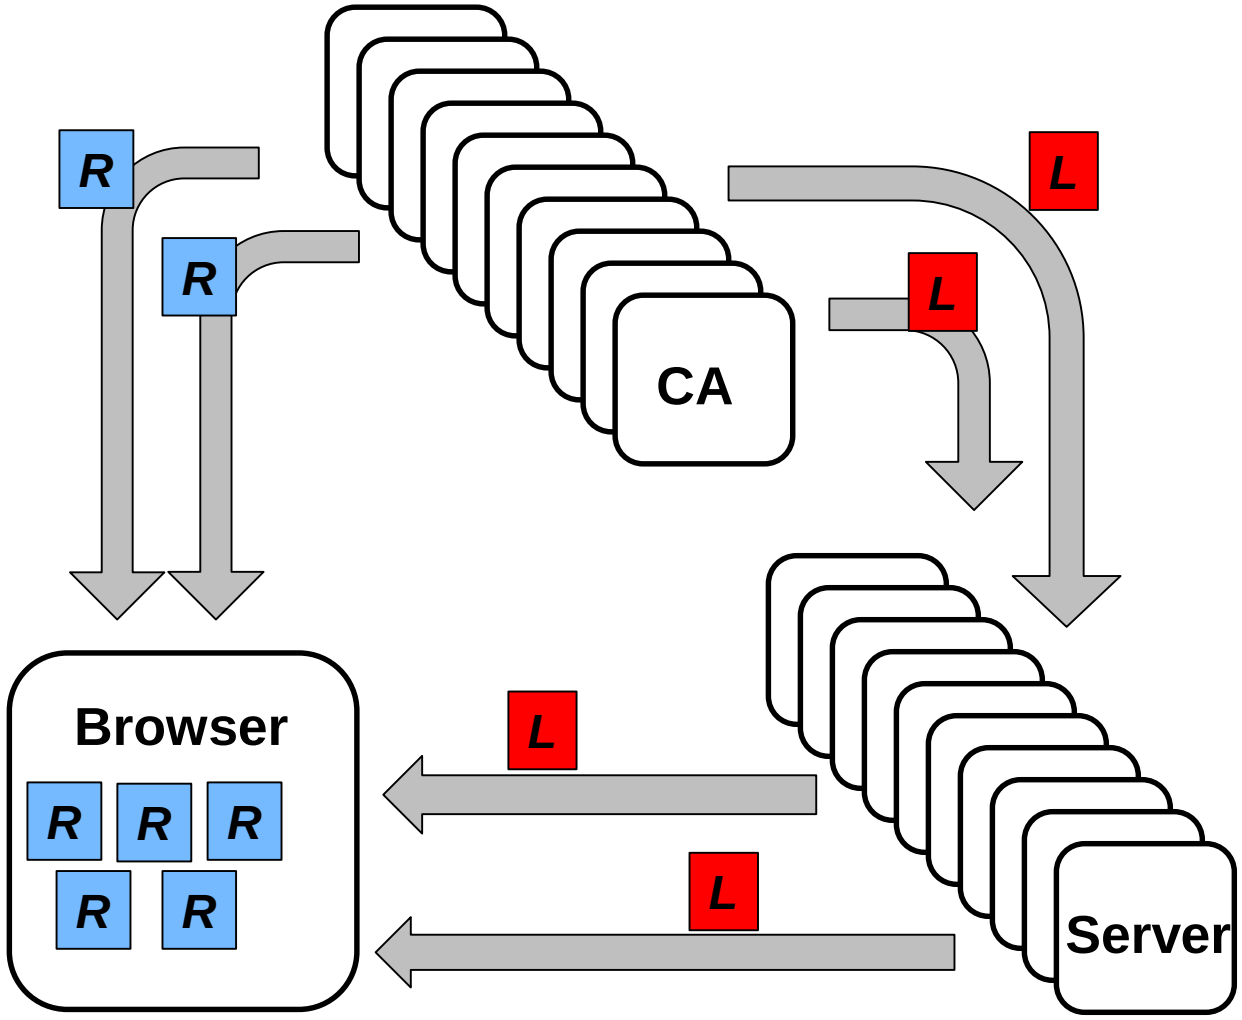

Certification Transparency (CT)

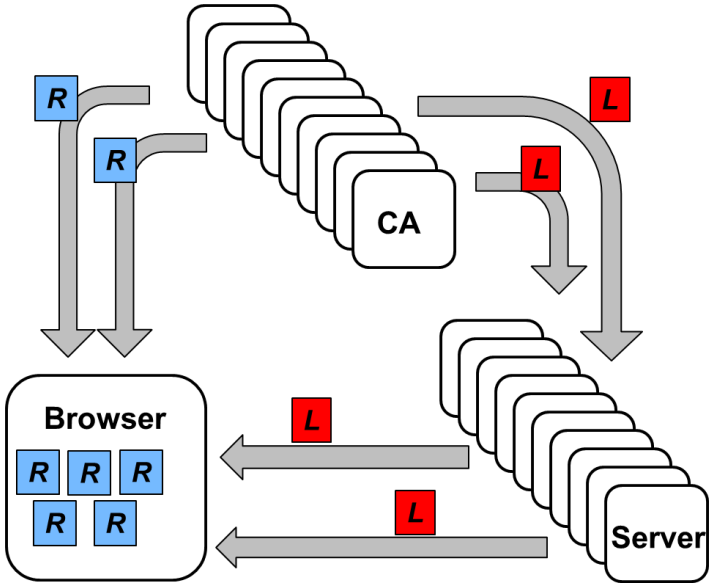

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

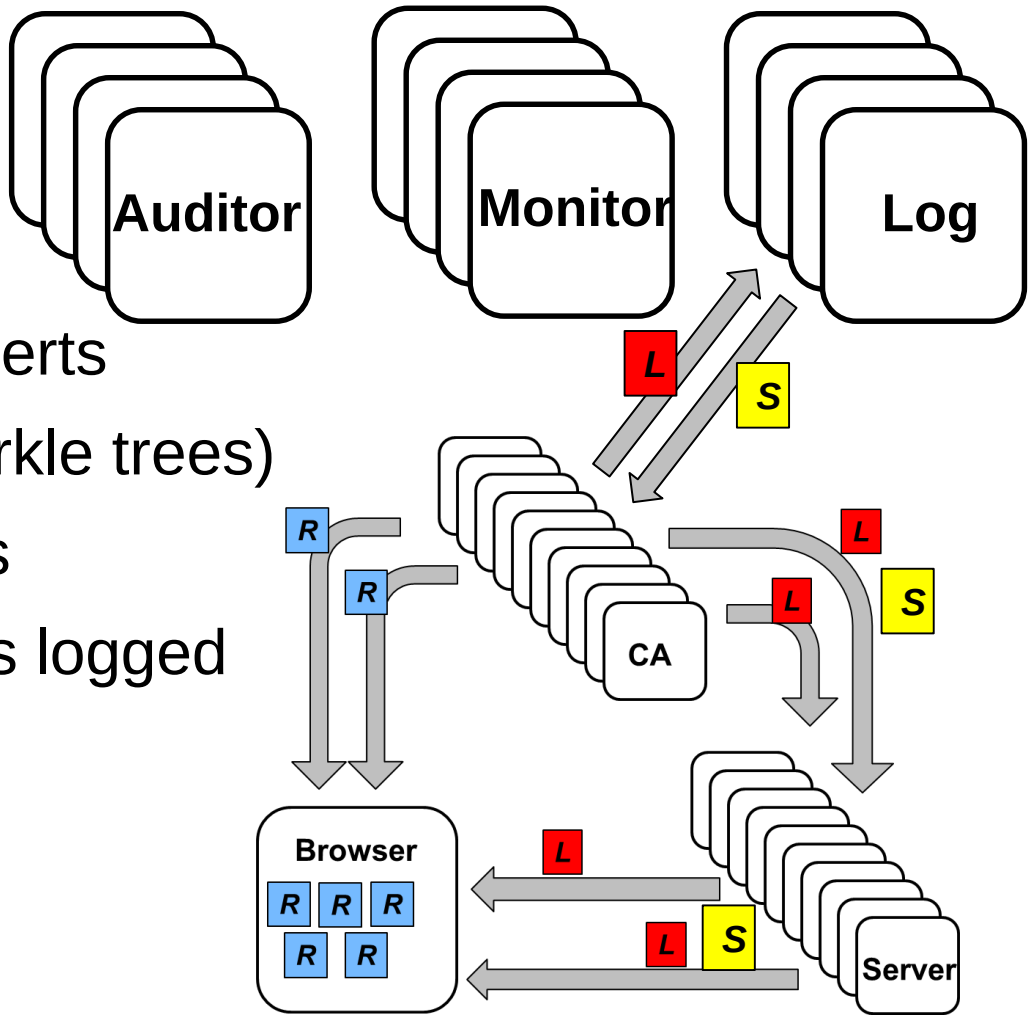

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

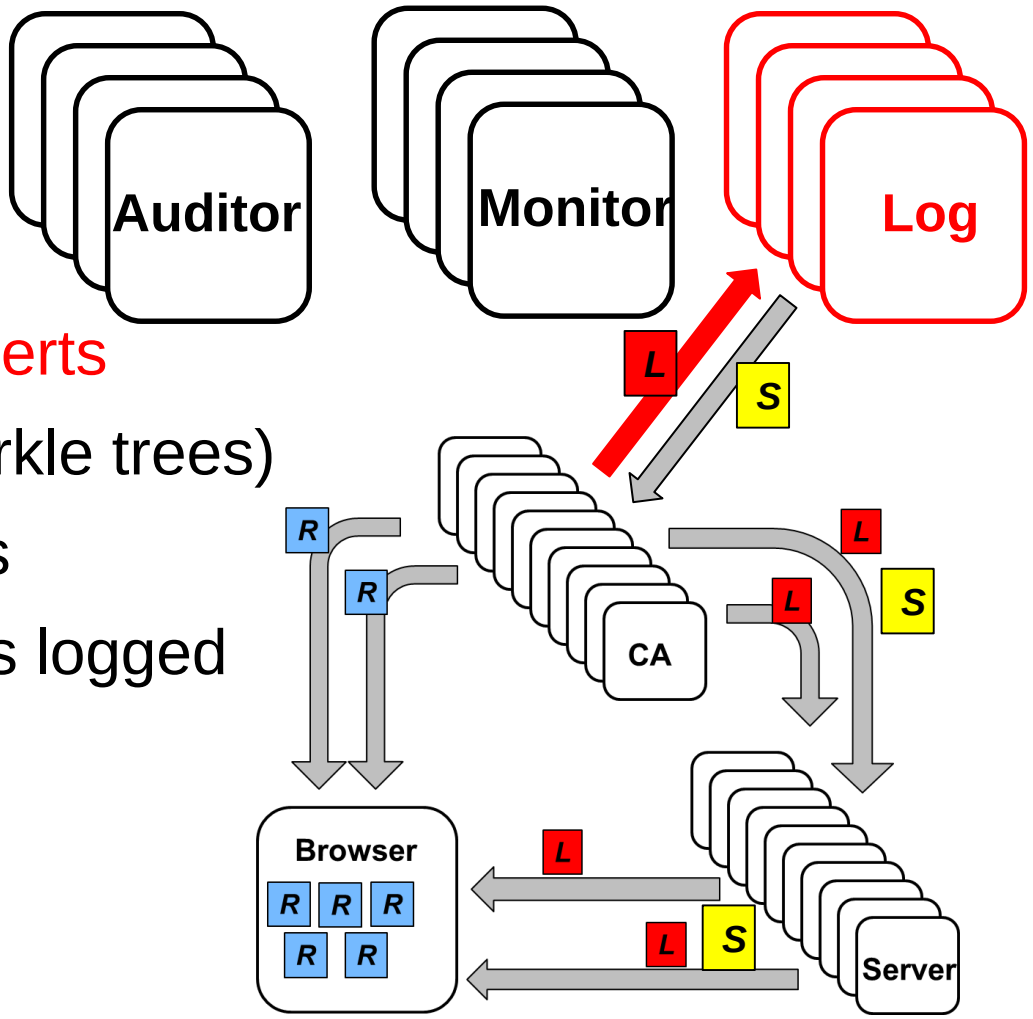

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

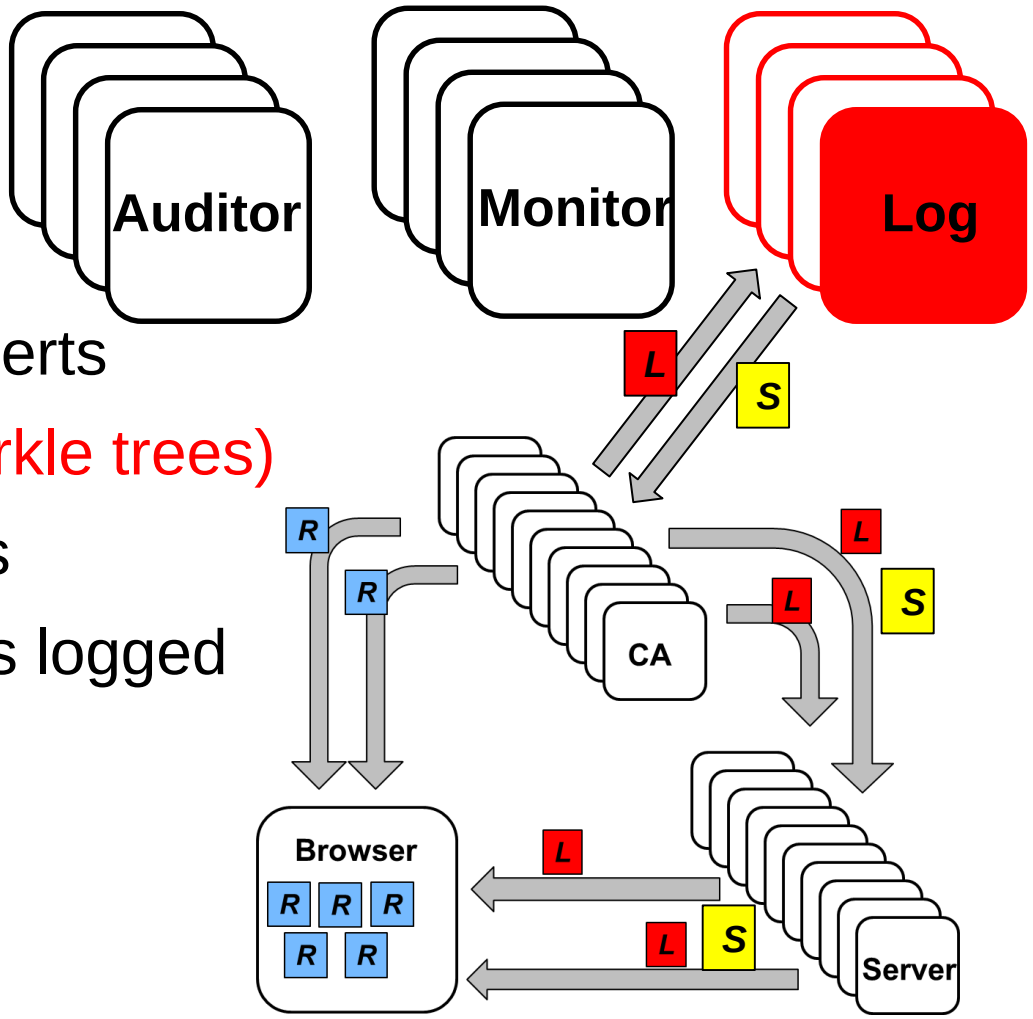

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

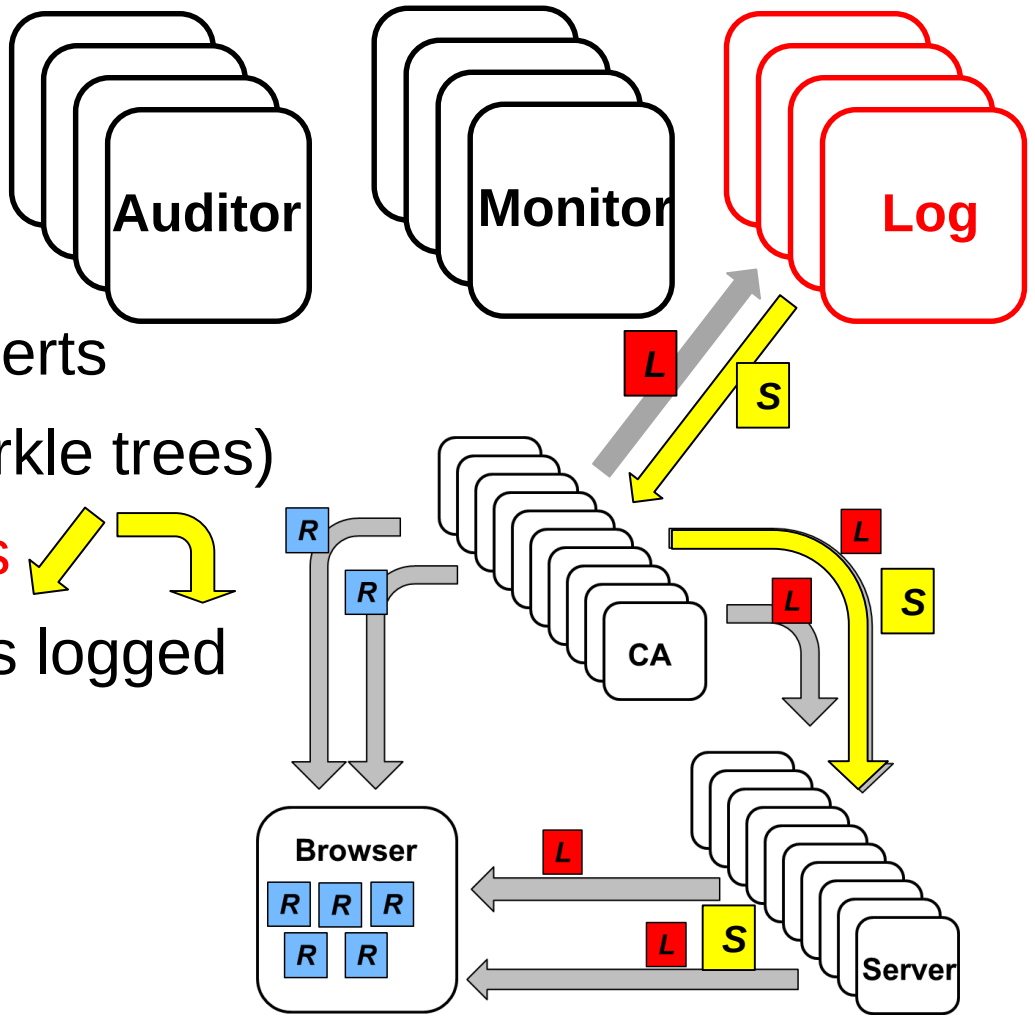

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

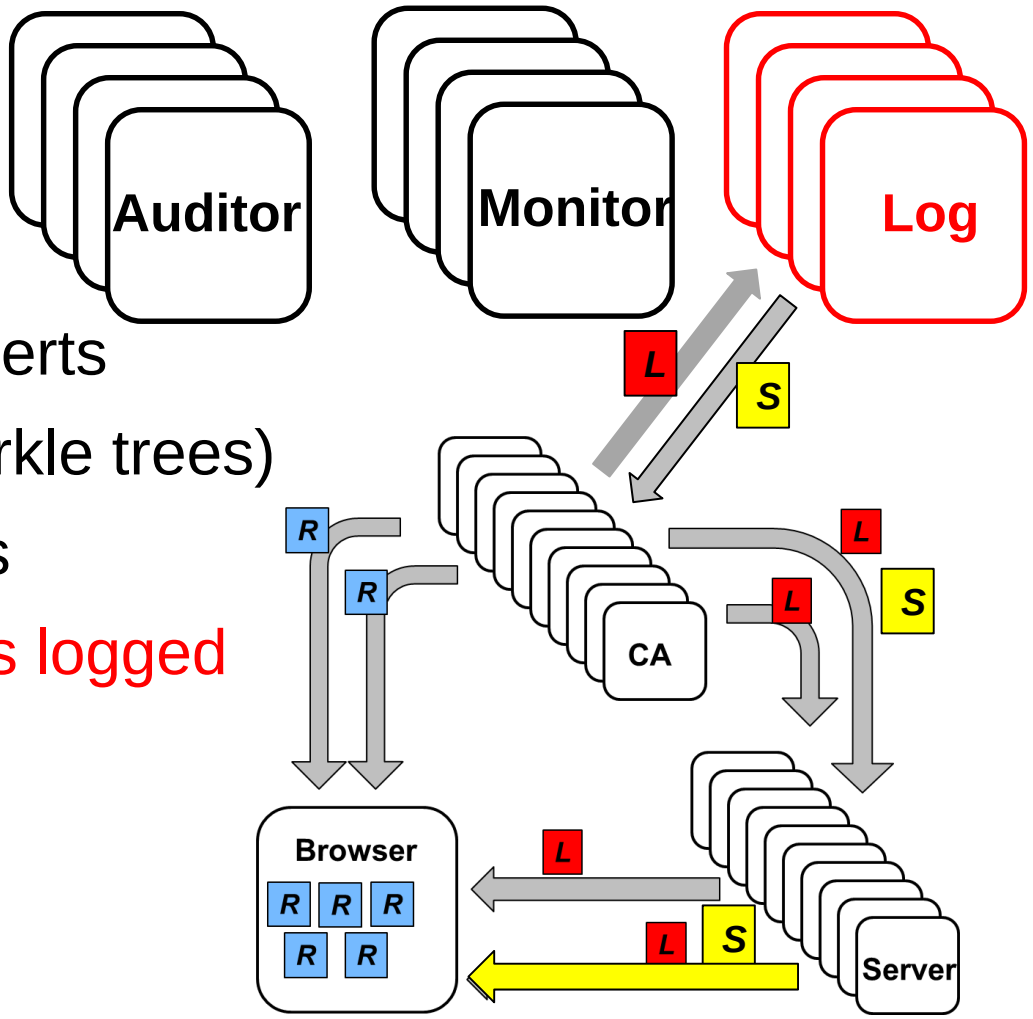

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

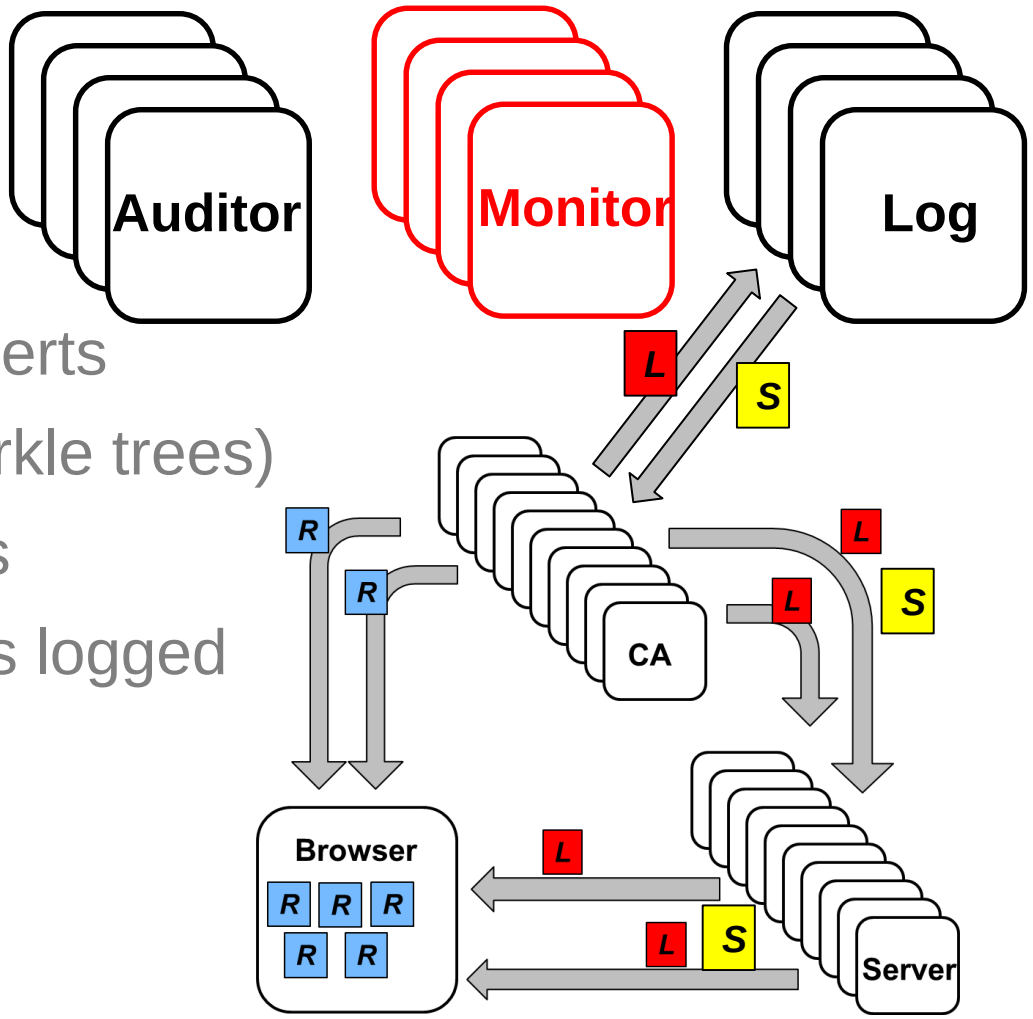

Certification Transparency (CT)

- Logs
  - Public record of certs
  - Append only (Merkle trees)
  - Servers get SCTs
  - SCTs proof cert is logged
- Monitors
  - Assert log content
- Auditors
  - Assert log behavior

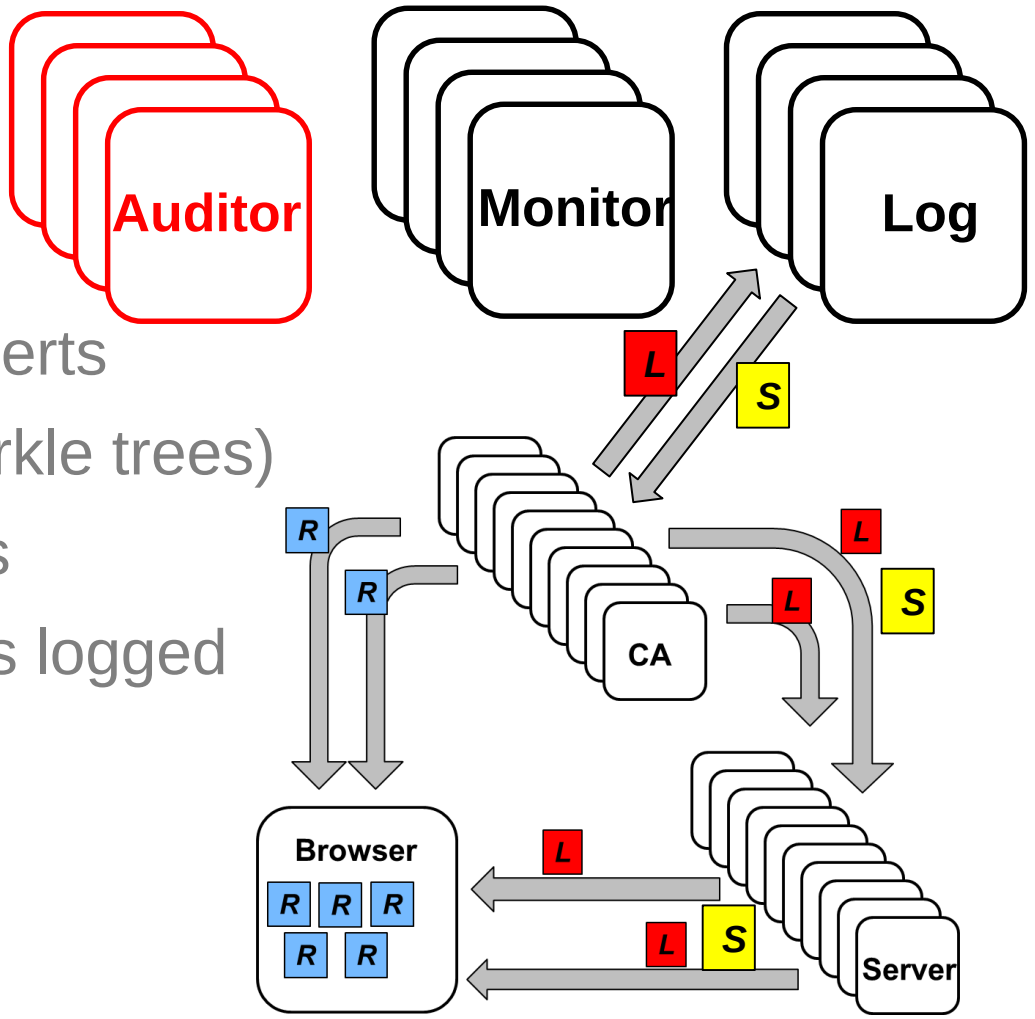

Methodology

Methodology

- Created CT monitor
- Monitored all public logs
  - 3 Google
  - 7 CA-based
  - Plausible (NORDUnet)
- Campus measurements
  - All HTTPS sessions for a week
  - 232 million HTTPS sessions

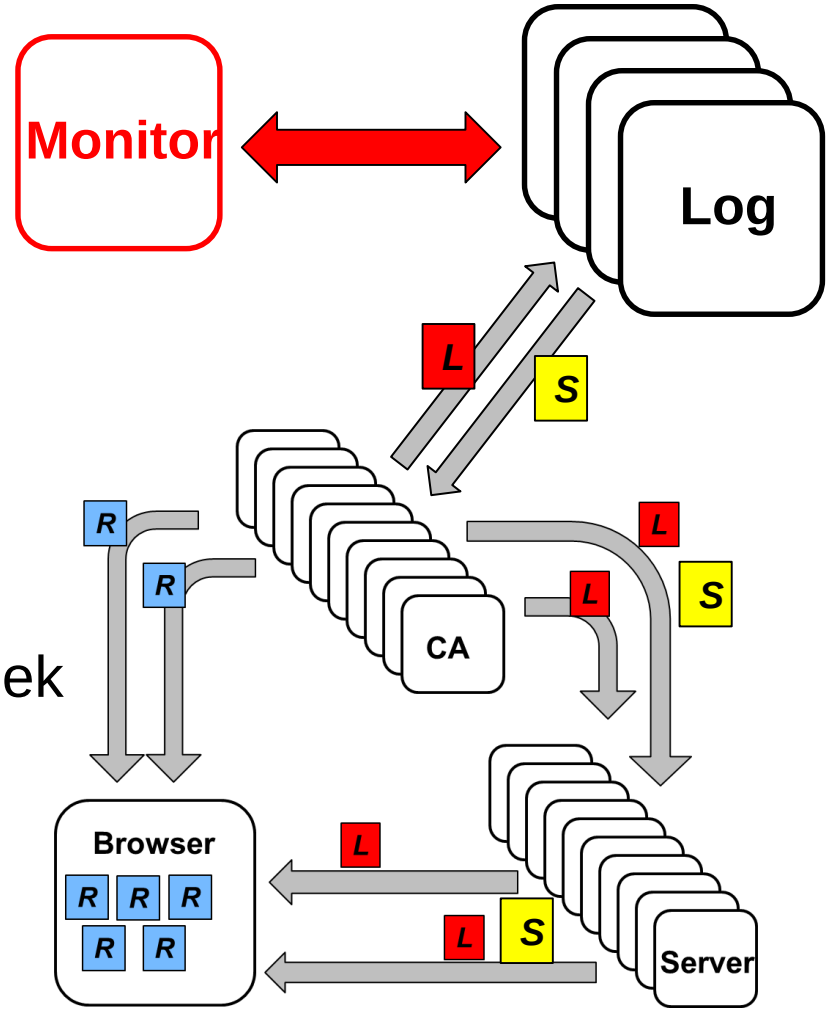

Methodology

- Created CT monitor
- Monitored all public logs
  - 3 Google
  - 7 CA-based
  - Plausible (NORDUnet)
- Campus measurements
  - All HTTPS sessions for a week
  - 232 million HTTPS sessions

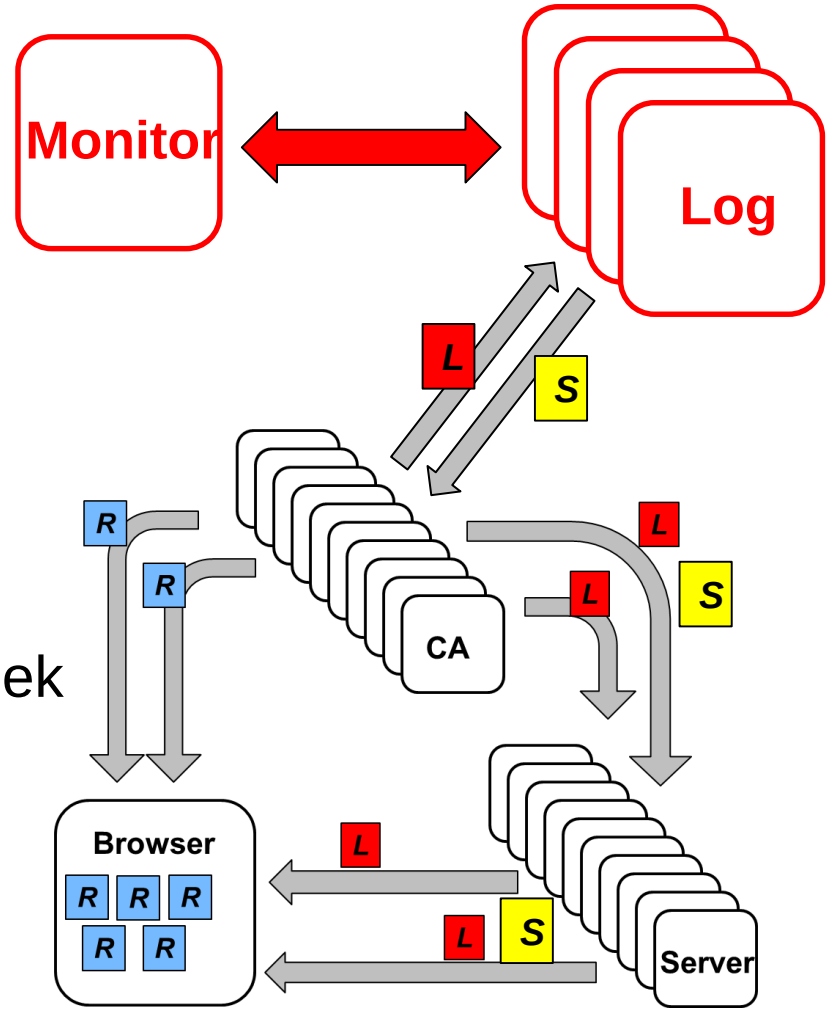

Methodology

- Created CT monitor
- Monitored all public logs
  - 3 Google
  - 7 CA-based
  - Plausible (NORDUnet)
- Campus measurements
  - All HTTPS traffic for a week
  - 232 million HTTPS sessions

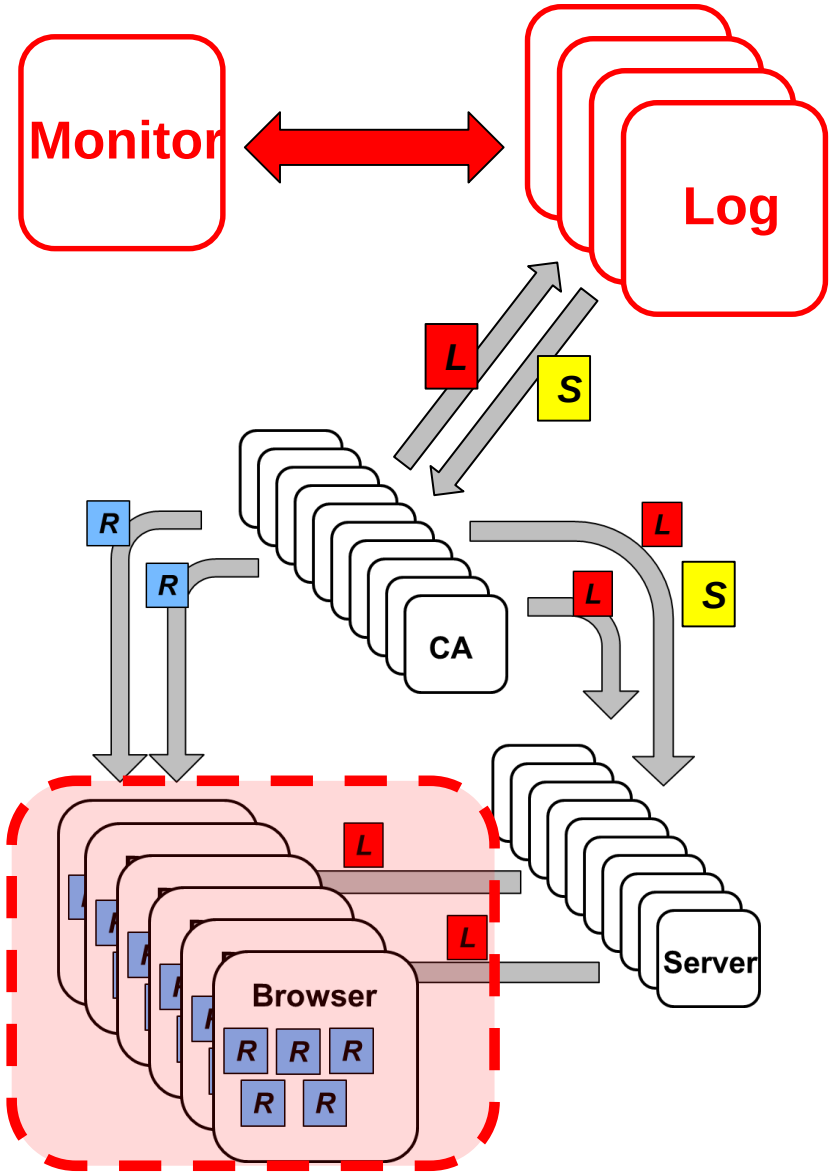

Basic log properties

- All logs allow use of HTTPS
-
-
-
-

Basic log properties

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
-
-
-

Basic log properties

Table 1. Basic properties of the CT logs.

Log name	Operated by	Submitted
Pilot	Google	2013-03-25
Aviator	Google	2013-09-30
Rocketeer	Google	2014-09-01
Digicert	Digicert	2014-09-30
Izenpe	Izenpe	2014-11-10
Certly	Certly	2014-12-14
Symantec	Symantec	2015-05-01
Venafi	Venafi	2015-06-11
WoSign	WoSign	2015-09-22
Vega	Symantec	2015-11-13
Plausible	NORDUnet	Not Subm.

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
-
-
-

Basic log properties

Table 1. Basic properties of the CT logs.

Log name	Operated by	Submitted	URL	Roots	MMD	UI	TTP
Pilot	Google	2013-03-25	ct.googleapis.com/pilot	474	24 hr	1 hr	22 min
Aviator	Google	2013-09-30	ct.googleapis.com/aviator	474	24 hr	1 hr	22 min
Rocketeer	Google	2014-09-01	ct.googleapis.com/rocketeer	474	24 hr	30 m	34 min
Digicert	Digicert	2014-09-30	ct1.digicert-ct.com/log	57	24 hr	1 hr	12 hr
Izenpe	Izenpe	2014-11-10	ct.izenpe.com	40	24 hr	1 min	< 1 min
Certly	Certly	2014-12-14	log.certly.io	183	24 hr	10 min	< 1 min
Symantec	Symantec	2015-05-01	ct.ws.symantec.com	19	24 hr	6 hr	< 1 min
Venafi	Venafi	2015-06-11	ctlog.api.venafi.com	357	24 hr	2 hr	3 min
WoSign	WoSign	2015-09-22	ct.wosign.com	12	24 hr	1 min	< 1 min
Vega	Symantec	2015-11-13	vega.ws.symantec.com	19	24 hr	6 hr	< 1 min
Plausible	NORDUnet	Not Subm.	plausible.ct.nordu.net	442	24 hr*	12 min	2 min

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots;
-
-

Basic log properties

Table 1. Basic properties of the CT logs.

Log name	Operated by	Submitted	URL	Roots	MMD	UI	TTP
Pilot	Google	2013-03-25	ct.googleapis.com/pilot	474	24 hr	1 hr	22 min
Aviator	Google	2013-09-30	ct.googleapis.com/aviator	474	24 hr	1 hr	22 min
Rocketeer	Google	2014-09-01	ct.googleapis.com/rocketeer	474	24 hr	30 m	34 min
Digicert	Digicert	2014-09-30	ct1.digicert-ct.com/log	57	24 hr	1 hr	12 hr
Izenpe	Izenpe	2014-11-10	ct.izenpe.com	40	24 hr	1 min	< 1 min
Certly	Certly	2014-12-14	log.certly.io	183	24 hr	10 min	< 1 min
Symantec	Symantec	2015-05-01	ct.ws.symantec.com	19	24 hr	6 hr	< 1 min
Venafi	Venafi	2015-06-11	ctlog.api.venafi.com	357	24 hr	2 hr	3 min
WoSign	WoSign	2015-09-22	ct.wosign.com	12	24 hr	1 min	< 1 min
Vega	Symantec	2015-11-13	vega.ws.symantec.com	19	24 hr	6 hr	< 1 min
Plausible	NORDUnet	Not Subm.	plausible.ct.nordu.net	442	24 hr*	12 min	2 min

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots; most CA-operated use few roots
-
-

Basic log properties

Table 1. Basic properties of the CT logs.

Log name	Operated by	Submitted	URL	Roots	MMD	UI	TTP
Pilot	Google	2013-03-25	ct.googleapis.com/pilot	474	24 hr	1 hr	22 min
Aviator	Google	2013-09-30	ct.googleapis.com/aviator	474	24 hr	1 hr	22 min
Rocketeer	Google	2014-09-01	ct.googleapis.com/rocketeer	474	24 hr	30 m	34 min
Digicert	Digicert	2014-09-30	ct1.digicert-ct.com/log	57	24 hr	1 hr	12 hr
Izenpe	Izenpe	2014-11-10	ct.izenpe.com	40	24 hr	1 min	< 1 min
Certly	Certly	2014-12-14	log.certly.io	183	24 hr	10 min	< 1 min
Symantec	Symantec	2015-05-01	ct.ws.symantec.com	19	24 hr	6 hr	< 1 min
Venafi	Venafi	2015-06-11	ctlog.api.venafi.com	357	24 hr	2 hr	3 min
WoSign	WoSign	2015-09-22	ct.wosign.com	12	24 hr	1 min	< 1 min
Vega	Symantec	2015-11-13	vega.ws.symantec.com	19	24 hr	6 hr	< 1 min
Plausible	NORDUnet	Not Subm.	plausible.ct.nordu.net	442	24 hr*	12 min	2 min

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots; most CA-operated use few roots
- Many roots included in many published logs
-

Basic log properties

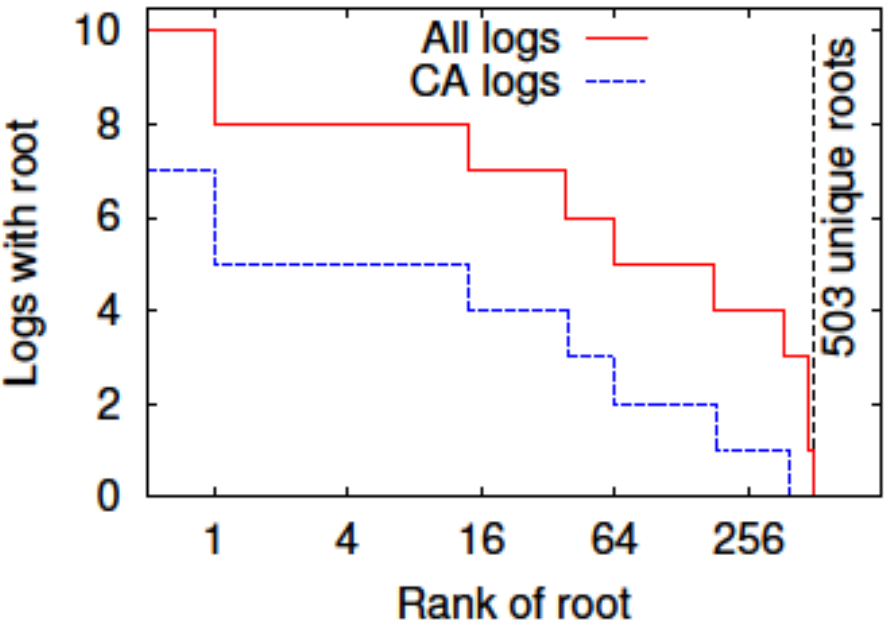

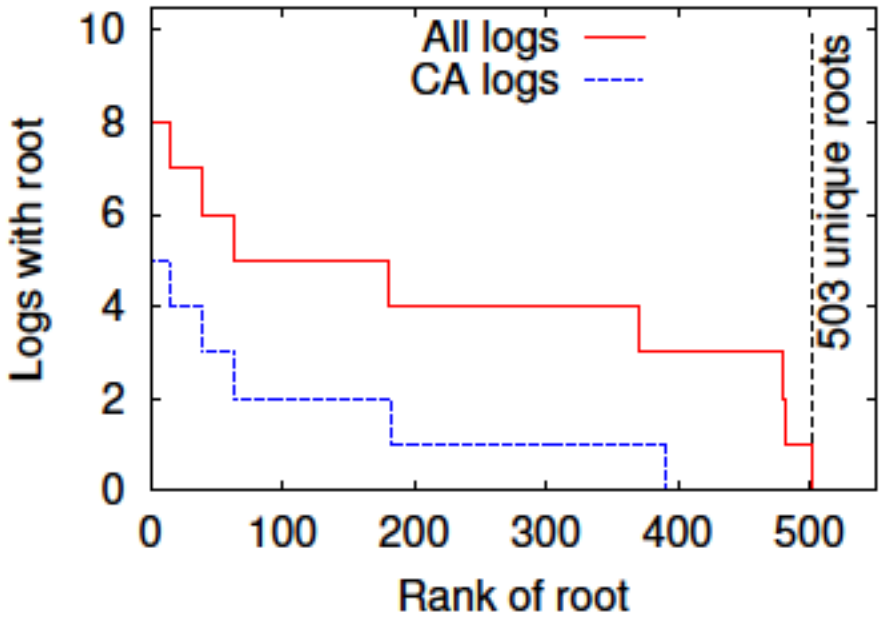

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots; most CA-operated use few roots
- Many roots included in many published logs
-

Basic log properties

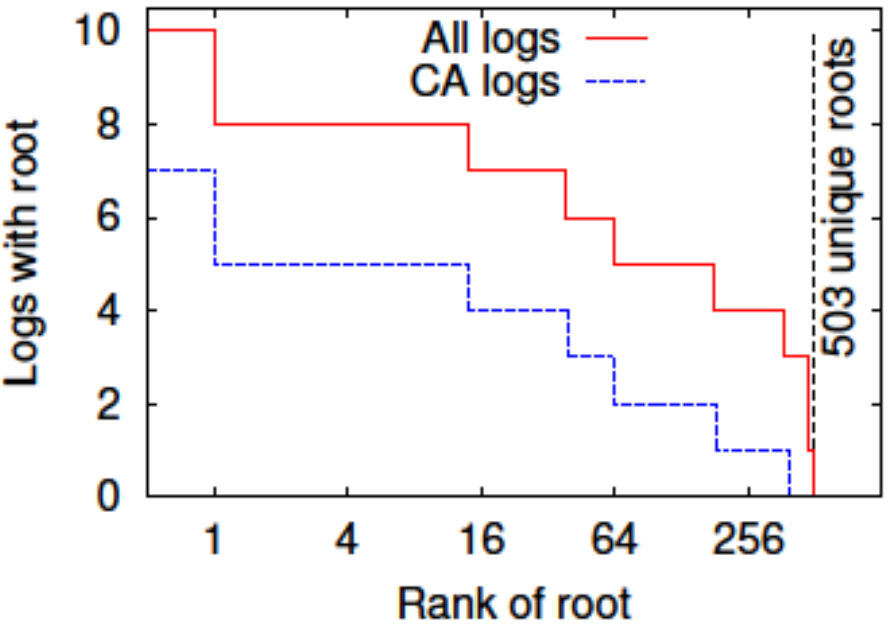

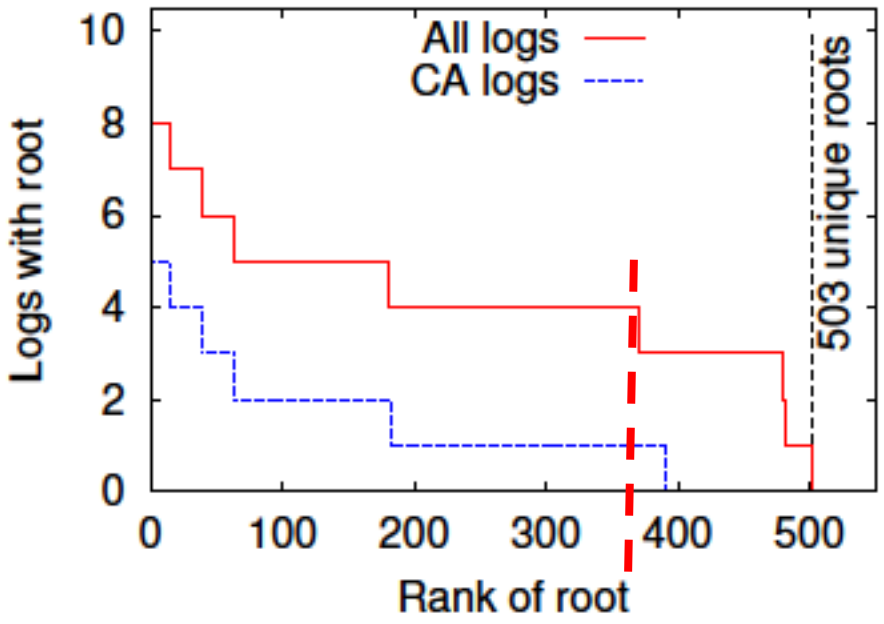

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots; most CA-operated use few roots
- Many roots included in many published logs
-

Basic log properties

Table 1. Basic properties of the CT logs.

Log name	Operated by	Submitted	URL	Roots	MMD	UI	TTP
Pilot	Google	2013-03-25	ct.googleapis.com/pilot	474	24 hr	1 hr	22 min
Aviator	Google	2013-09-30	ct.googleapis.com/aviator	474	24 hr	1 hr	22 min
Rocketeer	Google	2014-09-01	ct.googleapis.com/rocketeer	474	24 hr	30 m	34 min
Digicert	Digicert	2014-09-30	ct1.digicert-ct.com/log	57	24 hr	1 hr	12 hr
Izenpe	Izenpe	2014-11-10	ct.izenpe.com	40	24 hr	1 min	< 1 min
Certly	Certly	2014-12-14	log.certly.io	183	24 hr	10 min	< 1 min
Symantec	Symantec	2015-05-01	ct.ws.symantec.com	19	24 hr	6 hr	< 1 min
Venafi	Venafi	2015-06-11	ctlog.api.venafi.com	357	24 hr	2 hr	3 min
WoSign	WoSign	2015-09-22	ct.wosign.com	12	24 hr	1 min	< 1 min
Vega	Symantec	2015-11-13	vega.ws.symantec.com	19	24 hr	6 hr	< 1 min
Plausible	NORDUnet	Not Subm.	plausible.ct.nordu.net	442	24 hr*	12 min	2 min

* Plausible operates with an unofficial MMD of 24hr.

- All logs allow use of HTTPS
- Venafi uses RSA with SHA-256, rest use ECDSA over NIST P-256
- Google+Plausible many roots; most CA-operated use few roots
- Many roots included in many logs
- Most logs have significant compliance margin; i.e., UI+TTP << MMD

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

Log name	Operated by	Entries
Pilot	Google	10,831,024
Aviator	Google	10,069,865
Rocketeer	Google	8,140,991
Digicert	Digicert	229,858
Izenpe	Izenpe	65,812
Certly	Certly	161,740
Symantec	Symantec	113,674
Venafi	Venafi	4,626
WoSign	WoSign	11,188
Vega	Symantec	80
Plausible	NORDUnet	5,893,906

- Three classes: Large, medium (CA-based), small (CA-based)
-
  -
-
  -
-

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

			Validation		
Log name	**Operated by**	**Entries**	**DV**	**OV**	**EV**
Pilot	Google	10,831,024	87%	8%	5%
Aviator	Google	10,069,865	87%	8%	5%
Rocketeer	Google	8,140,991	87%	8%	5%
Digicert	Digicert	229,858	18%	5%	78%
Izenpe	Izenpe	65,812	31%	1%	68%
Certly	Certly	161,740	36%	3%	61%
Symantec	Symantec	113,674	21%	5%	74%
Venafi	Venafi	4,626	85%	10%	5%
WoSign	WoSign	11,188	97%	1%	2%
Vega	Symantec	80	95%	0%	5%
Plausible	NORDUnet	5,893,906	88%	7%	5%

- Three classes: Large, medium (CA-based), small (CA-based)
- Medium most EV certificates
  - Both Digicert (27.6%) and Symantec (56.2%) of EV sessions on campus
-
  -
-

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

			Validation		
Log name	Operated by	Entries	DV	OV	EV
Pilot	Google	10,831,024	87%	8%	5%
Aviator	Google	10,069,865	87%	8%	5%
Rocketeer	Google	8,140,991	87%	8%	5%
Digicert	Digicert	229,858	18%	5%	78%
Izenpe	Izenpe	65,812	31%	1%	68%
Certly	Certly	161,740	36%	3%	61%
Symantec	Symantec	113,674	21%	5%	74%
Venafi	Venafi	4,626	85%	10%	5%
WoSign	WoSign	11,188	97%	1%	2%
Vega	Symantec	80	95%	0%	5%
Plausible	NORDUnet	5,893,906	88%	7%	5%

- Three classes: Large, medium (CA-based), small (CA-based)
- Medium most EV certificates
  - Both Digicert (27.6%) and Symantec (56.2%) of EV sessions on campus
- Large (crawl-based) fairly representative of the wild
  - E.g., campus 4.9% EV, large logs all have 5% EV
-

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

			Validation		
Log name	Operated by	Entries	DV	OV	EV
Pilot	Google	10,831,024	87%	8%	5%
Aviator	Google	10,069,865	87%	8%	5%
Rocketeer	Google	8,140,991	87%	8%	5%
Digicert	Digicert	229,858	18%	5%	78%
Izenpe	Izenpe	65,812	31%	1%	68%
Certly	Certly	161,740	36%	3%	61%
Symantec	Symantec	113,674	21%	5%	74%
Venafi	Venafi	4,626	85%	10%	5%
WoSign	WoSign	11,188	97%	1%	2%
Vega	Symantec	80	95%	0%	5%
Plausible	NORDUnet	5,893,906	88%	7%	5%

- Three classes: Large, medium (CA-based), small (CA-based)
- Medium most EV certificates
  - Both Digicert (27.6%) and Symantec (56.2%) of EV sessions on campus
- Large (crawl-based) fairly representative of the wild
  - E.g., campus 4.9% EV, large logs all have 5% EV
- Small logs have large portion test certificates

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

			Validation			Encryption algorithm			
Log name	Operated by	Entries	DV	OV	EV	RSA (1024)	RSA (2048)	RSA (4096)	EC (256)
Pilot	Google	10,831,024	87%	8%	5%	2%	79%	3%	16%
Aviator	Google	10,069,865	87%	8%	5%	1%	78%	3%	17%
Rocketeer	Google	8,140,991	87%	8%	5%	1%	75%	4%	21%
Digicert	Digicert	229,858	18%	5%	78%	0%	96%	3%	0%
Izenpe	Izenpe	65,812	31%	1%	68%	0%	95%	5%	0%
Certly	Certly	161,740	36%	3%	61%	0%	94%	5%	0%
Symantec	Symantec	113,674	21%	5%	74%	0%	97%	2%	0%
Venafi	Venafi	4,626	85%	10%	5%	0%	93%	5%	1%
WoSign	WoSign	11,188	97%	1%	2%	0%	99%	1%	0%
Vega	Symantec	80	95%	0%	5%	0%	95%	0%	2%
Plausible	NORDUnet	5,893,906	88%	7%	5%	3%	90%	3%	4%

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  -
-
-

Certificate analysis

Table 2. Distribution of certificate validation types and signature hashes.

			Validation			Encryption algorithm			
Log name	Operated by	Entries	DV	OV	EV	RSA (1024)	RSA (2048)	RSA (4096)	EC (256)
Pilot	Google	10,831,024	87%	8%	5%	2%	79%	3%	16%
Aviator	Google	10,069,865	87%	8%	5%	1%	78%	3%	17%
Rocketeer	Google	8,140,991	87%	8%	5%	1%	75%	4%	21%
Digicert	Digicert	229,858	18%	5%	78%	0%	96%	3%	0%
Izenpe	Izenpe	65,812	31%	1%	68%	0%	95%	5%	0%
Certly	Certly	161,740	36%	3%	61%	0%	94%	5%	0%
Symantec	Symantec	113,674	21%	5%	74%	0%	97%	2%	0%
Venafi	Venafi	4,626	85%	10%	5%	0%	93%	5%	1%
WoSign	WoSign	11,188	97%	1%	2%	0%	99%	1%	0%
Vega	Symantec	80	95%	0%	5%	0%	95%	0%	2%
Plausible	NORDUnet	5,893,906	88%	7%	5%	3%	90%	3%	4%

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  - And, log more Elliptic Curve (EC) certificates
-
-

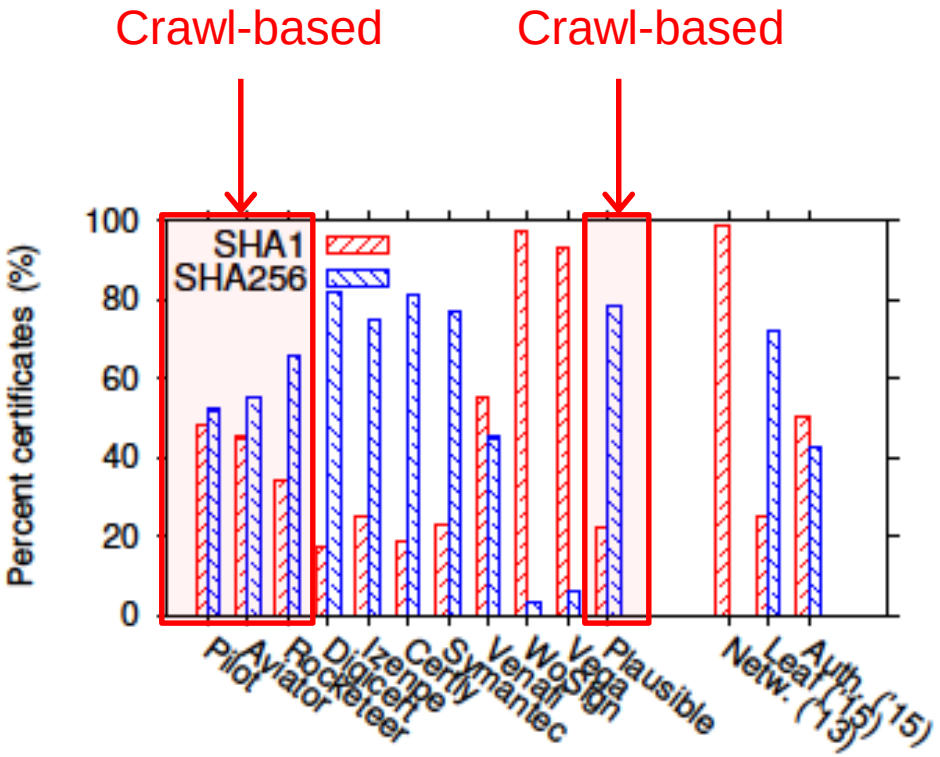

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  - And, log more Elliptic Curve (EC) certificates
- Crawl-based (large) logs see many weak signatures (SHA1)
  -
-

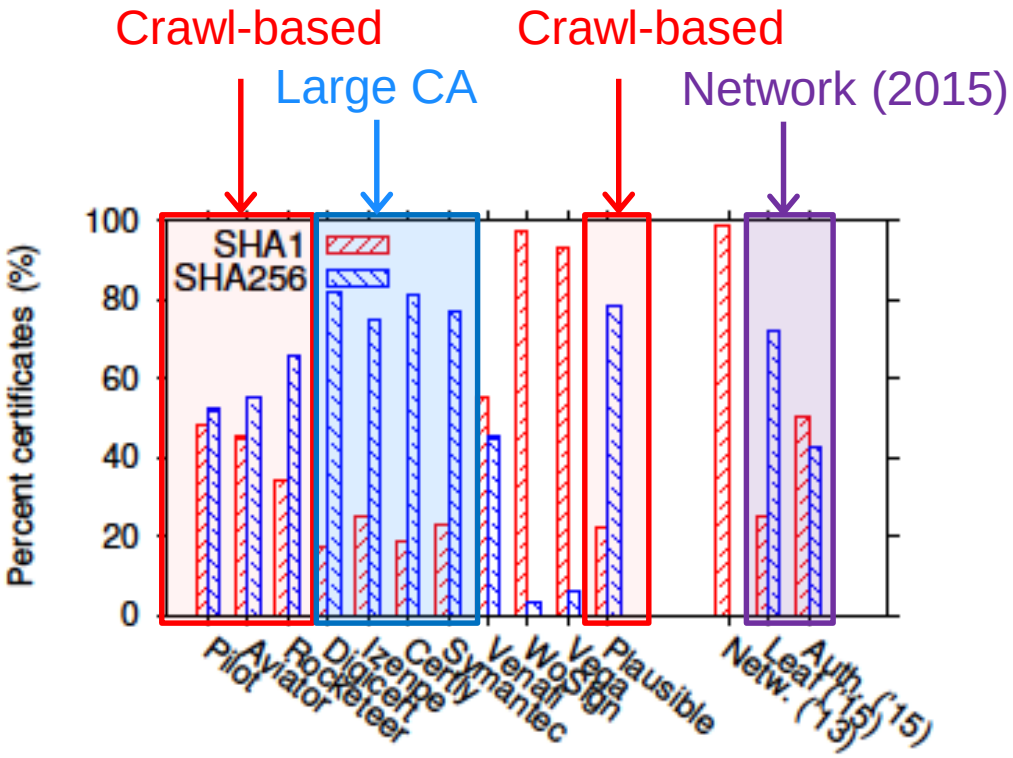

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  - And, log more Elliptic Curve (EC) certificates
- Crawl-based (large) logs see many weak signatures (SHA1)
  - Much more than large CA logs; Consistent with network numbers
  -
-

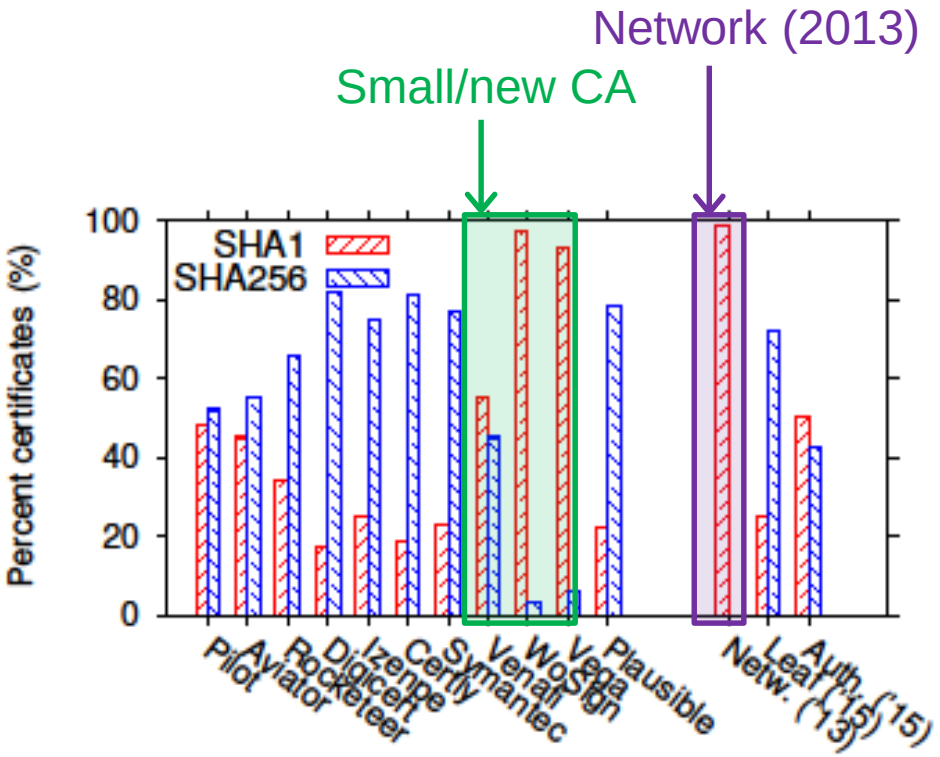

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  - And, log more Elliptic Curve (EC) certificates
- Crawl-based (large) logs see many weak signatures (SHA1)
  - Much more than large CA logs; Consistent with network numbers
  - Small/new CA logs (mostly test certificates!!) and old network even more
-

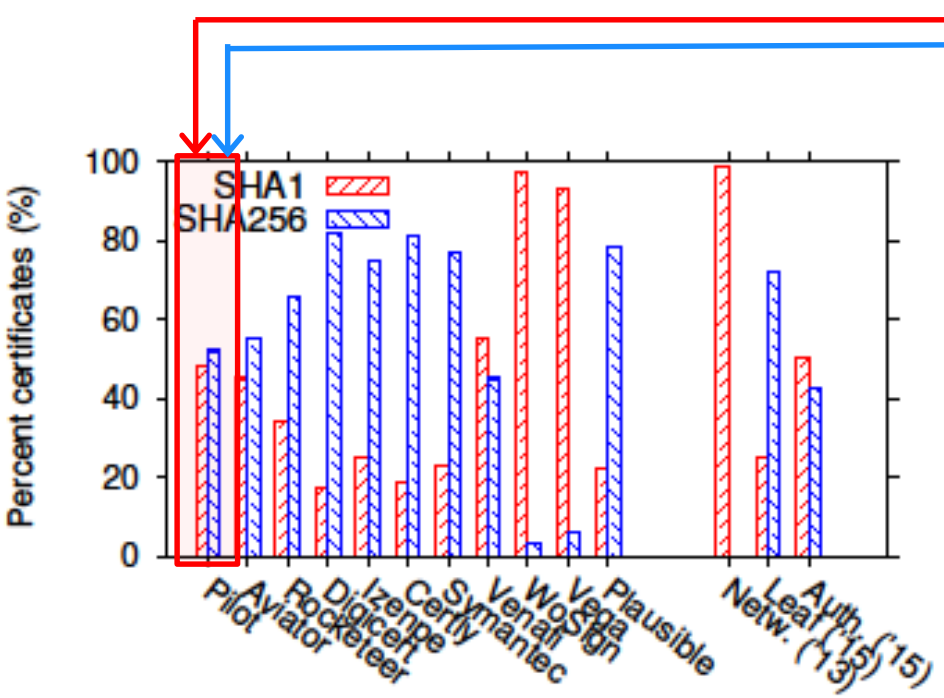

Example CDFs based on Pilot

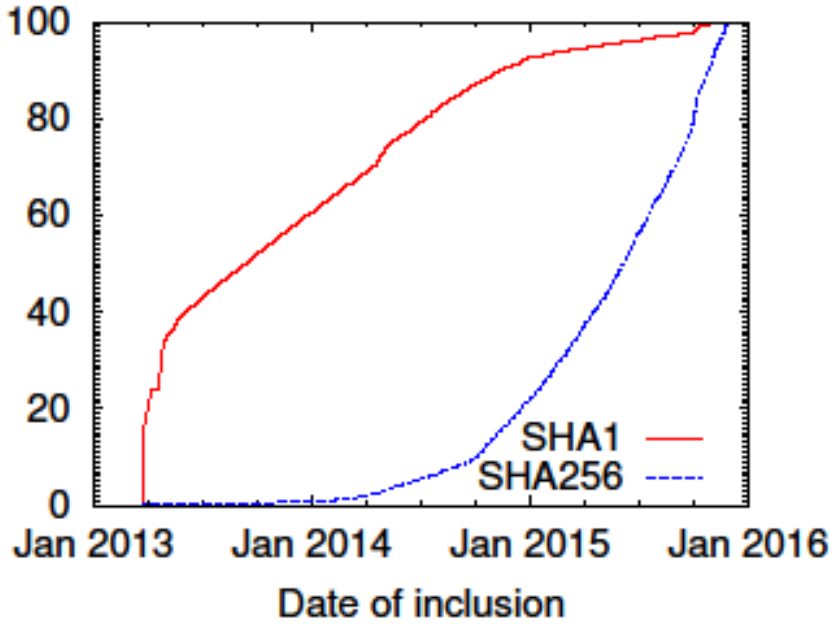

- Crawl-based (large) logs most weak key algorithms (1024 RSA)
  - And, log more Elliptic Curve (EC) certificates
- Crawl-based (large) logs see many weak signatures (SHA1)
  - Much more than large CA logs; Consistent with network numbers
  - Small/new CA logs (mostly test certificates!!) and old network even more
- SHA256 is taking over, but new SHA1 certificates are still being added to the logs

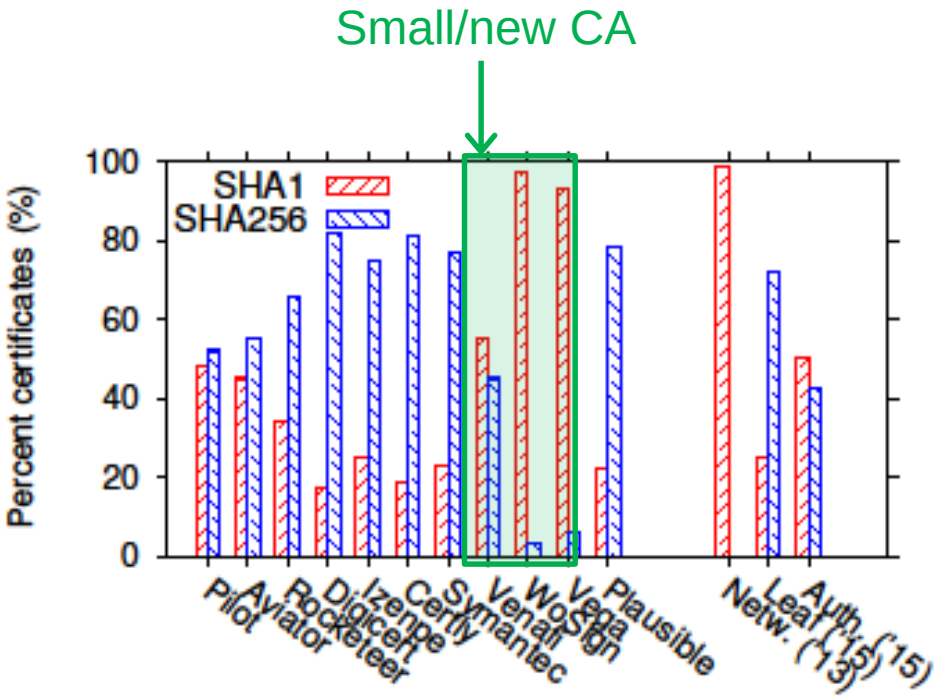

- As we have seen, the small CA-based logs really stick out
  - E.g., large number of SHA1 certs (mostly test certificates)

Validation test

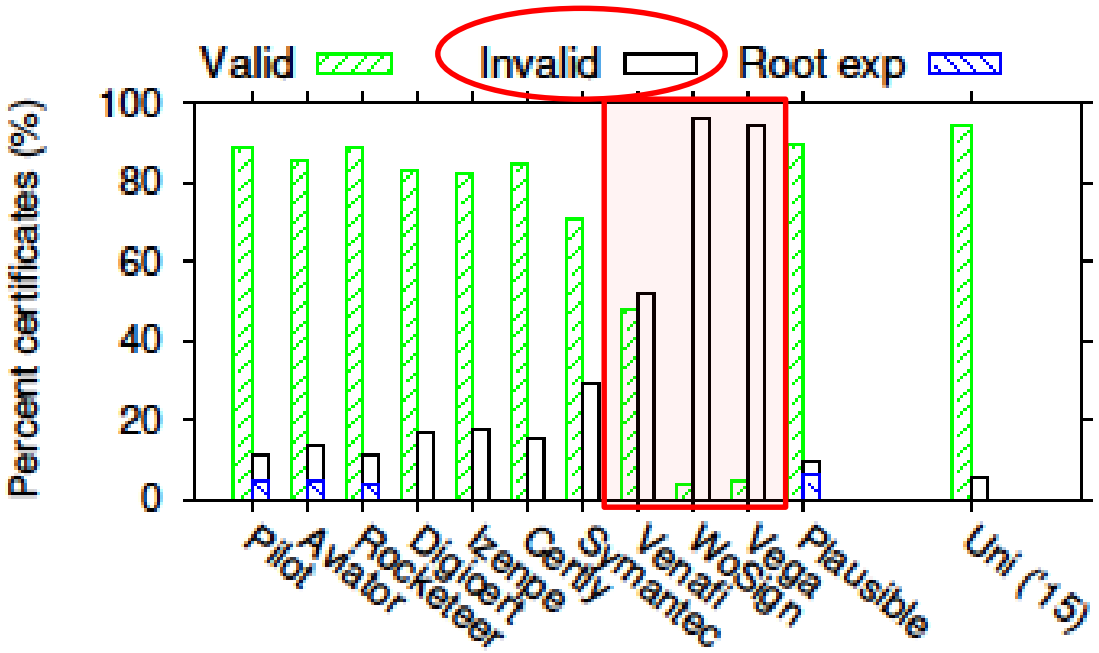

- As we have seen, the small CA-based logs really stick out
  - E.g., large number of SHA1 certs (mostly test certificates)
  - These logs have many invalid certs (do not validate using Mozilla root store)

Validation test

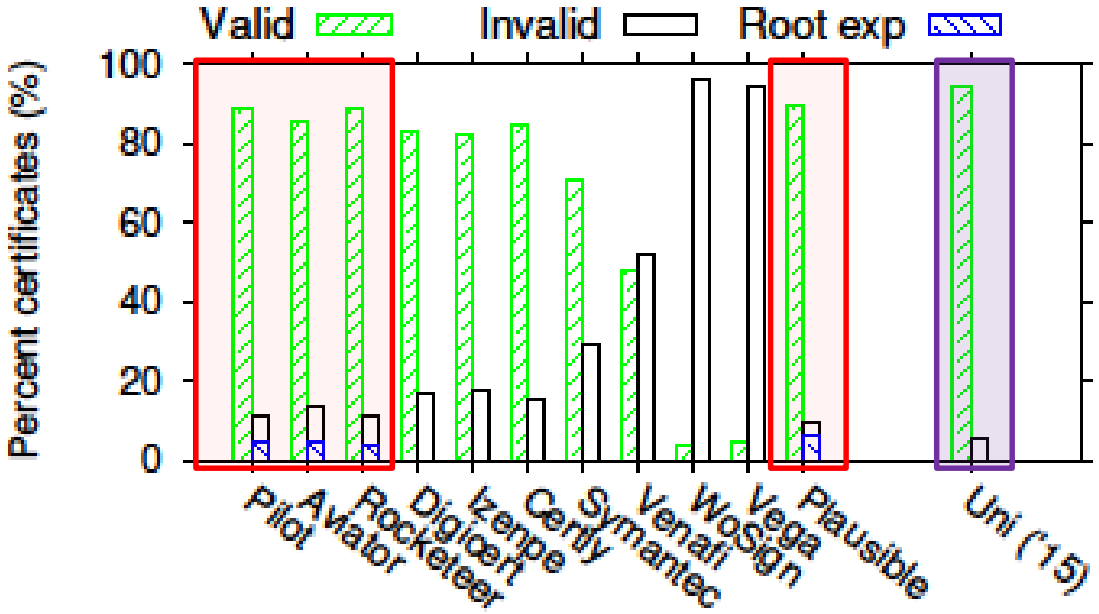

- As we have seen, the small CA-based logs really stick out
  - E.g., large number of SHA1 certs (mostly test certificates)
  - These logs have many invalid certs (do not validate using Mozilla root store)
- Crawl-based logs consistent with what seen on network
  -

Validation test

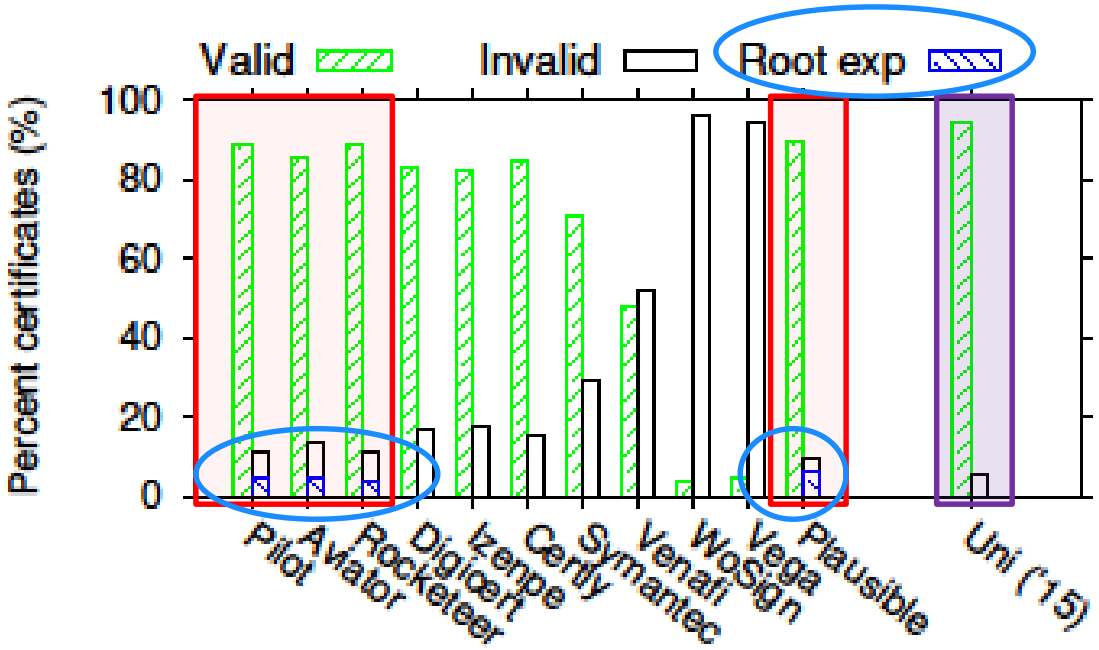

- As we have seen, the small CA-based logs really stick out
  - E.g., large number of SHA1 certs (mostly test certificates)
  - These logs have many invalid certs (do not validate using Mozilla root store)
- Crawl-based logs consistent with what seen on network
  - Subset of invalid certs have expired roots (comparison even more similar …)

Cross-log publication

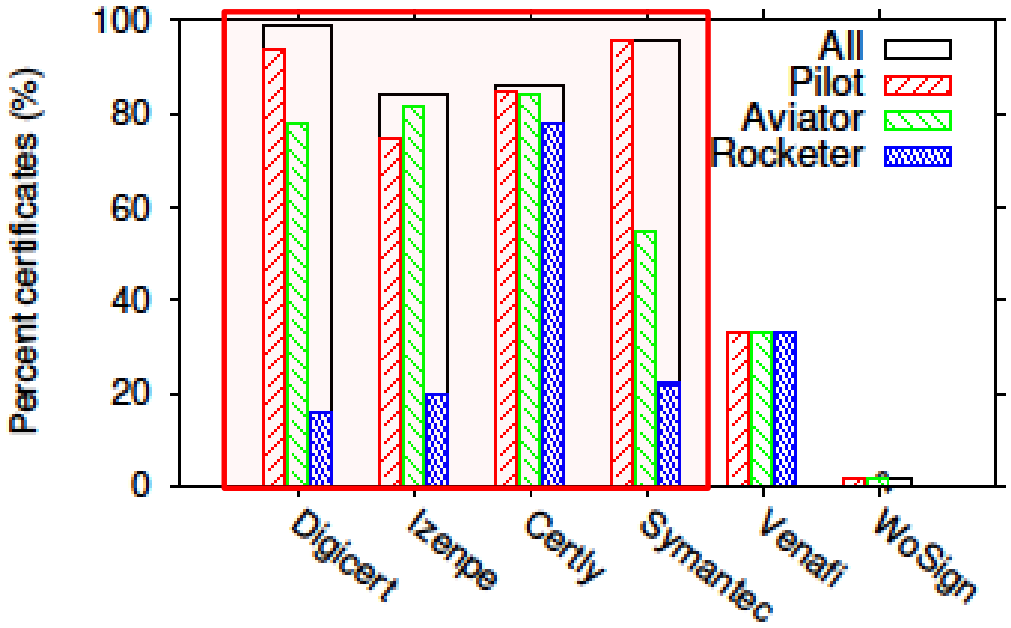

- Among the four large CA logs, most certs are also logged in Google logs ...
  - Remember Chromes 1+1 policy
  -

Cross-log publication

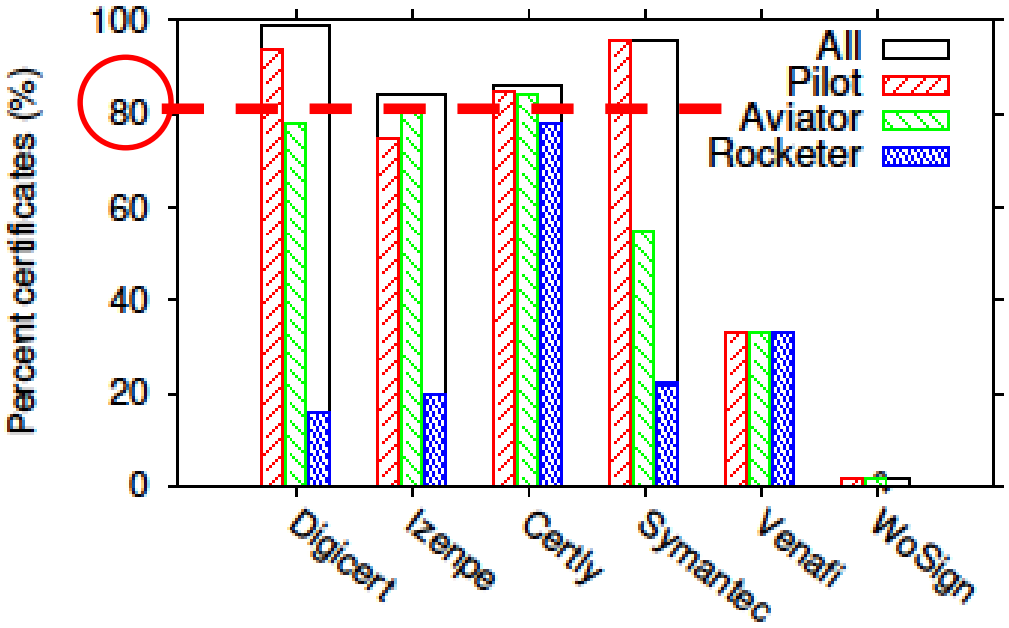

- Among the four large CA logs, most certs are also logged in Google logs ...
  - Remember Chromes 1+1 policy
  - More than 80% in at least one Google log

Cross-log publication

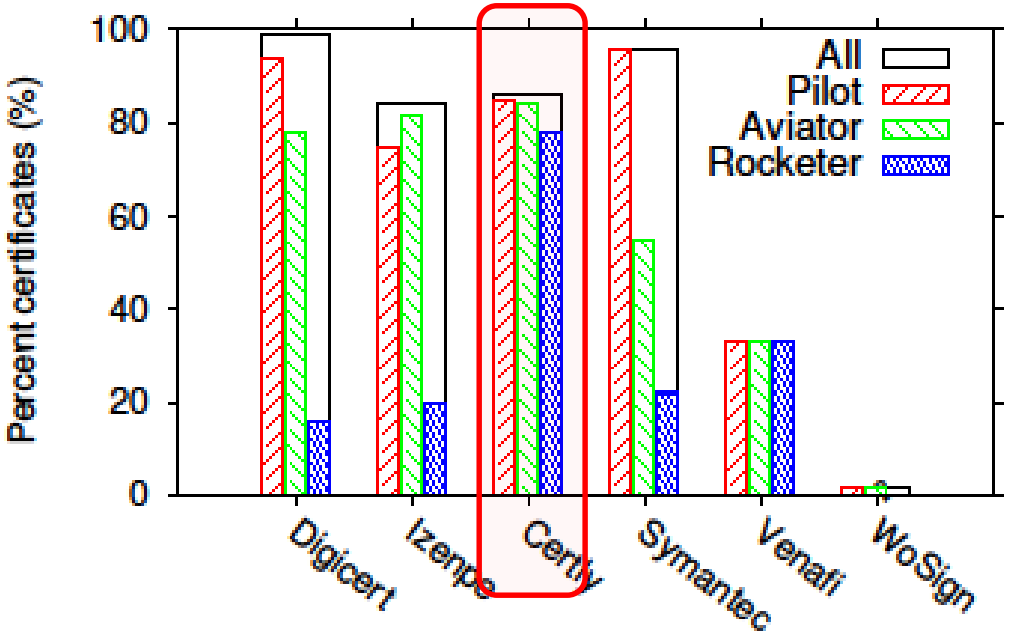

- Among the four large CA logs, most certs are also logged in Google logs ...
  - Remember Chromes 1+1 policy
  - More than 80% in at least one Google log
  - Certly use all three; the other three large CAs typically use 2 of 3
  -

Cross-log publication

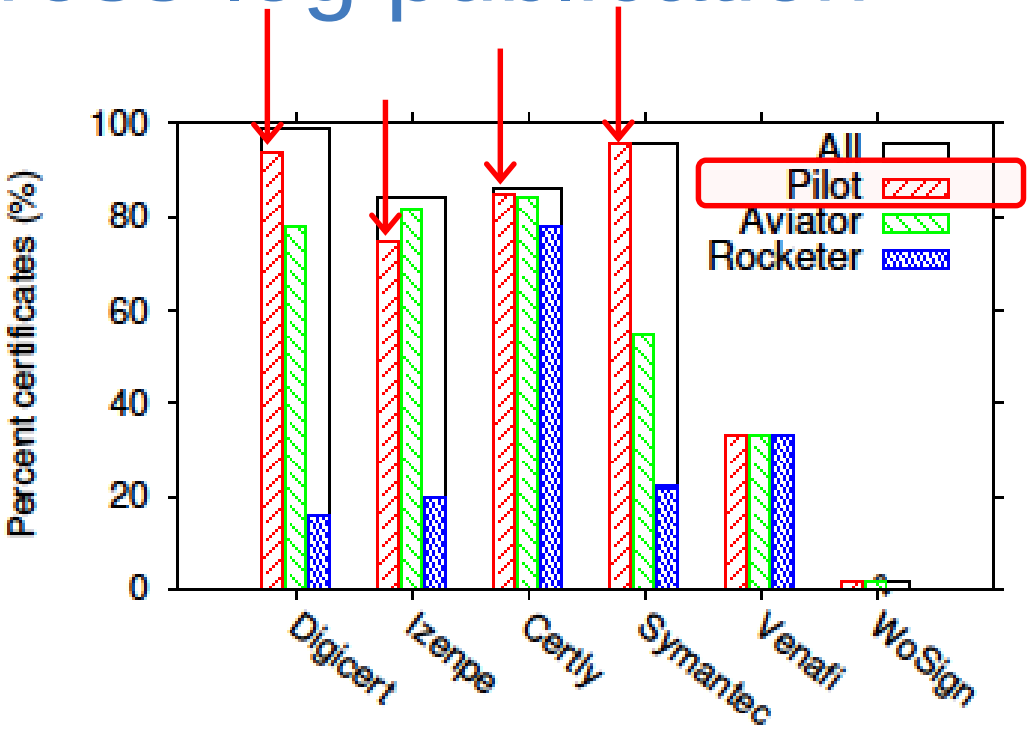

- Among the four large CA logs, most certs are also logged in Google logs ...
  - Remember Chromes 1+1 policy
  - More than 80% in at least one Google log
  - Certly use all three; the other three large CAs typically use 2 of 3
  - Bias towards Pilot partially age related

Temporal analysis

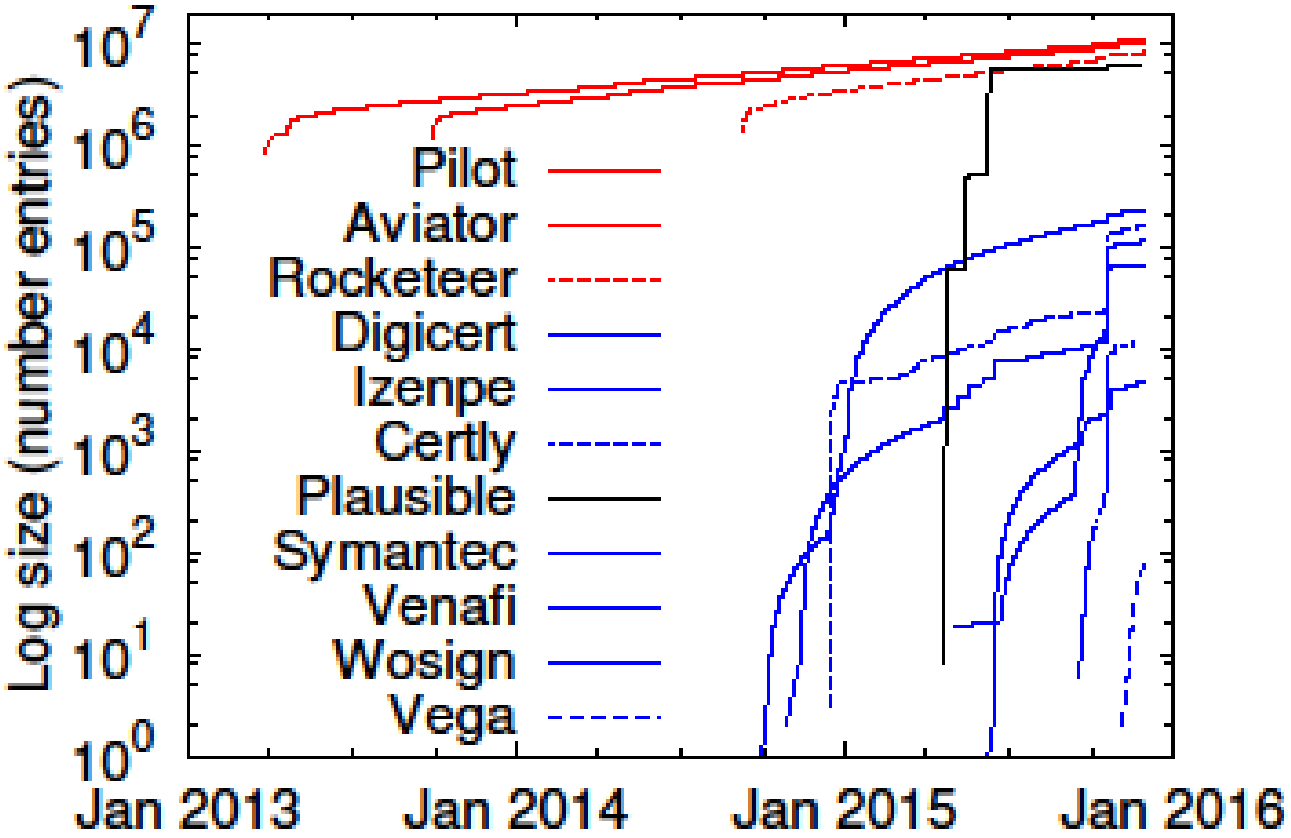

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
-
-

Temporal analysis

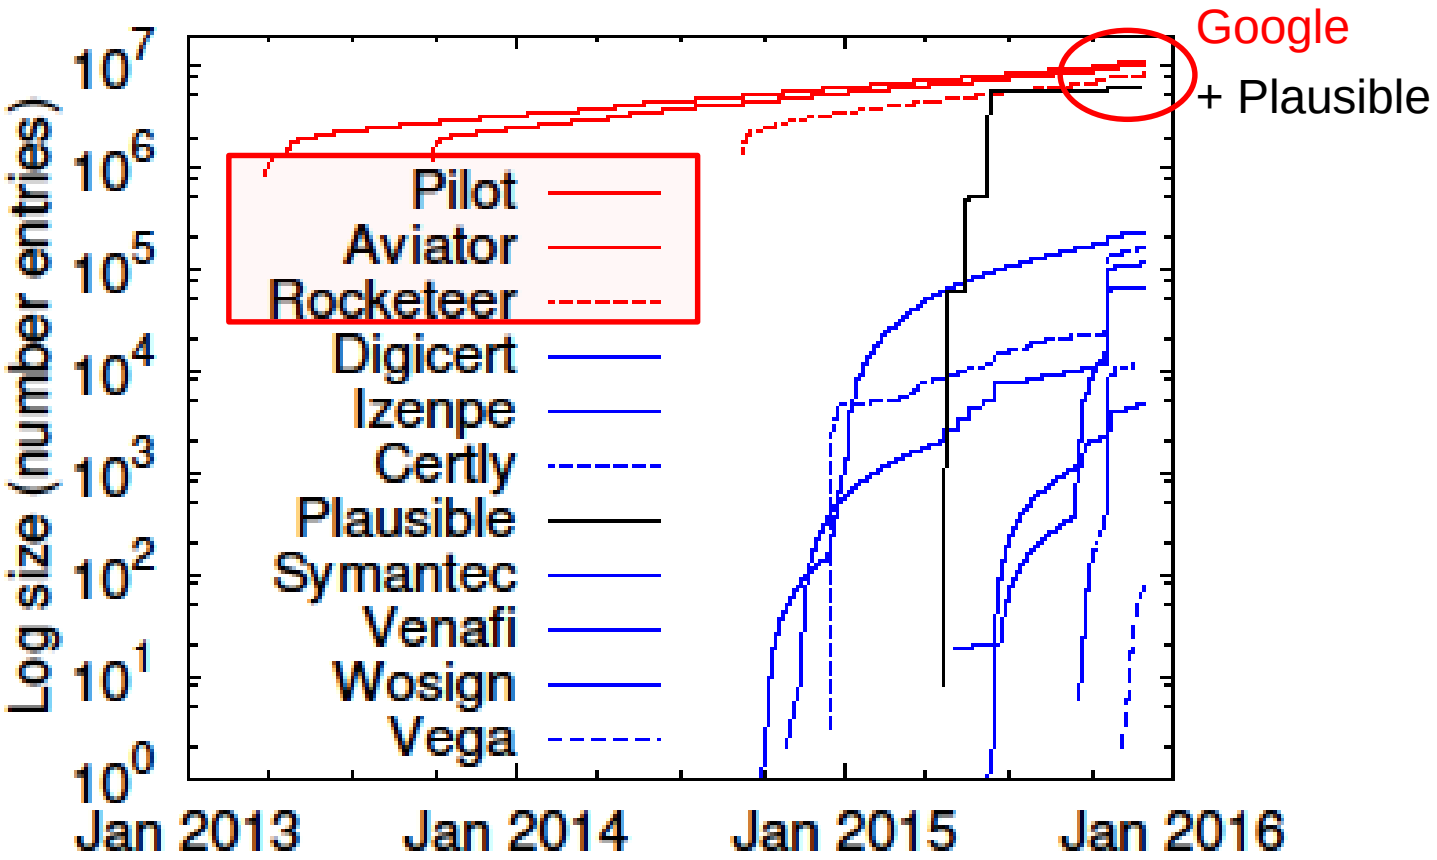

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
-
-

Temporal analysis

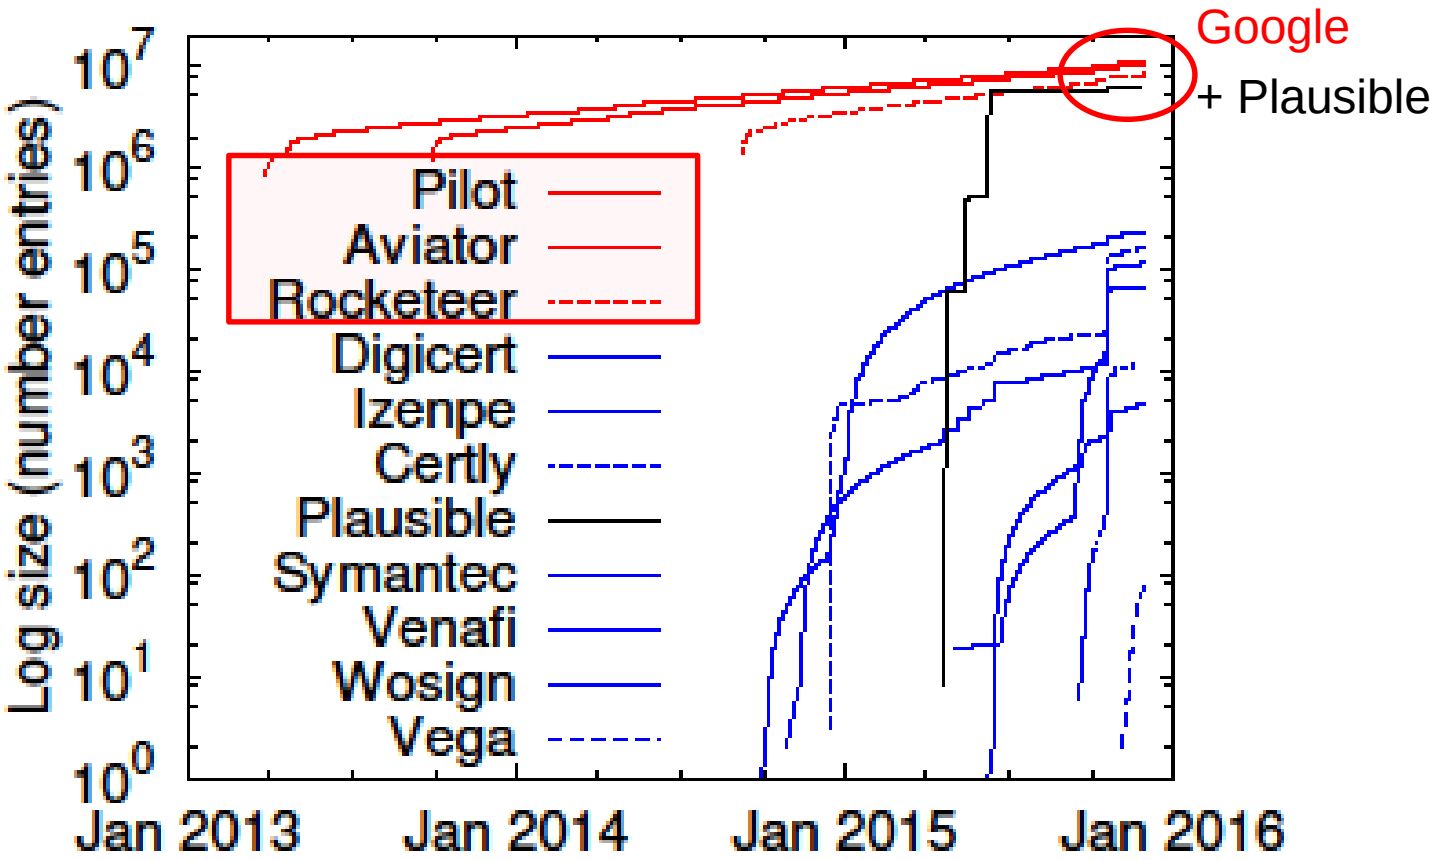

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
- Strict size ordering of Google logs
-

Temporal analysis

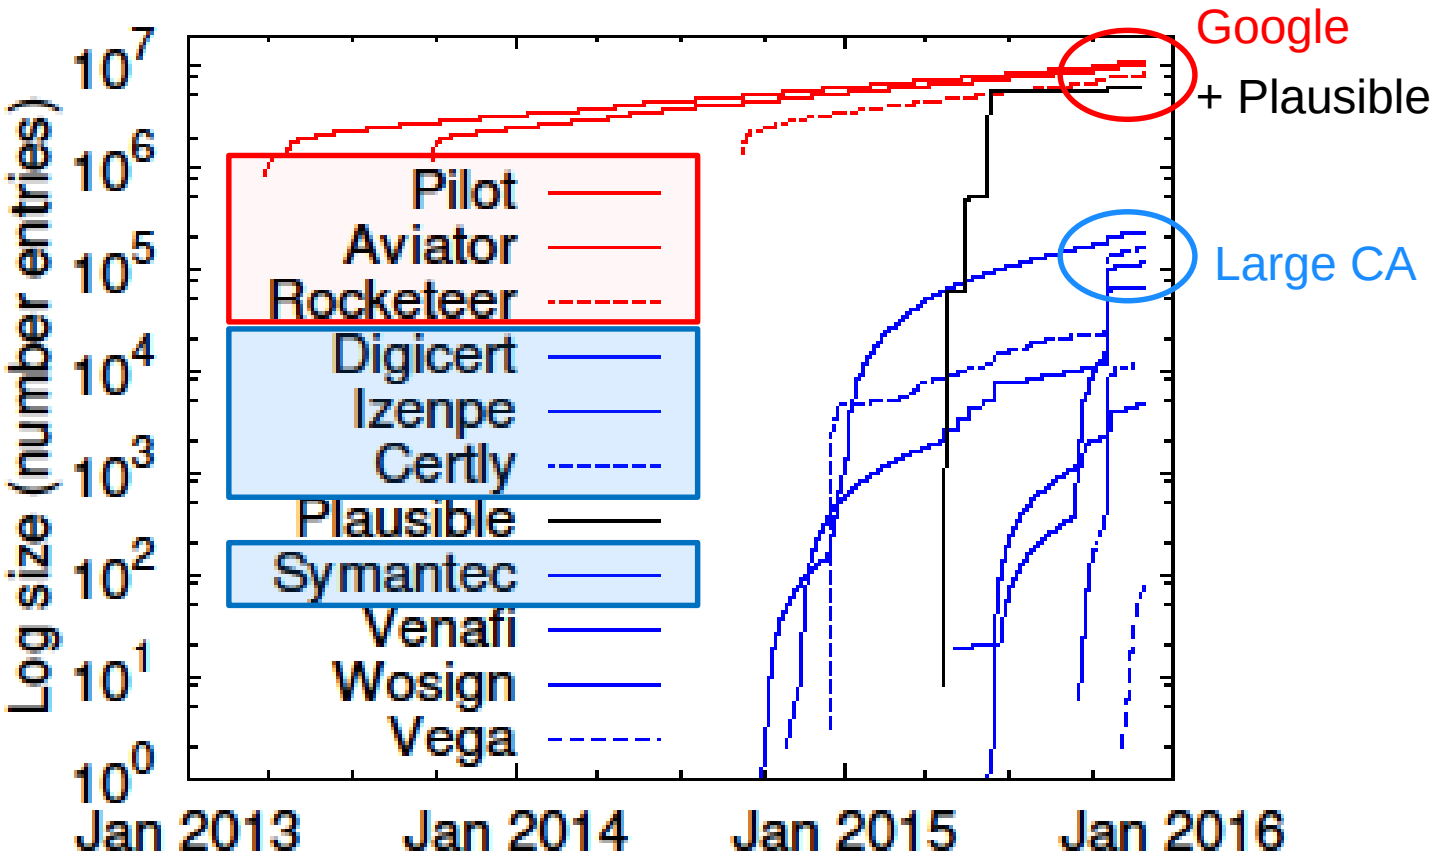

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
- Strict size ordering of Google logs
- Size ordering changes among CAs

Temporal analysis

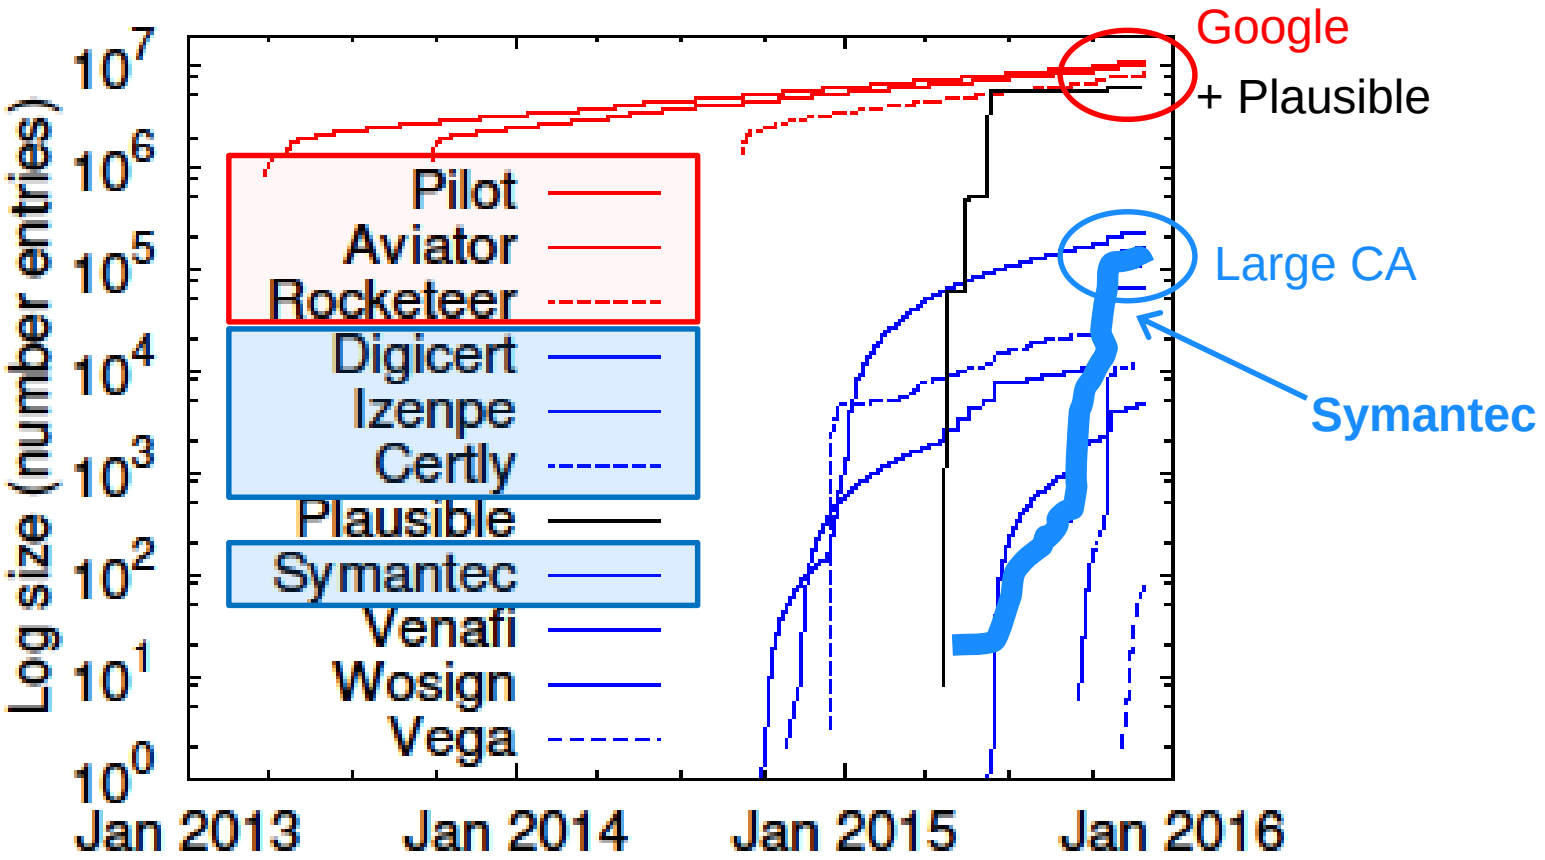

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
- Strict size ordering of Google logs
- Size ordering changes among CAs **(e.g., Symantec incident ...)**

Temporal analysis

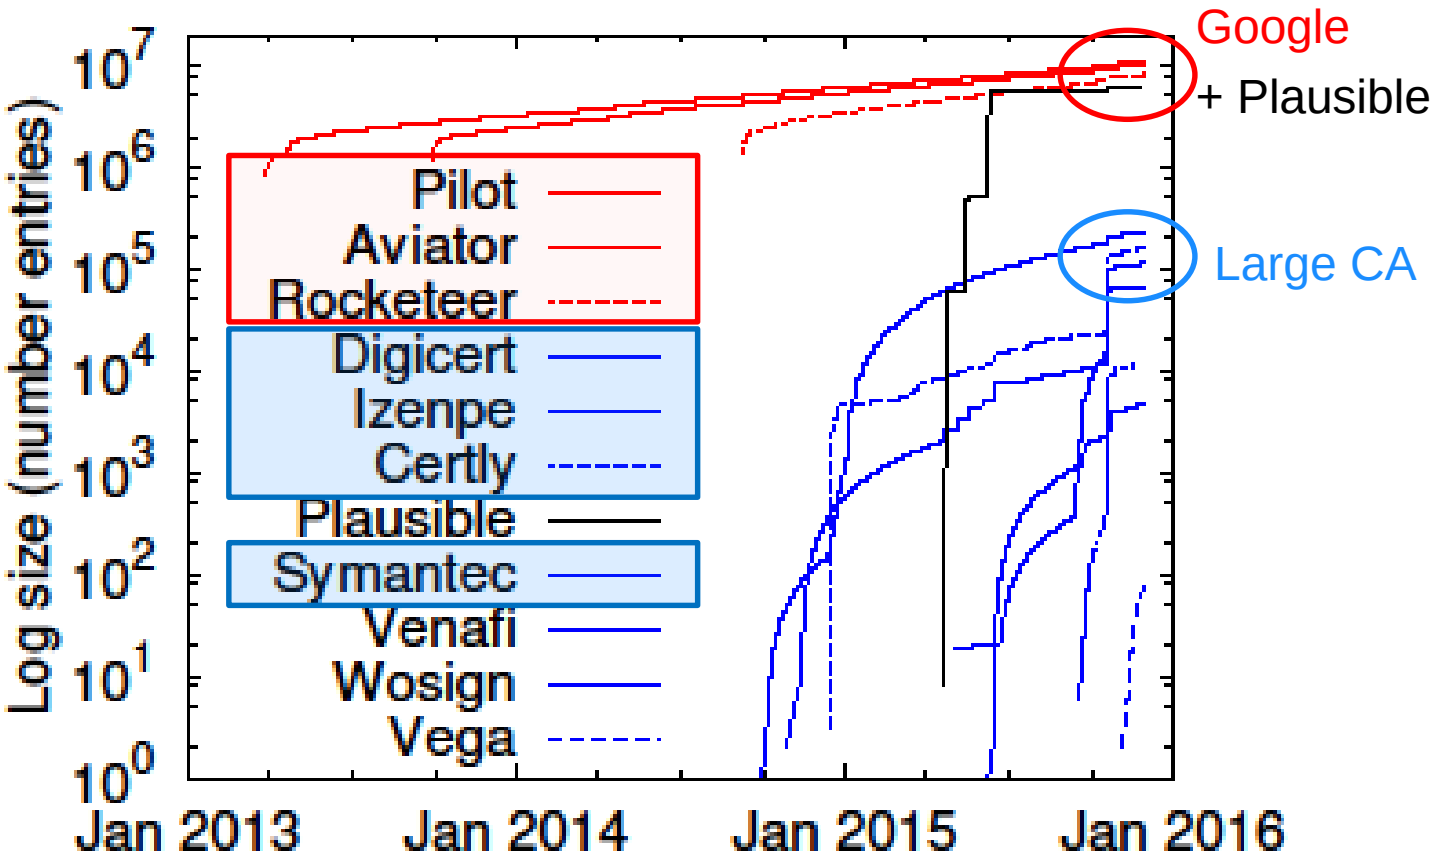

- CT logs are strictly append-only
  - Increasing use of short-lived certs and HTTPS
- Strict size ordering of Google logs
- Size ordering changes among CAs (e.g., Symantec incident ...)

Temporal analysis examples

- Pilot (first Google) log shows spike in EV around the time that Chromes EV policy took effect (Jan ‘15)

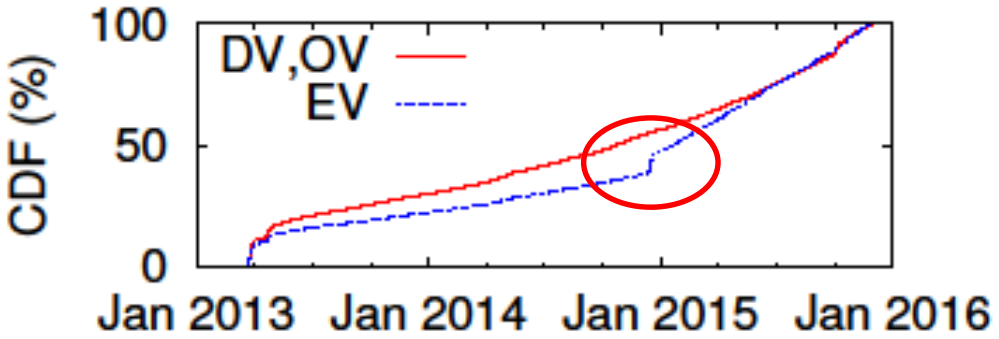

- Digicert started their log around the same time and have been adding EVs steadily ever since

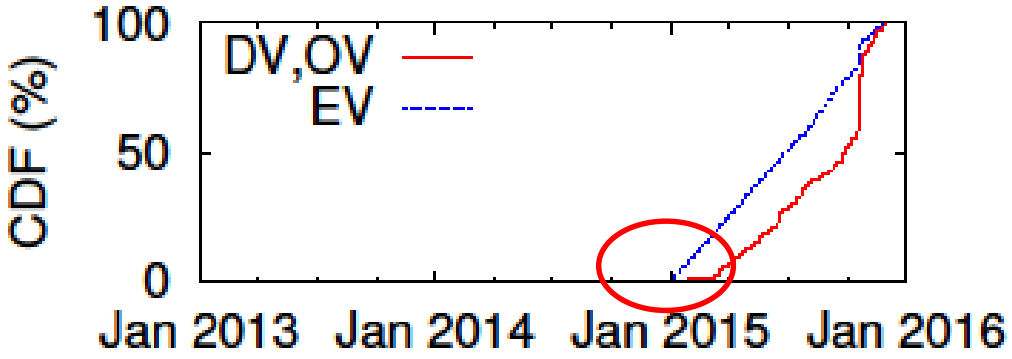

- Symantec: EV and DV goes more hand-in hand

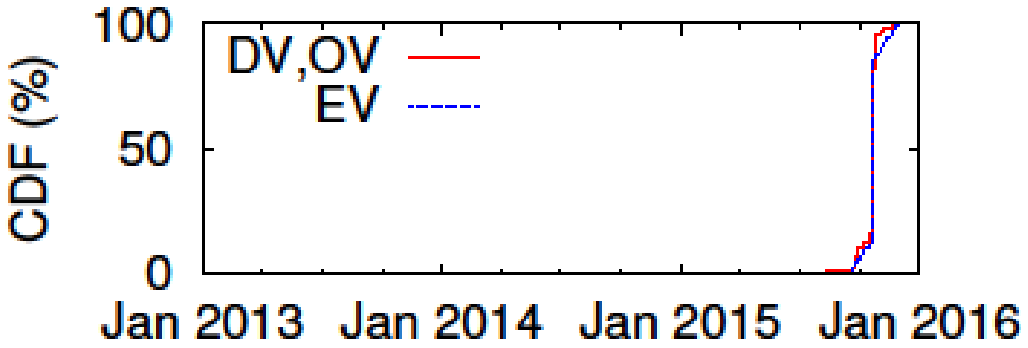

Temporal analysis examples

- Pilot (first Google) log shows spike in EV around the time that Chromes EV policy took effect (Jan ‘15)
- Digicert started their log around the same time and have been adding EVs steadily ever since (spike in DVs after Symantec incident)
- Symantec: EV and DV goes more hand-in hand (again, Google requires Symantec to log all certs, due to their 2015 incident)

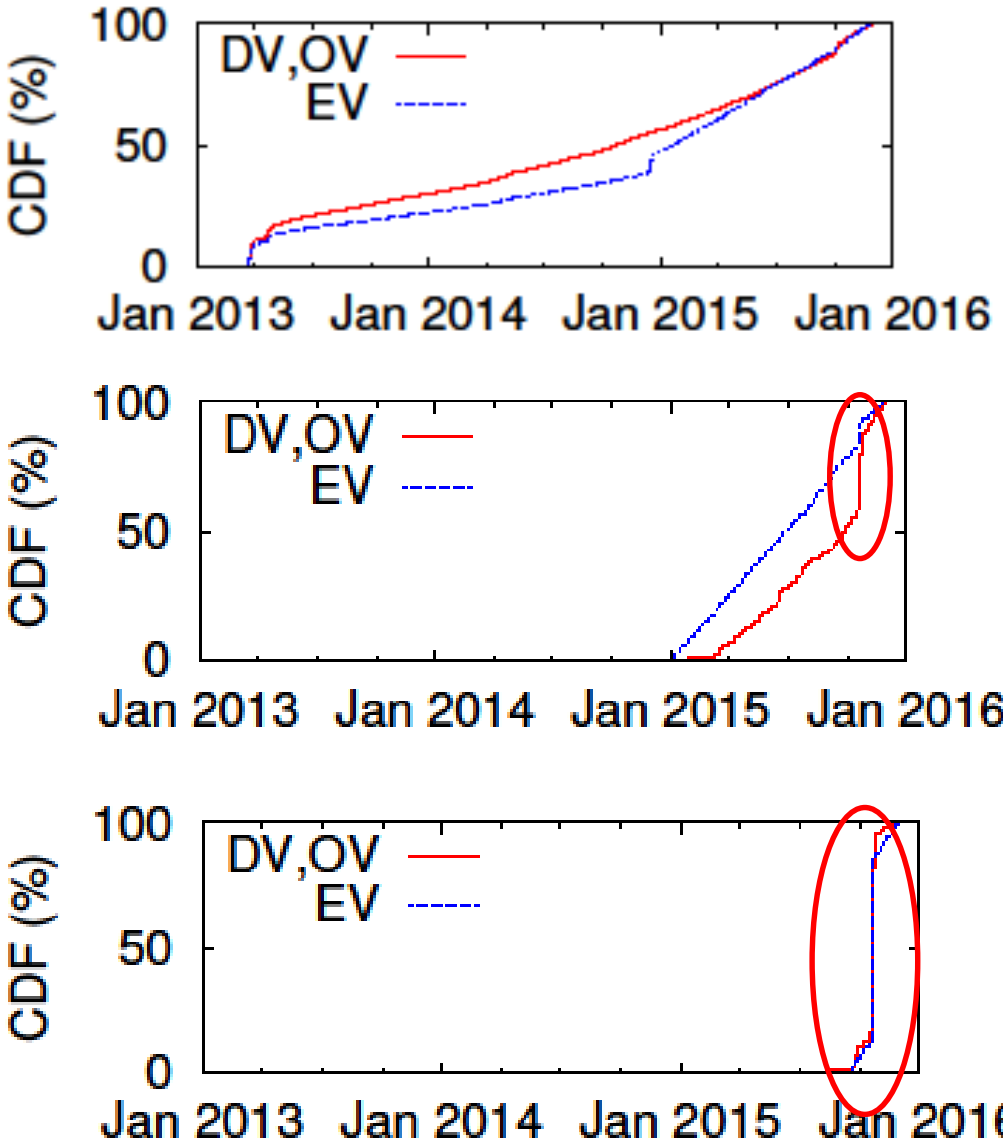

Popularity-based analysis

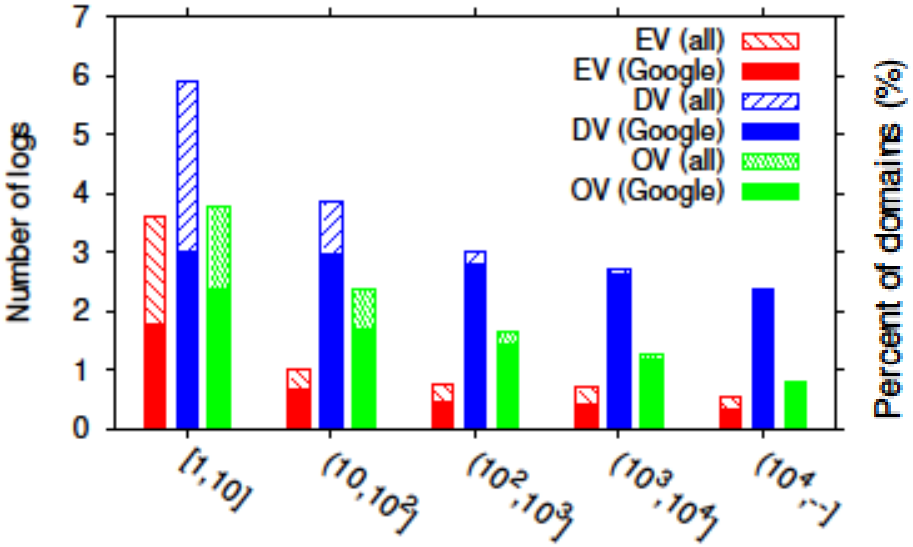

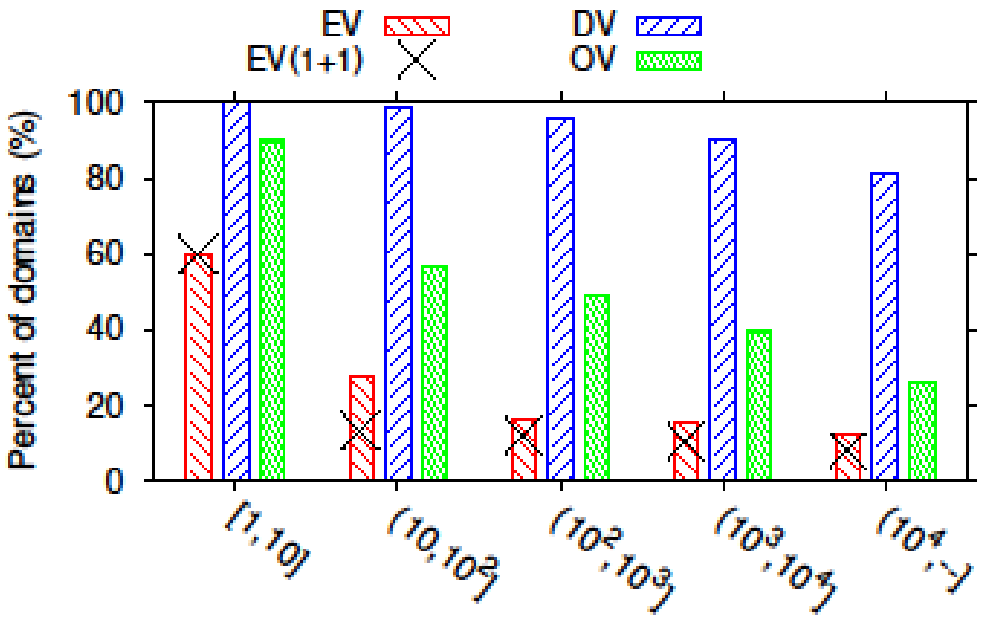

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
-
  -
  -
  -

Popularity-based analysis

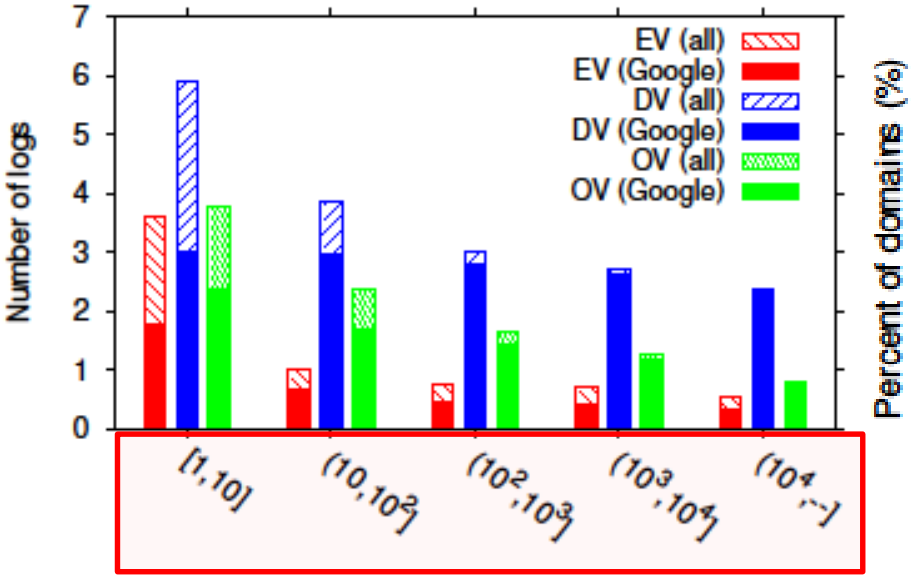

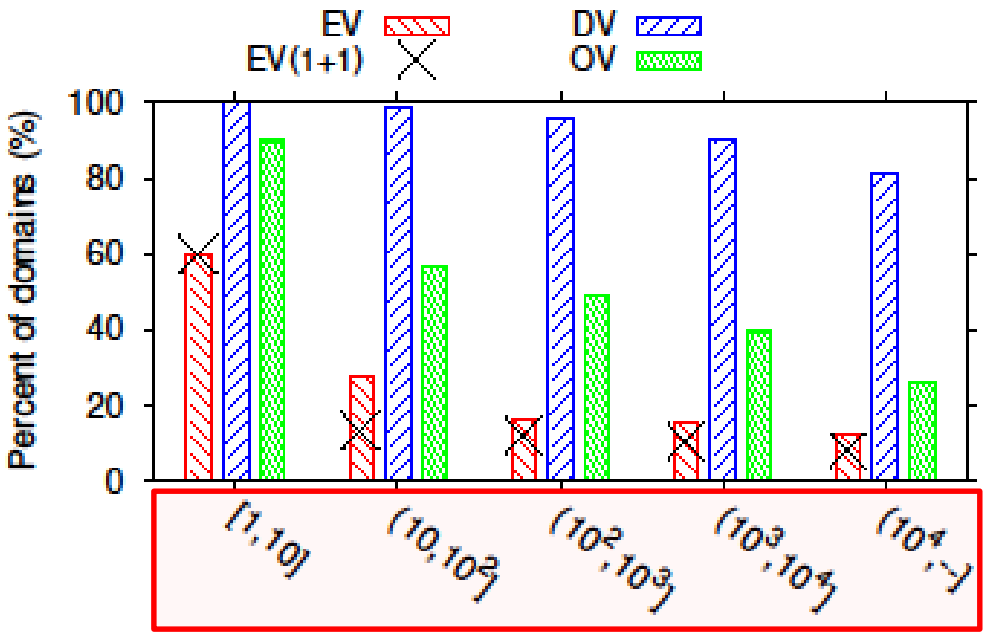

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
-
  -
  -
  -

Popularity-based analysis

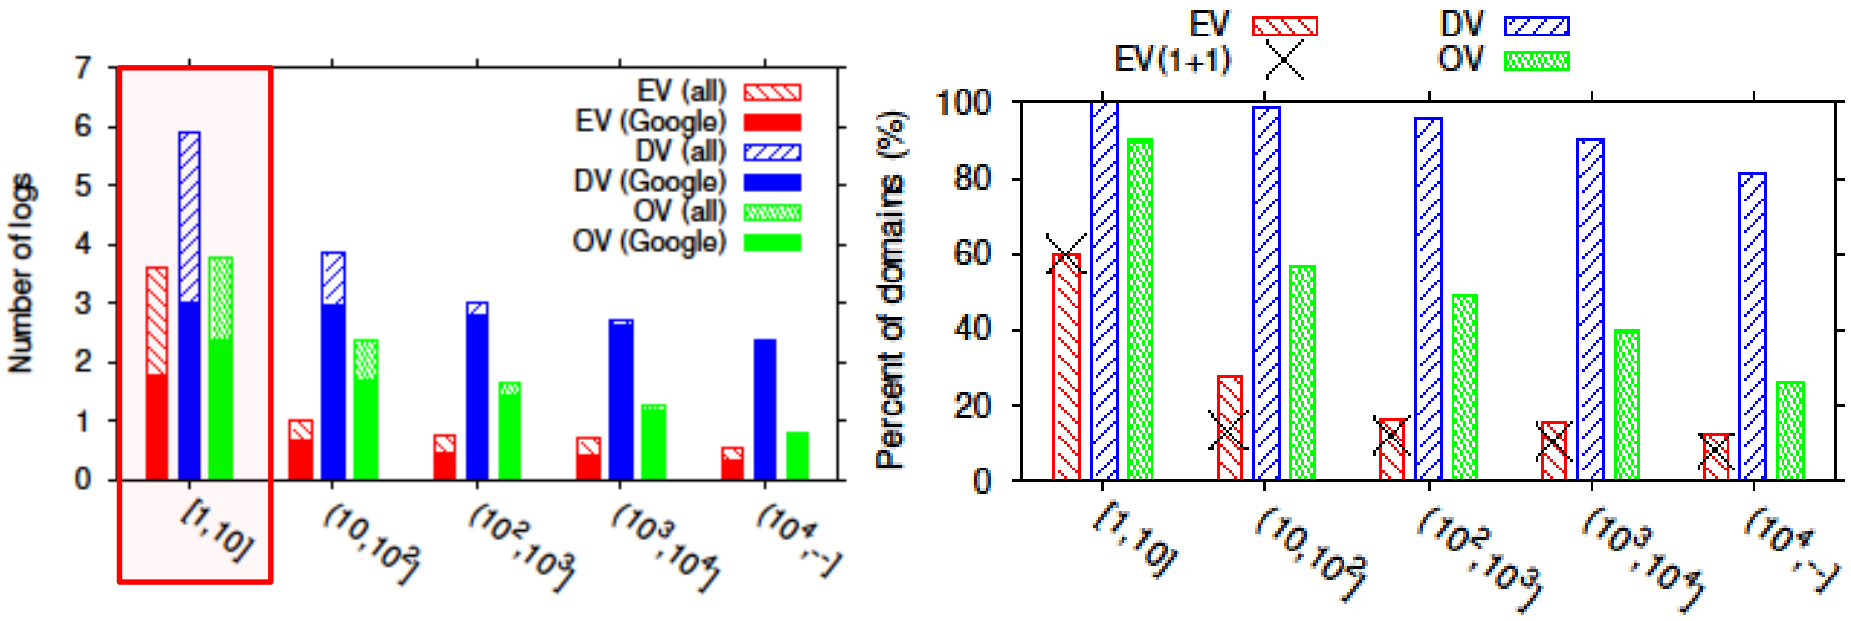

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
- Most popular (e.g.,top-10) domains best log coverage
  - Visible in most logs (across cert types)
  - Largest fraction of domains visible in at least one log (+DV high all types)
  - Largest fraction of domains that fulfill Chrome's EV policy (of course only applicable to domains with EV certs)

Popularity-based analysis

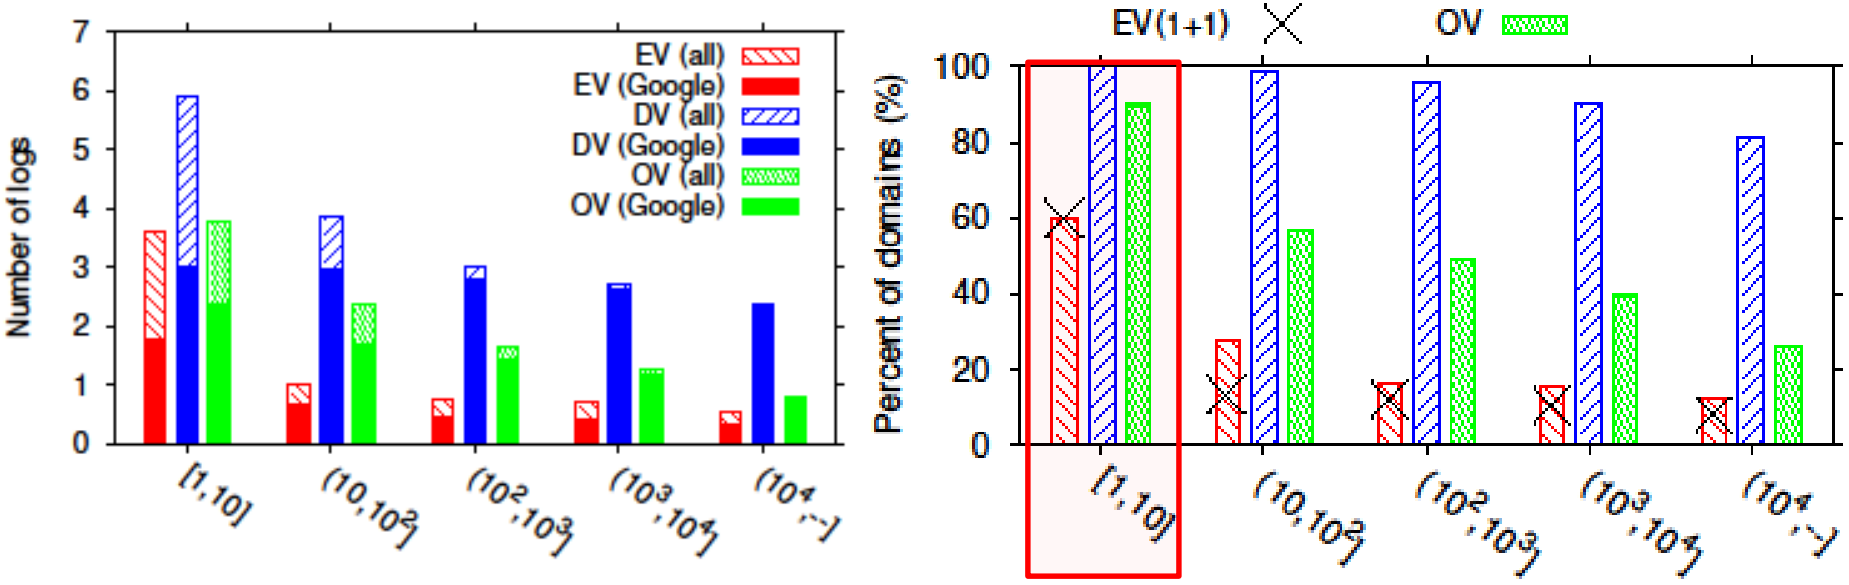

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
- Most popular (e.g.,top-10) domains best log coverage
  - Visible in most logs (across cert types)
  - Largest fraction of domains visible in at least one log (+DV high all types)
  - Largest fraction of domains that fulfill Chrome's EV policy (of course only applicable to domains with EV certs)

Popularity-based analysis

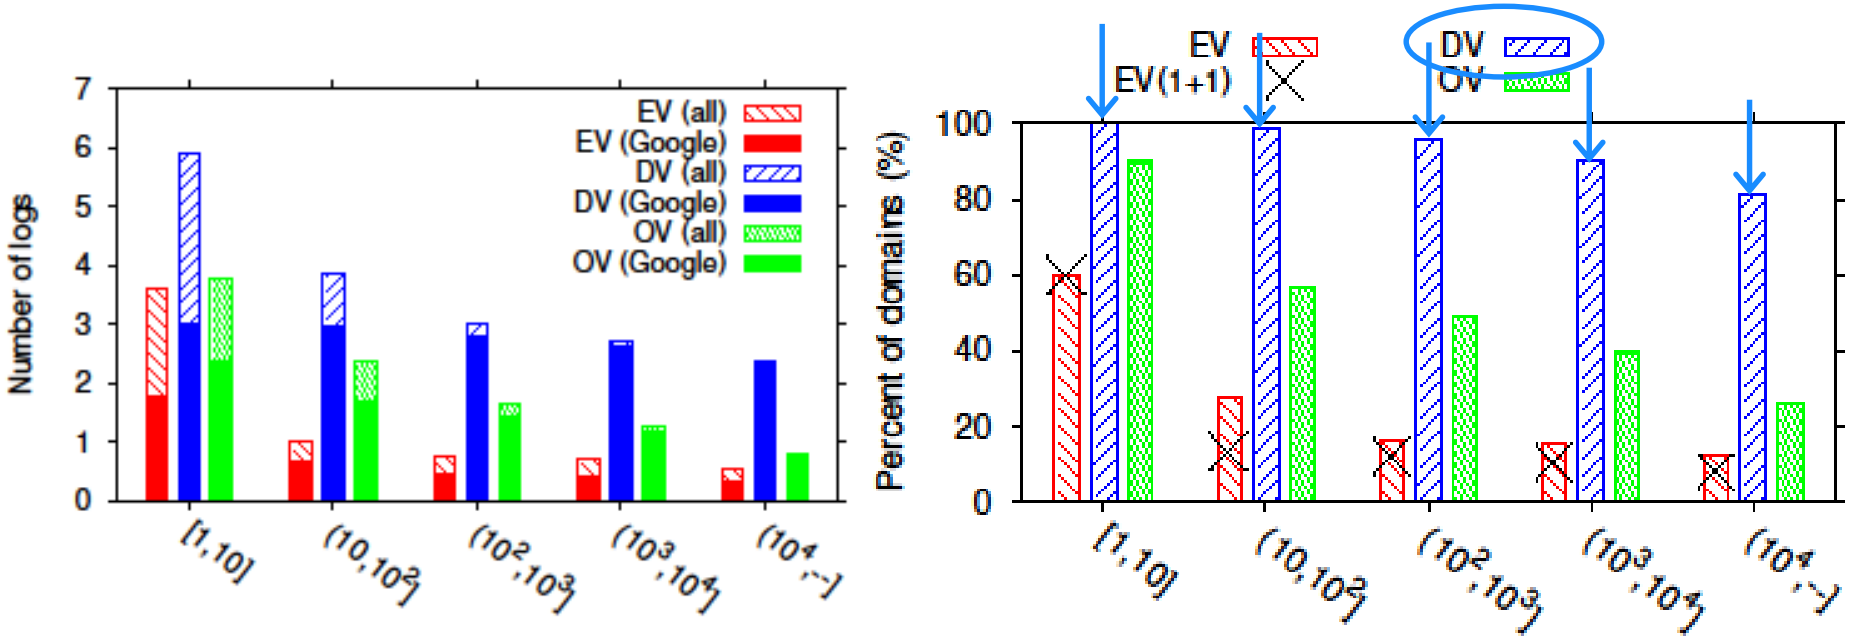

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
- Most popular (e.g.,top-10) domains best log coverage
  - Visible in most logs (across cert types)
  - Largest fraction of domains visible in at least one log (+DV high all types)
  - Largest fraction of domains that fulfill Chrome's EV policy (of course only applicable to domains with EV certs)

Popularity-based analysis

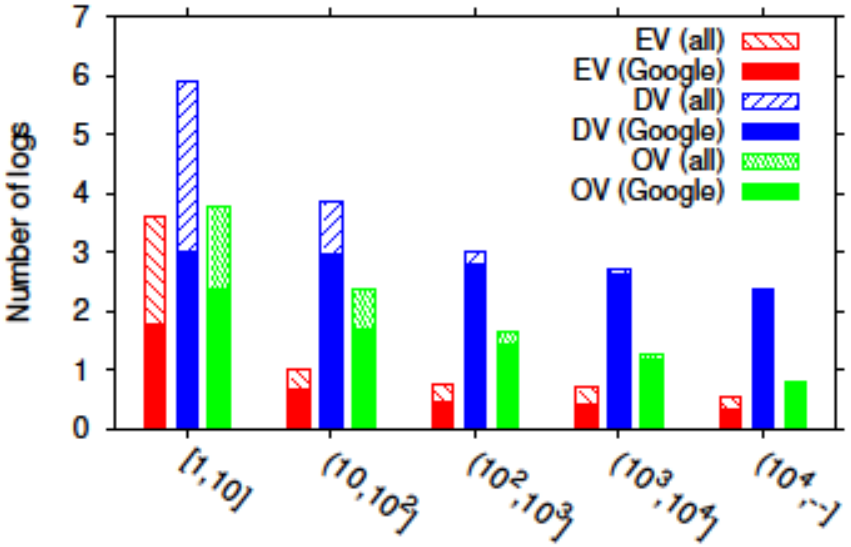

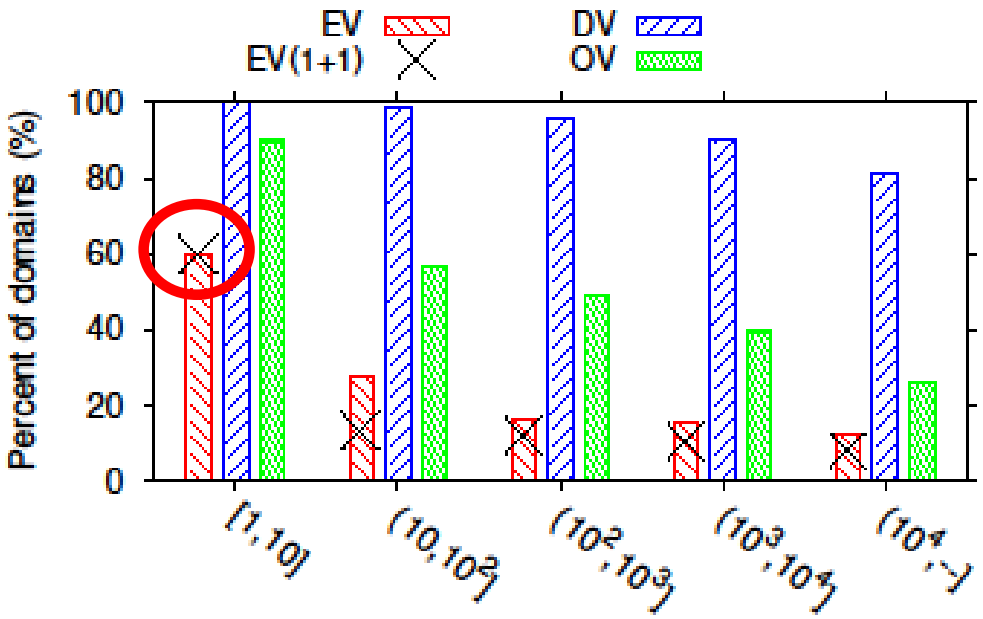

- Popularity of domains based on campus sessions
  - Rank of domains + logarithmic sized buckets
- Most popular (e.g.,top-10) domains best log coverage
  - Visible in most logs (across cert types)
  - Largest fraction of domains visible in at least one log (+DV high all types)
  - Largest fraction of domains that fulfill Chrome's EV policy (of course only applicable to domains with EV certs)

Conclusions

- Characterized eleven CT logs with basic monitor
  - All public at that time (3 Google, 7 CAs, Plausible)
  - Complemented with passive campus measurements
- Significant log differences based on operator; e.g.:
  - Google logs are crawl-based, use larger root stores, and are more representative of what is seen in the wild (e.g., by Chrome browser and campus users), including weaker keys, hashes, etc.
  - CA-based logs appear to be focused on helping CAs comply to Chrome's EV policy (and future DV extensions by Chrome and Firefox)

Conclusions

- Characterized eleven CT logs with basic monitor
  - All public at that time (3 Google, 7 CAs, Plausible)
  - Complemented with passive campus measurements
- Significant log differences based on operator; e.g.:
  - Google logs are crawl-based, use larger root stores, and are more representative of what is seen in the wild (e.g., by Chrome browser and campus users), including weaker keys, hashes, etc.
  - CA-based logs appear to be focused on helping CAs comply to Chrome's EV policy (and future DV extensions by Chrome and Firefox)

Thanks for listening!

A First Look at the CT Landscape: Certificate Transparency Logs in Practice

LINKÖPING UNIVERSITY

Niklas Carlsson (niklas.carlsson@liu.se)
www.ida.liu.se/~nikca/