Issues in Key Pinning

Chris Palmer and Ryan Sleevi,
Google, Inc.

Nondeterminism in Path Building

There are often multiple paths from an EE to a trust anchor, for various reasons. This can cause false negatives during pin validation.

Does enforcing pin validation for “known” or “system” roots solve enough of the problem?

Can we favor pinnable paths during path building?

Is this a problem for TACK, or just HPKP?

Pin Time-out and Revocation

An issue of control for site operators, and an ease-of-use issue. (Consider TACK's proliferation of keys, and increased granularity of control.)

TACK's "aging" (enforce a pin for no longer than it has been observed, with a minimum observe-before-enforce time and a maximum enforcement time) is clearly a good idea. Is it enough, or do we need explicit pin breaking?

How many keys can site operators really manage?

Are backup pins helpful enough to make a requirement?

Ease-of-Use

Does it matter, from a security perspective, how we communicate pins? No, **as long as pins are validated during TLS session setup and before application traffic**.

In TLS, in X.509, at application layer, elsewhere? The main distinction is ease-of-use and tooling for site operators. TACK, HPKP, and possibly other proposals make different application-independence/ease-of-use trade-offs.