HTTPS Token Binding & TLS Termination

Brian Campbell

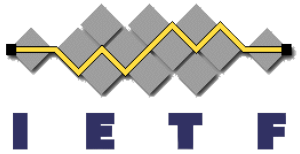

IETF 97
Seoul
November 2016

Situation

- Very common in HTTPS application deployments to have TLS ‘terminated’ by a reverse proxy sitting in front of the actual application
- For applications in such deployments to take advantage of token binding, some information needs to be communicated from the TLS layer to the application
- In the absence of a standard means of conveying the appropriate token binding information, different implementations will do it differently
  - Terrible for interoperability
  - A boon to unneeded complexity
  - Improved opportunity to get things wrong

Proposed Solution

- Work to standardize something in this WG!
- Hopefully not controversial

Two General Approaches

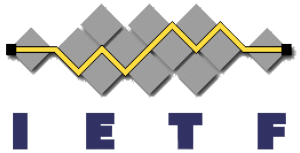

1. The TLS terminator validates the Token Binding Message and passes it (or some variation) along to the application
   - More work for the TLS layer
   - Easier reconciliation of supported key parameters
2. The application validates the Token Binding Message with sufficient info provided as headers by the TLS terminator
   - EKM, the negotiated key parameters
   - Hard to terminate the connection with the client
   - Not sure how renegotiation would work
- Miscellaneous thoughts
  - What about version?
  - TLS terminator must sanitize headers either way
  - Only one level of proxying supported
  - Applications likely need configuration

So...

- Does the WG think this is work worth pursuing?
- Feedback on the approach
- Write a draft
  - Me
  - You?
- IETF magic happens!