eduroam – a world-wide network access roaming consortium on the edge of preserving privacy vs. identifying users

Overview. eduroam is a world-wide roaming consortium for Wireless LAN access. It is a federated environment that spans more than 1.000 participating administrative domains in approx. 50 countries world-wide (see http://www.eduroam.org for details).
It is a publicly-funded consortium for users in the public sector (education and research institutions only). Privacy considerations are a high-profile topic: privacy of users is demanded by national and European regulations (EU Data Protection Directive, etc.); but at the same time, network abuse cases create a demand for identifying, blacklisting and maybe prosecuting users.

Basic operational model. eduroam's basic operation model is trying to meet these contradictory demands purely by using standardised technologies of IEEE and IETF. The combination of

- IEEE 802.11i and 802.1X: provides means to reveal identity and credentials only to the user's identity provider (IdP); also enables privacy of transmitted data on-the-air
- RADIUS proxy system: proxying over multiple hops conceals hotspot location from IdP
- EAP types w/ mutual authentication: prevents credential leakage to unauthorised third parties
- EAP types w/ support for anonymous outer identity: conceals user's identity from hotspot

provides minimum disclosure of any sensitive information to only the parties which are authorised to get that piece of information. At the same time

- rigid logging of all authentications
- defined contacts for hotspot and identity provider operators
- defined escalation procedures in case of abuse

enable a hotspot operator to pinpoint the exact user if need be. This requires the cooperation of the identity provider (which may involve authorities in the hotspot country and IdP country).

Recent refinements and their challenges. On operator's requests, the basic operational model has been augmented with extra information recently. These refinements have an impact on privacy, and efforts need to be taken to limit this impact.

- RFC4372: Chargeable User Identity (CUI)

  Thanks to the extensive privacy support in the basic operational model, it is possible for an eduroam user to appear as a different user to a hotspot every time he logs in: by changing his local MAC address and his EAP outer identity. Blacklisting of such a user in case of abuse becomes impossible for the hotspot operator since there is no invariant handle to identify the user. The operator needs to consult with the IdP, which is unlikely to happen in real time. The implementation of CUI at Identity Providers allows hotspot operators to recognise users on re-entry, while still not revealing the actual identity.

  This measure is seen as privacy-invading by some of our legal specialists, because a persistent identifier for the user can be seen as personally identifiable information, which is subject to the EU Data Protection Directive.

One possible unintended misuse of this attribute is that multiple hotspot providers could cooperate and create mobility profiles of users, since the CUI value would globally mark an individual. Eduroam is mitigating this threat by generating CUI values which are different for each hotspot (but of course stay persistently the same within any given hotspot). This is achieved by using unique identifiers for hotspots, see below, and using these identifiers as an input to the CUI generation algorithm. This gives the CUI similar properties as the eduPersonTargetedID attribute in web federations.

- RFC5580: Operator-Name

  The use of CUI required to vary the returned CUI value per hotspot. RFC5580's Operator-Name allows to tag authentication requests with the realm name of the operator. This enables per-hotspot CUIs and also helps in debugging connectivity problems because the Identity Provider can immediately see which path the authentication took in the proxy chain.

  Of course, sending Operator-Name reveals a user's visited location to the identity provider immediately, which reduces the user's privacy. Surprisingly, this is not commonly seen as a big issue. The reason being that the next generation of eduroam authentication routing is going to be based on RADIUS/TLS with dynamic discovery, which will anyway create a directly tracable link between hotspot operator and identity provider.

User Perception. Surprisingly, few end users seem to notice or care about privacy. There is a significant gap between what technology allows a user to do to conceal his identity, and what users actually do. (Example: while the identity *could* look different for every reauthentication with a different MAC address and outer identity, only a tiny fraction of users makes actual use of these features). As a result, much of the argumentation about privacy rests with the operational and research staff, and is not driven by end user demand.

On a more general note, any discussion about user privacy should be twofold: One, what are the theoretical limts in protocols regarding privacy support, Two: to what extent is user privacy actually achievable within these limits, given that end user's knowledge of protocols and options is usually rather limited.

One related, and particularly interesting question is how to make privacy-supporting elements work without requiring corresponding expertise from the end user.Example: an IEEE 802.1X supplicant could randomize a device's MAC address prior to every authentication, without user interaction. This would remove one element that could otherwise be abused to create mobility profiles about the user.