Short talk

SPORES: Transferring files as spores through the wind instead of up in the Cloud

Daniel Bosk and Sonja Buchegger, KTH

Abstract SPORES is a protocol that provides a file transfer service that does not rely on a central service but provides strong privacy guarantees.

Consider that Alice wants to share a file with Bob. She uploads the file to "the Cloud" (someone else's computer) and sends a link to Bob. Bob fetches the file from the third-party storage. This allows the third party to infer much information. SPORES changes this.

Alice has several devices (laptop, smartphone, etc.) with varying online–offline patterns, we call these devices her device squad. The squad knows which device has which files and each device of the squad has a model of the other devices' online–offline patterns. Alice's devices can coordinate the sending of the file. Similarly, Bob also has a device squad, and his squad can coordinate receiving the file. Instead of Alice giving Bob a link to the file in the Cloud, Bob gives Alice a route where the destination is his squad (the *set* of Bob's devices).

At the core of our proposal lies a stateless, probabilistic onion-routing protocol called SPOR. SPOR relies on a set-based mix-header format based on Sphinx [@Sphinx] (we dubb this extension SphinxES, for Sphinx Extended for Sets). Given an algorithm for secure peer sampling, a probabilistic model of node availability, SPOR will populate the layers of SphinxES with sets of devices to reliably deliver packets from one set of devices to another over a network of mix-nodes with high churn.

This yields several interesting properties, among others: SPOR can provide connections over a network of unreliable devices, unlike Tor's network of dedicated, always online devices; SPOR removes the weakest-link bandwidth problem present in Tor, sets of nodes at each layer provides a wider path and thus more bandwidth; at the same time, SPOR still provides privacy guarantees en par with Tor.