Factsheet

SamKnows Shield – Security

Security

Security

The SamKnows Internet Measurement Platform is relied upon by ISPs, internet companies, governments, consumers and academics worldwide. It is therefore imperative that the highest possible levels of security are maintained. Every aspect of the platform is designed with security in mind.

Firewalls

All SamKnows servers (infrastructure and test servers) utilise iptables or firewalld as firewalls to restrict access to only necessary services from necessary hosts. Internal infrastructure servers such as databases are completely inaccessible except from other SamKnows servers. Only servers and ports that require access from the outside world - for example TCP Port 80 on a webserver - would have that port fully open. We also actively monitor firewall ruleset integrity through our monitoring tools.

SamKnows access

Internal access to servers or data is restricted where only the infrastructure team and select other personnel have access to such systems; staff only have access to servers required for their role. Direct access to servers must be done through either a bastion server or through the secure office network and is done through password authentication or key authentication on a per-user basis managed centrally using LDAP.

Personally identifiable information

The amount of personally identifiable information stored by SamKnows is limited to:

- Full name
- Email address
- Address

All metadata and personally identifiable information is both encrypted at rest and in transit on SamKnows servers. It is never stored or transmitted to/from the Whitebox.

Server management

All servers are centrally managed by Puppet and are kept up to date with the latest security patches. Puppet is controlled by our secured puppet master servers and git repositories that require two-factor authentication or SSH authentication. Only authorised SamKnows personnel have access to make puppet changes.

Virtualisation

All SamKnows-hosted infrastructure servers utilise dedicated physical hardware only and are in no way shared with any other users. Where an ISP chooses to host all or part of the platform themselves internally, virtualised machines may be used. However, the virtual machines should be provisioned according to the SamKnows hosting requirements and other applications should not co-exist on these servers.

User password storage

Users' passwords are never stored or transmitted in plain text. Passwords are hashed and salted before being stored in the database. Passwords are hashed using the SHA-256 hashing algorithm.

Transfer of measurement data

All data transferred between SamKnows infrastructure is conducted over SSL, authenticated using verified public/private key pairs.

CPE agent management traffic encryption

Our data collection infrastructure exposes APIs over HTTPS to allow measurement agents (Whiteboxes and CPE) to report test results, check for software and configuration updates. This is secured over TLS using a private SamKnows CA. The SamKnows agent in the router trusts the private SamKnows root CA. No additional CAs are trusted. Of course, in the case of a router integration, the router itself may already trust other CAs.

SamKnows One traffic

The SamKnows One analytics platform is exposed over HTTPS. All communications to the web servers are secured by TLS using SSL certificates obtained from a public CA.

Data exports

Data extracts provided to clients are made available via a secure FTP & HTTPS server. Each individual client has their own secure area on this server (isolated from the rest of the operating system) where their data can be accessed via a complex password. No user has the capability to see anything other than their own data.

Data center security

SamKnows stores data on servers located in PCI DSS compliant datacentres. Data is encrypted both at rest and in-transit and is not stored outside of the EU or GDPR 'third-countries'.

All infrastructure servers that store data utilize dedicated physical hardware only and are in no way shared with any other users (we do not used shared virtual machines).

Whitebox placement in the home network

Whiteboxes are typically installed in-line (i.e. “man in the middle”) in the user’s home network. This is required in order to identify the presence of cross-traffic. Modern versions of the Whitebox utilize a hardware switch internally, with the SamKnows applications only having access to port counters via the switch’s API. User-generated cross-traffic does not pass across the CPU, so is unavailable to packet capture applications.

Router security (with physical access)

As with any hardware device, a sufficiently well-resourced and motivated attacker can compromise the device if they have physical access. Should this occur, it is important to stress that this:

- Does not grant the attacker elevated access to the SamKnows backend infrastructure. The sole interface the Whitebox/router has to SamKnows is via the HTTPS interface the DCS exposes to report test results or check for updates
- Only allows the user to modify the software on the individual router they have physically compromised.
- Does not make available any personally identifiable information or passwords as these are not stored in any form on the Whitebox/router

Router security (without physically access)

Due to the placement within the home network, there is no surface area with which to exploit direct access to the Whitebox remotely (because there is no direct access).

The only possible attack vector is via the DCS update procedure. All possible efforts have been taken to prevent such an attack. As highlighted previously, this is secured via TLS using a private SamKnows CA (so even a public CA breach and DNS hijack cannot compromise security here). The DCSes themselves are accessible to only a handful of SamKnows staff, and only via LDAP authenticated and authorized SSH connections. The DCSes are only accessible via SSH from a limited set of trusted SamKnows IP addresses.