F-SECURE SECURITY CLOUD

Purpose, function and benefits

October 2015

CONTENTS

F-Secure Security Cloud is a cloud-based digital threat analysis system operated by F-Secure Corporation. It consists of a constantly growing and evolving knowledge base of digital threats fed by data from client systems and automated threat analysis services.

Given that data gathered from users' devices is essential to the proper functioning of Security Cloud, we have taken careful steps to ensure that F-Secure's privacy principles are strictly upheld. This whitepaper provides an overview of the steps taken to ensure that Security Cloud's data collection and processing is done in a way that will not reveal sensitive information about an individual or device.

F-SECURE LABS
Security Technology
Whitepaper

F-Secure

F-SECURE SECURITY CLOUD IN BRIEF

F-Secure Security Cloud (*later: Security Cloud*) is a cloud-based digital threat analysis system operated by F-Secure Corporation. It consists of a constantly growing and evolving knowledge base of digital threats fed by data from client systems and automated threat analysis services.

The adoption of a cloud-based threat analysis service offers many benefits over traditional methods. Data accumulated from a large number of client nodes allows us to build an accurate, real-time picture of the global threat situation. This approach also allows us to faster utilize recently gathered threat intelligence to protect customers.

Products from F-Secure utilize Security Cloud to a varying degree. Traditional anti-malware products rely on a combination of local scanning engines and Security Cloud services. Products that are designed to be lightweight can rely solely on the Security Cloud for threat detection. The latter approach is used to reduce CPU, memory and network bandwidth usage on a client device.

Given that data gathered from users' devices is essential to the proper functioning of Security Cloud, we have taken careful steps to ensure that F-Secure's privacy principles are strictly upheld. All data collection and processing is done in a way that will not reveal sensitive information about an individual or device. Confidential user-generated content is not sent to the cloud. More information about privacy and security can be found later in this document, and in the Security Cloud Privacy Policy.

Version history

March 26th 2015	v1-00, MA
April 10th 2015	v1-01, MA
October 8th 2015	v1-02, VL
October 29th 2015	V1-03, VL

DISCLAIMERS

- **The purpose of this document is to help customers better understand how F-Secure products function and the benefits Security Cloud provides. This document is not designed to be a legally binding agreement that defines the content of products and services provided by F-Secure Corporation.**
- **Security Cloud is a constantly evolving set of systems and processes. This document may become partly inaccurate as this evolution takes place. F-Secure Corporation will update this document every time major changes are made to our systems or processes. The latest version will always be available on the F-Secure website.**
- **Any metrics presented in this document are valid at the time of publication. Metrics may change over time. Presented metrics should therefore be interpreted as approximate ballpark figures.**

SECURITY CLOUD BENEFITS

Global malware behavior tracking

Most security products from F-Secure utilize reputation services for threat recognition. Client software calculates a cryptographic hash for an object (typically an executable) and performs a network query (of the calculated hash) with Security Cloud. This process provides us with a very granular picture of the worldwide distribution of software and how malicious programs spread. Malware detection is based on the analysis of samples and the behavior of executables on client computers. Security Cloud enables us to examine how suspected malicious programs behave globally and are spread between computers, countries and continents.

Sharing threat data forms a protective network

New threats are constantly appearing in the wild. Any user with a Security Cloud-enabled product may be the first to encounter a new threat and provide our systems with scanning results, file metadata, malware behavior data, and a sample of the file. All users will subsequently benefit from this data and receive protection faster and more accurately than with traditional security products. Essentially, Security Cloud-enabled products form a network where they cooperate and share threat data.

Enhances traditional malware scanning

Traditional security products rely mainly on malware scanning engines installed locally on the device. Our latest product offerings combine traditional local scanning approaches coupled with lookups to the Security Cloud. Security Cloud can provide the latest information about new threats before they have been incorporated into definition databases commonly utilized by traditional scanners. Security Cloud can also factor in the global behavior of programs to speed up the detection of new malware. This enables traditional anti-malware products to utilize up-to-date information from the cloud while retaining the ability to operate in an off-line mode.

Autonomous malware scanning

Products can utilize Security Cloud as their main anti-malware scanner. This function may require the user's consent to upload certain types of non-executable files from the device to Security Cloud for scanning. Most files will not need to be uploaded because Security Cloud already has enough data about them to make a verdict. This way the user can receive sufficient protection even if only a very small fraction of all their files are uploaded for scanning. This behavior saves resources and is especially beneficial for mobile devices.

Sample collections

Security Cloud includes a comprehensive archive of objects (files, URLs, etc.). This collection is updated by a number of independently sourced sample feeds. Malicious and suspicious files uploaded for scanning by users of Security Cloud-enabled products also make up part of this collection. A limited collection of common clean files are also included in this collection. By re-scanning these sample collections periodically, Security Cloud is able to automatically deliver clean or malicious verdicts for a great deal of queries made by client machines.

Automated analysis

In order to keep up with today's rapidly changing threat landscape, automated processing and analysis of malware is required. Security Cloud sources a variety of automated analysis processes (all maintained by F-Secure) on samples in our backend. This analysis feed enables Security Cloud to classify a significant amount of new malware quickly and automatically. This automation enables us to provide fast responses to new threats.

Resource savings

Computing power

Security Cloud-enabled products can offload CPU-heavy scanning tasks to the cloud service. This saves CPU-time, and thus battery life, especially on mobile devices. A reputation query to Security Cloud is often faster than scanning an object on the local device.

Device storage

Security Cloud enables us to create products that completely omit local scanning engines. This makes the security product more lightweight to both install and run. Such products are developed for devices where storage space is scarce (such as mobile devices, appliances and tablets).

Bandwidth

Security Cloud can save significant amounts of bandwidth by reducing or eliminating the need to update definition databases on the local device. Products with local engines typically update definition databases rather infrequently. Products that do not utilize local engines can save significant network bandwidth usage (an estimated 80-90%) by skipping these database updates completely.

Detection speed improvement

The definition database update process used in traditional anti-malware products is often a heavy and lengthy operation. Each database needs to be compiled, tested and published with tight quality controls. Clients relying on traditional scanning engines must download and apply updates. Although products with local engines receive periodic updates to their definition databases, the frequency at which these updates arrive cannot be compared to the reaction speed of cloud-based systems. Because Security Cloud relies on a cloud-based reputation database, new detections are (in effect) instant for clients using Security Cloud services.

Flexibility

Security Cloud provides a wide range of independent services. This enables product designs that use all, or subsets of those services. This allows F-Secure to maintain a comprehensive portfolio of products on different operating systems that all benefit from our centralized threat intelligence.

HOW DOES SECURITY CLOUD WORK?

Overview

Security Cloud is a complex modular online service that protected devices connect to through the Internet. The picture below provides a coarse schematic overview of Security Cloud's most important modules.

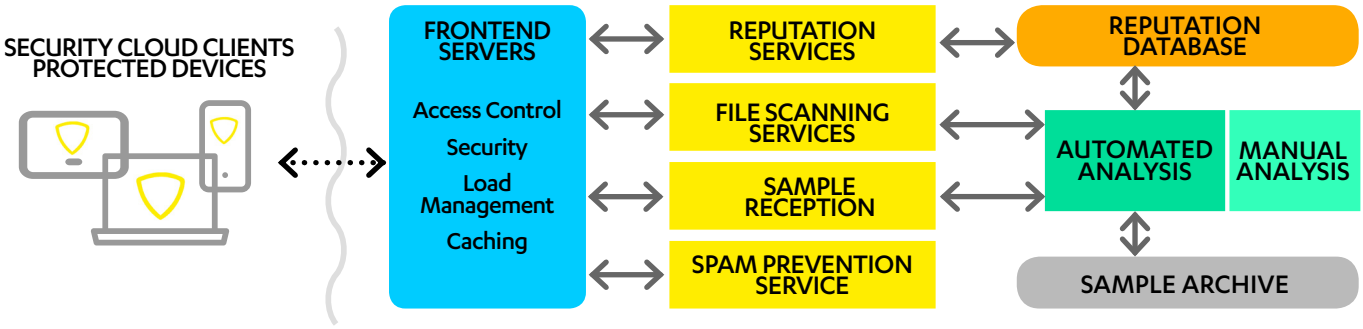

Security Cloud Clients – Protected devices

Security Cloud provides a set of services that individual client-side security components can connect to. These client-side components are developed by F-Secure. Products using Security Cloud are mainly developed by F-Secure Corporation, but some third parties may have an agreement with F-Secure to utilize one or more of these services in their own products.

Reputation services

One of the core services in Security Cloud is the reputation lookup service. This service enables clients to query the reputation of computer networks (such as home router and public Wi-Fi) and objects, (such as files or URLs).

Networks are verified by connecting devices with a service that identifies altered DNS settings, and verifies if these settings have been changed to expose people to attacks.

Objects are checked by calculating the object's cryptographic hash and sending the hash to the reputation service. The reputation service returns a verdict for the queried hash (trusted, malicious or unknown). Security Cloud may request metadata for unknown objects, or the object itself, for further analysis. Clients respond to such requests according to their settings and privacy policies.

File scanning services

Lightweight products rely mainly on Security Cloud's file scanning service for malware analysis. This service works in conjunction with our reputation service. If the reputation of a file is unknown, the client can be instructed to upload the file to our backend for malware analysis. The results of this analysis may cause the file to be flagged as suspicious and sent on for further processing in our automated analysis systems. This deeper analysis is a potentially time-consuming process, and hence it is not performed on every unknown file we receive.

Sample archive

Our sample archive contains files that Security Cloud has received from a variety of sources, including participating devices. Both malicious and suspicious files may be present in this collection. Malicious files are generally archived permanently, whereas suspicious files are removed as soon as they are not deemed malicious. A collection of very common clean files is also maintained. Clean files are not collected directly from customer devices.

Spam prevention service

The unsolicited messaging (spam) prevention service works by analyzing hashes of some header fields from email messages in an attempt to classify the message as spam. Message content itself is never uploaded to Security Cloud.

Analysis systems

Today's rapidly developing threat landscape requires a highly automated approach to malware analysis. Files classified as suspicious are received from many sources and go through several analysis steps. These steps can include, but are not limited to, metadata analysis, structural analysis, statistical analysis, and behavioral monitoring. F-Secure's backend AI-driven automation examines the file's metadata and analysis results, and either performs further analysis steps, or classifies the object as either clean or malicious. Unclear cases can be flagged for manual analysis and may be examined by a human.

SECURITY CLOUD METRICS

Security Cloud is a high-volume system. It is a critical component for most computers protected by F-Secure's technology. This section presents some data that gives a picture of the volumes Security Cloud is dealing with.

Number of queries per day received by Security Cloud's servers	About 5 billion per day (mid 2015)
Number of unique samples received per day by Security Cloud	About 1 000 000 (mid 2015)
Average number of items (files and URLs) checked per day by a client's local Security Cloud module	About 28 000 per day on average
Locally resolved Security Cloud queries for files (client cache hit rate)	About 60%
Locally resolved Security Cloud queries for URLs (client cache hit rate)	About 96%

SECURITY AND PRIVACY

Security Cloud is a system where strict security practices are always applied. Our servers house a massive collection of malicious software that could be harmful to both F-Secure's own systems and the rest of the Internet. Strict security practices are also applied when dealing with any data collected from client devices. All data is anonymized on the client before transmission to Security Cloud. Data that could be used to determine the identity of the device or the user of that device is never collected. All network connections between clients and our Security Cloud are encrypted.

PRIVACY PRINCIPLES

Privacy is one of F-Secure's core values. Privacy is carefully considered during all Security Cloud planning and development work. We collect only the minimum amount of data from clients necessary to provide our services. Every transferred bit must be justifiable from a threat fighting perspective, and data is never collected for presumed future needs. The following table documents our privacy principles in full detail.

Minimize upstream of technical data	Data about the customer's computer is not transferred and collected unless the data is essential for providing the protection service.
Do not upstream personal data	The system is designed to not send any information that can identify the person using the computer. Such data is not needed for the operation of Security Cloud. Security Cloud-enabled clients use several algorithms to prevent private data from being transmitted and filter out such data from, for example, URLs and file paths.
Use anonymous identifiers	Clients generate unique anonymous IDs that can't be connected to the user's, license owner's, or device owner's true identity. These kinds of IDs are used when repetitive connections from the same device need to be tracked.
Prevent backend data consolidation	Clients use several different unique anonymous IDs for different connections to Security Cloud. This makes it impossible for F-Secure Corporation to profile users by comparing user IDs from different systems.
Do not store IP-addresses	The customer's IP-address is never stored. City-level geo-mapping may be done and the result stored if it is beneficial for providing the service.
Do not trust the network	All network transfers are encrypted using strong crypto. Asymmetric encryption is used for authentication when appropriate.

SECURITY PRINCIPLES

Secure by design	A system is never secure unless it has been designed to be secure. Security can't be added as the last feature in a project. This is something that is acknowledged when developing Security Cloud and its related systems.
Encrypted network traffic	Data is never transferred in plaintext over the Internet. Encryption is, in addition, used to ensure the integrity of various objects. F-Secure utilizes a mixture of generally available crypto libraries and protocols, and customized crypto code.
Separated malware environments	Storing and testing malicious software is a challenging task that we have over 20 years of experience in. All malware handling is performed in networks separated from the Internet and other F-Secure networks. Storage and testing networks are separated from each other, and files are transferred using strictly controlled methods.
Professional monitoring	All critical systems in Security Cloud are monitored by F-Secure personnel. All systems storing or testing malware are hosted by F-Secure itself.
Controlled access	Only a limited number of F-Secure employees have access to Security Cloud's critical systems. Such access is granted, revoked and documented according to a documented and controlled process.
Open attitude	The most fundamental principle in all security work is an open and humble attitude. We have put a lot of effort into securing Security Cloud, but the work is never finished. A secure system can only be maintained by promoting an open attitude where problems in the system are reported, analyzed and fixed promptly. This attitude includes public openness if we encounter incidents that put customer security in jeopardy.

SEE ALSO

Security Cloud privacy policy

https://www.f-secure.com/en/web/legal_global/privacy/security-cloud

CONTACT INFORMATION

If you have any further questions about Security Cloud, please contact:

F-Secure Corporation
Tammasaarenkatu 7
PL 24
00181 Helsinki
Finland

https://www.f-secure.com/en/web/home_global/support/contact