16-11-2017

Deliverable D8.4
Certificate Transparency Log v2.0 Production Service

Deliverable D8.4

Contractual Date:	31-10-2017
Actual Date:	16-11-2017
Grant Agreement No.:	731122
Work Package/Activity:	8/JRA2
Task Item:	Task 6
Nature of Deliverable:	O (Other) – demo
Dissemination Level:	PU (Public)
Lead Partner:	DFN (LRZ)
Document ID:	GN4-2-17-16C404
Authors:	L. Nordberg (NORDUnet), M. Ahltorp (NORDUnet), I. Golub (PSNC), D. Schmitz (DFN (LRZ))

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 731122 (GN4-2).

Abstract

This document gives an overview of the GÉANT Certificate Transparency (CT) log service being developed by JRA2 T6, a new IETF standards-based approach to discover misissued certificates. It explains CT and describes the CT log service, its benefits and users, technology and architecture, and operational aspects. It also provides a roadmap for GN4-2 Period 2.

Table of Contents

Table of Figures

Executive Summary

The GÉANT Certificate Transparency (CT) log service is a new development in GN4-2, and continues work originally started by SUNET/NORDUnet. It is a new, Internet Engineering Task Force (IETF) standards-based approach to discover misissued certificates. By providing transparency of all issued certificates through the log, the CT log service enables its users – domain owners, certification authorities (CAs), browser vendors and web users – to monitor and audit the certificates for any given domain. This helps to strengthen the trust and security of the web space and thus of the Internet as a whole.

Additional benefits that this service brings to the community are direct control over service availability; low latency for accessing the service; the potential for retrieval of research data, taking privacy concerns into account; increased legitimacy for CT in general, resulting in a more secure web; and increased trust among entities participating in the CT log service.

The GÉANT Certificate Transparency log service has three main aspects: installation of the CT log server, and providing two types of support to users: in using the service and/or in contributing to the GÉANT CT log service by providing CT architecture components in their own domain.

This document gives an overview of the Certificate Transparency log service that is being developed in GN4-2. It describes how it is structured, the benefits of the service, and to whom it is offered. It covers the technical background of the service, as well as the operational aspects, namely, supporting infrastructure, service policies, service metrics, and operations and support teams. The document also provides a service roadmap for GN4-2 Period 2, and concludes with an overall assessment of the service.

The GÉANT Certificate Transparency log service is planned for production in spring 2018.

1 Introduction

Certificate Transparency (CT) is a new, Internet Engineering Task Force (IETF) standards-based approach to discover the misissuance of domain certificates by using a public, append-only log. Malicious or undependable certification authorities (CAs) can create certificates for a domain without the knowledge, instruction or approval of the domain owner. CT provides a tool to detect such actions and can thus help in their fast and successful resolution.

The development of a CT log server by Joint Research Activity 2 Network Services Development, Task 6 Network Security (JRA2 T6) during Period 1 is a continuation of work begun by SUNET/NORDUnet prior to GN4-2. The development is shaped towards supporting and delivering the GÉANT Certificate Transparency log service, which is scheduled to transition into production during Period 2.

This document gives an overview of the CT log service that GÉANT will provide to its users; a demonstration of the production service[1] will be scheduled once that transition has taken place.

Section 2 describes the service, beginning with an explanation of CT and including the service's benefits and users.

Section 3 provides a technical description and Section 4 considers service operations, covering supporting infrastructure, operations and support team, policies and metrics.

Section 5 explains the Certificate Transparency service roadmap for GN4-2 Period 2. Section 6 presents an overall assessment of the service.

[1] "Other: demo" was the Type stated in the Technical Annex for this deliverable.

2 Certificate Transparency Service

2.1 Certificate Transparency

Certificate Transparency is a system for storing website certificates in public, append-only untrusted logs, as described in IETF RFC 6962 [RFC6962]. CT is recognised by browser vendors and certification authorities as an important and innovative technology that increases trust in the general Internet, including R&E networking.

An important effect of CT is that domain owners are able to monitor logs for certificates issued for their domains. They can thus see the certificates that they are using, but also potential certificates that have been issued without their knowledge. This provides a novel and powerful tool for domain owners to control the certificates issued for their domain.

Web browsers and other HTTPS clients that use CT require certificates to be present in a number of public logs in order to trust and use them. This makes it possible to efficiently detect and prove misissuance of certificates. Every entity that handles a certificate can submit it to a log, together with the certificate chain from a known certification authority (CA). This means that CAs, web browsers, web server operators and other entities that see a certificate can submit a certificate to CT logs.

The append-only property of a CT log means that an entry in a log never changes, is never removed and is never moved to another position in the log. This fundamental property of a log is easily checked by asking the log for proofs which can be verified without the need for any trust in the log itself. By adding a hash of each log entry to a Merkle tree and signing the tree head with the log key, log servers make it possible for auditors to efficiently verify the consistency of the log by retrieving inclusion and consistency proofs from the log and making the same calculations themselves.

2.2 Service Description

The GÉANT Certificate Transparency log service has several work areas:

- To establish, provide, maintain and manage a central CT log server, open for anyone to submit certificates to it or to check existing certificates within it. In the production service, the log will be configured to accept certificates signed by CAs in all the common trust stores.
- To provide support to the GÉANT community in understanding how CT logs can help them to protect their network and contribute to strengthening Internet security, as well as to provide

help in using the log to protect the web space of the domain owners, primarily from the GÉANT community.

- To provide support to organisations, primarily from the GÉANT community, that want to contribute to the GÉANT CT log service by establishing and running a Certificate Transparency architecture component themselves.

GÉANT is a vendor-independent party that supports the interests of all targeted users of the CT log service, which might not be the case with logs that are run by CAs or browser vendors. Currently, not all logs accept all certificates. For example, CAs typically run their own log instance accepting certificates signed by themselves in order to comply with browser requirements. They might (although not very often seen in practice) or might not accept certificates published by other CAs, for whatever reason they choose - technical, economic, political...

The GÉANT CT log service equally and equitably supports any web-related party, entity, organisation or person. It also accepts all certificates from all well-known trust stores. By being open to any organisation or user group, the GÉANT CT log service helps in reducing the number of logs one has to search through. Considering that the provided service is available and reliable, with all the aforementioned aspects, GÉANT CT log service will help to improve the certificates verification process.

However, even though the service is open and available for all Internet users, the support in how to use the service and how to contribute to CT infrastructure is aimed primarily for the GÉANT community. Apart from maintaining the central CT server and providing help in using the CT log to check the certificates for their domain, GÉANT also provides support to institutions to run a part of the infrastructure themselves.

For example, an NREN component of the CT log will provide higher availability and performance of the CT service, especially to the domain owners associated with that NREN, as they are – from the routing perspective – closer to the log. It will thus enable more direct access to, exploitation and leverage of the benefits.

2.3 Service Benefits

The provision of a central GÉANT CT log service has multiple benefits for all parties involved in its provision and use:

1. Increasing legitimacy for CT in general, resulting in a more secure web.
2. Increasing trust among entities participating in the CT log service
3. Direct control over service availability.
4. Low latency for traffic towards the service.
5. Knowledge-building in the field of CT.
6. Potential retrieval of research data, taking privacy concerns into account.

Currently, the only provider of CT logs open to all CAs is Google. However, in order to enhance the transparency and trustworthiness of assigned certificates, which leads to stronger security for each domain, there should be more logs available than just Google's. GÉANT, as an independent and non-

biased provider, can help all involved parties in strengthening overall web security, thus leading to a more secure Internet.

2.4 Service Users

CT users range from web browsers, domain owners, CAs and browser vendors to all Internet end users. Any of them can at any time be in any of the following roles: log auditor, log monitor, or log contributor.

Web browsers and domain owners do not use CT log servers directly but rather through CT log monitors, which periodically fetch all new certificates from one or more CT logs and alert domain owners if a certificate has been issued for their domain.

CAs are encouraged to submit new certificates to logs directly. They can benefit from having access to logs that accept their certificates, as well as from predictable availability.

Currently, the only **browser** that uses CT is Chrome, which requires a certificate to be logged in at least two logs in order to consider it trustworthy, one of which has to be a non-Google log. Having GÉANT as such a second log places GÉANT in a strong strategic position, especially taking into consideration that this is, at the moment, pioneering work.

In the GÉANT community, most NRENs are in the role of domain owners and web users. Some NRENs also have the role of a CA. It is also expected that some NRENs will set up their own CT log monitoring systems and point them at the GÉANT CT log service as a service for the organisations they support – universities, research institutes and their other domain owners. NRENs interested in participating in operating parts of the GÉANT CT log service can do so by providing the necessary infrastructure and NOC resources.

3 Technical Description

The GÉANT CT log service consists of one CT log server [RFC6962] as defined by the architecture of the catlfish [catlfish] software, originally developed by NORDUnet. The catlfish architecture defines three distinct types of nodes: Frontend, Merge and Sign, as presented in Figure 3.1. In the interest of security, the service design minimises the number of components that must be trusted.

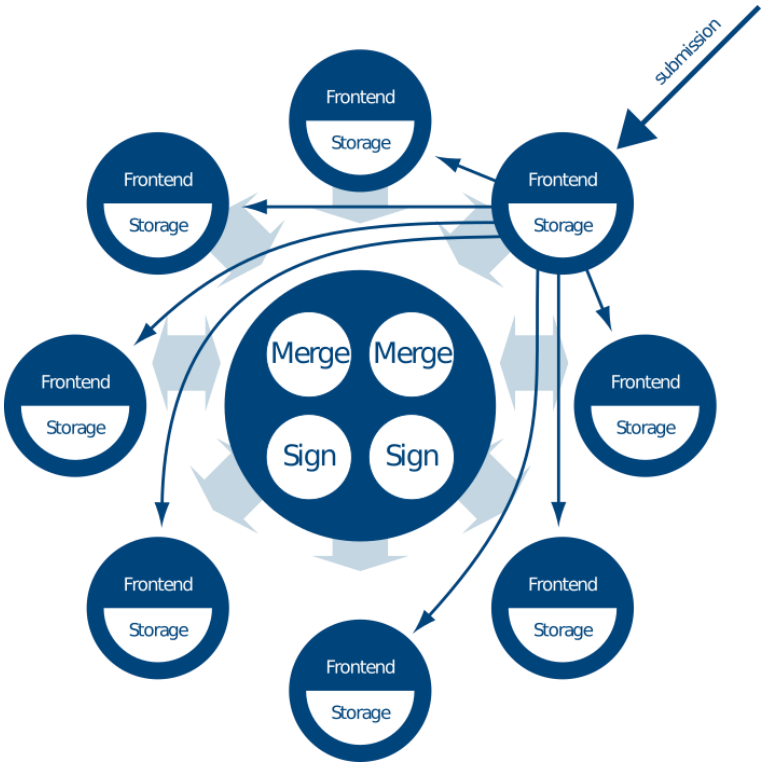

Figure 3.1: CT log server architecture

The GÉANT CT log service will need more than one of every type of node in order to fulfil the redundancy and performance requirements.

Frontend nodes are the public-facing nodes. Each Frontend node has a complete copy of the published log, which enables them to answer any read queries (read operations) without referring to other nodes. They are also responsible for coordinating submissions (write requests), forwarding the data to all other Frontend nodes and finally requesting a promise of publication from a Sign node. Frontend nodes are designed to be run at sites that are not fully trusted, and a single Frontend node cannot by

itself compromise the integrity of the log. Each additional Frontend node provides both added redundancy and performance resilience.

Merge nodes are at any one time divided into one Primary Merge Node and the Secondary Merge Nodes. The Primary Merge Node periodically collects the submissions gathered at the Frontend nodes and decides the order in which the submissions should be published. After having sent the order to the Secondary Merge Nodes, the Primary Merge Node then requests a publication signature from a Sign node and distributes the signature and the order to all Frontend nodes. Any submission missing from a Frontend node is also distributed to that Frontend node by the Primary Merge Node. If the Primary Merge Node fails, a new Primary Merge Node is selected from among the Secondary Merge Nodes.

Sign nodes have the cryptographic keys that are the basis of the cryptographic trust of the log. They issue two types of signature: promises of publication (Signed Certificate Timestamp (SCT)) and publication signatures (Signed Tree Head (STH)). The Sign nodes must be well secured, but they have no state, and can therefore be run on a completely read-only system.

GÉANT CT log service differs in design when compared with CT solution from Google [https://github.com/google/certificate-transparency]. GÉANT CT log solution architecture consists of several simple components with well-defined interfaces (APIs). This architecture enables more controlled and easier to develop architecture, as each component can evolve as needed, not necessarily making an imperative for the other components to change and thus contributing to the correctness, robustness and overall security of the solution.

With such design, GÉANT CT log can be established as a multi-domain solution. In this way, individual components can be placed in different organisations/geographically distributed places and thus share the burden and increase availability and reliability of the service.

In addition, before GÉANT involvement, there were no scalable open-source solutions, and even though the situation is changing lately, multi-domain solution is provided only by GÉANT.

The technology infrastructure used to provide the CT log service consists of free and open source software developed in-house and available here [https://www.ct.nordu.net/]. It is run on Unix-based systems capable of hosting Erlang virtual machines (VMs) and executing Python code.

4 Service Operations

Operations of the CT log service are defined from the fit-for-purpose perspective and supported through the underlying technology to ensure its fit-for-use aspect. It also requires establishment of the supporting infrastructure, as well as the definition of operations and support teams, service policies, and the service metrics according to which the service will be evaluated and its success assessed. The operational aspects described in this section form part of the service definition and will be put in place in Period 2.

4.1 Supporting Infrastructure

In addition to the architecture of the Certificate Transparency log server described in Section 3, the service is supported by the following components:

- Systems and service operations framework for service deployment and orchestration.
- Monitoring system for checking the health of systems and services and raising alarms.
- Hardware security module (HSM) infrastructure for keeping sensitive keys off general-purpose computers.
- Issue-tracking system for operational issues.

All of these components will be established, validated and tested before the service transitions into production.

4.2 Operations and Support Teams

There are three teams that support the GÉANT CT log service:

- **Operations team**, which takes care of the service in production, its central and distributed components, and manages and maintains the installed infrastructure as well as providing support for the users. The team's focus is existing users and installations. It will be reachable at ct-ops@lists.geant.org.

- **Business development and service management team**, which tracks the usage of the established service, gathers feedback and input from the user groups and actively takes part in the future-service roadmaps. The team's focus is potential and new users. It will be available at ct-team@lists.geant.org.

- **CT developers team**, which translates the input from the operations, business development and service management teams into future software and service versions. The team will be reachable at ct-dev@lists.geant.org, primarily for technical and implementation-specific topics.

4.3 Service Policies

The Certificate Transparency log service, as explained in Section 2.2, is open to all web browser users, domain owners, browser vendors and/or CAs to submit certificates and/or query, monitor or audit existing certificates in the log.

Additional support to any of the users – e.g. how to use the service and how to get the most benefit out of it, and/or how to participate in the GÉANT CT log service – is primarily offered to the GÉANT community.

Usage policies and expectations exist between CAs and log providers (e.g. fast access to a log), or between browser vendors and log providers (e.g. Google's expectation of at least 99% availability, as measured by Google). In addition to this, the log is required to sign and publish a new tree head at least once every 24 hours. In practice, a new tree head should be published much more often, typically several times per hour. Such known requirements are being taken into consideration in the design of the operational and production service.

4.4 Service Metrics

Availability is a key metric for measuring the success of a CT log service. Availability is actively being measured by a monitoring system. The monitoring system constantly submits test certificates to the log server and queries the server in order to establish confidence in log consistency and response times. Target availability is 99%.

Other service metrics, such as the number of certificates in the log, growth in log size and number of read or write requests can also be measured to assess the load on the infrastructure, but cannot be considered as service performance indicators as they do not fall within the area of influence of the CT development or operations team.

5 Service Roadmap

During Period 2 of GN4-2, the Certificate Transparency log service is scheduled to pass the pilot gate (as defined by GÉANT's product lifecycle management (PLM) process) and enter the service transition lifecycle phase in which the production infrastructure will be prepared with clean installations of all components, validated and tested from the functionality, performance and security perspective. The service is planned to be ready for production in spring 2018.

The CT software will eventually have to be updated to reflect the changes in the protocol as a result of the standardisation efforts in IETF (e.g. work on [RFC6962-bis]). The software will also have to be continuously updated with security fixes.

The major planned software releases will focus on documentation (1.0, 2018 Q1), performance (1.1, 2018 Q2) and the new protocol mentioned in the previous paragraph (2.0, 2018).

In parallel with the preparations for production, the strategy and planning phase will start for future versions of CT.

6 Conclusions

Development of the GÉANT Certificate Transparency log service is a continuation of work by SUNET/NORDUnet. The service should transition into production during GN4-2 Period 2.

Based on IETF standards, it enables discovery of misissued certificates. By providing transparency of all issued certificates through the log, the CT log service enables monitoring and audit of the certificates issued for a domain.

The service delivers clear trust and security benefits for its user community and for CT, the web space and the Internet as a whole.

The GÉANT Certificate Transparency log service has three main aspects: installation of the CT log server, and providing two types of support to users: in using the service and/or in participating in the GÉANT CT log service by running a service architecture component in their own domain. The primary GÉANT CT log service users are domain owners, certification authorities and web users.

The operations and support framework is designed primarily to support the GÉANT community in using the service and/or providing a similar CT log service in their domains. However, use of the CT log service is open to all CT users.

The service is planned for production in spring 2018, based on the infrastructure, policies, metrics and operational and support teams described in this document.

References

[catlfish] https://www.ct.nordu.net/

[RFC6962] B. Laurie, A. Langley, E. Kasper, *IETF RFC 6962: Certificate Transparency*, June 2013
https://tools.ietf.org/html/rfc6962

Glossary

CT	Certificate Transparency
CA	Certification Authorities
HSM	Hardware Security Module
HTTPS	Hyper Text Transfer Protocol Secure
IETF	Internet Engineering Task Force
JRA	Joint Research Authority
JRA2	Network Services Development
JRA2 T6	JRA2 Task 6, Network Security
NOC	Network Operations Centre
NREN	National Research and Education Network
PLM	Product Lifecycle Management
R&E	Research and Education
RFC	Request for Comments
SCT	Signed Certificate Timestamp
STH	Signed Tree Head
VM	Virtual Machine