SAVANT: Aggregated Feedback and Accountability Framework for Named Data Networking

Diarmuid Ó Coileáin
CTVR, School of Computer Science and Statistics,
Trinity College Dublin,
Dublin 2, Ireland.
collindi@tcd.ie

Donal O'Mahony
CTVR, School of Computer Science and Statistics,
Trinity College Dublin,
Dublin 2, Ireland.
Donal.OMahony@cs.tcd.ie

ABSTRACT

Content providers (i.e., entities that own or are licensed to sell or distribute content e.g., BBC, Netflix) are looking for more efficient, secure, cheaper, accountable and scalable mechanisms for the delivery of content to end-users. The *Information Centric Networking* (ICN) paradigm offers solutions to some of these challenges by decoupling a user's trust in content from where it is obtained by enabling the content to *self-verify* (i.e., the user can establish integrity, trust and provenance in content received from trusted or untrusted infrastructure). However, there are still associated challenges with using ICN architectures related to content feedback and accountability. In this paper, we propose an ICN architecture extension for content feedback and accountability called the Savant framework, which we apply to the *Named Data Networking* (NDN) architecture.

Categories and Subject Descriptors

C.2.4 [**COMPUTER-COMMUNICATION NETWORKS**]: Distributed Systems–Distributed applications

General Terms

Management, Measurement, Design

Keywords

Information Centric Networking, Content Feedback and Accountability, Analytics

1. INTRODUCTION

Content feedback and accountability information enables a content provider and a content distributor (e.g. Content Distribution Network (CDN)) to track the content delivered to consumers. We define *feedback* as the information returned to a content provider so that future or in-progress operations can be monitored, supported, altered or corrected. This information can help a content provider to determine what content is popular, the geographical location of users and what kind of performance clients are receiving from the network [1]. In comparison to feedback, we define *accountability* as the willingness of trusted or untrusted communicating entities to produce accurate and verifiable information about the content distribution process. The primary difference between a feedback architecture and an accountability architecture is that when there is a problem the latter has the tools to pinpoint (with non-repudiation) the responsible entity [5]. In the remainder of this paper we briefly outline Savant's framework for providing content feedback and accountability (depicted in Figure 1), which we apply to the NDN architecture.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).
ICN'14, September 24–26, 2014, Paris, France.
ACM 978-1-4503-3206-4/14/09.
http://dx.doi.org/10.1145/2660129.2660165.

2. THE SAVANT FRAMEWORK

The main components in the Savant architecture include content providers, content ingestion, accountability engine, NDN caches and NDN clients (collectively referred to as *NDN agents*), configuration manager and aggregator functions. The *content ingestion* process prepares content for distribution to many different users, devices and networks, performing tasks such as transcoding, resolution conversion, encryption and adding *metadata*. The *metadata* specified during ingestion identifies the content provider. This helps the NDN agent establish contact with the closest available *accountability engine* responsible for collecting information for that content provider.

The accountability engine is composed of geographically dispersed infrastructure located close to the end-user (similar to CDN infrastructure). It has primary responsibility for collecting, aggregating and validating published feedback and accountability information collected from NDN agents. Once two-way communication has been established between the NDN agent and accountability engine, the accountability engine will issue a continuous stream of NDN *Interests* for feedback and accountability information.

Both the accountability engine and NDN agents run a configuration manager and aggregator functions. At their simplest level, *aggregator functions* work to filter, summarise and publish feedback data containing specified attributes. Moreover, they determine what information is collected and the frequency of collection. Due to the nature of the NDN publish/subscribe architecture, only requested information is pulled for aggregation (i.e., there is no duplicate or unnecessary information requested for aggregation).

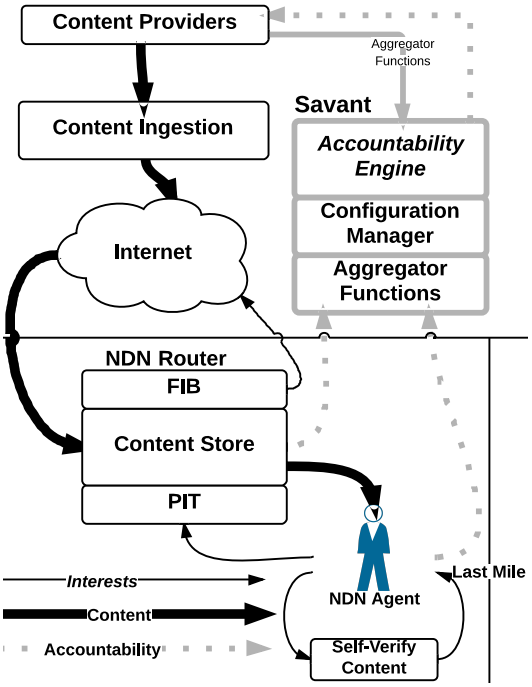

Figure 1: The NDN architecture with the Savant framework.

Finally, aggregator functions are installed (and uninstalled) on distributed accountability engines and NDN agents by their local *configuration manager*. The configuration manager also manages the publication of key/value pair attributes and aggregator function output for collection by accountability engine infrastructure.

2.1 Savant Accountability

A major challenge in distributing content from untrusted infrastructure is the inability to directly observe or trust the interactions between communicating peers [1]. However, NDN (and ICN architectures) offer natural support for accountability due to published content's ability to self-verify. This is underpinned by public/private key pairs, deterministic inputs and outputs (i.e., *Interest*/Data), hash functions and digital signatures [5][3]. The Savant framework provides the remaining tools for accountability such as log auditing and providing certified public/private key pairs to agents. To achieve this, we follow similar principles to those already described by PeerReview (which provides a secure record of accountability in distributed systems [3]) and Reliable Client Accounting (RCA) (a non-repudiable accounting system for hybrid CDN-P2P systems [1]). These systems use a hash chain of log entries in order to detect inconsistencies between the logs of communicating peers. This is achieved by getting both sending and receiving peer infrastructure to maintain tamper evident logs for all communication-taking place between peers [1][3].

3. IMPLEMENTATION

In 2011, a team at Conviva performed an analysis of data collected from over 1 million unique viewers, content providers, Internet service providers and CDNs in order to determine how video quality impacts end-user engagement [2]. In order to prove the viability of the Savant architecture we attempted to partially replicate that experiment by gathering and comparing (with Conviva's analysis) a subset of these metrics for a small number of ICN users.

We achieved this objective using an implementation of the Savant framework in conjunction with small modifications to the NDN Video [4] software. First, both the NDN Video server and Savant Server run on independent Amazon Elastic Compute Cloud (EC2) micro-instance machines (i.e., Ubuntu-12.04.3-64 bit; memory: 613MB and disk: 8GB). Second, we published a video-on-demand (VoD) content file using the NDN Video software to the NDN Video Server. This machine acted as the data source for NDN clients requesting content from EC2. Third, we modified the NDN Video Client to publish accountability events as it receives and renders video content from the NDN Video Server. Fourth, this published feedback and accountability information is collected by the Savant Server (i.e., the accountability engine) by requesting chunks of published information from the NDN Clients. Finally, the Savant Server processes, aggregates and republishes/persists the data collected based on aggregator functions.

4. CONCLUSION

Based on this implementation and the Conviva's analysis of end-user engagement we were able to demonstrate the viability of an ICN architecture collecting real-time analytic information supporting the ICN content distribution process. For example, Savant consistently measured NDN Video's join-time at just over 10 seconds for VoD content. In comparison, Conviva measured 95% of users in their analysis having a join-time of less than 10 seconds for CDN content distribution [2]. These and similar comparisons focus attention on the importance of feedback and accountability metrics for ICN content providers, content distributors and end-users.

5. ACKNOWLEDGEMENTS

This research, a part of the Telecommunications Graduate Initiative, is funded by the Irish Higher Education Authority under the Programme for Research in Third-Level Institutes (PRTLI) and co-funded under the European Regional Development fund.

6. REFERENCES

[1] P. Aditya, M. Zhao, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, and B. Wishon. Reliable client accounting for p2p-infrastructure hybrids. In *Proc of NSDI'12*, pages 8–8, 2012.

[2] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang. Understanding the impact of video quality on user engagement. In *Proc of the ACM SIGCOMM*, 2011.

[3] A. Haeberlen, P. Kouznetsov, and P. Druschel. Peerreview: Practical accountability for distributed systems. In *Proc of SOSP*, 2007.

[4] D. Kulinski and J. Burke. Ndn video: Live and prerecorded streaming over ndn. Technical report, Technical Report, The NDN Project Team, 2012.

[5] A. R. Yumerefendi and J. S. Chase. The role of accountability in dependable distributed systems. In *Proc of HotDep*, volume 5, pages 3–3, 2005.