Savant: Aggregated Feedback and Accountability Framework for Named Data Networking

Diarmuid Ó Coileáin (collindi@tcd.ie) & Donal O'Mahony (Donal.OMahony@cs.tcd.ie)
Trinity College Dublin, Dublin 2, Ireland.

Introduction

Problems:

- No accounting or accountability in ICN architectures.
- Content distributed from **trusted** and **untrusted** infrastructure
- Cannot determine ISP performance
- Cannot adapt to failures
- Cannot detect misbehavior
- Cannot bill customers
- Cannot track audiences

Feedback: is the information returned to a content provider so that future or in-progress operations can be monitored, supported, altered or corrected e.g., QoE, buffering time, buffering ratio, rendering quality, end-user engagement (e.g., content views, advertising impressions), user demographics (e.g., geographic location and device type), etc.

Accountability: is the willingness of trusted or untrusted communicating entities to produce accurate and verifiable information about the content distribution process. Accountability is the ability to establish integrity, provenance, auditability and non-repudiation in the accounting information received.

The Savant Framework

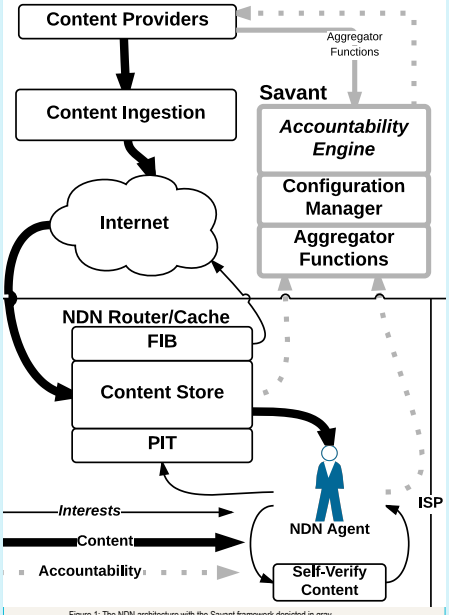

Figure 1: The NDN architecture with the Savant framework depicted in gray.

- Savant pushes **primary responsibility for accounting and accountability** out to the NDN caches and NDN clients (collectively referred to as **NDN agents**)
- 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.
- The **configuration manager** installs/uninstalls aggregator functions, manages the publication of key/value pair attributes and aggregator function output for collection by accountability engine infrastructure. Both the accountability engine and NDN agents run a configuration manager and aggregator functions.

Aggregator Functions [2]

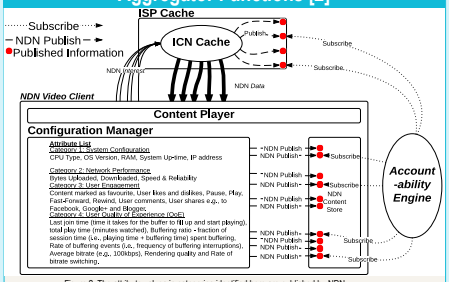

Figure 2: The attribute values in categories identified here are published by NDN agents. The accountability engine subscribes to published data at NDN agents.

- Based on **event processing languages** [2],
- Run queries continuously as new data arrives
- They filter, summarize and publish feedback data containing specified attributes
- Monitor for patterns of **complex events** notifying interested parties on occurrence e.g., Turn off/on accountability
- What feedback information to collect, the frequency to collect it, how to aggregate data returned.
- Use ***Authenticated Interests*** [7] add commands and authentication tags (i.e., digital signatures or message authentication codes (MACs)) to NDN *Interests* to support applications efficiently running commands on remote fixtures.

Named Data Networking: Natural Support for Accountability [4]

- Strong public π_n & private key σ_n pairs to node: n
- Trust, integrity and provenance can be established in published NDN content
- Hash function e.g., SHA-256
- Digital signatures
- Deterministic inputs and outputs (*Interests* and Data)

Savant Accountability Framework [1][3][4]

- Log entry $e_i = (h_i, s_i, t_i, c_i)$
- Recursive hash value h_i, sequence number s_i always increasing, type t_i (e.g., *Interest* or Data) and type specific content C_i (e.g., <node-id><message>)
- Authenticator $\alpha_i^n = (s_i, h_i, \sigma_n(s_i||h_i))$ attached to each message sent – signed by n's private key σ_n
- Accountability engine collects authenticators for both agents
- Separate sub hash chain maintained for each communicating node. Also a separate set of authenticators.
- Authenticators are cumulative (similar to RCA [1]). Consequently, a certain amount of privacy can be maintained between NDN agents
- Logs are tamper evident
- C_i contains accounting information e.g., QoE, bitrate, etc.

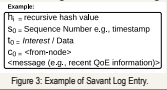

Figure 3: Example of Savant Log Entry.

Client Hash Chain [1][3]

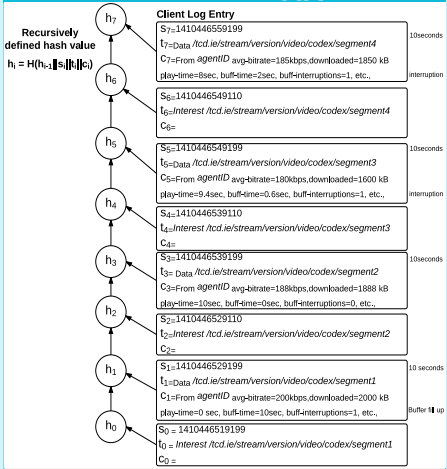

Figure 4: Savant Log Entry: Deterministic *Interest* (request) and Data (response). The recursively defined hash value forms a sub hash chain for each peer it communicates with linking previous log entries making the log tamper-evident. Consequently, one pair of authenticators is maintained for each node an NDN agent communicates with.

Savant: Accounting and Accountability Collection from ingress/egress/ISP cache or NDN cache

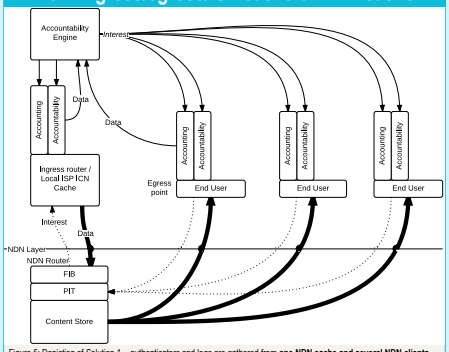

Figure 5: Depiction of Solution 1 - authenticators and logs are gathered from **one NDN cache and several NDN clients**, which need to be correlated correctly by the accountability engine (due to time-shifted content and in-network caching).

The Savant framework propose collecting accounting and accountability information from:

- Solution 1: **ingress point** (or from a local ISP NDN cache) and **egress point** (i.e., NDN client) in a network
- Solution 2: for **all content distributed** from the NDN router to NDN clients in the network.

Work-in-progress: Preliminary Results

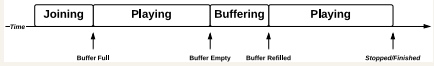

Figure 6: NDN Video State Machine. Adapted from [5]

- Amazon Elastic Compute Cloud (**EC2**) micro-instance machine (i.e., Ubuntu-12.04.3-64 bit; memory: 613MB and disk: 8GB).
- Modified NDN Video Client to publish accounting events as it receives and renders video content from the NDN Video Server [6]
- Measured the following accounting metrics (i.e., a subset of metrics from Conviva [5] study on the impact of video quality on end-user engagement): **join-time**, **session duration**, **buffering**, **rate of buffering events** and **average bitrate**.
- Scalability: we extrapolated that Savant could support 104 busy concurrent NDN agents for each Amazon EC2 micro-instance using only about 61% of network resources used by one NDN Video [6] content distribution session.

Future Work

- Enhance accountability
- Enhance accounting
- Analyze in large environment

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 *USENIX*. Available: http://dl.acm.org/citation.cfm?id=2228298.2228309
2. G. Cugola and A. Margara, "Processing flows of information: From data stream to complex event processing". Available: http://doi.acm.org/10.1145/2187671.2187677
3. A. Haeberlen, P. Kouznetsov, and P. Druschel. Peerreview: Practical accountability for distributed systems. In Proc of SOSP, 2007.
4. 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.
5. 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 ACM SIGCOMM, 2011.
6. D. Kulinski and J. Burke, "Ndn video: Live and prerecorded streaming over ndn," Technical Report, The NDN Project Team, Tech. Rep., 2012. [Online]. Available: http://http://named- data.net/publications/techreports/trstreaming/
7. Jeffery Burke, Paolo Gasti, Naveen Nathan, and Gene Tsudik. 2012. Securing Instrumented Environments over Content-Centric Networking: the Case of Lighting Control. *CoRR* abs/1208.1336 (2012). http:// arxiv.org/pdf/1208.1336.pdf

Acknowledgements

Ireland's EU Structural Funds Programmes 2007 - 2013
Co-funded by the Irish Government and the European Union

EUROPEAN REGIONAL DEVELOPMENT FUND

An Roinn Post, Fiontar agus Nuálaíochta
Department of Jobs, Enterprise and Innovation

Investing in your future

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.