

# Runtime Analysis of Whole-System Provenance

Thomas Pasquier\*  
University of Bristol

Xueyuan Han  
Harvard University

Thomas Moyer  
University of North  
Carolina at Charlotte

Adam Bates  
University of Illinois at  
Urbana-Champaign

Olivier Hermant  
MINES ParisTech  
PSL Research University

David Eyers  
University of Otago

Jean Bacon  
University of Cambridge

Margo Seltzer  
University of  
British Columbia

## ABSTRACT

Identifying the root cause and impact of a system intrusion remains a foundational challenge in computer security. *Digital provenance* provides a detailed history of the flow of information within a computing system, connecting suspicious events to their root causes. Although existing provenance-based auditing techniques provide value in forensic analysis, they assume that such analysis takes place only retrospectively. Such post-hoc analysis is insufficient for realtime security applications; moreover, even for forensic tasks, prior provenance collection systems exhibited poor performance and scalability, jeopardizing the timeliness of query responses.

We present CamQuery, which provides inline, realtime provenance analysis, making it suitable for implementing security applications. CamQuery is a Linux Security Module that offers support for both userspace and in-kernel execution of analysis applications. We demonstrate the applicability of CamQuery to a variety of runtime security applications including data loss prevention, intrusion detection, and regulatory compliance. In evaluation, we demonstrate that CamQuery reduces the latency of realtime query mechanisms, while imposing minimal overheads on system execution. CamQuery thus enables the further deployment of provenance-based technologies to address central challenges in computer security.

## CCS CONCEPTS

• **Security and privacy** → **Operating systems security**; *Information flow control*; Intrusion detection systems;

## KEYWORDS

Whole-system Provenance; Information Flow Tracking; Graph Processing; Linux Kernel

## ACM Reference Format:

Thomas Pasquier, Xueyuan Han, Thomas Moyer, Adam Bates, Olivier Hermant, David Eyers, Jean Bacon, and Margo Seltzer. 2018. Runtime Analysis

\*Part of this work was completed at Harvard University and at the University of Cambridge.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CCS '18, October 15–19, 2018, Toronto, ON, Canada

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5693-0/18/10...\$15.00

<https://doi.org/10.1145/3243734.3243776>

of Whole-System Provenance. In *2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, October 15–19, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3243734.3243776>

## 1 INTRODUCTION

Timely investigation of system intrusions remains a notoriously difficult challenge [66, 94, 96]. While security monitoring tools provide an initial notification of foul play [13, 41, 86, 91, 95, 97], these indicators are rarely sufficient in and of themselves. Instead, crafting an appropriate response to a security incident often requires scouring terabytes of audit logs to determine an adversary's method of entry, how their reach spread through the system, and their ultimate mission objective. Such investigations not only require a human-in-the-loop, but are excruciatingly slow, at times requiring months of investigation and thousands of employee hours [56]. This delay between an event's occurrence and its diagnosis represents a tremendous window of opportunity for attackers – as they continue to exploit the system, defenders are still just getting their bearings.

*Digital provenance* (or *provenance* for short) refers to the data being used in a variety of ways to address the challenges of forensic audits. By parsing individual records into causal relationship graphs that describe a system's execution, provenance enables defenders to leverage the full historical context of a system and to reason about the interrelationships between different events and objects. With provenance, forensic investigations can trace back a given security indicator (e.g., a port scan) to the attacker's point of entry (e.g., a malicious email attachment) [53] and then trace forward from the entry point to determine what other actions the attacker has taken on the system.

Unfortunately, provenance-based auditing's growing popularity has uncovered significant limitations in its performance and scalability. Early efforts to integrate provenance querying into production systems indicated that, even for modestly small organisations (e.g., 150 workstations), forensic queries can take on the order of hours or days to complete [61]. In an actual attack scenario, where a timely incident response could make the difference between victory and defeat, such delays are unacceptable. Moreover, to date, provenance-aware systems have supported causal reasoning only as an after-the-fact forensic activity [54]; this is unfortunate, because provenance is also invaluable to a variety of runtime security tasks such as access control [76, 77], integrity measurement [92], and regulatory compliance [8, 15, 68, 81]. To date, the design of low latency mechanisms for realtime provenance analysis has not been given adequate consideration in the literature.































