

SYM-AM-16-042



PROCEEDINGS OF THE THIRTEENTH ANNUAL ACQUISITION RESEARCH SYMPOSIUM

WEDNESDAY SESSIONS VOLUME I

Improving Security in Software Acquisition and Runtime Integration With Data Retention Specifications

Daniel Smullen, Research Assistant, Carnegie Mellon University
Travis Breau, Assistant Professor, Carnegie Mellon University

Published April 30, 2016

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Panel 9. The Operational and Developmental Dimensions of Cybersecurity

Wednesday, May 4, 2016	
3:30 p.m. – 5:00 p.m.	<p>Chair: Rear Admiral David H. Lewis, USN, Commander, Space and Naval Warfare Systems Command</p> <p><i>The Cybersecurity Challenge in Acquisition</i> Sonia Kaestner, Adjunct Professor, McDonough School of Business, Georgetown University Craig Arndt, Professor, Defense Acquisition University Robin Dillon-Merrill, Professor, Georgetown University</p> <p><i>Improving Security in Software Acquisition and Runtime Integration With Data Retention Specifications</i> Daniel Smullen, Research Assistant, Carnegie Mellon University Travis Breaux, Assistant Professor, Carnegie Mellon University</p> <p><i>Cybersecurity Figure of Merit</i> CAPT Brian Erickson, USN, SPAWAR</p>



Improving Security in Software Acquisition and Runtime Integration With Data Retention Specifications

Daniel Smullen—is a Research Assistant enrolled in the software engineering PhD program at Carnegie Mellon University. His research interests include privacy, security, software architecture, and regulatory compliance. [dsmullen@cs.cmu.edu]

Travis Breaux—is an Assistant Professor of Computer Science in the Institute for Software Research at Carnegie Mellon University (CMU). His research program searches for new methods and tools for developing correct software specifications and ensuring that software systems conform to those specifications in a transparent, reliable, and trustworthy manner. This includes compliance with privacy and security regulations, standards, and policies. Dr. Breaux is the Director of the CMU Requirements Engineering Lab and co-founder of the Requirements Engineering and Law Workshop, and he has several publications in ACM- and IEEE-sponsored journals and conference proceedings. [breaux@cs.cmu.edu]

Abstract

The Department of Defense (DoD) Risk Management Framework (RMF) for IT systems is aligned with the National Institute for Standards and Technology (NIST) guidance for federal IT architectures, including emergent mobile and cloud-based platforms. This guidance serves as a prescriptive lifecycle for IT engineers to recognize, understand, and mitigate security risks. However, integrators are left with the challenge—during acquisition and during runtime integration with external services—to reason about the actions on data inherent in their system designs that may have confidentiality risks. These risks may lead to data spills, loss of confidentiality for mission data, and/or revelations about private data related to service members and their families. Solutions are needed to assist acquisition professionals to align system data practices with the RMF and NIST guidance, as well as DoD IA directives—particularly with respect to the collection, usage, transfer, and retention of data. To provide support to this end, we extended our initial automation framework to support reasoning over data retention actions using a formal language. We propose an evaluation method for these extensions, carried out through simulations of real-world IT systems using imitation but statistically accurate synthetic data. Our language aims to address dynamically composable, multi-party systems that preserve security properties and address incipient data privacy concerns. Software developers and certification authorities can use these profiles expressed in first-order logic with an inference engine to advance the RMF, express data retention actions that promote confidentiality, and re-evaluate risk mitigation and compliance as IT systems evolve over time.





ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net

Improving Security in Software Acquisition with Data Retention Specifications

Daniel Smullen,

Travis Breaux

Navy Postgraduate School Acquisition Research Symposium

May 4, 2016

Motivation

- DoD and the Defense Industrial Base (DIB) leverage 3rd party service compositions to outsource infrastructure and services.
 - What sort of data do they want?
 - What will they do with my data once they have it?
 - What am I willing to give them?
- In prior work: we studied actions for collection, use and sharing [RE'15]
- What about data retention?

[RE'15] T. Breaux, D. Smullen, H. Hibshi, "Detecting Repurposing and Over-collection in Multi-Party Privacy Requirements Specifications." *IEEE 23rd International Requirements Engineering Conference*, Ottawa, Canada, pp. 166-175, Sep. 2015.

Data Retention Actions


- Redaction
 - Directly remove some type of data from a collection.
- Perturbation
 - Don't directly remove a specific type of data, just reduce the quality of the data, or remove specific data points.
- Data Append
 - Combine disjoint data to infer something, create a new type of data.

Running Example

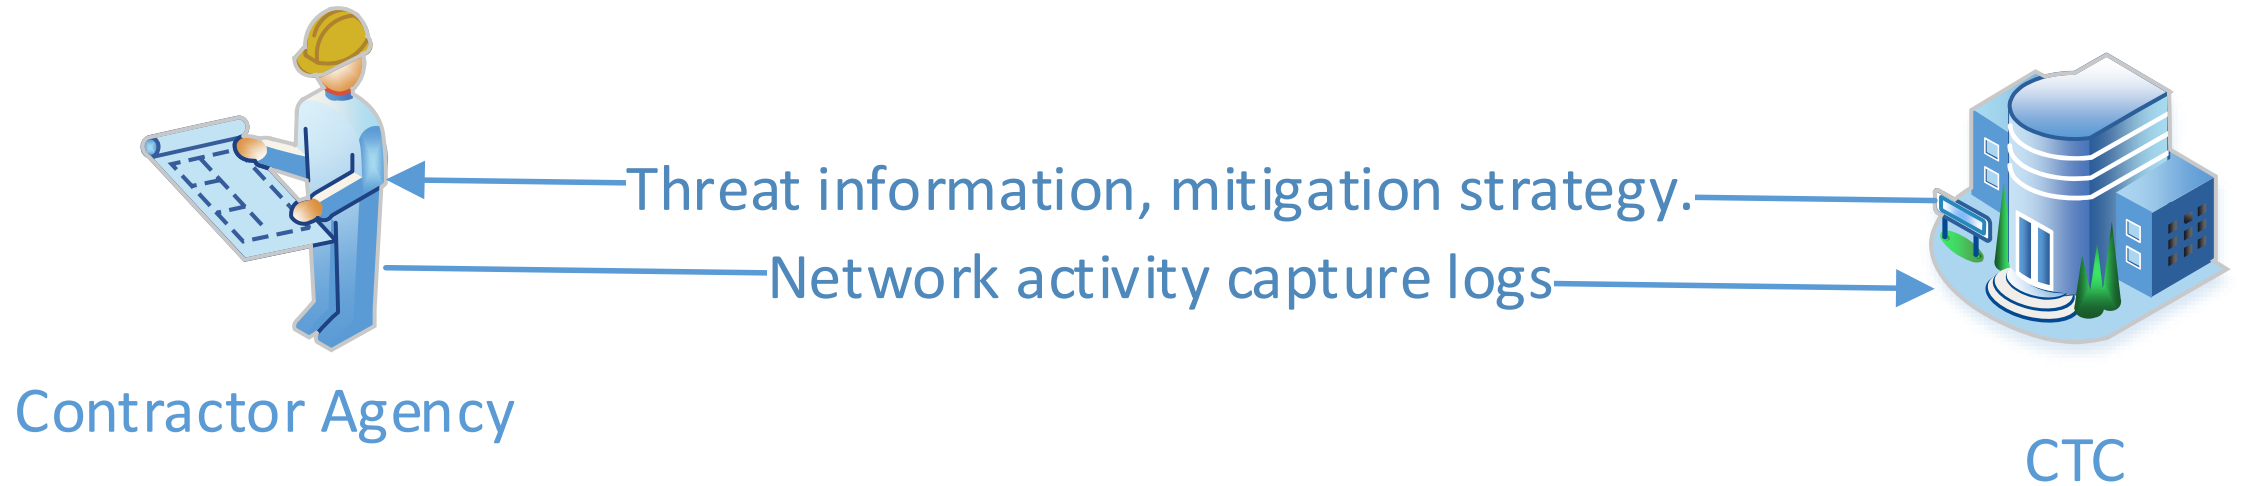
- Cyber threat information sharing portal.
- A **cyber threat clearinghouse** and **DoD contractor** have recently crafted a ***data sharing agreement*** that enables them to collaborate to share **cyber threat information** via this portal.



Running Example




Running Example



Technical Approach

- Eddy; formal language with semantics based on Description Logic (DL).
- Specify and analyze data flows.
- Use this as a tool to measure how specified actions affect *unanticipated disclosure*.



Eddy Language Structure

- Two sections; header section, policy section.
- Define data, actors, purposes in header.
 - All concepts can have sub-concepts described through DL *subsumption*.
 - Can have equivalent concepts described through DL *equivalence*.
- Define actions (with modality – permission, prohibition, obligation) in policy section.

Using Eddy

1. Analyze policy text to extract requirements, code into Eddy:

Model **Open** **Purpose** **specified** **obligation**.

*“All O-2 or higher service members **must collect** the **name** and **rank** from **network analysts** in order to complete a **network analysis authorization form**. This must be done to **assure compliance** with standing orders.”*

SPEC HEADER

D authorization_form > name, rank

SPEC POLICY

O **COLLECT** authorization_form FROM **network_analysts** FOR **assure_compliance**

Using Eddy

2. Tool compiles Eddy into Description Logic:

- $\text{name} \sqsubseteq \text{authorization_form}$
- $\text{rank} \sqsubseteq \text{authorization_form}$
- $p_1 \equiv \text{COLLECT} \sqcap \exists \text{hasObject.authorization_form} \sqcap \exists \text{hasSource.network_analysts} \sqcap \exists \text{hasPurpose.assure_compliance}$
- $p_1 \sqsubseteq \text{Obligation}$

SPEC HEADER

ATTR NAMESPACE "http://gaius.isri.cmu.edu/example.owl"

ATTR DESC "This is an example policy for medical data, written in Eddy."

P treatment > diagnosis, prescription, blood-tests

D patient-labs > bloodwork

A medical-professional > phlebotomist, doctor, nurse

A laboratory > phlebotomist

SPEC POLICY

P COLLECT bloodwork FROM phlebotomist FOR treatment

P COLLECT bloodwork FROM laboratory FOR treatment

P USE bloodwork FROM phlebotomist FOR marketing

P USE patient-labs FROM phlebotomist FOR anything


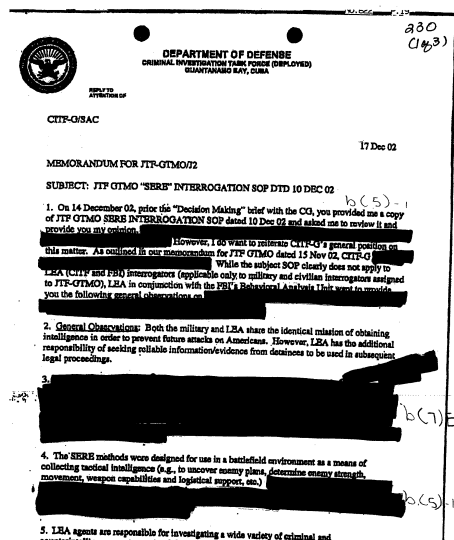
P USE bloodwork FROM medical-professional FOR diagnosis

R USE bloodwork FROM medical-professional FOR anything



Further Extending Eddy for Data Retention

- **Redaction** means to remove data elements from a dataset.
- Redaction is useful as a data minimization strategy when data cannot be shared, or when it can be pared down to the minimum necessary.



These are paper redactions, but we can redact data from digital sharing, too.

Further Extending Eddy for Data Retention

- ***Data append*** refers to a general class of methods that link two or more data elements together.
- By *prohibiting* data append, downstream parties are bound to limit the use of a redacted dataset; cannot combine to recreate original.
- Fixed requirement for the data prior to sharing, assuring disjointness from other datasets post-transfer to a third party.

Further Extending Eddy for Data Retention

- ***Perturbation*** refers to a general class of methods that introduces statistical inaccuracies into data; conforms to statistical profile of original data;
 - Changing data values.
 - Removing values.
 - Adding new values.
- Eddy language does not assume that data perturbation is implemented by any particular method.

Experimental Design

- *How do data append, redaction and perturbation systemically affect data subject unanticipated disclosure?*
- Microsimulation [Lov16]; a technique for analyzing real-world situations based on synthetic data.
- Combine with data sharing agreement profile (Eddy specification).
- Perform analysis.

[Lov16] R. Lovlace, “Spatial Microsimulation in R”, CRC Press, 2016.

Why Generate Synthetic Datasets?


- Third-party suppliers want to see the technology works before sharing their data.
- De-anonymized datasets for public release lack the sensitive data we want to protect in our analysis.

Our Approach:

- Empirically motivated, micro-simulation based on Monte Carlo to achieve technical realism.


How to Sample Data Retention Specifications?

- Based on randomly sampling from the population of data retention agreements expressed in the extended Eddy.



Sampling Data Retention Specifications

- Requires expert analysis to perform data segmentation; 3 steps.
 1. Examine captured data.
 2. Determine what data is associated with a threat.
 3. Determine what data is extraneous.
- Experts use this analysis to refine the data attributes described in Eddy.



Proposed Evaluation Metrics

- Categorical re-identification approach proposed by [Swe02].
- Number of possible results (threats) proportional to uniqueness of threat with respect to data attributes.
- Threat employed by one organization easier to identify compared to threat used by many organizations.
- Threat associated with one individual easier to identify compared to association with many individuals.

[Swe02] L. Sweeney, “k-Anonymity: a model for protecting privacy”. *International Journal for Uncertainty in Fuzzy Knowledge Based Systems*, 557-567, 2002.

Proposed Evaluation Metrics

- Dataset is queried to match threat information with likely threat.
- The results of this query will be *possible* threats.
- Probability of correct identification, given threat is identified;
 - $P(\text{correct identification}|\text{threat identified}) = \frac{1}{\text{count}(\text{ident. query results})}$

Conclusions, Lessons Learned

- Synthetic data generation process requires aggregate data, deep knowledge of data characteristics.
- Designing evaluation functions to determine data utility is bound to queries on data that are being executed.
- Tension between *business value* derived from data and *risk of unanticipated disclosure* of confidential information.

Conclusions, Lessons Learned

- Impossible to measure or predict impact of data retention actions without knowing how data will be used.
- System integrators/acquisition specialists must recognize that both data **and** queries on data have confidentiality risks built in.
- *We have proposed a method to calculate and engineer confidentiality impact; analysts must have intuition that system has inherent confidentiality risk.*

Future Work

- Use machine learning to augment and/or replace expert judgement in data segmentation process.
- Reduces the need for personnel to analyze the data, instead it can be broken up into categories automatically.

Future Work

- Specification enforcement mechanisms;
 - Prevent data from being used in unspecified way, rather than check conformance.
- Seed data with predictable data points that show evidence of unauthorized uses.
- Feedback mechanisms;
 - Collect and report/deny/redact data at runtime.

Future Work

- Specification enforcement mechanisms;
Would allow integrators/analysts to:
 - Take control of how 3rd party services use downstream data.
 - Evaluate whether to use a 3rd party service based on whether it is truthful about specified data practices.
- Allows enforcement of new requirements for 3rd party based on analysis;
 - Can force redaction of certain sensitive data,
 - Prevent data mixture with respect to classification levels, or data types,
 - Require data to always be combined.