W.R. Simpson & K.E. Foltz, Int. J. of Design & Nature and Ecodynamics. Vol. 11, No. 3 (2016) 295–305

ACCESS AND PRIVILEGE IN SECURE BIG DATA ANALYSIS

W.R. SIMPSON & K.E. FOLTZ
Institute for Defense Analyses, USA

ABSTRACT

The distributed data sources and strict security controls of the Enterprise Level Security (ELS) architecture present challenges for data mining. The ELS architecture is a secure enterprise system that enforces strict security controls in a uniform way across an enterprise. It includes end-to-end bilateral authentication for all human as well as machine interactions and verifiable claims-based access controls. Claims provisioning is automated and centrally managed based on authoritative attributes of active entities in the enterprise. While these security provisions are necessary for secure systems, they present some unique challenges to big data analyses. Key among these are non-standard schemas, non-standard access and privilege, restricted access to analysis outcomes, and overall privilege handling. Some of the distributed data sets may be fully or partially accessible, or even not accessible. Users with limited access may compute different results than those with broad access. We discuss the problems encountered for data mining in an ELS architecture and possible solutions.

Keywords: access control, big data tools, escalation, privilege, security, standardized roles, standardized schemas.

1 INTRODUCTION

This paper describes techniques to provide data mining services to individual entities based on the level of access they have to the data. This allows a wider use of data mining services across the enterprise by allowing more open access to the data mining tools while providing automated access controls at the data sources themselves. This also allows data mining of highly sensitive data sets where no one entity is allowed full access or where data mining across the full data set is prohibited.

The primary method of providing access to data using Enterprise Level Security (ELS) is through web services and content access tagging. The ultimate goal is to provide standardized interfaces to the data itself, but current database systems are not yet prepared to provide this level of access. The problem of standardized schemas and taxonomies for big data was recently addressed [1], but it only addressed privacy concerns while not addressing security and access restrictions. With ELS, the database interfaces are locked down to a single, highly restricted interface for a single web service. All access is provided through this web service. The ELS model is used at the web service to provide end-to-end authentication and claims-based authorization through a standard set of enterprise-approved protocols and enterprise-provided claims validation and verification code.

Data mining services are provided with a web application front end that provides the same controls on access, to all enterprise users. The data mining services rely on the requesters to gain access to data. The data mining application itself does not have access to any data sources. This is in contrast to many systems for which the intermediary not only has full access but also limits access to data based on user accounts, roles, or other internal controls. The web service fronting the

This paper is part of the Proceedings of the International Conference on Big Data (Big Data 2016)
www.witconferences.com

ISSN: 1755-7437 (paper format), ISSN: 1755-7445 (online), http://www.witpress.com/journals
DOI: 10.2495/DNE-V11-N3-295-305

database requires authentication and claims from the data mining service as well as verifiable claims proving that the requester is requesting access to the data. This is done through a combination of authorization credentials by a Security Token Server (STS) with trust relationships in the enterprise.

Federation allows data to be shared across enterprises. This presents special challenges for data mining since enterprise trust agreements must be extended in limited ways to federation partners. This paper discusses some of these challenges and outlines potential solutions that maintain the ELS model.

2 RELATED WORK

Big data work has exploded, partly due to an increase in the amount of data available for analysis, and partly due to its own popularity. Big data is a broad term for data sets that are so large or complex that traditional data processing applications are inadequate [2]. Challenges include analysis, capture, data curation, search, sharing, storage, transfer, visualization, and information privacy. For this work, we treat big data as data sets that are larger than what can be analyzed using the common or standard techniques that are in practice [3]. This means that big data is a moving target based on available technology. The main concept is that new ways of doing the analysis must be developed, which assume an abundance of data rather than rely on old methods that attempt to be precise or make use of all of the data, but which involve extensive time and resources to do such detailed work. Instead of replacing the old methods, big data methods are often a feeder or supplement to them.

Big data analysis can focus on the bigger picture to inform further, deeper analysis. Big data analysis can be used to determine what data is available so as to choose the best method for further analysis. It can provide a roadmap of the data, which allows drilling down to key subsets. It can provide a rough estimate of a global parameter, which can be adjusted locally. Alternatively, it can be used just to show what is in the data, such as a literary analysis of a large collection of unknown texts.

A lot of the work in data mining focuses on how to use big data sets to do new things. Much of the work uses data that is either publicly available or generated in-house. In either case, access control rules are uniform across the data. Public data is available to anyone. In-house data is available to those owning the data. As a result, the problems focus on data quality, data heterogeneity, new applications for the data, combining different data sets, or other issues with the complete data sets.

In cases where data is not available publicly or in-house data is supplemented with restricted data from other sources, access and use rules come into play. For example, by using extensive medical records and genome sequencing, advances in personalized medicine may be possible. However, medical records and genome sequences may each have their own separate access rules that prevent a general-purpose big data analysis across a population. Some solutions involve creation of “cleansed” data sets for use in analysis, which preserve the basic properties of interest in the data while removing identifying attributes. Others involve an all-powerful entity that is given blanket access to all records, and then providing strict controls on what this entity can release, such that end-to-end access rules are preserved while allowing aggregation at the powerful intermediary.

Using cleansed data requires the initial effort to create the cleansed data set, as well as the ongoing effort to maintain this data set. This can be a significant effort, and even cleansed data can inadvertently leak sensitive information. Releasing data to a powerful intermediary may not be desirable where sensitive information is concerned, as it provides a central point of attack.

This work proposes the ELS architecture as a way to manage access for big data analysis across an enterprise. This does not solve the problem of inaccessible data, but does make it easy to set enterprise rules for access so that all data that need to be accessible for analysis purposes is accessible.

3 ELS BACKGROUND

ELS is an architecture that provides secure access to data and services for all active entities in an enterprise. It does not address passive entities, such as information packages, static files, and/or reference data structures. Passive entities are the target of activities and do not initiate activities and cannot assume the role of a requester or provider. Active entities are those entities that change or modify passive entities, request or provide services, or participate in communication flows. Active entities are users, hardware, and services. All active entities in the enterprise have enterprise X.509 Public Key Infrastructure (PKI) certificates [4], and their private keys are stored in tamper proof, threat mitigating storage. Communication between active entities in the enterprise requires bi-lateral PKI end-to-end authentication. Active entities must be named in accordance with the enterprise naming instruction. Authorization in the operational environment is implemented by a verifiable access control claims-based process.

Claims are part of an authorization credential issued by a trusted Security Token Server (STS) and signed by that entity to preserve integrity. A claims-based credential is sent to the provider in a Simple Object Access Protocol (SOAP) [5] envelope containing a Security Assertion Markup Language (SAML) token, which includes issuance time and expiration time. Figure 1 displays Active Entity B performing authorization of Active Entity A, and Active Entity B retrieving content from a passive entity.

For access control, the required credential is the SAML Token, which is constructed at runtime by an STS that has access to an Enterprise Attribute Store (EAS) as described below. The SAML may also be created by a trusted federation partner in accordance with federation agreements. In each case, the SAML is provided directly to the provider after authentication. The provider verifies and validates the SAML and extracts the claims as needed.

3.1 Core tenets

Each component of every enterprise solution should be tested against a set of fundamental evaluation criteria or tenets. These tenets are the core philosophical drivers of all architectural decisions. The ELS tenets are as follows:

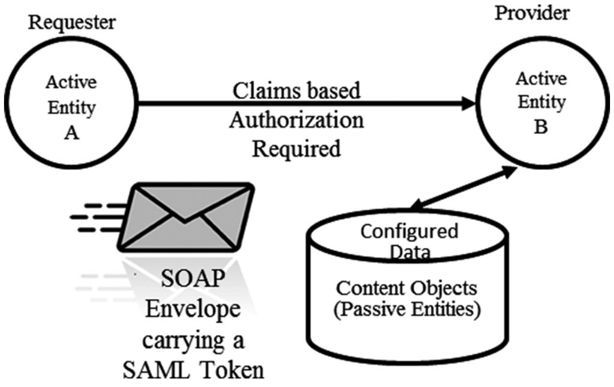
Requester
Provider
Active Entity A
Claims based Authorization Required
Active Entity B
Configured Data
Content Objects (Passive Entities)
SOAP Envelope carrying a SAML Token

Figure 1: Entity communication.

1. Malicious entities are present and our systems need to function with these embedded threats rather than rely on filtering them out.
2. Simplicity. Added features come at the cost of greater complexity, less understandability, greater difficulty in administration, higher cost, and/or lower adoption rates that may be unacceptable to the organization.
3. Extensibility. Any construct should be extensible to the domain and the enterprise, and ultimately to cross-enterprise and coalition.
4. Information hiding. This involves only revealing the minimum set of information to the requester and the outside world needed for making effective, authorized use of a capability.
5. Accountability. This means being able to unambiguously identify and track which active entity in the enterprise performed each operation.
6. Minimal detail. Add detail to the solution to only the required level. This preserves flexibility of implementation at lower levels.
7. Service-driven rather than a product-driven solution.
8. Lines of authority are preserved and information assurance decisions are made by policy and/or agreement at the appropriate level.
9. Need-to-share as over-riding the need-to-know.
10. Separation of function. Sometimes referred to as atomicity, this allows for fewer interfaces, easier updates, maintenance of least privilege, reduced and more easily identified vulnerabilities, and improved forensics.
11. Reliability. Security works even when adversaries know how it works.
12. Trust but verify (and validate). Trust should be given out sparingly, and even then, trusted outputs need checking.
13. Minimum attack surface. The fewer the interfaces and the less the functionality in the interfaces, the smaller the exposure to threats.
14. Handle exceptions and errors. Exception handling involves logging, alerting the Enterprise Support Desk (ESD), and notifying the user.
15. Use proven solutions. Select products, technologies, techniques, and algorithms that have sufficient evidence of maturity in their intended use.
16. Do not repeat old mistakes. Eliminate known vulnerabilities and exploits, use a flaw remediation system, and use patching and repairing.

Key to the big data challenge are (1) simplicity, (2) extensibility, (3) information hiding, (6) service driven, and (8) need-to-share. Simplicity and extensibility allow easy scale-up to large and diverse data sets. Information hiding allows aggregation of data without revealing all of the data itself, which is a key function of big data analyses, on sensitive data sets. Being service-driven allows general access based on a single simple interface rather than requiring a number of different access protocols, accounts, or product suites. This makes it easier to set up and maintain data access for analysis. Need-to-share is a policy that, as a default, makes data available to those who are allowed to view it instead of requiring additional and often burdensome access policies for access to data.

In addition, the idea of end-to-end security is a derived requirement [6]. End-to-end security passes original requester credentials all the way to the data store itself. Access policies are determined by the data owner and applied to the actual requester, not the intermediary. The intermediaries may provide a standardized interface for the data, but they act only as relays and have no elevated privileges beyond what the requester provides.

4 BIG DATA WITH ELS

This section describes how to merge big data analysis and ELS. This includes some changes to the way big data analysis is performed to accommodate strict ELS security controls. It also describes how to shape an ELS instantiation to better prepare it for big data analysis.

4.1 Basic ELS preparations

The first step in a secure environment is to issue strong identity credentials. Without these, nothing else works since verification of identity is at the heart of all ELS security. The recommended approach is PKI using X.509 certificates, but any sufficiently strong authentication for the purposes of the enterprise will suffice.

Next, an attribute store must be established. It must either import or natively manage entity attributes, where entities include users, machines, services, and applications. A central access control registry describes the available enterprise services and their access control rules, as specified by the data owners associated with each service. A claims engine applies entity attributes to the access control rules to generate claims, which are tokens proving that an entity is allowed to access a service or application. The claims engine places generated claims into a claims repository, which has an associated “Provide Claims” service.

An STS, or collection of STSs, is used by requestors to generate authorization credentials. The requester authenticates to the STS and the STS passes this identity to the “Provide Claims” service to retrieve associated access claims. These are packaged into an SAML token and provided back to the requester, signed by the STS for integrity. The requester forwards this to the desired service or application endpoint, which performs security checks on the SAML and authenticated requester identity.

The key to providing analysis services across the secure enterprise is consistency across the enterprise. All data providers must front their data with a web service interface, and possibly a web application if direct user interactions are desired. The web service/application interface, combined with universal strong authentication and SAML-based authorization from trusted STSs provides the basis to control access to all data across the enterprise.

4.2 Big data analysis with ELS

After the basic components are in place, big data relies on providing appropriate controls at the data repositories. Data owners can set their own policies depending on the type of data they own, its sensitivity, and any organizational or legal restrictions in place. The EAS can be configured to contain attributes relevant to big data analysis since these are applicable across an enterprise and apply to many data sets.

The actual big data analysis can be conducted in different ways:

- Local analysis,
- Enterprise level services for data analysis,
- Data owner-provided tailored analysis services,
- Federated services for multiple data sources.

For local analysis, the individual conducting the analysis simply downloads all applicable data to a local system and then performs the analysis using this system. This may stress networks and local computational resources if the data sets are especially large, but it provides a straightforward way to do the analysis using the ELS infrastructure for access controls.

For enterprise level data analysis, tools are made available as enterprise services. In this case, access to these services is open to all, but the access to data is based on the requester of the service. Different entities with different levels of access receive different results from the same query since they have access to different underlying data. This provides a more scalable solution since anyone can perform analyses without requiring a dedicated infrastructure. However, personal data sets, where each person has access only to his own data, would be difficult to perform analysis on, and these are often interesting data sources for broad analyses.

For data-owner-provided services, the data owners provide tailored big data services that involve elevated privileges for certain functions, such as averages, totals, or regional statistics. In this case, the data owner can provide internal mechanisms to do broad analysis of data sets while preventing sensitive data from leaking. Since the data owner can manage access to the data, the data owner's services can be provided specific privileges that are denied to any other entity. This allows escalation without risking an outside entity's abusing this escalation since the data owner maintains control of the external access endpoints.

For federated services, groups of data owners create analysis services that have full access to all data sets. The services only provide aggregated statistics from these data sets. This is similar to a single data owner providing such services, except this can scale to the full enterprise and allow analyses across larger data sets.

Local analysis requires the fewest changes to existing data stores, but offers the least scalability and functionality. Use of enterprise services provides more scalability by sharing computation resources and algorithms, but still lacks aggregation services over full data sets. Data owner services provide greater aggregation across a single data set, and with proper enterprise, planning may offer enterprise-wide scalability. Only federated service models provide full scalability and full aggregation capabilities, but these can introduce security vulnerabilities as a single all-powerful entity that can access many data sources. The best solution for an enterprise depends on the goals, resources, and data protection requirements of the enterprise.

4.3 Data-driven access controls

Today, databases have access controls based on user accounts [7] and [8]. These limit access to particular rows, columns, views, and actions. However, in an enterprise in which data is shared by many different owners, coordinating the database schemas, accounts, views, and other information is prohibitively costly. In ELS enterprises, the approach is to create views dynamically based on claims. Instead of user accounts or roles, which are managed by the database management system, the data owner maps claims to dynamic views of the data. Each view contains the data that the individual with the associated claim is allowed to access.

The main challenge associated with this approach is coordinating the claims and views. For a static set of access privileges, a static set of views with associated claims may be sufficient. However, in many cases the content is generated dynamically. In this case, claims must describe the set of data a requester is allowed to access, and may refer to internal schemas to allow proper creation of dynamic views. Coordination of enterprise-wide claims with local database schemas requires careful attention to maintain proper mappings.

As an example of this approach, we consider an enterprise financial database. This database has many predefined roles. These are determined by the data owner and placed in the format of an Access Control requirement (ACR) for storage in the enterprise service registry. The roles may be

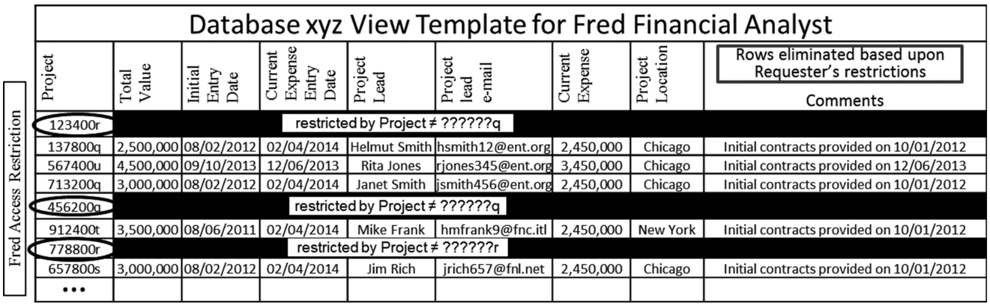

Database xyz View Template for Fred Financial Analyst

Fred Access Restriction

Project	Total Value	Initial Entry Date	Current Expense Entry Date	Project Lead	Project lead e-mail	Current Expense	Project Location	Rows eliminated based upon Requester's restrictions Comments
123400r	restricted by Project ≠ ??????q							
137800q	2,500,000	08/02/2012	02/04/2014	Helmut Smith	hsmith12@ent.org	2,450,000	Chicago	Initial contracts provided on 10/01/2012
567400u	4,500,000	09/10/2013	12/06/2013	Rita Jones	rjones345@ent.org	3,450,000	Chicago	Initial contracts provided on 12/06/2013
713200q	3,000,000	08/02/2012	02/04/2014	Janet Smith	jsmith456@ent.org	2,450,000	Chicago	Initial contracts provided on 10/01/2012
456200g	restricted by Project ≠ ??????q							
912400t	3,500,000	08/06/2011	02/04/2014	Mike Frank	hmfrank9@fnc.itl	2,450,000	New York	Initial contracts provided on 10/01/2012
778800r	restricted by Project ≠ ??????r							
657800s	3,000,000	08/02/2012	02/04/2014	Jim Rich	jrich657@fnl.net	2,450,000	Chicago	Initial contracts provided on 10/01/2012
•••								

Figure 2: Access tailored data view.

arbitrarily complex since the claims engine will compute whether or not they are satisfied and provide any variables or restrictions requested. The result is shown in Fig. 2.

One issue with this approach is that the claims, which are generally assigned and managed by the data owners within the enterprise, must be used within the database for access controls associated with views [9] and [10]. Changes in the EAS may affect low-level access policies of many databases within the enterprise. Database schemas and data elements must be provided with standard definitions to facilitate the analyses, but even under those circumstances the data may not be equivalent, making joins and other queries that coordinate across databases difficult.

4.4 Escalation of privilege

In order to have a complete analysis, the requester may have to use some form of escalation of privilege. Such escalation can be for aggregation or exposure.

Aggregated data is often made available to requesters even though the underlying raw data is not. This occurs when access to aggregated data does not violate the least privilege concept. Examples include statistical values such as the mean, median, and sum; summary geographic distribution data such as value ranges and geographic coincidence; and type sequence summaries for included data. One way to provide this access is for data owners to allow privilege to aggregated data through their access and privilege requirement documentation in the EAS registry. Specific rules determine which aggregations are permitted and who can receive them. In this case, the data owner provides and enforces the access rules. Another option is to escalate privilege in the analysis software in order to view and aggregate the data for restricted users. In this case, the aggregation service has access control responsibilities. The first method is preferable because it gives the data owner control over the data that requesters receive, but the second has the advantage that it can be layered on existing data stores.

Escalation for exposure of data is performed to provide special access privileges to an individual in order to perform a specific analysis task. This is a sensitive process since any mechanism that creates such an escalation may be misused by nefarious entities. Escalation of privilege is a prominent part of many attack vectors. For ELS systems, it is recommended that individuals assigned to the analysis tasks be pre-screened for least restriction. However, the most talented individual or most available individuals may not meet these requirements. In these cases, analysts should be provided delegated claims. In the delegation process, temporary claims are created for individuals by individuals who have those claims and are willing to be accountable for that delegation. A multi-party delegation by all data owners could allow a trusted entity to perform an analysis across all data sets with restrictions and access expirations that are agreeable to all data owners. In contrast to aggregation escalation, which is often permanent and based on attributes, exposure escalation is temporary and based on exactly the claims required for the task. This limits direct access to the data. In addition, unlike aggregation escalation, in which the aggregator is called by many other entities, the exposure escalation is limited to a single entity, reducing the attack surface.

4.5 Big data analysis using federation data

Big data analysis can be challenging when sharing information across enterprise boundaries. Within an enterprise, organizational boundaries can be addressed using the single common EAS. Such mechanisms do not exist across enterprises. Solutions with ELS include the following:

- Incorporate outside entities into the EAS.
- Provide separate credentials to vetted outsiders.
- Delegate claims to credentialed outsiders.
- Require a basic ELS setup within the partner organization.

When incorporating foreign entities into the EAS, their existing authentication credentials are trusted. This essentially brings the outsiders into the enterprise. This creates the obvious problem that these entities are not part of the enterprise but are treated as if they were. In some cases, this might be an acceptable solution, but in general, it is not.

With enterprise-provided credentials for vetted outsiders, they are more like enterprise members since they have received similar vetting. Credentials can reflect that they are not full members if this distinction is important. However, this can be costly and difficult to implement, and again it brings outsiders inside the system.

With delegation, outsiders have existing authentication credentials. Claims are delegated directly based on these credentials instead of assigning native attributes from which these claims are computed. In this case, all access and privilege is granted explicitly through a delegation assignment by an authorized individual within the organization. Delegation rules can limit who is allowed, and under what circumstances they are allowed, to delegate.

With a partner ELS setup, identity can be established through PKI or a similar method, and claims can be shared through SAMLs issued internally and validated in the partner organization using federation agreements. These federation agreements translate partner identities, claims, and attributes to local identities, claims, and attributes, such that the two partners need not change their internal security infrastructure. They need only establish a new federation agreement.

Federated ELS provides the cleanest and most secure access, but it requires the most work to set up. Other solutions offer trade-offs between security and ease of setup.

4.6 Data leakage

Data leakage may occur in an ELS environment based on differing access levels. For example, a certain data element may be sensitive, so averages are provided only for a sample size of at least N values in order to help mask the individual data values. In this case, if one active entity has access to the same data as another active entity, plus a small number of additional values, this could leak information. By both entities computing averages and comparing the values, they can compute the average of the difference, which is a small sample size with less than N items.

For ELS, the same person can often have different roles, which allow different levels of access. By logging in as different roles, a single person can perform such analyses and compute statistics that may reveal sensitive information. This is a potential security vulnerability associated with providing very broad access to data analysis. The finer the access controls, the more potential for data leakage. The enterprise must decide where the proper balance is between accessibility and confidentiality and determine access controls and data analysis functions accordingly.

This is an interesting challenge because fine-grained access controls generally provide higher security, but in this case, they can hurt it. Fine-grained access controls make the most data available to all entities that are allowed to access it, which is generally an improvement in the availability-versus-confidentiality trade-off space. However, if these controls are fine enough, individual data elements can be extracted through comparison of aggregated results. One solution is reducing the fidelity of results so that exact values are harder to compute accurately. Another is restricting the analyses that can be performed. These solutions can address the security problem, but they reduce functionality.

In some cases, the access problems may be more fundamental, such as involving an inherent conflict between needing to make some data available while hiding other data across a large population with different access rights. In this case, the ELS and big data analysis combination can serve to bring these inconsistencies to the surface so they can be addressed directly. Addressing these issues moves enterprise security from unknown and ad hoc access rules to calculated and conscious choices about data access.

5 CONCLUSION

The ELS architecture enables data owners to leverage trusted enterprise services to provide access controls for data in a standard way. This breaks the standoff between functionality and security, offering a middle ground where data owners maintain control of data and users who satisfy the rules for access automatically gain access to data. ELS does not address inherent limitations in sharing. Instead, ELS focuses on the cases in which sharing can be enabled and provides the mechanism by which rules can be established and enforced to enable such sharing. Unlike traditional methods, which would either restrict access or release data that is sensitive, ELS allows a way to provide the maximum level of access to each individual based on his individual attributes.

ELS enables big data analyses across different data sets with different security restrictions. With non-ELS systems, a choice between security or functionality is required. Data sets must be public or made available through special agreements despite rules governing use. The restrictions on public data or special agreements are too restrictive for a general purpose method of analysis. With ELS, data owners maintain control over their assets by using enterprise services to determine access based on data-owner-defined rules and trusted EASs. For big data analysis of sensitive data, this can make the difference between running an analysis and not running it. The flexibility of ELS allows trade-offs to match the security requirements of the data owners and the enterprise.

Federated sharing poses additional challenges since there may be no common language for access rules, no common trusted identity issuer, and no common trusted attribute stores. Without these bases for common security policies, agreements must be put in place to either trust other partners' security infrastructure or integrate their users into an existing ELS infrastructure. Different options exist according to the partner's existing infrastructure and willingness to share and trust security information and functions.

Analysis by different users can result in different results since the data-driven model of ELS provides access based on the requester's credentials. There may be no complete picture of the data if access policies prevent full data access by any individual. However, in these cases the whole picture is not achievable without special security considerations, so ELS provides what is available to each user, which improves availability while maintaining confidentiality. One area of concern specific to ELS and big data analysis is the idea that some data may be leaked by different users' using comparative analysis. Although the data sets analyzed may all be big enough to hide individual data values, differences may be small enough to extract details about parts of the overall data sets. Care must be used in coordinating ELS access policies and big data analysis parameters.

For ELS to work, some low-level details must be configured properly. Distributed databases must use standard schemas, standard views, and standard approaches to escalation of privilege where appropriate. This may require internal changes to existing data stores. Within the ELS architecture, other options, such as aggregation services, can provide a quick fix with minimal changes to existing systems.

This research is part of a body of work for high assurance enterprise computing using web services. Elements of this work include bi-lateral end-to-end authentication using PKI credentials for all person and non-person entities, a separate SAML credential for claims-based authorization, full encryption at the transport layer, and a defined federation process. Many of the elements of this work are described by Simpson *et al*. [11–14].

ACKNOWLEDGMENT

This work was supported in part by the U.S. Secretary of the Air Force and The Institute for Defense Analyses (IDA). The publication of this paper does not indicate endorsement by any organization in the Department of Defense or IDA, nor should the contents be construed as reflecting the official position of these organizations.

REFERENCES

[1] NIST, *NIST Big Data Interoperability Framework (seven volumes),* NIST Special Publication 1500-1, NIST Big Data Public Working Group (NBD-PWG) Definitions and Taxonomies Subgroup, Information Technology Laboratory, 2015.

[2] Magoulas, R. & Lorica, B., *Introduction to Big Data.* Release 2.0, O'Reilly Media: Sebastopol, CA, 2009.

[3] Siwach, G. & Esmailpour, A., Encrypted Search & Cluster Formation in Big Data. *ASEE 2014 Zone I Conference. University of Bridgeport*, Bridgeport, CT, March 2014.

[4] IETF RFC 2585, *Internet X.509 Public Key Infrastructure Operational Protocols: FTP and HTTP*, May 1999.

[5] Mishra, P. Maler, E., Cahill, C.P., Hughes, A.J., Beach, M., Metz, B.R., Randall, R., Wisniewski, T., Reid, E.I., Austel, P. & Hondo, M., *Conformance Requirements for the OASIS Security Assertion Markup Language (SAML) V2.0*. OASIS Standard, March 2005.

[6] Foltz, K. & Simpson, W., Enterprise level security – basic security model. *19th World Multi-Conference on Systemics, Cybernetics and Informatics, WMSCI 2016*, Orlando, FL, March 2016, In Publication.

[7] Ullman, J., *First Course in Database Systems*, Prentice–Hall Inc., Simon & Schuster, 1997.

[8] Hershey, W. & Easthope, C., A set theoretic data structure and retrieval language. *Spring Joint Computer Conference, May 1972 in ACM SIGIR Forum*, **7(4)**, pp. 45–55, 1972.

[9] Oracle, *My SQL Stored Programs and Views*, available at http://docs.oracle.com/cd/E19078-01/mysql/mysql-refman-5.0/stored-programs-views.html#stored-routines-syntax (accessed October 2014).

[10] Purdue, *Using Stored Procedures to Set Views*, available at https://www.cs.purdue.edu/homes/ninghui/projects/Topics/DB_FineGrained.html (accessed October 2014).

[11] Simpson, W. & Chandersekaran, C., A SAML framework for delegation, attribution and least privilege*, 3rd International Multi-Conference on Engineering and Technological Innovation, IMETI 2010*, Vol. 2, pp. 303–308, Orlando, FL, 2010.

[12] Simpson, W. & Foltz, K., Lecture notes in engineering and computer science. *World Congress on Engineering 2015*, *Wide Area Network Acceleration in a High Assurance Enterprise,* pp. 502–507, London, July 2015.

[13] Simpson, W., Chandersekaran, C. & Foltz, K., Lecture notes in engineering and computer science. *World Congress on Engineering and Computer Science 2014, Distributed versus Centralized Protection Schema for the Enterprise,* pp. 68–73, Berkeley, CA, October 2015.

[14] Simpson, W. & Foltz, K., Lecture notes in engineering and computer science. *Proceedings World Congress on Engineering and Computer Science 2015*, *Maintaining High Assurance in Asynchronous Messaging,* Vol. 1, pp. 178–183, Berkeley, CA, October 2015.