



Synchronization of Unique Identifiers Across Security Domains

by Mark R. Mittrick and Gary S. Moss



ARL-MR-712

February 2009

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-MR-712**February 2009**

Synchronization of Unique Identifiers Across Security Domains

Mark R. Mittrick and Gary S. Moss
Computational and Information Sciences Directorate, ARL

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.				
1. REPORT DATE (DD-MM-YYYY) February 2009		2. REPORT TYPE Final		3. DATES COVERED (From - To) October 2005–Present
4. TITLE AND SUBTITLE Synchronization of Unique Identifiers Across Security Domains		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Mark R. Mittrick and Gary S. Moss		5d. PROJECT NUMBER 9TV0VC		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory AMSRD-ARL-CI-IC Aberdeen Proving Ground, MD 21005-5067		8. PERFORMING ORGANIZATION REPORT NUMBER ARL-MR-712		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT The main thrust of the Global Force Management Data Initiative (GFM DI) is the standardization and exposure of force structure information across the Department of Defense (DOD) components. The GFM Community of Interest developed the GFM extensible markup language schema definition to facilitate the exchange of information between and within the DOD enterprise. To implement the GFM DI, each component must instantiate an Organizational Server (Org Server) that can provide data conforming to this published schema. Key to this data standardization effort is the ability to refer to any GFM DI element via an identifier that is unique across the enterprise, called a Force Management Identifier. This capability is being implemented using a construct called Enterprise-wide Identifiers (EwID), developed by the U.S. Army Research Laboratory in 2001. A Web application called an EwID Enterprise Seed Server (ESS) manages the allocation of EwID prefixes, which Org Servers use to generate EwIDs. Maintaining a single-point control of this process assures that all EwID prefixes, and therefore EwIDs, are unique. However, for the GFM DI to be truly global, the information must be consistent across security domains. Thus, EwIDs will be unique regardless of the classification level of the associated data. The GFM DI is addressing this requirement by implementing a cross domain solution that enables a classified instance of the ESS to delegate an EwID Seed request to an unclassified instance of the ESS.				
15. SUBJECT TERMS GFM, cross domain info exchange, EwID, radiant mercury				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 20
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED		
				19b. TELEPHONE NUMBER (Include area code) 410-278-4148

Contents

List of Figures	iv
Acknowledgments	v
1. Introduction	1
2. The Enterprise-wide Identifier (EwID)	1
3. Mission Need	2
4. Radiant Mercury Guard Selection	2
5. Communication	3
5.1 Low-to-High Message Traffic.....	5
5.1.1 1A, 1B: UpdateDatabase WS Request	5
5.1.2 1C, 1D: UpdateDatabase WS Response.....	5
5.2 High-to-Low Message Traffic.....	6
5.2.1 2A, 2B: GetNewSeed WS Request	6
5.2.2 2C, 2D: GetNewSeed WS Response.....	6
6. Status	7
7. Issues	8
8. References	9
Distribution List	10

List of Figures

Figure 1. EwID structure (2).....	2
Figure 2. GFM cross-domain flow chart.....	6
Figure 3. Multilevel ESS allocation (nested structures indicate data replication from lower to higher security domains).....	7

Acknowledgments

This report would not have been possible without the efforts of the Global Force Management team and the Tactical Information Fusion Branch.

INTENTIONALLY LEFT BLANK.

1. Introduction

The Joint Staff, Force Structure, Resources, and Assessment Directorate and the Office of the Under Secretary of Defense for Personnel and Readiness established the Global Force Management Community of Interest (GFM-COI) in the summer of 2003. An excerpt from a conference paper (1) on the GFM-COI states:

The objective for Global Force Management (GFM) is to establish a transparent and universal process to manage, assess and display the worldwide disposition of US forces. This includes US force availability, readiness and capability in order to assess the risks associated with proposed allocation, assignment and apportionment options. Fundamental to GFM and foundational to transformation is the GFM Data Initiative (GFM DI), which addresses organizing force structure data in a joint hierarchical way for integration across Service lines.

Force structure data is instrumental to integration because it ties everything in the Department of Defense (DOD) together. Prior to the GFM DI, each service had a unique means of representing and documenting force structure. This information was in different formats and distributed using many different sources, which made it hard to locate and use. The GFM DI is addressing this issue by exploiting the Net-Centric Data Strategy and current data representation standards and tools.

One of the challenges is that the process must be usable across service, coalition, and security classification boundaries. This requires a system for generating identifiers that are unique across these boundaries. The U.S. Army Research Laboratory (ARL) has developed a generalized construct called Enterprise-wide Identifier (EwID) that is suitable for this purpose.

2. The Enterprise-wide Identifier (EwID)

EwIDs are 64-bit sequences consisting of a 32-bit (4-byte) prefix concatenated with a 32-bit suffix (see figure 1). The Enterprise Seed Server (ESS) allocates the prefix referred to as an EwID Seed. Anyone with a legitimate need can request an account on an ESS and can obtain EwID Seeds as needed. The ESS ensures that each EwID Seed is unique; it is the responsibility of the recipient of an EwID Seed to provide the means to generate unique 32-bit suffixes for their particular application. The generation of EwIDs is fundamental to the implementation of an organizational server because the server uses them to implement force management identifiers that identify all data in the XML schema definition.




Figure 1. EwID structure (2).

EwID Seeds are currently available from the NIPRnet instance of the ESS located at <https://ess.arl.army.mil>. At the time of this writing, ARL was in phase 2 of the cross domain solution (CDS) process, which is a DOD level accreditation. ARL will instantiate a SIPRnet instance of the ESS once it has officially obtained the interim authority to operate the CDS.

3. Mission Need

An ESS delivers a prefix which is used to build EwIDs by data sources such as the GFM organization servers and related authoritative data sources (ADS). The GFM ADS will have a cross-domain implementation to pass the GFM data from lower to higher (i.e., more restrictive) security domains, and many operational systems will use the data from these GFM ADS. The ESS's allow the GFM ADS to maintain tracking information through point of contact (POC) information or a link to an EwID-tracking Web service (see section 5). The ESS tracking service implements an iterative search for the information by following consecutive links provided by the ADS maintainer. With many operational systems using the GFM DI data and passing data tagged with EwIDs between security domains, the EwID tracking service must be capable of immediately locating data from the high side, even when it resides on or is created and transferred from the low-side network. The ability to generate and track EwIDs in real time to support operational systems is critical and requires the development of a cross-domain solution.

4. Radiant Mercury Guard Selection

A variety of factors determined which cross-domain solution was ultimately chosen. First, the GFM/ESS team investigated the currently available cross-domain guards with regard to the services that they supported. After the initial evaluation, the team narrowed the selection down to three choices—Radiant Mercury (RM) guard, Information Support Server Environment guard, and U.S. Naval Research Laboratory pump. The next step focused on more in-depth criteria involving conversations with vendors to select the one guard that would best fit the needs of GFM. From those conversations, it was determined that the Radiant Mercury high-assurance guard would be the most cost-effective solution for GFM because it supported bidirectional

transmission control protocol and internet protocol (TCP/IP) socket capability and extensible markup language (XML). In addition, RM was secret and below interoperability and top secret and below interoperability certified and provided on-site support and training as well as help with the approval process.

5. Communication

The ESS was originally implemented in Macromedia* ColdFusion* (ColdFusion has since been acquired by Adobe†). Given that the ESS is a Web application, the ARL developer decided that the most robust means of developing a distributed capability was via a Web service (WS). WS's are standards based and agnostic with regard to the underlying software or hardware platform. This flexibility is vital to scalability and portability of the software. In addition, Macromedia developed the ColdFusion MX release based on the Java 2 Enterprise Edition (J2EE). Therefore, ColdFusion applications can be hosted on most popular operating systems due to the availability of J2EE compliant application servers. Ultimate control of the ESS has not been decided, so the final software/hardware platform is yet to be determined.

Implementing WS's across an RM guard is not technically difficult. However, the RM guard does not have any proven means of facilitating the transfer from an information assurance perspective. WS's are remote procedure calls encoded using the XML and transmitted over a full duplex TCP/IP socket using the hypertext transfer protocol (HTTP) packets containing simple object access protocol (SOAP) payloads. The RM guard supports TCP/IP sockets but does not come with any support for the parsing of HTTP packets. Furthermore, although RM has an XML parsing capability, it does not use a standard XML parser, so it is not a simple thing to validate XML/SOAP (3) against a published XML schema (4). This does not appear to be a problem for RM technicians as they have been able to program the RM message analysis and generation parser to pass the WS calls. However, at the time of this writing, the GFM team has yet to undergo phase 2 of the CDS process that will test for vulnerabilities as part of security testing and evaluation.

A client application initiates a WS call by sending an HTTP request packet to the Web server hosting the WS. The server processes the request and then sends an HTTP response packet back to the client. Therefore, each WS call requires a full-duplex channel through the RM guard (effectively one channel in each direction). Communication between the NIPRnet (unclassified) instance of the ESS and the SIPRnet (secret) instance of the ESS requires the following two WS's:

* Macromedia and ColdFusion are registered trademarks of Macromedia, Inc.

† Adobe is a registered trademark of Adobe Systems, Inc.

1. UpdateDatabase – replicate specified database operations.
2. GetNewSeed – retrieve the next available EwID Seed.

The unclassified ESS is the “master” server and is responsible for handing out all EwID Seeds, regardless of security domain. The unclassified (NIPRnet) network is referred to as the “low-side” and the secret (SIPRnet) network as the “high-side” to indicate the relative security level. For the low-side server to maintain a single point of control over the allocation of EwID Seeds, the high-side server must request seeds from the low-side server through the RM guard. This is the purpose of the GetNewSeed WS. The high-side server invokes the GetNewSeed WS when a user asks for a new seed. The low-side server response contains the new seed value.

The other requirement for cross-domain communication is for EwID tracking via its seed. EwID tracking is a service provided by the ESS that returns the POC information from the ESS account to which the seed was assigned. The ESS maintains records in a database of all EwID Seed allocations and assignments. Seed allocation guarantees that a specified number of seeds can be obtained, whereas seed assignment is the reception of a seed. EwID tracking is a service provided by the ESS that ultimately retrieves information about the item tagged with a specified EwID. However, since the ESS does not manage the generation of EwIDs from EwID Seeds, it is the responsibility of the ESS user (e.g., an ADS owner) to maintain POC information in the corresponding ESS account as well as a reference to a tracking service. The ESS will delegate the task of information retrieval and display it to the provided tracking service. To this end, ARL will publish a tracking service Web service description language (5). The high-side user must be able to track EwID Seeds or EwIDs allocated from the local security domain and lower. Therefore, all information necessary for tracking EwID seed allocation needs to be replicated to higher security domains.

To address this, the high-side ESS hosts an UpdateDatabase WS. Whenever the low-side ESS generates a structured query language (SQL) statement associated with EwID Seed management, it calls the high-side UpdateDatabase WS, via the RM, with an XML-encoded representation of that SQL. The high-side server decodes the XML back into SQL and applies it to a shadow copy of the low-side database.

Given that data is replicated to higher security domains, there is still an issue that would arise if a user would attempt to track an EwID that was requested from a lower security domain (from where the tracking request is originated). The problem is that the uniform resource locator (URL) of a tracking service is in the address space of the network where it is hosted and that network is not reachable from a higher security domain. Therefore, the ESS would not be able to delegate the tracking request to the specified tracking service. The solution is for each tracking service to be registered in a Universal Description Discovery and Integration (UDDI) (6) repository. The ADS owner would specify the uniform resource name (URN), which is basically a unique string that will identify the UDDI entry for the tracking service in a UDDI registry yet to be identified. Each security domain will have one UDDI registry for this purpose. The ADS

owner will register the tracking service URL for each security domain in the respective UDDI registry, but the URN associated with a particular ADS owner will be the same across security domains. This level of indirection, using a URN to look up the tracking service URL, solves the issue. The UDDI entry can also store a URL, pointing to contact information on the current network.

Implementing a WS call from the low-side network to the high-side network (or vice versa) requires a cross-domain device (guard) because one network is not reachable from the other. The RM guard has two Ethernet interfaces so it can act as a proxy to facilitate the transmission of packets from one network to the other. WS calls on either network must use the RM guard's IP address as the WS endpoint; then, the guard performs network address translation (NAT). NAT is the process of modifying network address information for the purpose of remapping a given address space into another (7).

Figure 2 shows the two WS requests that result in four distinct data channels through the RM guard. Message types are labeled with a number: **1** for the **UpdateDatabase** WS and **2** for the **GetNewSeed** WS. A label also contains a letter indicating handling of the message as it goes through the guard. For instance, the UpdateDatabase WS request starts out as **1A** on the low side of the guard and then continues as **1B** on the high-side network. The RM guard inspects each packet to ensure that it adheres to the schema. If it does, the guard passes it, providing NAT in the process. If any part of the message is not in the expected format, the RM guard grounds the packet (blocks it) and logs the action.

5.1 Low-to-High Message Traffic

5.1.1 1A, 1B: UpdateDatabase WS Request

The request keeps the required tables mirrored from the unclassified side to the classified side by, in effect, implementing remote SQL transactions in real time.

5.1.2 1C, 1D: UpdateDatabase WS Response

The WS returns success or failure of that update, saving any failed transactions to a log file and alerting ESS via e-mail. Later, the administrator can process failed transactions by running a software utility.




Figure 2. GFM cross-domain flow chart.

5.2 High-to-Low Message Traffic

5.2.1 2A, 2B: GetNewSeed WS Request

The high-side server sends a seed request via the RM guard to the low side. The high-side seed server has an account on the low-side seed server; the request payload need only contain the account number of the high-side server. This account number is an EwID. The ESS reserves one EwID Seed for each server instance and uses it to generate primary keys for all its database tables.

5.2.2 2C, 2D: GetNewSeed WS Response

The WS returns the newly allocated EwID Seed to the high-side server, which maintains all tracking information for the user requesting the seed.

Figure 3 shows how the GetNewSeed WS can facilitate EwID Seed allocation given two or more security domains via a cascading mechanism (left-pointing arrows represent the GetNewSeed WS request). Also depicted is the chained replication of data upward via the UpdateDatabase WS request so that at any classification level, there is access to the tracking information for the current classification level and below.




Figure 3. Multilevel ESS allocation (nested structures indicate data replication from lower to higher security domains).

6. Status

- The RM guard was installed at ARL in May 2006. It consisted of a Sun Blade 150 computer with Trusted Solaris 8 as its operating system. It also contained the RM software, which performed the guard functions.
- The ESS software was developed and tested in November 2006.
- The Cross Domain Validation and Approval Request was completed and submitted in January 2007.
- The Cross-Domain Appendix was completed and submitted in January 2007.
- Phase 1 of the CDS process was completed in February 2008. This approval allowed the process to go to phase 2.
- The Data Owner's Guidance was completed and signed in June 2008.
- Currently, the GFM team is working on phase 2.

7. Issues

A major issue for the GFM team was that as a computer scientist, there was limited background and knowledge in information assurance subjects. There were also significant cultural differences between the computer science discipline and information assurance, which led to difficulties in communicating concepts such as data representation, software architecture, and programming implementation details. This communication gap was responsible for multiple rewrites, which resulted in severe setbacks within the Defense IA/Security Accreditation Working Group review process.

Another issue the team encountered was that the Cross Domain Information Exchange (CDIX) process was reengineered by the newly formed Unified Cross Domain Office (UCDMO).^{*} The UCDMO has been tasked with taking charge of the CDIX process, has recognized many of these gaps, and is working with the cross-domain community to address and correct these issues as well as streamlining the entire process.

^{*} Located at the U.S. Army Adelphi Laboratory Center, Adelphi, MD.

8. References

1. Chamberlain, S.; Boller, M.; Sprung, G.; Badami, V. Establishing a Community of Interest (COI) for Global Force Management. *Proceedings of the 10th International Command and Control Research and Technology Symposium*, Ritz-Carlton Hotel, McLean, VA; 13–16 June 2005; also see http://www.dodccrp.org/events/10th_ICCRTS/CD/track12.htm.
2. Chamberlain, S. An Enterprise Identifier Strategy for Global Naming Across Arbitrary C4I Systems. *Proceedings of the 6th International Command and Control Research and Technology Symposium*, U.S. Naval Academy, Annapolis, MD; 19–21 June 2001; also see http://www.dodccrp.org/events/6th_ICCRTS/Tracks/Papers/Track2/059_tr2.pdf.
3. Mitra, N.; Lafon, Y. *Soap Version 1.2 part 0: Primer (second edition)*; 27 April 2007; <http://www.w3.org/TR/2007/REC-soap12-part0-20070427/> (accessed 7 October 2008).
4. Lafon, Y. *Soap Version 1.2 part 1*; May 2003; <http://www.w3.org/2003/05/soap-envelope/> (accessed 7 October 2008).
5. Christensen, E.; Curbera, F.; Meredith, G.; Weerawarana, S. *Web Services Description Language (WSDL)*; 15 March 2001; <http://www.w3.org/TR/wsdl> (accessed 7 October 2008).
6. Universal Description Discovery and Integration Standard. <http://uddi.xml.org/> (accessed 7 October 2008).
7. Rekhter, Y.; Moskowitz, B.; Karrenberg, D.; de Groot, G. J.; Lear, E. Network Working Group Request for Comments: 1918; network address translation; <http://tools.ietf.org/html/rfc1918> (accessed 7 October 2008).

NO. OF
COPIES ORGANIZATION

1 DEFENSE TECHNICAL
(PDF INFORMATION CTR
only) DTIC OCA
8725 JOHN J KINGMAN RD
STE 0944
FORT BELVOIR VA 22060-6218

1 DIRECTOR
US ARMY RESEARCH LAB
IMNE ALC HR
2800 POWDER MILL RD
ADELPHI MD 20783-1197

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CI OK TL
2800 POWDER MILL RD
ADELPHI MD 20783-1197

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CI OK PE
2800 POWDER MILL RD
ADELPHI MD 20783-1197

ABERDEEN PROVING GROUND

1 DIR USARL
AMSRD ARL CI OK TP (BLDG 4600)

NO. OF
COPIES ORGANIZATION

1 DIR USARL
AMSRD ARL CI I
B BROOM
2800 POWDER MILL RD
ADELPHI MD 20783-1197

2 DIR USARL
AMSRD ARL CI IB
R WINKLER
L TOKARCIK
2800 POWDER MILL RD
ADELPHI MD 20783-1197

1 DIR USARL
AMSRD ARL CI NT
G CIRINCIONE
2800 POWDER MILL RD
ADELPHI MD 20783-1197

ABERDEEN PROVING GROUND

22 DIR USARL
AMSRD ARL CI IC
F BRUNDICK
S CHAMBERLAIN
T HANRATTY
H INGHAM
M MITTRICK (15 CPS)
G MOSS
J RICHARDSON
M THOMAS

INTENTIONALLY LEFT BLANK.