

Secure Data Services Platform (SDSP)

A Web Services Based Data Distribution Capability And Architecture

24 March 2008

Prepared by:



LVI-OakHill

1320 Old Chain Bridge Road, Suite 250

McLean VA 22101

703-761-3060

<http://www.lvi-usa.com>

This document is unclassified.

Table of Contents

| | |
|---|----|
| 1. Introduction..... | 3 |
| 2. Background..... | 3 |
| 3. Reference Architecture Overview..... | 3 |
| 4. Reference Architecture Core Features | 4 |
| a. Service Based Design..... | 4 |
| b. Embedded Management and Control Functionality | 5 |
| c. Legacy/Heritage Transition Enabler | 6 |
| d. Robust Scalable Hardware/Software Infrastructure | 7 |
| 5. Detailed System Description..... | 9 |
| a. Baseline System Build (SDSP-Core)..... | 9 |
| b. Cross Domain Enhancement (SDSP-CDE) | 10 |
| c. Application Services Interface (SDSP-ASI) | 11 |
| d. Extended Data Storage Layer (SDSP-EDL)..... | 11 |
| 6. Capabilities Review | 12 |
| 7. Framework Analysis | 13 |
| a. Pros/Cons | 13 |
| b. Expected Efficiencies | 14 |
| c. Programmatic/Technical Challenges | 14 |
| 8. Conclusion | 15 |

List of Figures

| | |
|--|----|
| Figure 1: SDSP-Core High Level Architecture | 9 |
| Figure 2: SDSP-CDE Enhancement High Level Architecture | 10 |

List of Tables

| | |
|---|----|
| Table 1: Concept Development Paths..... | 12 |
|---|----|

Executive Summary

LVI has designed a reference architecture for secure web services based data distribution that can be used to augment and extend existing infrastructures. The capability is an integrated set of systems collectively called the Secure Data Services Platform (SDSP). The SDSP is an infrastructure capability which resides in a data center and provides services for multiple applications and end users through a series of data movement services.

Functionally, the SDSP ingests high volume data in any format and then provides the data to consumers via standard web service calls. The initial design of the SDSP is based on an Oracle database system running in a highly scalable clustered computing environment on the Linux operating system. However, the platform was designed from the ground up as a non-proprietary web services solution so other database products could be substituted depending upon the client needs. The SDSP can be deployed into any environment with structured data and can be rapidly configured to provide a consumer-oriented web service. Extending the system to include new data feeds can be accomplished using either web service calls to other services, or data ingest routines on the SDSP. The time frame for including a new “feed” is driven by the complexity of the data source and the degree to which the relationships and meaning between the data elements is documented. In addition to moving data, the SDSP provides an exhaustive set of metrics to track data entering, being processed and exiting the system. These metrics are available through the web services interface, allowing data producers and consumers to have insight into the operation of all aspects of the system. To use a simple analogy the SDSP is like an efficient and scalable valet service for structured data. Data arrives and is “parked” in the system until it is removed and used by other consumers. We have experience parking various types of DOD and Intelligence Community data and thus far have been able to rapidly incorporate each data feed presented to us using the Commercial Off-The-Shelf (COTS) tools on which the SDSP is built.

LVI’s current customers are using the SDSP as a transition enabler as they move from legacy systems to a new service based paradigm. The reference architecture provides a complete transition road map for providing consistent data distribution across both legacy and newer service based tools within and between organizations.

This white paper outlines the history of the SDSP development, explains the current and future generations of the SDSP, and analyzes the architecture from the perspective of stated DOD and Intelligence Community requirements.

1. Introduction

LVI developed this reference architecture for a data and information distribution solution that combines data and web services tiers creating a self contained service based infrastructure. This reference architecture is implemented as an integrated set of systems collectively called the Secure Data Services Platform (SDSP). The architecture provides for a phased deployment that we call generational cycles. These “generations” contain iterative cycles that are managed and developed using agile development processes. LVI has currently defined six generations of the architecture, each of which adds increasing levels of capability and support for data and information handling needs. Development requires this type of deliberate growth to account for the maturity of emerging technologies. While these technologies contain great promise, their adoption first by commercial organizations and then by US Government agencies will require both time and cultural acceptance. Additionally, the implementation plan provides for rapid deployment of initial capability in operational environments, with iterative updates to complete the entire generational change.

A version of the SDSP is set for testing and accreditation within the Defense community and is scheduled for fielding starting in the last quarter of calendar year 2008 (CY2008).

2. Background

The first generation of the SDSP architecture originated from a research & development (R&D) effort by the Under Secretary of Defense for Intelligence (USDI). The results identified that a system combining a back-end data tier with a front-end web services tier could provide the level of flexibility, security, scalability and stability necessary to support internal and external data customers. The second generation was green lighted to support the data infrastructure needs of the Defense Intelligence Agency (DIA)'s Collection Management Mission Applications (CMMMA) Project. Along the way, LVI has extended the architecture to provide several additional generations that can expand to support environments beyond that of the original DOD prototype.

3. Reference Architecture Overview

At its core, the current SDSP design combines an Oracle 10gR2 relational database system paired with a service oriented mid-tier running on a scalable, clustered operating environment. The current design supports both Solaris 10 and Linux operating system environments. The platform includes all the capabilities required for data ingest, security labeling, data storage, user authentication, subscription access, and data dissemination. The web services tier provides the typical REST, SOAP, ATOM, and RSS style services seen in commercial solutions. In addition, the system employs Public Key Infrastructure (PKI) certificates and Secure Socket Layer (SSL) technologies to address the secure authenticated delivery of data to consumers.

LVI has outlined a series of extensions that can be deployed on top of the core architecture to extend the capability of the SDSP for specific environment. One of these extensions is the ability to provide a cross domain database capability using Oracle's Cross Domain Security Solution (CDSS). There are a number of infrastructure enhancements required to support this extension, but for customers needing a DOD/IC

mandated PL4 solution, this extension does allow the SDSP to provide this capability. A second extension expands the web services capabilities to support interfaces that are more consistent with the Web 2.0 philosophy. This enhancement adds capabilities based on various W3C standards to provide the type of rich data sharing environment found in the commercial Internet space. These include support for syndicated distribution of content, and peer to peer services for connecting consumers. It also allows the SDSP to be a repository of application produced data and introduces a limited capability for unstructured data. A third extension offers semantic services utilizing triple store technology and expands support for unstructured and semi-structured data.

From the database tier, the current architecture supports a combination of technologies including Oracle Real Application Clusters (RAC), data warehouse style relational databases, ontological triple store databases and non-database collections. On the web tier, the architecture supports traditional Simple Object Access Protocol (SOAP) and Representational State Transfer (REST) services as well as Atom/RSS feed and publication services. All of these combine to offer a flexible, scalable solution for handling a variety of data types, data rates and data roles. As technologies advance, new hardware and software can be included in the overall architecture to address specific mission needs. While the initial design uses Oracle technology, there is nothing that precludes the use of other technologies, including Open Source database products like MySQL.

In keeping with the commercial practices and standards, the SDSP relies heavily on the eXtensible Markup Language (XML). The SDSP is capable of processing XML data from external sources using a variety of ingest methods and can store and manage the data in its raw XML form or in a processed form more suitable for relational database storage. On the output side, the SDSP provides outbound subscription data as XML documents. In addition, the system conforms to the Atom/RSS model definition for collections and entities for both feed services and publishing services.

4. Reference Architecture Core Features

The reference architecture contains five core features which guide the overall design and implementation:

- Service Based Design
- Embedded Management and Control Functionality
- Legacy/Heritage Transition Enabler
- Robust Scalable Hardware/Software Infrastructure
- Phased Implementation Methodology

The following paragraphs discuss these features in detail.

a. Service Based Design

Starting with the original prototype, the SDSP was created following the Service Oriented Architecture (SOA) paradigm. While the term SOA has become passé, the concepts behind SOA provide a viable method for developing service based architectures for enterprise level systems. From the

SDSP perspective, we fully support the original spirit and concepts behind SOA and use this not as a technology buzzword but as a true paradigm for how to build scalable, flexible architectures supporting rapid growth in delivered service capability and quality. From day one, the SDSP architecture has incorporated a clear distinction between the data and web service tiers. Standard web services technologies, such as SOAP and REST have been included in the design to insure we can fully support today's commercial best practices. Additional technologies will be added through each successive update to complete the desired capability set and to maintain pace with changes in the commercial and government communities.

The proper use of these technologies drives a requirement for fundamental changes in the way every component of an application is built. Legacy database applications share the characteristic of tight coupling between the application and the data because they were designed for the user to be sitting at or very close to where the actual hardware and storage existed. Even the advent of server based applications did not really change this paradigm. They were developed in environments where 100Mbps connectivity was the norm. The underlying database schema did not require significant optimization because it was easy to post-process volumes of data at the desktop. However, this operational design does not work in a distributed global web environment. The entire underlying database structure has to change to truly support a service based environment. It simply does not work to 'add SOA services' to most legacy/heritage applications.

The SDSP does much of the intelligent processing of data at ingest time and allows the consumer -- whether a system or a person -- to request only the data required. Standards based service technologies are fully implemented and the database tier supports application development in a service based environment. This allows us to expose all the services of the SDSP architecture for use by any application. This includes capabilities like the embedded Lightweight Directory Access Protocol (LDAP) repository or the typical create, read, update, and delete (CRUD) database services.

b. Embedded Management and Control Functionality

One of LVI's core competencies is system and application management. Our principal engineers have direct experience with the design and implementation of management and control systems for programs in the commercial and government space (UUNET [now Verizon], AT&T, Sprint, Defense Intelligence Agency (DIA), Naval Research Lab (NRL), Federal Bureau of Investigation (FBI) and the Defense Information Systems Agency (DISA)). Through our collective experience, we've employed the best practices of commercial industry and incorporated the needs of government agencies into the SDSP.

The SDSP collects metrics during all phases of the data distribution cycle, from ingest through delivery. It collects metrics for the infrastructure

hardware and the data processing layer where data is ingested or distributed. At the data layer, usage metrics provide data producers, data consumers, security administrators, and system administrators with visibility into all aspects of the system. They allow data producers, those who provide data to the system, to see when and how their data was received, how it was processed, stored and managed on the SDSP, and how it was distributed to the various data consumers. It allows system and security administrators to audit and trace the flow of data from producer to consumer to verify the security and integrity of all data flowing through the system. It provides consumers with a clear picture of how the system is operating and provides them with the ability to assess the quality of the data they are receiving.

Collected metrics are available through the various subscription services. This allows the metrics to be used to drive external monitoring tools ranging from simple status web pages to sophisticated manager of manager (MOM) systems. Since the metrics are exposed in real-time through the web services layer they can be used by the SDSP and other systems to drive automated changes to the operational and security posture of the system.

Control at the infrastructure level is provided by standard COTS tools integrated into the SDSP. Management metrics provides control at the application layer by allowing each interfacing application to retrieve metrics related to the data and services the application is using.

c. Legacy/Heritage Transition Enabler

The SDSP supports transitioning critical legacy systems into a service based environment. During our analysis of several such systems for the defense community, we discovered several areas where the SDSP can help transition legacy systems.

The first scenario is to have the SDSP provide a customer/consumer-facing web services layer than can be an intermediary service during the transition process. This applies to legacy systems which cannot support the load requirements of additional user processing or where retrofitting a web services capability is too complex and/or too costly. The current SDSP design uses a COTS ingest tool called Oracle Warehouse Builder (OWB) which is used to rapidly build ingest scripts. The data can be ingested in various methods ranging from flat files to direct database calls to a legacy system. Other ingest methods are also possible either within OWB or by using other tools to receive, processes and load incoming data. Whether this comes from in-house generated parsing tools based on C, PERL, .NET, Java, or other programming/scripting languages or through other commercial data and text processing tools will depend on the specific situation. The key principle is that the SDSP doesn't constrain the user to just having OWB, but offers multiple ways to handle the ingest process either with legacy tools or newly defined methods. In the case of simple flat files, the development of ingest logic and the setup of appropriate data schema could be completed within a

single day. For one of the more complex cases, we performed a direct login and extracted 25 tables of data, merged the data on the SDSP based on consumer requirements, and then provided the data back as a service. This capability was significantly less expensive and much faster than converting the legacy system. Using these techniques, the cost and implementation time can be reduced exponentially, and they can be used to provide web based services where it would otherwise be impractical.

A second transition scenario is to move commonly used reference data onto the SDSP. Applications in the same portfolio frequently ingest and store the exact same reference data from an external data source. In some cases this results in infrequently used data driving entire system architecture processing and storing requirements. Shared reference data can be brought into the SDSP and applications can make web services calls when the data is required. This can significantly reduce storage and processing requirements depending on the type of data. A variation is to bring any new data into the SDSP requiring applications to make web calls to the SDSP to obtain the data. This approach can be used when the incoming data is not available through an existing web service. The combination of these two approaches can extend the life of a critical legacy application at a relatively low cost, even during development of newer capabilities.

The third scenario is to develop new web based applications. The SDSP provides the capability to both read and write data to the data service tier via standard web service calls. The benefit is that developers have a standards based, commercial style web services platform to build upon when integrating with legacy systems. This allows the SDSP to function within an organization as transition point into the web services development environment without requiring fundamental changes to the legacy system. One of the key tenets of web services development is true de-coupling of data from the application. Many 'web based' applications being developed today are not truly web services or SOA compliant because the application and data are still tightly coupled. This is the case with most legacy or stovepipe systems within the DOD and Intelligence Community.

d. Robust Scalable Hardware/Software Infrastructure

From a hardware standpoint, the SDSP has been built and tested across a wide range of platforms. LVI has tested against various hardware combinations ranging from a single commodity PC to multi-node Oracle RAC systems. For our largest DOD customer, the production SDSP is a clustered computing solution which uses a ten node Oracle RAC solution. In addition to the hardware design, LVI has also built the SDSP on Solaris 10, RedHat Linux, SUSE Linux and Windows 2003 Server operating systems.

Our experience has shown that combining various hardware and software components can produce a more scalable and flexible architecture for hosting the SDSP. In particular, the Oracle RAC design on Linux servers has some

significant built in fault tolerance and redundancy capabilities. It also offers some unique scalability options, especially for legacy database transitions. Additionally, the use of clustered MySQL databases provides a compelling solution for those seeking to reduce software licensing costs and extend the Open Source model. In both cases, the compute clusters can be created with various sized Intel/AMD processor based servers, or by using larger systems like those provided by Sun, IBM and specialty manufacturers. As more processing capability is required the underlying server hardware can be adjusted to expand the cluster capacity. Servers can in some cases be added or removed without impacting the operation of the overall platform.

The robustness of the SDSP when running on a clustered computing environment creates another opportunity for cost savings when transitioning into a data center environment. The SDSP can be used as a system consolidation platform. The production infrastructure is robust enough to support multiple databases so it is possible to co-host multiple systems on the hardware infrastructure thus reducing database licensing costs and overall hardware footprint in the data center. Scaling the system by altering the underlying hardware infrastructure can be addressed by extending the current baseline, or by replacing hardware components as the technology advances.

One of the biggest knocks against Oracle RAC is the inter-process communication impact that occurs with increasing number of nodes. In some cases this can be addressed by altering the hardware infrastructure. The SDSP architecture allows both for changes in hardware and the addition of multiple database instances (some RAC, some not) to support the overall client need. This same logic extends to the software components that support the web services tier. New capabilities can be added by using the current software set, or by adding new software products to the overall mix. The integration is simplified by maintaining a clear delineation of duties between the web and data services tiers and using standards based COTS products to build each tier.

LVI is involved with several R&D efforts aimed at exploring and extending the hardware and network communications bounds related to large scale, distributed database solutions. Working with government and industry partners on these leading edge efforts allows LVI to plan for these technologies now, thereby ensuring that the SDSP architecture can more easily accept technology improvements for hardware, software and communications as they enter main stream use.

5. Detailed System Description

The following sections describe the key features and capabilities of each component of the SDSP architecture.

a. Baseline System Build (SDSP-Core)

The SDSP-Core is the starting point for the reference architecture. It provides base level capabilities to support ingesting structured data sources and allows for subscription access to the processed data. Beginning with the baseline system and continuing throughout all successive enhancements, the SDSP services model is expanded and evolved to keep pace with client needs and technology advances. **Figure 1** provides a high level view of the SDSP-Core architecture.

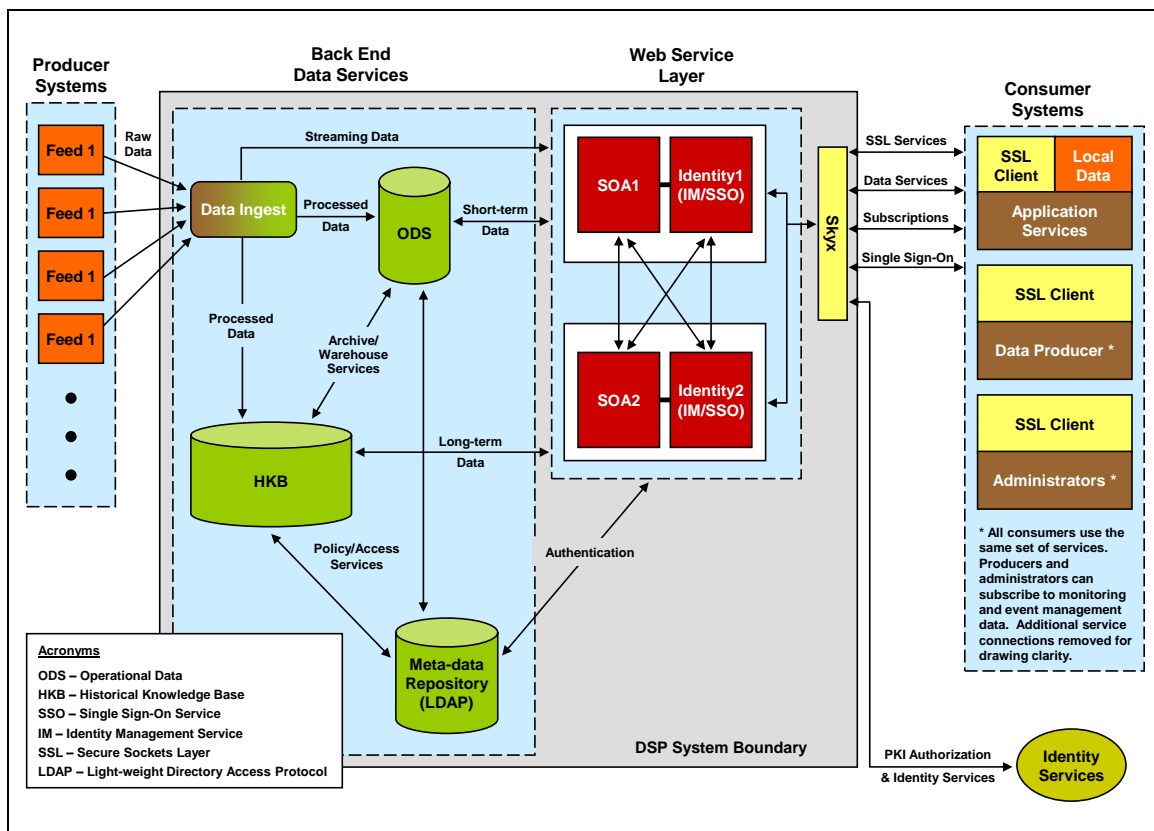


Figure 1: SDSP-Core High Level Architecture

Capabilities included in the SDSP-Core:

- generic data ingest and processing methods
- structured data storage
- security labeling of database content (Controlled Access Program Coordinating Office (CAPCO) compliant) – as supported by the database technology

- subscription services using Secure Socket Layer (SSL) and Public Key Infrastructure (PKI) technologies to secure individual sessions and authenticate data consumers
- extensive metrics and audit collection
- single sign-on (within the SDSP realm)
- service based XML document delivery

b. Cross Domain Enhancement (SDSP-CDE)

The cross domain enhancements extend upon the SDSP-Core by adding a PL4 database structure to support sharing structured data across multiple security domains. This particular solution is designed to support our DOD and Intelligence Community customers, but is applicable to any organization that is looking to bridge multiple network domains that require different levels of access and data control. Additional services and data feeds are added to the SDSP-CDE to support the unique demands of cross domain aware applications. Unlike the SDSP-Core which established a broad set of basic capabilities, the SDSP-CDE focuses on solving specific aspects of the cross domain data transfer problem. Since the system was originally developed to support our Department of Defense (DoD) customers, the primary focus is

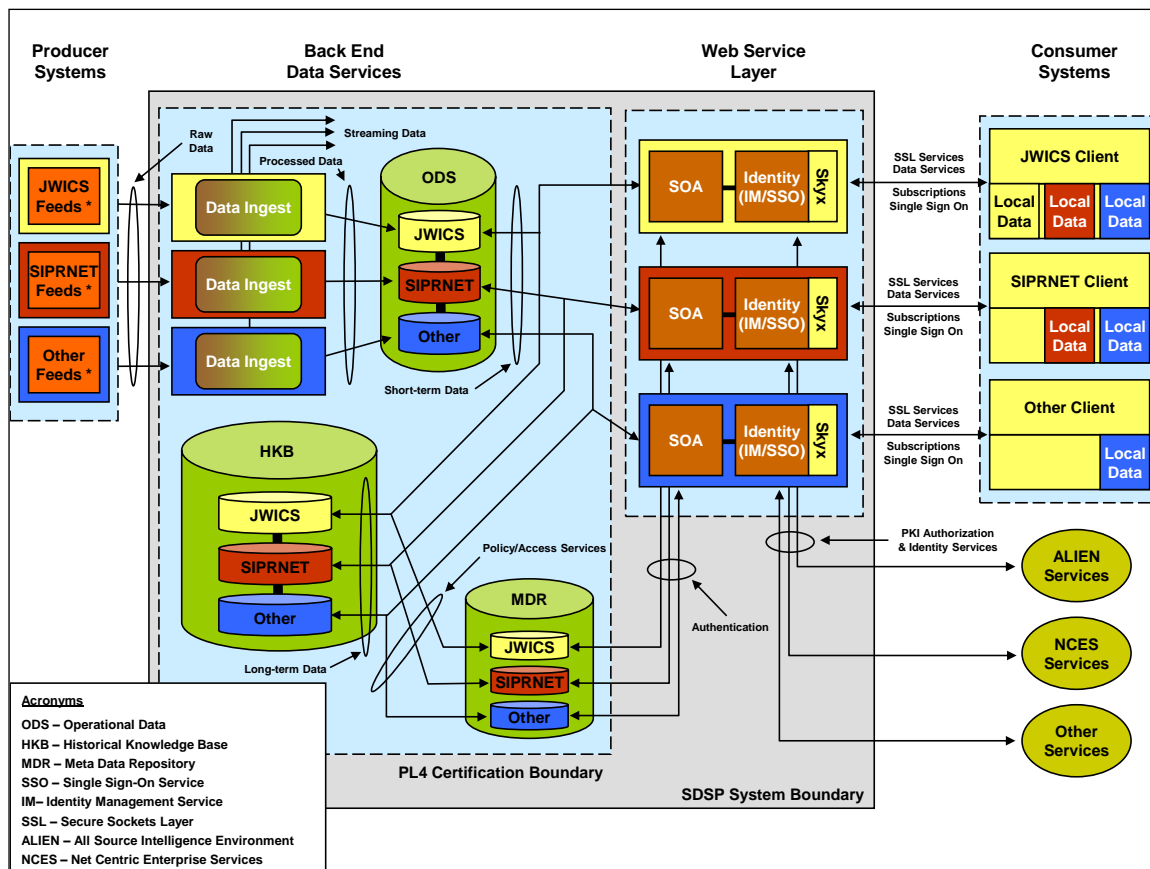


Figure 2: SDSP-CDE Enhancement High Level Architecture

cross domain support for secure networks (i.e. NIRP ↔ SIPR). However, the system is also capable of supporting separate organizational domains within a single security domain, each having distinct boundaries that might otherwise impede data sharing. For example the separation of law enforcement data and DOD data within the FBI's Secret network environment. **Figure 2** provides a high level view of the SDSP-CDE architecture.

c. Application Services Interface (SDSP-ASI)

The SDSP-ASI adds services and data storage support for applications to use the SDSP as a repository of record for reference data, meta-data and identity management data supporting individual applications. This includes a schema to store data created by the applications, eliminating the need for heavy weight applications that require embedded database services. In a truly service based architecture the database and data storage mechanisms are part of the core infrastructure of the organization and should not be considered as components of the application space. This requires that applications divest themselves of data retention functions and rely on the data services provided by the core infrastructure.

The SDSP-ASI provides a series of web services that enable the applications to take advantage of the data services tier without requiring direct database connections which would introduce brittleness to the architecture. Applications then become collections of services, which allow them to exist at multiple points on different domains using the same structure and code base.

Other features added with the SDSP-API include:

- Support for Business Process Execution Language (BPEL) – *allows for modeling interactions between business processes*
- LDAP services – *allows for storing and associating application specific attributes to embedded PKI attributes*
- Auto-generated formatting of data – *current includes KML/KMZ to support Google Earth*
- Enterprise Java Beans (EJB) and the Java Persistence API (JPA) – *allows for persistent storage in the web services layer through*

d. Extended Data Storage Layer (SDSP-EDL)

The SDSP-EDL adds support for new data storage methodologies including unstructured data and semantic data stores. Much of the work for this enhancement includes merging in other systems and solutions that are focused on solving some of these issues. Since the SDSP architecture is based on commercial standards, it leaves open the inclusion of tools built and deployed under legacy/heritage programs or through other “nextgen” acquisition efforts. By merging these capabilities, this generation completes the SDSP

support for structured, unstructured and semi-structured data sources. It includes new services to address search and discovery across the multiple data sources and fully supports advanced collaboration efforts for rigorous analytic data models. While the SDSP-Core laid the ground work by making services the principle component for exposing functionality, the SDSP-EDL adds the final pieces that drive semantic modeling and use of the data and services across all the systems and tools comprising the enterprise.

6. Capabilities Review

A recent Defense Request for Information (RFI) defined 19 key capability requirements which the agency was particularly interested in seeing in solutions. We felt these capabilities were an excellent representation of community wide requirements. We further divided these into six core areas and added an additional set of data capabilities. **Table 1** shows the relationship between these capabilities and the current development path for the SDSP. The concepts are grouped by capabilities and a check mark (✓) indicates which component of the SDSP begins the support for each concept. After their introduction, the concepts gain further refinement through successive iterations and enhancements to each component.

Table 1: Concept Development Paths

| ✓ = Starting point for concept | SDSP Components | | | |
|--|-----------------|-----|-----|-----|
| | Core | CDE | ASI | EDL |
| Transport Capabilities | | | | |
| High volume, low latency delivery | ✓ | | | |
| Varied size data elements from bytes to gigabytes | ✓ | | | |
| Intelligent routing and dissemination | | ✓ | | |
| Seamless dynamic routing of data across logical protection boundaries | | ✓ | | |
| Control Capabilities | | | | |
| Orchestration of routing and workflow | | | ✓ | |
| Extensible data management and transformation | ✓ | | | |
| Centralized management of distributed integration configurations | ✓ | | | |
| Policy-driven, access-based confidentiality solutions | | ✓ | | |
| Efficient management tools to reduce complexity and costs of training and Operations and Maintenance (O&M) | ✓ | | | |

| ✓ = Starting point for concept | SDSP Components | | | |
|---|-----------------|-----|-----|-----|
| | Core | CDE | ASI | EDL |
| Role-based, hierarchical management and resource monitoring of federated components | | | ✓ | |
| Distributed configuration and caching of deployment resources and routing rules (i.e. no single point of failure) | ✓ | | | |
| Availability Capabilities | | | | |
| Assured access from any point | ✓ | | | |
| Diverse client connectivity and support for multiple protocols | ✓ | | | |
| Highly distributed and scalable at both the service and application layer | ✓ | | | |
| Event and access driven service invocations | ✓ | | | |
| Adaptable to a wide range of data formats and protocols | ✓ | | | |
| Availability Capabilities (continued) | | | | |
| Scriptable and declarative environments based on configuration rules, instead of compiling behavior into code | | ✓ | | |
| Security Capabilities | | | | |
| Enterprise security and access control models compatible with PKI | ✓ | | | |
| End-to-end data auditing and traceability across a diverse infrastructure | ✓ | | | |
| Data Capabilities | | | | |
| Unstructured data sources | | | | ✓ |
| Support for semantic web and analytical database technologies | | | | ✓ |

7. Framework Analysis

a. Pros/Cons

The SDSP reference architecture provides a solution for the core capabilities as presented in this paper. The architecture was not developed specifically for a single organization; nonetheless, some adjustments may be required to match the unique nature and mission of each organization seeking to deploy the SDSP. One instance of the SDSP-Core is nearing

accreditation within the DOD environment, with the SDSP-CDE scheduled to begin accreditation in FY09. The baseline system provides the key capabilities required to start using the system and the planned enhancements can be easily expanded to support other unique requirements. The deployment plan (functions, schedule, cost) for each component has already been developed and design changes addressing specific client requirements can be readily added. Portions of the ASI component are already being tested and prepped for deployment. The bulk of the EDL component is through the preliminary design stage and product updates and capabilities are being identified and surveyed for best fit.

b. Expected Efficiencies

Beyond the expected efficiencies of deploying a service based environment, LVI can offer efficiencies in time and cost based on the level of work already completed for the SDSP architecture. The current design represents the culmination of nearly four years of research and development at a cost approaching \$10M in government and private R&D funding.

c. Programmatic/Technical Challenges

The primary challenge is modifying the current baseline to address the unique requirements of a specific organization. The choice of hardware and software components for the DOD build may not be suitable for other environments. In which case, new products may require investigation and testing before introduction into the architecture. One obvious example would be replacing the Oracle components of the current DOD specific implementation with another database product. While it doesn't require a complete overhaul of the architecture, it would require time to test and validate correct operation. Furthermore, while the SDSP-Core meets many of the capabilities expressed throughout the government and commercial communities, the manner in which they are met and presented may not be sufficient for a specific organization. The policy, processes or unique mission needs of each organization will ultimately drive the necessary changes to allow a successful operational deployment of the SDSP. LVI is confident that the SDSP architecture has the flexibility and robustness to meet these challenges and fully address each issue.

Within the Defense industry, LVI engineers have worked with national level organizations (NRO, NSA, NGA, ODNI) to access data feeds provided by various legacy programs. LVI understands their missions and is already a consumer of many of their products. Our partnership largely has shaped the design of the SDSP, especially in terms of the ingest metrics and subscription services. Building the SDSP-Core has given LVI a better understanding of the technical and programmatic challenges associated with expanding the SDSP architecture to meet new requirements and include the new components. Interestingly, information sharing

initiatives of the other IC organizations show the same set of traits -- successes and issues-- already investigated by LVI in building the SDSP.

8. Conclusion

The Secure Data Services Platform developed by LVI provides a practical, sustainable solution to fundamental data sharing requirements. It represents a careful study of information sharing issues and successes. While many of its features and functions are closely aligned with requirements openly discussed by the security-focused Defense and Intelligence Community, they are also applicable to other industry segments where the integrity, pedigree and security of data on the move is important. One can easily envision the use of the SDSP in a medical or legal environment where careful control of personal data that is governed by specific legal requirements (i.e. Sarbanes-Oxley, Health Insurance Portability and Accountability Act (HIPAA)). The SDSP offers an infrastructure starting point that embodies the concepts required to meet the critical data security and accountability requirements driving today's information sharing needs. The SDSP is already moving through certification toward a broader deployment in the Defense community, and its capabilities are available today for demonstration at the LVI facilities. We recommend the SDSP to any commercial or government organization looking for a practical means to start their transformation toward ubiquitous data sharing.

For more information contact us at:

info@lvi-usa.com

List of Acronyms

| | |
|--------|--|
| CAPCO | Controlled Access Program Coordinating Office |
| CDSS | Cross Domain Security Solution |
| CMMA | Collection Management Mission Applications |
| COOP | Continuity of Operations Planning |
| COTS | Commercial-Off-The-Shelf |
| CY | Calendar Year |
| DDMS | Department of Defense Discovery Metadata Specification |
| DIA | Defense Intelligence Agency |
| DISA | Defense Information Systems Agency |
| DOD | Department of Defense |
| DOD/IC | Department of Defense/Intel Community |
| HIPAA | Health Insurance Portability and Accountability Act |
| IC | Intelligence Community |
| IC-ISM | Intelligence Community Information Security Marking |
| MOM | Manager of Mangers |
| NGA | National Geospatial-Intelligence Agency |
| NRO | National Reconnaissance Office |
| NSA | National Security Agency |
| ODNI | Office, Director of Naval Intelligence |
| O&M | Operations & Maintenance |
| OWB | Oracle Warehouse Builder |
| PKI | Public Key |
| R&D | Research and Development |
| RAC | Real Application Clusters |
| RFI | Request for Information |
| REST | Representational State Transfer |
| SDSP | Secure Data Services Platform |
| SOA | Service Oriented Architecture |
| SOAP | Simple Object Access Protocol |
| SSL | Secure Socket Layer |
| USDI | Under Secretary of Defense for Intelligence |
| WAN | Wide Area Network |
| XML | eXtensible Markup Language |