


The SOA Magazine

Feature Article



SOA in Government: A Law Enforcement Use Case

by Girish Juneja, Blake Dournaee, Joe Natoli and Steve Birkel

Published: February 9, 2008 (SOA Magazine Issue XV: February 2008)

[Download this article as a PDF document.](#)

Abstract: IT culture within the public sector has long been known to be unique. The responsibilities of managing a wide range of often critical public services establish a distinct set of priorities that can't be compromised, especially when it comes to a reliance on technology. SOA adoption has grown substantially in government agencies in federal, state, and local sectors. There is an increasing realization that the strategic benefits of service-orientation can help overcome many of the traditional cost and efficiency-related IT problems.

This article (comprised of excerpts from the book "Service Oriented Architecture Demystified" [REF-1]) introduces the recently developed Federal Enterprise Architecture reference models and explores how they have been applied in support of a major SOA initiative within a law enforcement agency. Topics covered range from service identification, service taxonomy, security issues, and common SOA infrastructure components required for government IT enterprises, including a specific enterprise service bus configuration termed as the "hyper-ESB". Throughout this study, the importance of standardization and canonical data standards are constantly highlighted as a key success factor.

Introduction: From Silos to FEA

*"An elephant: A mouse built to government specifications."
- Lazarus Long, from Robert Heinlein's "Time Enough for Love"*

Government institutions across the world, at national, regional, and local levels, are significant consumers of technology. Governmental services affect us all. They can range from areas of defense and national security, to health, taxation, law enforcement, judiciary, environment, energy, social services, disaster management, and land use management.

Technology clearly needs to play an important role for any organization responsible for any one of these areas. In the public sector, automated systems perform a common, fundamental function: getting information to and from the "users." In other words, systems need to be in place to effectively share data between agencies and the public community comprised of citizens and businesses.

However, as with large private corporations, information technology in government institutions has been built in silos, where the interoperability and exchangeability of information is only an afterthought, leaving service-orientation concerns by the wayside. For example, government agencies tend to structure information silos in terms of agency-specific objectives. The more agency-specific silos created, the greater the integration problems that will have to be overcome down the road, when agencies need to communicate with each other (or between agency user communities).

We can draw a parallel between the big bus and little bus metaphor, where communication with different agencies falls into the big bus category and communication with citizens and businesses falls into the little bus category. Admittedly, integrating the information silos of various governmental institutions is a tall order for reasons that have little to do with

technology.

There are some encouraging signs that this is changing. In the United States, for example, the Office of Management and Budgets (OMB), the Office of E-Government (E-Gov), and Information Technology (IT) have established the Federal Enterprise Architecture (FEA) program, an initiative that is producing reference models in support of realizing SOA in government environments. These models essentially equip government institutions with frameworks and tools that can enhance collaboration and help analyze investments in order to reduce redundancies, encourage reuse, and improve performance and service quality.

Here is the official definition of the FEA, as published on the E-Gov portal [REF-2]:

The FEA is being constructed through a collection of interrelated "reference models" designed to facilitate cross-agency analysis and the identification of duplicative investments, gaps, and opportunities for collaboration within and across Federal Agencies. Figure 1 displays the five FEA reference models. The next few sections focus on the Service Component Reference Model (SRM) and its associated Data Reference Model (DRM). The SRM provides a framework for services to be derived from business objectives, while the DRM deals with the issue of canonical data representation, context, and exchange (one of the most important problems in realizing SOA in government).

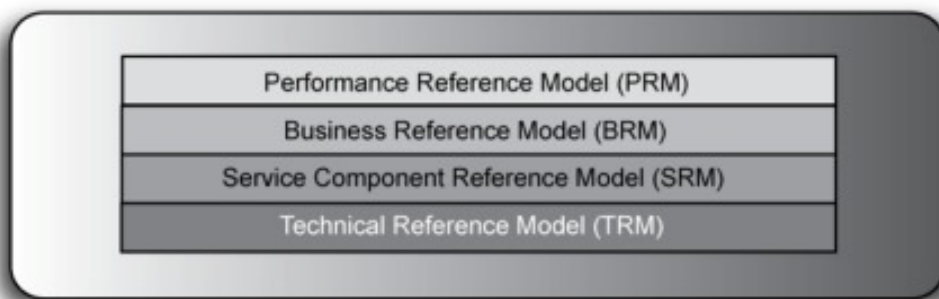


Figure 1: The four fundamental FEA reference models.

Service Component Reference Model (SRM)

The FEA Reference Model defines SRM as follows:

The SRM is a business-driven, functional framework classifying Service Components according to how they support business and performance objectives. It serves to identify and classify horizontal and vertical Service Components supporting federal agencies and their IT investments and assets. The model aids in recommending service capabilities to support the reuse of business components and services across the federal government. The SRM is organized across horizontal service areas, independent of the business functions, providing a leverage-able foundation for reuse of applications, application capabilities, components, and business services.

The SRM's objectives appear to be in-line with the common goals of an SOA business architecture in that they attempt to map business process requirements to SOA capabilities, and so can be thought of as the blueprint that decomposes and translates business requirements into pieces that can be re-assembled as services.

Using the SRM, you can perform complete assessments of current and future states to clarify what functions of the applications and IT infrastructure can and should be decomposed into reusable services. You can also determine where the standardization of functionality is required and also when customization is really needed.

Data Reference Model (DRM)

In a typical government environment, data, its representation, its context, and the means of its exchange all represent critical foundations for an SOA implementation. However, when introducing SOA into an already silo-based environment, the underlying (also silo-based) data architecture can be incompatible and may therefore require a new model. The DRM attempts to provide guidance on how to establish a suitable data model to support of the goals of SOA.

Here's the official definition for the DRM:

It is the FEA mechanism for identifying what data the federal government has and how that data can be shared in response to business/mission requirements. The DRM provides a frame of reference to:

- Facilitate Communities of Interest (which may be aligned with the LoBs (Lines of businesses) delineated in the FEA Business Reference Model) in establishing common language.
- Enable needed conversations to reach credible cross-agency agreements around: governance, data architecture and an information sharing architecture.

The DRM provides guidance to enterprise architects and data architects for implementing repeatable processes to enable data sharing... [REF-3]

The DRM specification goes on to define three key concepts that require standardization by the specific government community of interest (COI):

- The *data context*, to be agreed on by a government COI, is the purpose for which the data assets are being defined; what subject area should the data address, who needs it, who maintains it, how it is accessed, stored and secured (left for later versions of DRM) and how it ties into the FEA Business Reference Model (BRM).
- *Data description* is the specific structure and semantics of data required based on the data context.
- *Data sharing* is defined as the set of services that are standardized for data sharing for the data context area with the specific data description.

While data context, data description, and data sharing are abstract terms, in most if not all applications of the DRM, the data description is provided as an XML schema, the data context is defined by XML namespaces, and data sharing is expressed via XML-based request-response exchange patterns that can be used within a Web services framework.

In the upcoming SOA law enforcement use case we will discuss how concepts outlined in the SRM and DRM models come together to address a real world problem.

SOA in Law Enforcement: A Use Case

Law enforcement agencies are focused on preventing crime and providing reliable internal security. This sets the stage for an effective proving ground for SOA and the FEA models. The following use case documents how these models can be leveraged in support of establishing SOA.

In the United States, many federal, state, and local agencies are engaged in crime prevention and internal security. At the federal level this starts with the Department of Transportation and the Department of Homeland Security (within which can be found the US Immigration and Customs Enforcement and US Customs and Border Protection branches). Extending to state and local agencies, it continues with the state departments of corrections, local and state law enforcement agencies, courts, and any number of additional, related agencies.

In most law enforcement agencies, the personnel use different networks and systems to search for information on suspects, each with their own authentication and search mechanisms, and each with their own result data formats.

Searching through these independent information silos is inefficient. There are missed linkages across the silos, and ultimately, poorer law enforcement due to a low quality data architecture. For example, if an Illinois state police officer requests a criminal history record from Iowa, the search might have to be carried out through specific terminals that accept field-value pair-based queries, while the results coming back may be in an Iowa law enforcement-specific data format.

In this example the system used to provide the message switching capability is NLETS, or the National Law Enforcement Telecommunications System. NLETS provides a message switching solution that links together state, local, and federal law enforcement and justice agencies for the purpose of information exchange. This service, which does not host any information itself, is based on a frame relay. It is used widely by all of the states and federal agencies related to justice. The data exchanged includes driver records, state criminal records, license information, immigration records, AMBER alerts (an alert system designed to help locate and protect children), hazardous material (Hazmat) warnings, weather bulletins, terrorist alerts, and so on. NLETS records over 40 million transactions every month.

Figure 2 illustrates the information flow for someone from a state law enforcement agency performing a query similar to the one described previously. This query is being issued against the NLETS service, the National Crime Information

Center database (NCIC), the state department of motor vehicle records, and other data sources.

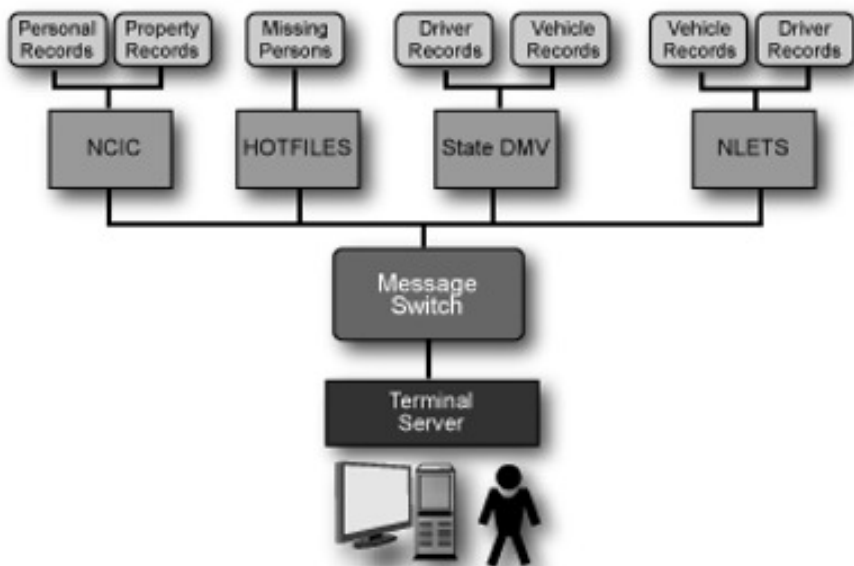


Figure 2: A typical state justice environment for accessing law enforcement information.

Authorized law enforcement personnel use the state law enforcement portal interface to request information based on a driver's license. The terminal server receiving the request then forwards it to the state message switch over TCP/IP. This switch acknowledges the message, and then, based on the query parameters, sends it to state systems.

Examples of these systems are the state department of motor vehicles, NLETS, NCIC (with access to gang and violent activity records, stolen gun, license plate, and boat records, and so on), and state hot files (for missing people). The responses from these systems are then sent back to the terminal server, which then converts the data into a normalized format to whatever extent possible for the law enforcement agent on the terminal.

However, to provide a uniform interface for search and results in the above example, a point-to-point mapping at the terminal server was deemed the only viable option.

In order to create a canonical data format that the state departments of justice could use (thereby reducing the n times $(n-1)$ mapping problem to $2n$, as shown in Figure 3), the U.S. Department of Justice's Office of Justice Programs (OJP) and the Global Justice Information Sharing Initiative (Global) came up with the Global Justice XML Data Model (GJXDM) [REF-4]. The GJXDM includes a data model, a data dictionary, and an XML schema.

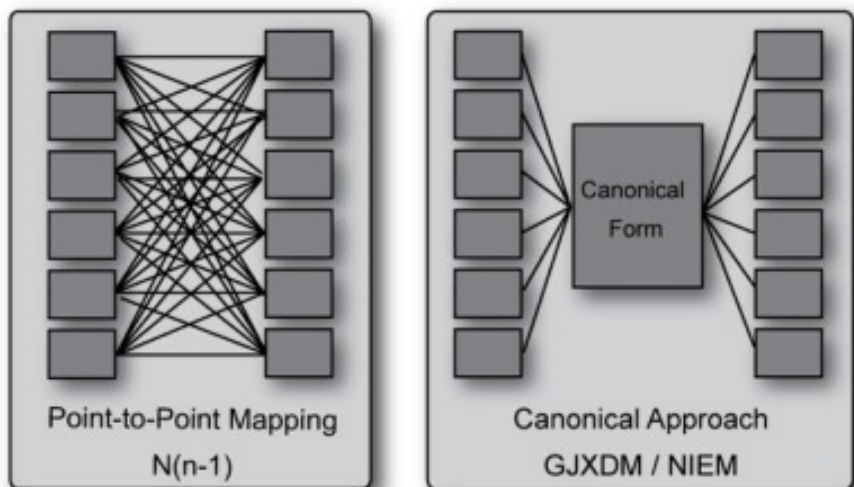


Figure 3: Point-to-point mapping versus the canonical approach used for justice data exchange.

With NLETS adopting GJXDM and offering a Web service-based interface for querying justice data across states, the point-to-point mapping problem is reduced to a canonical data exchange. The criminal justice information captured through this model includes rap sheets, court case records and juvenile files.

With GJXDM as the canonical XML data format, the data flow can be simplified in that the terminal server just converts to and from GJXDM. The one downside is that GJXDM exists as a fairly complex set of XML schemas and processing actual workflows based on GJXDM can be computationally intensive (unless XML processing is accelerated in some manner).

Although this solves the data standardization issue, we do not have any reusable service at this point. If an authorized homeland defense application needs access to state justice data there is no service visibility or interaction model available. Conforming to such a model would allow more direct access to the state department of motor vehicles data. Using GJXDM, the homeland defense agency must have a terminal server deployed. This translates into costs both in terms of time and real investments.

Several other systems are also available with complementary and overlapping purposes. A few examples include:

- HSIN (Homeland Security Information Network), which makes real-time threat information available to law enforcement and first-responders through a Web-based system.
- The Regional Information Sharing System (RISS), which links law enforcement agencies throughout the nation for multi-jurisdictional crime fighting and anti-terrorism. RISS is a shared intranet service.
- FBI's Law Enforcement Online (LEO), another secure intranet service providing Web-based access to law enforcement-related information.

All of these are examples of systems with the same core problem: they are conceived for a single agency purpose and built as a silo. Each was delivered independently, with different authentication schemes and usage models.

If we think in terms of the ultimate goals of justice data sharing, there is a clear benefit in introducing SOA with the support of canonical data standards. Left as they are, these types of silos will continue to suffer from security lapses due to disparate authentication schemes, duplicate data entry, maintenance problems, and general inefficiency from the user's perspective.

Common Vocabulary: Data Standards

Extending data interoperability for public safety throughout government bodies is a pressing need: while GJXDM is a commendable effort, the need for wider data interoperability across public safety related government agencies is immediate.

The National Information Exchange Model (NIEM) is a partnership of the U.S. Department of Justice and the Department of Homeland Security designed to develop, disseminate, and support enterprise-wide information exchange standards and processes. Their purpose is to enable jurisdictions to effectively share critical information in emergency situations, as well as support the day-to-day operations of agencies throughout the nation.

In the NIEM specifications, GJXDM is one of multiple data contexts, each of which has its own namespace with a common data description and data sharing guidelines. A NIEM-based data exchange allows for services across the Department of Justice, public safety, health and Homeland Security to exchange data in a canonical form.

Common Service Interaction Profiles

The universal communication network established by SOA is based on accepted protocols and standards, but it requires more than an interoperable data format to be successful. At a state agency level, this might translate into the need to agree to either a common transport (such as HTTP, HTTPS, MQ Series, etc.), or deploy systems that can interoperate between different transports and different interaction models. It is critical, therefore, to look at the patterns supported by the enterprise service bus (ESB) which can be deployed in support of these environments. Such patterns should also be supported by the service composition and business process execution environments.

Services: SOAP and (not or) REST

Department of Justice-related information, such as Amber Alerts, needs to be shared widely with the public. The service concept therefore does not need to be tied down to a specific technology artifact. A service can be made available via SOAP or REST, based on its requirements and its supporting infrastructure components.

Ideally, its infrastructure should not be restrictive as to how the services are made available. It is quite likely that

portions of the Department of Justice applications will become service-oriented from the ground up, while others will start off as service facades. For example, the FBI's LEO service could be made available via RSS feeds based on XML without re-engineering the entire application. These feeds could then integrate into the service bus at a state department of justice, which, in turn, should be able to provide interoperability between SOAP and REST-based services.

Security on the Network

Provincial and national governments are cautious adopters of technology, especially in the areas of public safety, for obvious reasons. Security is a top concern and to ensure at least a minimal level of network security, public safety environments are kept on encrypted private networks. The good news is that within state agency IT environments, the SOA security requirements, while stringent, are addressable with technology available today:

- *SSL Authentication* – Verifies the domain associated with the specific agency by parsing and validating the session-level certificate.
- *Canonical Data Validation* – Enforces the GJXDM/NIEM-specific content model, with the ability to support disparate schemas across agencies.
- *XML Content Attack Prevention* - Protects against XML context attacks, including semantic threats, denial of service threats, and parser exploits.
- *Authentication and Authorization* – Authorizes the requesting person (not the system) for the specific operations using the credentials presented. This is usually performed by matching the user and group information in an LDAP-accessible X.500 directory and ensures that attributes match a specific operation.
- *Security Conformance* – Verify the WS-I Basic Security Profile conformant WS-Security header on the SOAP message containing the OTA payload.

Service Taxonomy

An example of how tightly coupled information can be exposed as services in law enforcement is illustrated in Figure 4, which is one representation of the types of business services that can be found in a state department of justice. These and similar services have broader usage in state government agencies and can be orchestrated as part of a workflow in a "big bus" environment.

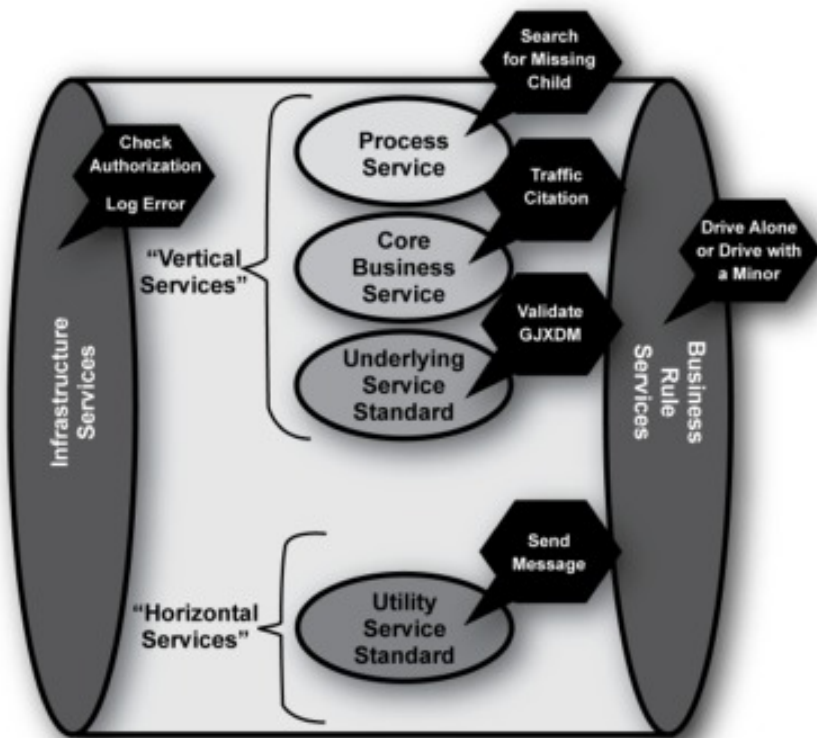


Figure 4: How various services, both business and utility-centric, can be exposed specifically for law enforcement.

Here are some examples of core, entity-centric business services that collectively establish a service taxonomy for law enforcement:

- Citing Officer (organization & affiliation, relationships, plus roles/groups)
- Subject/Driver (summary or full person record, driver license information)
- Appearance (court case details)
- Citation (issued to records, issued from records)
- Incident (road conditions, intoxication, speed, charge, etc.)
- Vehicle (identification, registration)
- Document (pictures)
- Location

In order to create a set of reusable services that can be composed at the business process level, the IT environment has to build a layered stack. The lowest part of the stack performs horizontal, reusable technical functions such as content routing, messaging, and security.

The services higher up in the stack represent the key elements of the taxonomy because they encapsulate the core business logic. These services are coarse-grained and provide highly shared data relevant to most law enforcement and public safety business processes. Examples of the shared data include lists of law enforcement officers, subjects, agencies, vehicles, as well as key transaction data types such as citations, incidents, appearance, and so on.

With these (and other) business services, applications can be composed to deliver key services such as searching for missing person, or searching for a match on a stolen vehicle in a manner that is easily changeable and extendable. Further, as these services need to be exposed externally to the state justice agencies, they can be composed to conform to data exchange patterns required to interoperate with external agencies using GJXDM or NIEM-based canonical data exchange.

Service Infrastructure

Infrastructure components required specifically for an SOA deployment in support of this law enforcement use case

include the following.

Service Containers

In law enforcement agencies IT adoption by necessity is cautious in nature. This means that a tremendous amount of data remains in legacy formats. In order to create a flexible service-oriented environment without massive reengineering, the service container solutions need to be able to provide service façades that wrap legacy systems.

Enterprise Service Bus

Once the core business services are defined and can be made available by using appropriate service containers, a service bus environment is required to enable the composition of services and define business process workflows associated with the composite services that are based on business use cases. Reliability and auditing are key requirements for executing business process workflows in law enforcement. So, utility services that provide messaging and security functions need to be carefully architected in a reusable manner.

When assessing an ESB, it is critical that its features support the requirements of law enforcement and government agencies in general. Here is some common criteria:

- The ESB must have high-performance support for local and remote service composition and orchestration environments.
- The ESB should support XML workflows, SOAP, or REST.
- The ESB should offer core utility services such as transformation, sophisticated content validation, security, content routing, Web service firewalls, and non-XML adaptation.
- The ESB should not be tied down to a vendor-specific messaging implementation.
- The ESB should be vertically and horizontally scalable and, to whatever extent possible, centrally managed from a single console.

We'll refer to an ESB with the above features as a "Hyper-ESB."

Figure 5 illustrates an SOA implementation in law enforcement. Comparing Figures 2 and 5, the law enforcement personnel's requests are now accommodated by a service bus that invokes a relevant business process that then composes underlying business services in order to execute the process workflow. These business services use horizontal utility services that are available transparently via the service bus.

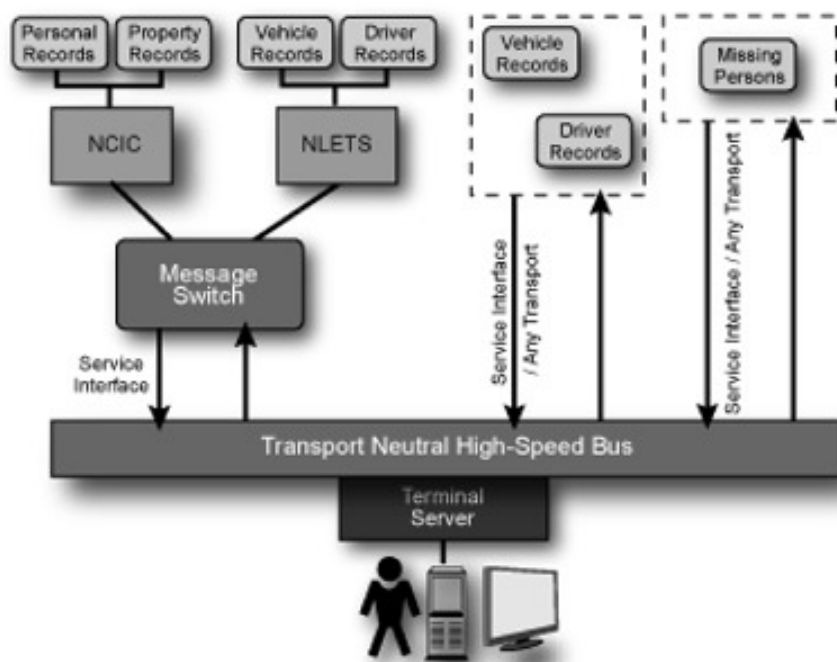


Figure 5: The SOA Environment in the State Justice.

There are today still barriers to full SOA realization in law enforcement environments that require reliability in message delivery with the flexibility of SOA. Traditional point-to-point, named queue based, proprietary messaging mechanisms do not allow for intermediaries that can dynamically bind to arbitrary queues based on request and response content. And SOAP-based reliable messaging (WS-RM) protocols are not yet mature enough to offer true interoperability.

So the pragmatic solution is to encapsulate and contain messaging services so that when better solutions are available, the underlying queuing mechanisms can be replaced without major impacts to system availability.

External Gateway

An enterprise service bus with service containers can make the previously described SOA implementation attainable. However, there is one more requirement that needs to be fulfilled. As we discussed earlier, law enforcement business environments have changed. Better information access is needed to and from the state law enforcement agencies, closer interoperability is needed with other public safety agencies, and the ability to disseminate information to the public is needed for emergency purposes.

With these expanded requirements, creating point-to-point interfaces to each segment of an external user base is not viable. What is missing is a secure, reliable means of exchanging information with external agencies. This must allow for network-based isolation and SOA-friendly security with content awareness to protect and then send/receive messages to appropriate service consumers.

The exchange patterns with external agencies are going to be based on canonical data (as per GJXDM/NIEM). However, because different agencies are at different stages of XML adoption for data interchange, some of the key authentication and request information is embedded as part of payload itself and not in the SOAP wrapper. Gateways are therefore needed to execute process flows that include what we identified earlier as security requirements, with the additional option to be able to take action based on the specifics of the payload.

Note: In the case that a Hyper-ESB type solution is adopted, the ESB and the external gateway can be collapsed into one solution environment, thereby reducing the total cost of the SOA migration.

Conclusion

While this article has focused on one specific use case, the reality is that state law enforcement agencies within the United States are at various stages of SOA adoption. While the cautious nature of governments may result in slower transitions toward SOA and related technologies, our experience has been that government agencies in the United States and across the world that have taken the lead with SOA have seen immediate benefits in reusability and reduction of redundancy in software and hardware.

References

- [REF-1] "Service Oriented Architecture Demystified" by Girish Juneja, Blake Dournaee, Joe Natoli, and Steve Birkel, Intel Press, http://www.intel.com/intelpress/sum_soa.htm
- [REF-2] E-Gov, "FEA Reference Models", <http://www.whitehouse.gov/omb/egov/a-2-EAModelsNEW2.html>
- [REF-3] The Federal Enterprise Architecture Program, "The Data Reference Model Version 2.0", http://www.whitehouse.gov/omb/egov/documents/DRM_2_0_Final.pdf
- [REF-4] US Department of Justice - Information Technology Initiatives, "Global Justice XML Data Model (Global JXDM)", http://www.it.ojp.gov/topic.jsp?topic_id=43

This article was excerpted from Service Oriented Architecture Demystified: A Pragmatic Approach to SOA for the IT Executive by Girish Juneja, Blake Dournaee, Joe Natoli, and Steve Birkel. Copyright © 2007 Intel Corporation. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4744. Requests to the Publisher for permission

should be addressed to the Publisher, Intel Press, Intel Corporation, 2111 NE 25 Avenue, JF3-330, Hillsboro, OR 97124-5961. E-Mail: intelpress@intel.com. Intel is a registered trademark of Intel Corporation. Other names and brands may be claimed as the property of others.

THE PRENTICE HALL SERVICE-ORIENTED COMPUTING SERIES FROM THOMAS ERL



[Home](#) [Past Issues](#) [Contributors](#) [What is SOA?](#) [SOA Glossary](#) [SOA School](#) [SOA Books](#) [About](#) [Legal](#)

Copyright © 2006-2008
SOA Systems Inc.