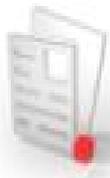




California Enterprise Architecture Program  
Office of the State CIO



# California

*Service-Oriented Architecture (SOA)  
And Federated Identity Management  
Vision*

August 27, 2007

**Preliminary Draft**

Lee Macklin, Acting Director  
California Enterprise Architecture Program  
Lee.Macklin@ceap.ca.gov

**Foreword by the State CIO .....3**

**Executive Summary .....4**

**Current Environment .....5**

**A Better Environment .....6**

**Governance Model .....7**

**Enterprise SOA Infrastructure .....10**

**Enterprise Identity Management.....12**

*Security and Authentication..... 13*

*The Security/Authentication Process ..... 14*

*Identity Management..... 14*

*Trust Model ..... 14*

**Enterprise SOA & IDM Roadmap .....16**

## Foreword by the State CIO

My Friends,

For three years now, we have been steadily working to establish a framework for enterprise-wide solutions and systems in California government. That framework includes major elements of our California State Information Technology Strategic Plan, the steps already taken to establish a statewide enterprise architecture, and the creation of IT governance bodies that encompass decision-making both from business and IT perspectives (i.e., the Enterprise Leadership Council, the Technology Services Board and the IT Council). The framework is largely in place (for more information on each of these topics, see [www.cio.ca.gov](http://www.cio.ca.gov)).

Now, we must engage ourselves in the even more arduous task of actually building enterprise systems within the context of that framework. This will not be easy work. We have already teed up a broad portfolio of projects that will propel us towards our enterprise goals. For example, we have major infrastructure modernization efforts now underway at Department of Justice, Department of Motor Vehicles, Employment Development Department, Department of Corrections and Rehabilitation, Department of Health Care Services and Department of Technology Services. We are pursuing enterprise-wide business management systems through the Fi\$Cal Project and in major ERP implementations underway at the State Controller's Office, Department of Corrections and Rehabilitation, Department of Transportation and the trial courts. We are building or rebuilding major case management systems all across government, including systems within Employment Development Department, Department of Motor Vehicles, Department of General Services, Department of Transportation, Department of Corrections and Rehabilitation, and the trial courts, as well as systems that deliver health and welfare benefits to millions of Californians.

But for these projects and systems to be truly "enterprise" in scope and operation -- as opposed to a series of siloed initiatives and efforts -- we must now commit ourselves to the cooperative development of enterprise standards that will meaningfully tie these systems together. The vehicle by which those standards will be developed is our enterprise architecture program. After prolonged consultation and deliberation, a consensus has developed that adopting a Service Oriented Architecture along with Federated Identity Management is the best course forward, the course most likely to lead to enterprise systems and capacity.

This document sets forth in the most basic terms possible our vision for that architecture and identity management approach. It also describes the nature of the issues that lie in the road before us and the type of decisions we will be making as we move down that road.

We seek your reactions, your support and your participation in this historic effort.

Clark Kelso  
Chief Information Officer  
State of California

## Executive Summary

It is time for an IT infrastructure in California that supports all the diverse lines of business from an *enterprise* perspective. It should be a standards-based environment that accommodates all stakeholders – different levels of government, as well as private industry partners. To build and maintain that environment over time, we will need to establish an inclusive governance model where all stakeholders can share their views and participate in creating standards.

There are two foundational components that make up this new IT infrastructure. First, an *Enterprise Service-Oriented Architecture (SOA)* infrastructure will host and manage new business services. Many of these business services will be implemented as shared web services. Second, an *Enterprise Identity Management (IDM)* system will manage all types of users in a consistent way. It will allow for various security policies to be applied as set by various security and privacy policy organizations.

The Roadmap section of this document lists the many steps that need to be accomplished to achieve success. It is anticipated that the roadmaps will be regularly updated resulting in many detailed projects. In order to accomplish the Roadmap, appropriate decision makers must provide strong policy on the project's sustainability, as well as identification of standards and processes and their enforcement. For example, to what degree will utilizing the Enterprise SOA and IDM infrastructures be mandatory? How will these Enterprise infrastructures be funded? Will standard language for SOA and IDM be enforced by the control agencies?

At some point, we may want to consider establishing an SOA Center of Excellence to manage and share policy, as well as serve as an exchange forum for all stakeholders. This approach is highly recommended by Gartner. In the meantime, several departments (DMV, EDD, DHCS, and OSI) are initial members of the newly formed SOA Governance Group and they have agreed to work together to get the enterprise SOA and IDM infrastructure in place. Each of these departments has new projects underway which will utilize this shared, services-based infrastructure. Additionally, the Identity Management Workgroup was formed several months ago to lead the effort in defining federated identity management details.

As we move forward into the new world of SOA and federated identity management, we need to pay particular attention to the sustaining financial model. We want to encourage stakeholders of all types to use this new enterprise environment. So, it is important that the financial model not be an impediment for providers and consumers of shared services.

## Current Environment

In the past, most business applications were built in a “silo” fashion. That is, they were built for a specific purpose and they were “hard-wired” to other systems which made them difficult and expensive to change. Some of these systems have been modified to allow them to exchange information in more common formats (XML). However, for the most part the interfaces are still very platform and language dependent. There are very few standards set.

Additionally, there is no consistency of how user identities are managed. Policies vary even within the same department. This means a user must have multiple accounts where rules and policies may be applied differently. Thus it is difficult to manage changes, and worse, inconsistencies could result.

This problem will be magnified as more users move to conducting business online, particularly via the web but also via automated call centers. The State Portal has been completely redesigned and now provides a much better user experience. However, the infrastructure behind the portal also needs a major upgrade to an SOA environment.

## A Better Environment

As we migrate from a silo environment to a business services based environment, we need a new infrastructure. The new business services will be implemented as web services – many will be shared services. For example, why not provide a single social security verification service and a single address verification service that all line of business applications can use instead of each building their own?

The new SOA environment provides opportunity to offer new business services that don't exist today. For example, the Department of Health Care Services has put into production a new set of shared web services that verify SSN and retrieve social security information, provide MEDS eligibility information, and provides a single view of medications history – integrating information from both government and private provider sources. All information is available real time, and the users access the information via standard web service interfaces (SOAP messages) and data is in XML format.

This does present some new challenges. First, we must have governance around the initial funding and project mechanisms for the SOA enterprise and IDM enterprise infrastructures, as well as a model for sustainability. Second, there must be governance around the management of the services. What is the process for deploying and certifying a service? What is the process for modifying a service? How will access to services be controlled? How do we deal with different security policies for different services? Or, for different classes of users?

It is obvious that we must find ways to simplify and standardize user access to web services. Due to the complexity of California government, a federated approach will most likely be required to satisfy all the stakeholders.

Additionally, we need to put a stake in the ground regarding interoperability standards. Which interface and protocol standards will we choose? Which standards will be set for how data is exchanged?

This is an even bigger challenge when one considers that it is not just state government. To achieve the biggest bang for the buck, different levels of government as well as partners in private industry must be included. So, a flexible standards model must be chosen.

The bottom line is we need to shift our thinking to an “Enterprise mentality” and leverage IT assets by reusable web services across lines of businesses and security domains.

This does not imply that we need to “pull the plug on the mainframe”, or move from COBOL to .NET (for example). It does require that we agree on standard interfaces and standard policies regardless of how the services are actually implemented.

## Governance Model

Since Enterprise SOA and Federated Identity Management is a highly shared environment, it is critical that a governance model be defined and implemented to manage an environment that is acceptable to all stakeholders. Standards must be set and policies determined for how services will be created, certified, modified, and retired as well as how security will be defined and enforced. Additionally, initial and sustaining funding models and the mechanisms to implement enterprise infrastructure must be agreed upon.

California is a very complex government primarily based on a peer relationship model. Therefore, the governance model consists of many groups each with a specific focus. Here is a brief description of some of the groups and how they relate:

The **Enterprise Leadership Council** (ELC) is a top-level council that was formed to address enterprise issues. This includes providing executive support for ERP consolidation, service oriented architecture (SOA), and federated identity management efforts underway. Details of the ELC can be viewed at <http://cio.ca.gov/stateIT/governance/leadership.html>.

**Enterprise Process Advisory Committee** (EPAC) was formed specifically to manage the consolidation of the state's financial, procurement, and HR systems. Recently, their charter was expanded to include statewide enterprise system projects. EPAC is a working committee under ELC. EPAC details can be viewed at <http://cio.ca.gov/stateIT/ITcouncil/committees/epac.html>.

**Information Technology Council** (ITC) was formed to develop a statewide IT strategic plan and adopt enterprise-wide IT standards and policies. Membership is primarily state CIO's. There are approximately 10 committees, 2 sub-committees, and 8 working groups under the ITC. For more information on the ITC, see <http://cio.ca.gov/stateIT/ITcouncil/index.html>.

One of the ITC committees is Enterprise Architecture. This group is focused on developing a statewide business reference model (BRM) and data reference model (DRM).

**The California State Information Security Office** is currently part of the Office of Technology Review, Oversight, and Security (OTROS) group within the Department of Finance. One of the responsibilities is to set statewide security policy, both physical and computer online. As a result of SB90, this office will be moved to the new Office of Information Security and Privacy Protection effective January 2008. For details on the current office, see <http://www.infosecurity.ca.gov/>.

**Office of Privacy Protection (OPP)** is primarily focused on the protection of consumer's privacy information. The office is currently part of the Department of Consumer Affairs. As a result of SB90, this office will be moved to the new Office of Information Security and Privacy Protection effective January 2008. For information on the current office, see <http://www.privacy.ca.gov/>.

### **Office of Information Security and Privacy Protection**

SB90 (Aug 2007) created a new Office of Information Security and Privacy Protection in the State and Consumer Services Agency effective January 2008. The existing “S” portion of OTROS, that is the State Information Security Office at DOF, and the Office of Privacy Protection will be moved into the newly created agency-level office.

**CalOHI** is a department within the Health and Human Services Agency that implements HIPAA (Health Insurance Portability and Accountability Act) and represents California in the federally sponsored nationwide HISPC (Health Information Security Privacy Collaboration) project. They are included in this document because they set privacy and security policy within the health line of businesses. The SOA and identity management infrastructure must support and enforce these policies. For more details see <http://www.ohi.ca.gov/state/calohi/ohiHome.jsp>

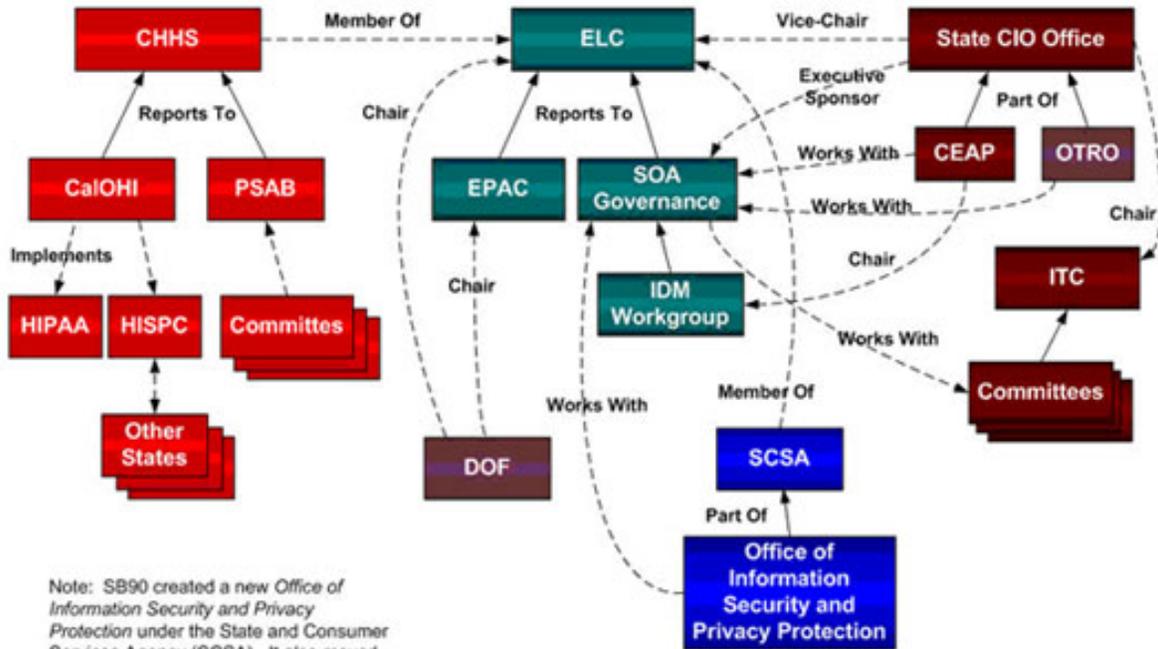
**Privacy and Security Advisory Board (PSAB)**. This new board reports to the California Health and Human Services Agency (CHHS) and was recently formed as part of HISPC Phase II. There are four committees under this board: Privacy, Security, Legal, and Education. PSAB will make policy recommendations to CHHS. The SOA Governance Group will need to understand the HIT policies recommended by PSAB to ensure they can be properly managed within the SOA and Identity Management infrastructure.

**The California Enterprise Architecture Program** is part of the State Office of the CIO. It is responsible for creating the vision and blueprint for enterprise SOA, enterprise federated identity management (IDM), the Service Reference Model (SRM), and Technical Reference Model (TRM). CEAP provides reference architectures, researches and makes recommendations on SOA policy, management, and interoperability among levels of government and public/private stakeholders. CEAP works with all lines of business and rolls up individual requirements and issues into a cohesive enterprise. CEAP coordinates with most of the organizations described in this document.

The **SOA Governance Group** was formed to determine how to initially get the enterprise SOA and federated identity management infrastructure in place. This group deals with funding and process issues, and it reports to the ELC. CEAP and the Identity Management Workgroup frame issues and make recommendations to the SOA Governance Group. They will also provide policies for shared service lifecycle management. Additionally, they will make decisions on how to best implement federated identity management.

For more information on California Enterprise Architectures efforts see <http://cio.ca.gov/stateIT/enterpriseArch.html>.

## SOA & IDM Governance



### Glossary of above Terms:

<b>CHHS:</b>	California Health and Human Services Agency
<b>CalOHI:</b>	California Office of HIPAA Implementation
<b>CEAP:</b>	California Enterprise Architecture Program
<b>DCA:</b>	Department of Consumer Affairs
<b>DOF:</b>	Department of Finance
<b>EPAC:</b>	Enterprise Process Advisory Committee
<b>ELC:</b>	Enterprise Leadership Council
<b>HIPAA:</b>	Health Information Privacy Accountability Act
<b>HISPC:</b>	Health Information Security and Privacy Collaboration (Federal)
<b>IDM Workgroup:</b>	Identity Management Workgroup
<b>ITC:</b>	Information Technology Council
<b>OTROS:</b>	Office of Technology Review, Oversight, and Security (until Jan 2008)
<b>OTRO:</b>	Office of Technology Review, Oversight (effective Jan 2008)
<b>PSAB:</b>	Privacy and Security Advisory Board
<b>SCSA:</b>	State and Consumer Services Agency
<b>SOA:</b>	Service Oriented Architecture
<b>State CIO:</b>	State Chief Information Officer

## Enterprise SOA Infrastructure

Will there be one SOA environment? No, there will be many. However, there should be a primary *enterprise* SOA environment within the Executive Branch and it is anticipated that it will reside at the Department of Technology Services.

So, what are the major components of an SOA environment? The primary components are a service bus, a service registry, and the processes to manage the service environment. Other components are also important such as proxy servers, databases to store service and policy information, and application servers. We are referring to the components in the primary SOA environment as an *Enterprise Service Bus* and *Enterprise Service Registry*.

It is anticipated that at least half of the costs for the enterprise SOA infrastructure will be in consultant fees to identify, establish, and manage processes.

It is very important that interoperability standards be set to ensure that the enterprise service bus and registry can communicate with other SOA environments that may reside in such departments as CalPERS, Franchise Tax Board, Hawkins Data Center, Counties, private industry partners, etc.

Note, it is the combination of the Enterprise Service Bus and the Enterprise Service Registry that produces *transparency* of web services. That is, the users (consumers) of services should not need to know where the services are actually located. All consumer applications point to the ESB and it figures out (via consulting with the Enterprise Service Registry) where the services are actually located. This provides a lot of flexibility as the web services can be updated and moved without affecting the users.

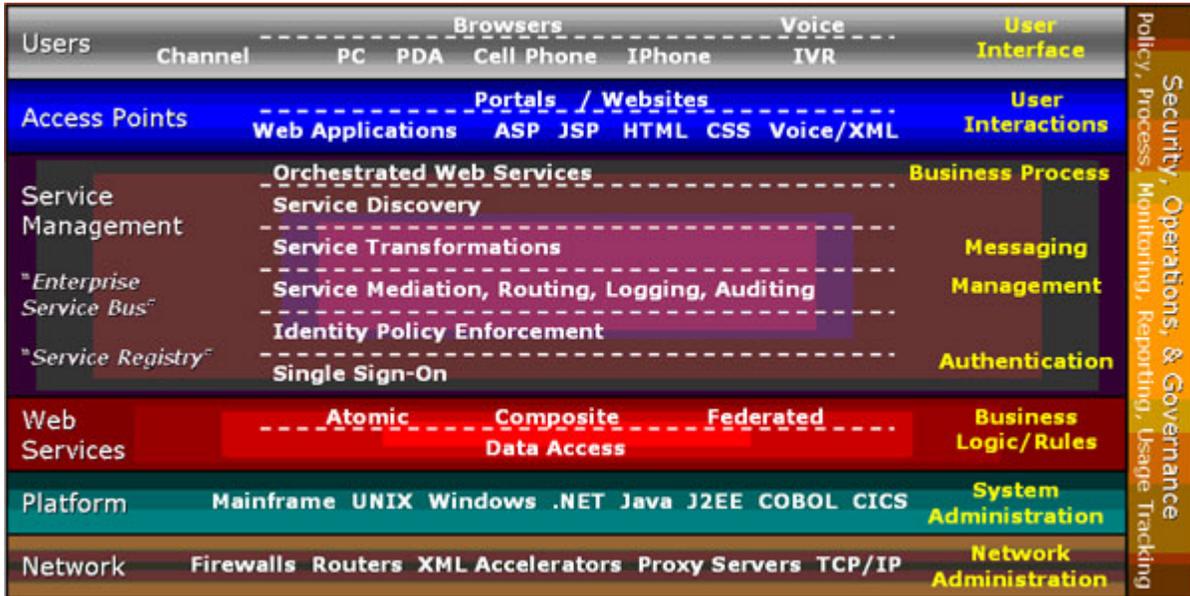
Shared services come in two categories: *shared external* services are consumed outside of your department; *shared internal* services are consumed across projects or lines of business within your department or agency. A policy needs to be determined whether either or both of the shared service types will use the *enterprise* SOA service registry on a mandatory or optional basis.

In addition to the State portal, many departments will have line of business portals. They are free to choose the portal vendor and platform that the portal is built on (subject to State portal user interface, accessibility, and usability requirements as determined by the Office of eServices). However, if a portal is going to host shared external services, then a policy needs to be set as to whether or not these services will be registered with the *enterprise* service registry and access to these services managed by the *enterprise* identity management system.

Funding, spending authority, and project mechanisms must be determined for DTS to procure install, configure, and manage the initial Enterprise SOA infrastructure.

An SOA Reference Architecture is provided as a visual illustration to show how all the pieces fit into the enterprise SOA puzzle. It is important to note that the scope of this vision is limited to online services accessed via the web, voice through an IVR system or call center, or other smart clients. It does not address the broader physical security issues or paper processes.

California SOA Reference Architecture



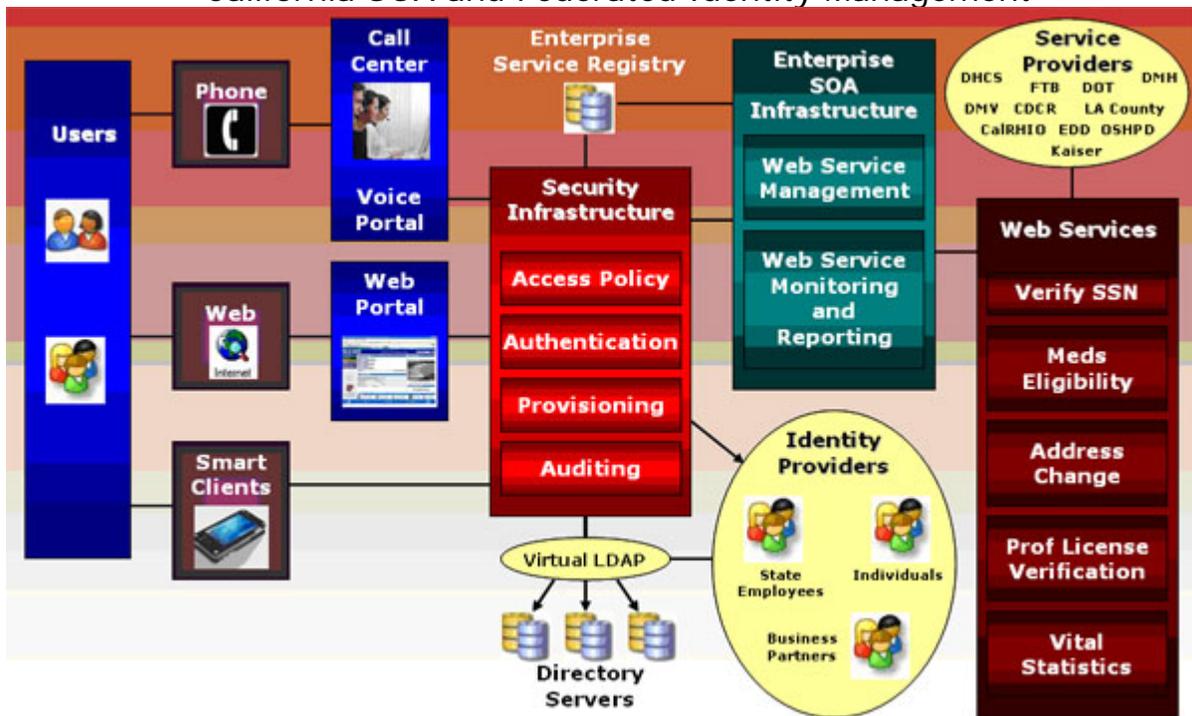
A key part of the enterprise SOA infrastructure will be creating a certification environment separate from the production environment. New or modified web services can be tested to ensure that they play nicely before moving them to the production environment. Of course, the certification process must be defined and published so developers understand the environment and the process to get their new services certified.

# Enterprise Identity Management

As we move to an online, shared services environment it is imperative that we implement a security system that ensures only authorized users get access to services that contain sensitive information. However, because security policies will vary widely across services, the identity management system must be flexible enough to determine which policy to apply, ensure that it is enforced, and meets auditing requirements.

In California, we do not have a centralized model. Therefore, we will most likely implement a federated identity management system. The goal is to establish multiple identity service providers, one for each class of user. An identity service provider will have the sole responsibility for authenticating its user class based on the security policy associated with the specific interaction (that is, the service they are trying to access). The service providers will trust the identity providers and not re-authenticate the user. However, the service provider will inspect the credentials created by the identity provider for the user to ensure that this user has been properly identified and the credential actually came from the identity provider.

California SOA and Federated Identity Management

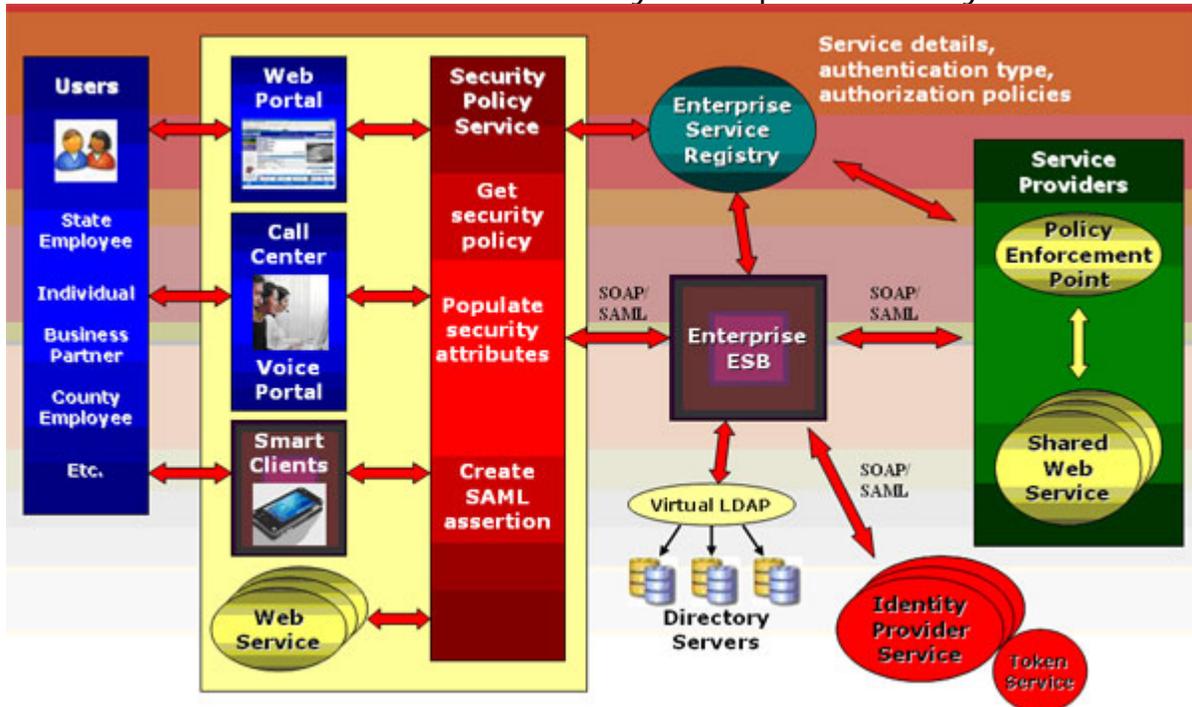


The above diagram illustrates how SOA and federated identity management work together. Regardless of how the user request was initiated (phone, web, smart client), the interaction must first successfully pass through the security infrastructure before being routed to the appropriate web service.

## Security and Authentication

The following diagram is a more detailed view of the security components. From the online perspective, a request for a shared web service could come from four possible sources (“channels”); web browser, voice, smart client, or another web service (for example, an eligibility service could invoke a verify SSN service). Regardless of which channel initiates the request, it is responsible for 1) getting the appropriate security policy, 2) populating the attributes for this specific user as specified in the security policy, and 3) creating the SAML assertion (a section in the SOAP message that deals with security).

CA SOA and Federated Identity Conceptual Security Model



One of the big decisions to be made is how to define and where to store the security *policies*. There seem to be three possibilities; 1) centralize all policies in a single (logical) repository, 2) leave all policies in local (line of business) repositories and query the appropriate repository at run time, or 3) provide meta information about the policies in a centralized (“Enterprise Service Registry”) repository. The current thinking is use the Enterprise Service Registry as a single place to retrieve enough information about the target web service to create the SAML assertion.

This same challenge exists regarding how to best architect where the security *attributes* are stored and managed. These would typically be the roles that a user is a member of, but attributes are not limited to just role-based information. It is anticipated that these user attributes will be maintained in local repositories (“LDAPs”), but a “Virtual LDAP” would “front-end” all the individual local ones. This means the channel applications would need to go to only one place (via the Enterprise ESB) to get the information necessary to populate the security attributes. In reality, the Virtual LDAP simply federates to all the local repositories.

## ***The Security/Authentication Process***

Web services communicate via industry standard SOAP messages. SOAP messages are a type of XML and are supported by practically every vendor. However, there are multiple choices for handling the security part of the SOAP message. The two leading standards are SAML and WS-SX. While there are several variations of these standards and under certain conditions they work together, for the purpose of this vision document we need to choose which one will be the primary. Both of these standards are evolving and hopefully, they will eventually merge into one.

Additionally, X.509 certificates are widely used for message authentication, message integrity, and message encryption. SAML, in conjunction with XML Encryption and XML Digital Signatures can accomplish the same thing. There is even a variation where X.509 certificates can work in concert with SAML. It is anticipated that California's SOA and Federated Identity Management infrastructure will need to support both X.509 certificates and SAML. A related issue is whether or not the State should provide a managed PKI service. This would enable a "single chain of authority" from local-state-federal government. The counties have asked the State to provide this service. The federal government is also encouraging this approach and they provide a bridge where state's can get certified.

One more potential problem needs to be addressed. The policy for digital signatures currently resides with the Secretary of State office. There are restrictions on digitally signing documents. However, it is unclear on whether or not this applies to digitally signing SOAP messages.

## ***Identity Management***

In order to arrive at the best policy and architecture decisions, CEAP and the Identity Management Workgroup will facilitate working sessions to determine the classes of users and the details of exactly how they are authenticated and authorized. The first all-day session will be focused on the Individual class of users. Some examples of this class will likely be citizens, residents of California, legal aliens, plus others. It is also anticipated that different security policies for authentication and authorization will be required. At the end of the day, we need to determine if we can agree on whether or not we can establish a single Individual Identity Service Provider and apply the different policies. The Identity Management Workgroup will frame the results in the form of a recommendation to the SOA Governance Group.

## ***Trust Model***

Another decision area will be how do we define and implement a trust model? A trust model states the conditions that consumers and providers of shared services agree upon. This includes restrictions based on user profile or role as well as broader sharing rules among providers. The model also states liabilities that the parties agree to. A decision will need to be made whether a single trust agreement can be drafted that covers all the different scenarios (perhaps via appendices) or will multiple trust agreements be needed (perhaps one for government-to-government and a separate one for government-to-private industry).

The bottom line is the enterprise identity management system must ultimately accommodate all stakeholders (levels of government, public/private). The governance for this very broad model will very likely need to be revisited as more classes of users are added to the infrastructure.

For a detailed description of federated identity standards see  
[http://cio.ca.gov/caIT/pdf/IDM\\_Standards.pdf](http://cio.ca.gov/caIT/pdf/IDM_Standards.pdf)

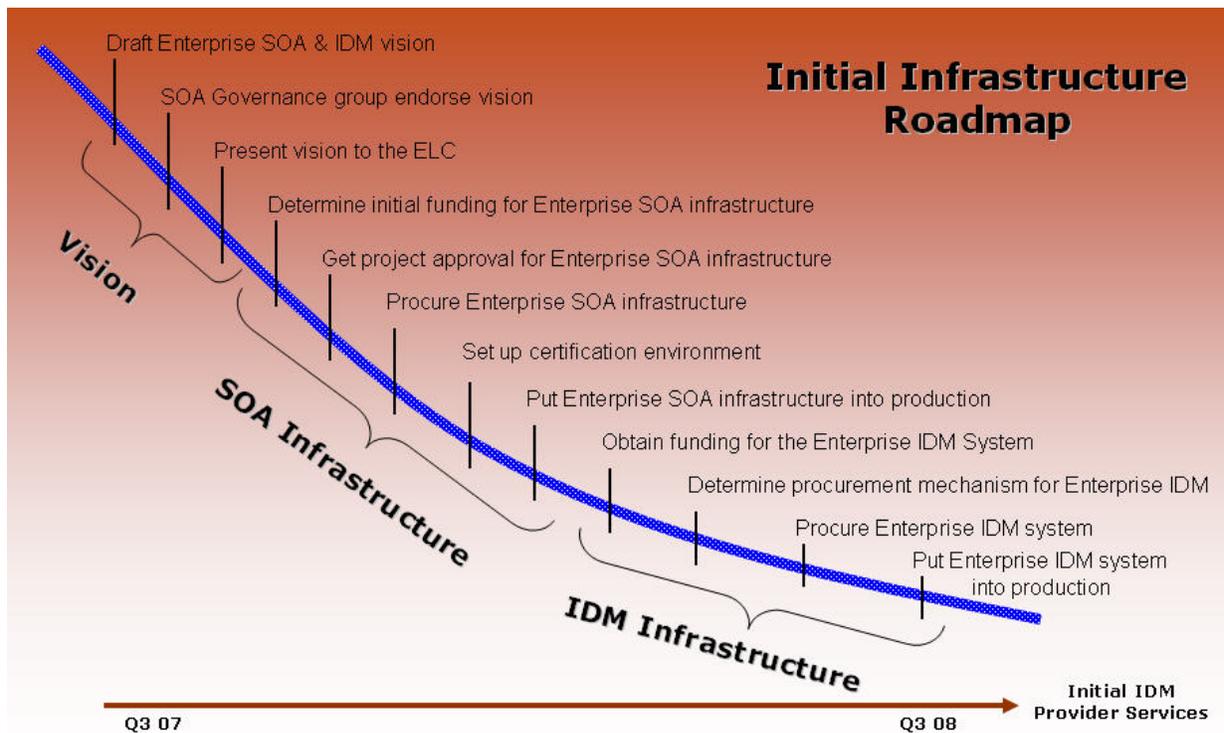
# Enterprise SOA & IDM Roadmap

To implement the vision stated in this document, this roadmap lists the key activities and decisions that need to be made and their approximate timeframes. It is anticipated that this roadmap will continually be updated as decisions are made and implementation progresses. More detailed timeframes will follow within each of the major areas (for example, SOA Infrastructure, IDM Infrastructure).

## Step 1. Articulate the SOA and IDM Vision

The general strategy is to first state the vision followed by two parallel phases: 1) get the SOA and IDM infrastructures in place, and 2) determine how Individual users and State Employee users will be defined and managed.

In the Initial Infrastructure phase illustrated below, the SOA Governance Group will debate this vision followed by a recommendation to the Enterprise Leadership Council (ELC) that it be adopted as the general direction for moving to a new enterprise information technology infrastructure.



Next, we must sort out the details of how to fund and procure the enterprise SOA infrastructure.

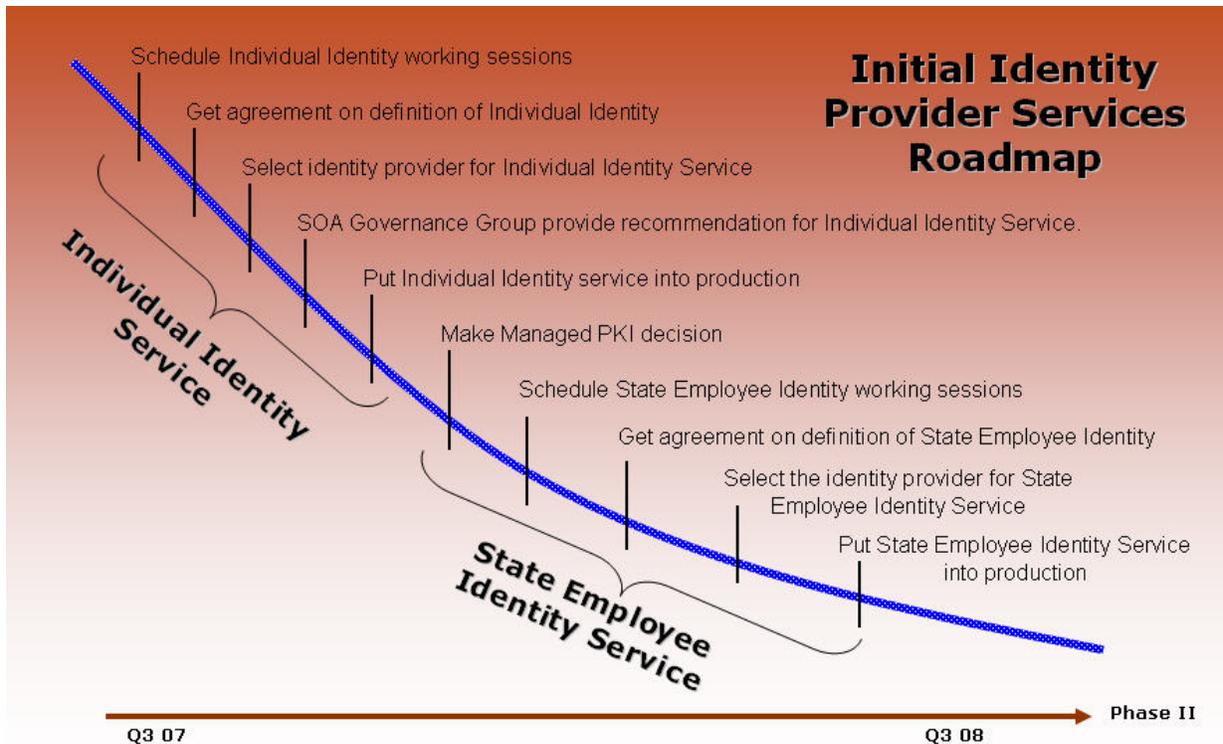
Step 2. Create Sustainable Funding Model:

In parallel with the procurement effort, we must determine how to translate the EDD/DOL funding into a DTS procurement for an enterprise identity management (IDM) system.

Step 3. Planning- First Implementation- Identity Provider Services:

Since we plan to migrate to a federated identity model, we must determine which user classes we plan to implement first. Based on initial discussions, it is recommended that we first define the Individual user and build the Individual Identity Provider Service. The Individual user could be a citizen, resident of California, legal alien, etc. CEAP and the Identity Management Workgroup will facilitate a working session where subject matter experts from state entities (and county representation) will be invited to describe exactly how they define this user class as well the details of authentication and authorization. The goal will be to agree on a common definition for this user class as well as how the policies can be best architected and enforced. The Identity Management Workgroup will frame this in a recommendation to the SOA Governance Group.

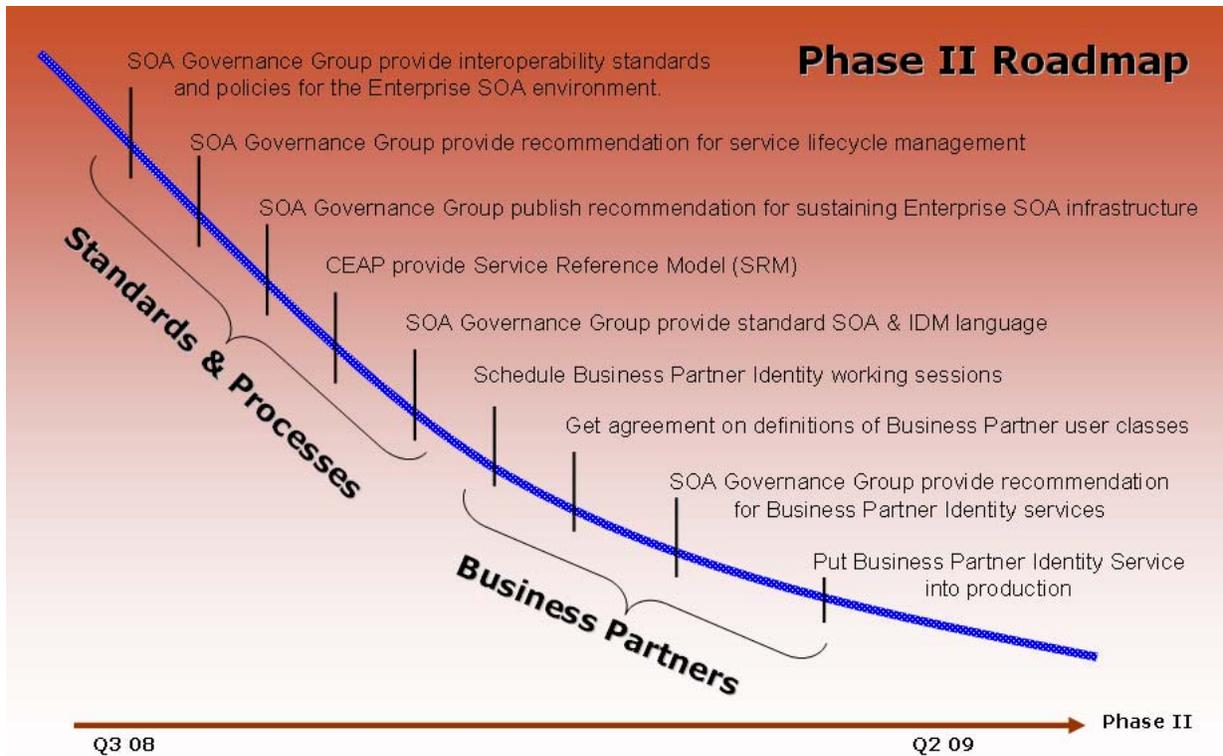
Assuming that X.509 certificates will be used in the federated identity process, as the Counties have requested a state-managed process, we need to make a decision on whether or not the state will provide a managed PKI solution. If the decision is to proceed, then a PKI program will have to be established. The State of Illinois has managed PKI in production and we could learn a lot and perhaps use their documents and processes as a starting point. We have had initial discussions and they have already provided some documents for our review. A managed PKI program includes establishing state-wide policies for all government stakeholders regarding how certificates will be structured, issued, and maintained as well as a certification policy and technical environment. Once we have sorted this out at the state level, we would then begin the process of becoming certified with the federal bridge. This would result in a single chain of authority from local to federal government meaning communication could occur at any level based on the trust policies and managed PKI certificate structure.



#### Step 4: Phase II Creates Policy - Standards and Processes for SOA & IDM access

Phase II of the roadmap focuses on creating documents that detail the standards and the processes for using the enterprise SOA and IDM infrastructure. These would be in the form of a series of recommendations provided by the SOA Governance Group to the Enterprise Leadership Council for adoption which would result in Executive Branch policy. This includes providing standard language that would be used in all funding request, project request, and procurement documents. Of course, the controlling agencies would need to also approve this language and agree to enforce it. This will provide vendors with a consist set of definitions so future services will be built in a standardized way.

Also in Phase II is the determination of how business partner users will be handled. There are many types of users in this category, e.g., healthcare providers, tax accountants, and employer representatives. Defining this class of user will be very complex and may very well require multiple identity provider services.



Current State of California’s Enterprise SOA and IDM

So, where are we now? At this point we have reference SOA and conceptual IDM architectures. DTS is in the process of sorting out how to fund the enterprise SOA infrastructure. EDD has applied for a Department of Labor grant to help fund the enterprise IDM system. The Department of Health Care Services has already built the first set of shared web services and they went into production in their environment. They will be moved into the enterprise infrastructure as soon as it is ready.