TRUSTED INFORMATION PROCESSES IN B2B NETWORKS

Chintan Doshi and Liam Peyton
S.I.T.E,University of Ottawa,Ottawa,Ontario,Canada

Keywords: Business process, circle of trust, business to business networks.

Abstract: The design, implementation and management of inter-organizational business processes that operate across the Internet have to address a number of issues that do not normally arise for business processes that operate solely within an organization. A framework is needed which supports traditional business process management and which also has the technical infrastructure in place to address federated identity management, privacy compliance and performance management. In this paper, we examine how this can be accomplished in an architecture with built in event logging and privacy auditing that deploys processes defined in the Business Process Execution Language standard (BPEL) into a "Circle of Trust" (CoT) architecture as specified by the Liberty Alliance standard for federated identity management. A sample business process scenario is implemented in the proposed framework and evaluated.

1 INTRODUCTION

The rapid proliferation of the Internet has transformed the way that organizations do business. E-business has opened up opportunities for business processes to be defined that integrate an organization's information systems with those of its external business partners in order to provide value-added services to consumers. To be successful though, the design, implementation and management of inter-organizational business processes that operate across the Internet have to address a number of issues that do not normally arise for business processes that operate solely within an organization. These issues include:

- Identify Management to safeguard consumer and employee identities during inter-organizational transactions over the Internet.
- Privacy Compliance to ensure that consumer personal data is safeguarded and handled appropriately in compliance with applicable government laws, industry regulations, and organizational policy.
- Performance Monitoring to ensure that processes are providing appropriate quality of service across organizations as well as meeting business performance objectives and Service Level Agreements.

In order to address these issues in a systematic manner, a framework is needed which supports traditional business process management and which also has the technical infrastructure in place to address federated identity management, privacy compliance and performance management. In this paper, we examine how this can be accomplished in an architecture with built in event logging and privacy auditing that deploys processes defined in the Business Process Execution Language standard (BPEL) into a "Circle of Trust" (CoT) architecture as specified by the Liberty Alliance standard for federated identity management.

2 BACKGROUND

A business process is defined in (Hammer & Champy, 2003) as "a collection of activities that takes one or more kinds of input and creates an output that is of value to the customer". In particular, a business process is a "flow or action made up of an organized sequence of tasks or activities that create, deploy and exchange artifacts. It has a measurable performance value such as rate of return. Finally, it conforms to laws, legislations and regulations". In this paper, we focus on processes executed by software in which the artifacts are in the form of electronic data.

Doshi C. and Peyton L. (2008).
TRUSTED INFORMATION PROCESSES IN B2B NETWORKS.
In *Proceedings of the Tenth International Conference on Enterprise Information Systems - SAIC*, pages 37-44
DOI: 10.5220/0001714500370044

The importance of business process automaton as a key enabler for e-business is emphasized in (Casati, F. and Shan, M. 2000). Business Process Execution Language (BPEL4WS) is emerging as a leading standard for business process automation. BPEL-WS works over a Service Oriented Architecture (SOA) (Huhns, M. N. and Singh, M. P. 2005). In SOA, different service providers publish a loosely-coupled service interface to expose their functionality in a platform and programming language independent manner. The web service technology stack consisting of SOAP, WSDL and UDDI (Curbera, F. et al. 2002) is typically used to implement a SOA.

The benefits and challenges of organizations collaborating in a B2B network to provide value-added services to customers was analyzed in detail in (Frichman, R.G., and Cronin, M.J., 2003). In particular, it highlighted the challenge of developing an information-rich, service-oriented, trust infrastructure to ensure transactions are conducted privately, securely, and in accordance with consumer preferences.

A conceptual framework for eliciting high level trust requirements in e-business is presented in (Jones et al. 2000). (Pavlou and Ratnasingam 2003, Ratnasingam 2002) analyze the importance of technology trust in B2B e-commerce and web services where technological trust is defined as "the subjective probability by which organizations believe that the underlying technological infrastructure is capable of facilitating transactions according to their confident expectations". The European Dependability initiative (Wilikens, M., Morris, P. and Masera, M., Eds. 1998) identified the four new drivers for trust in e-business systems as globalization, complexity of large-scale open systems, transition to virtual digital environments and rapidly evolving systems. A conceptual framework and survey of consumer trust online is presented in (Venkatesh Shankar, Fareena Sultan and Glen L. Urban 2002).

A Circle of Trust (CoT) is a federated identity management system in which an individual's identity and personal information is protected by a designated identity provider, while cooperating enterprises within the CoT can still share the individual's personal information as long as the individual's permission is obtained and their identity protected (Shin et al, 2004). The distinction between anonymous identity and pseudonymous identity was made in (Koch et al., 2005).

The Liberty Alliance Project was established in 2001 as a consortium of technology enterprises to create an open standard and set of specifications for federated identity management. The key objective of a Liberty Alliance Circle of Trust is to enable organizations to share data while protecting privacy of consumers. The Liberty Alliance identity federation framework (ID-FF) (Wason, T., 2003) is based on the OASIS specification for Security Assertion Markup Language (SAML) (Cantor et al, 2004). The web service framework (ID-WSF) specification (Kemp, Y., 2004) defines the creation, discovery, and invocation of interoperable identity web-services and permission-based attribute sharing. An overview of security and privacy in ID-WSF is given in (Landau et al, 2003).

In recent years, governments have passed significant legislation to regulate the mechanisms that organizations must put in place to protect privacy and document their compliance with the law. In particular, most legislation indicates that end users should have the right to know who their personal data is being shared with. In (Peyton, L., Doshi, C., and Seguin, P. 2007) it was shown that to support transparency and full compliance with privacy legislation, an Audit Trail Service was needed that tracks data-sharing events across a Circle of .Trust.

3 APPROACH

Our approach is to take BPEL definitions of a B2B process and extend them to fully integrate with and conform to the architecture and specifications defined by the Liberty Alliance for a Circle of Trust (CoT) as well as integrating an audit trail and event logging service across the B2B network.

For a given process, there are a number of special requirements that must be addressed by our approach. They can best be understood in terms of three stakeholders as shown in Table1.

Table 1: Trusted Information Process Requirements.

Consumer	Business	Administrators
(a)Protect Identity (b)Control over Data-sharing (c)Single-sign on convenience	(a)Visibility through performance monitoring (b)Document compliance (c)Transparency to build consumer confidence	(a)Verify compliance (b) Transparency

The consumer's foremost goal is to ensure their privacy. Their identity should be protected when their personal data is being shared by different service providers across B2B network. The consumer should be able to define access-control policies for the sharing of their data. Additionally, consumers may like to have the convenience of single-sign on (SSO) wherein, once authenticated, they can access different service providers in the same trust domain without having to login again.

Collaborating businesses want to manage process execution by monitoring key performance indicators and service level agreements. They want to document their use of personal data to ensure compliance with privacy legislations. Additionally, they may need means to build transparency into data-sharing activities beyond the minimum requirements of law to build consumer's confidence.

Government administrators or industry regulators seek to ensure and verify compliance of processes with laws and legislations.

Table 2 below lists the framework components in our approach which help meet these requirements. Each of these issues and the manner in which the framework component addresses them is explained in more depth in the next section through an example B2B process scenario.

Table 2: Trusted Policies in a Trusted Information Process.

B2B Issue	Framework Component	Use
Federated Identity	Identity Provider	Single-sign on , Protects Identity through pseudonyms
Privacy preserving data-sharing	Discovery Service and Policy Decision Points	Discovery Service enables any two service providers to exchange data while protecting identity. PDP enforce user's access-control policies.
Privacy Compliance	Audit Trail Service	Document and Verify Privacy Compliance, Transparency into Data-sharing
Performance Monitoring	Event Stats Service	Measuring performance

4 SCENARIO

4.1 The Business Process

Figure 1 below depicts a high level overview of a business process scenario we have implemented in our proposed architecture. We will use it to illustrate how B2B issues can be addressed. We first describe the business scenario in high level terms without any concern for how federated identity management, privacy compliance, or performance monitoring are addressed. Once the process is well understood and defined in BPEL, we then show how to extend the process and integrate it into a Liberty Alliance Circle of Trust in order to address those issues.

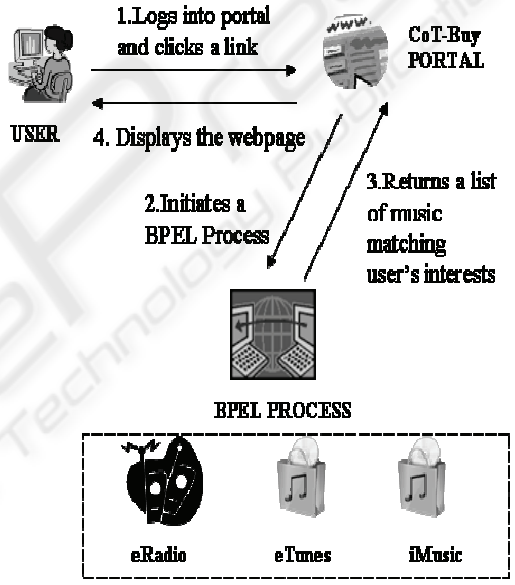

Figure 1: Scenario.

CoT-Buy, a service provider, runs a portal that allows users to search and buy music that matches their interests (such as favorite genres, bands, albums etc).

The users log on to the CoT-Buy portal and click on a link to search for music. This initiates the execution of a BPEL process that first gets user's music interests at an online radio service, eRadio, and then fetches matching music from vendors, eTunes and iMusic, across the B2B network. Finally, the BPEL process returns a list of music items along with their prices, vendor and download information to CoT-Buy portal which accordingly generates an on-the-fly webpage to the user displaying that personalized list.

The user subscribes to an online radio service called eRadio. It keeps track of user's listening patterns through descriptive music "tags" based on album, song, artist or genre such as "Comfortably Numb", "Rock" or "Pink Floyd". eRadio offers access to these user tags to third-party service providers such as CoT-Buy through an eRadio web-service.

eTunes and iMusic are online music vendors in the B2B network. They both provide an online catalogue service with an operation that takes a list of "descriptive music tags" as input and returns a list of matching or similar music tracks, their price and download information.

The Activity Diagram for the BPEL process in Figure 2 illustrates the detailed steps of the B2B process as follows:

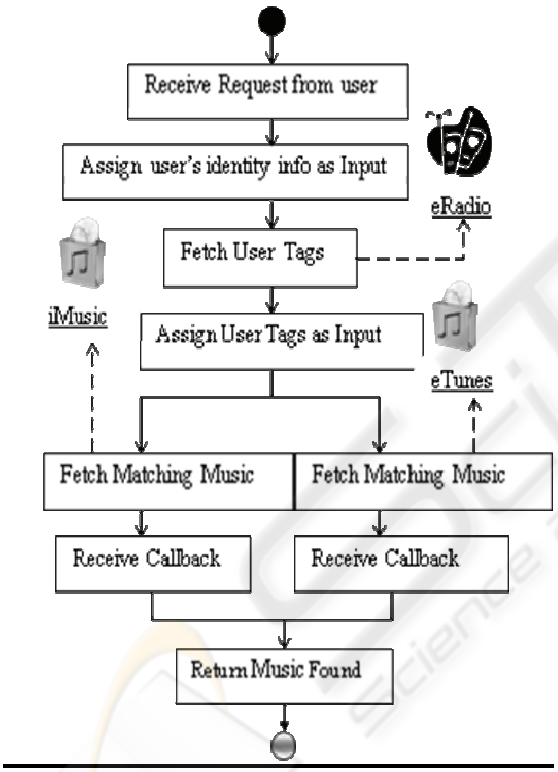

Figure 2: BPEL Activity Diagram for the scenario.

(a) The BPEL Process receives initial request from CoT-Buy (on behalf of a user "Bob"). This corresponds to a receive construct in BPEL syntax.

(b) It then constructs the input for calling the eRadio Service. In order to communicate directly with eRadio on Bob's behalf, CotBuy needs to reveal Bob's identity information at CoTBuy to eRadio, so that it can request tags for the correct user. This corresponds to the Assign construct (perhaps with XPATH queries) used in BPEL syntax.

(c) The BPEL Process synchronously invokes an operation on the eRadio webservice to fetch Bob's tags (based on the music he listens to at eRadio online service). eRadio web service returns a list of music tags that best reflect Bob's music interests at eRadio.

(d) The BPEL process sets the input for invocation of eTunes and iMusic catalogue services to the list of Bob's music tags obtained in previous step.

(e) BPEL asynchronously invokes eTunes and iMusic catalogue web-services in parallel to obtain a list of music items that matches or are similar to Bob's music tags.

(f) The BPEL process then receives a callback from the services invoked in previous step.

(g) The BPEL process returns the combined list of music items obtained in the previous step to CoT-Buy.

Notice that the above BPEL definition for the process:

(a) Needs to reveal Bob's identity information stored at CoTBuy in order to request his music tags from eRadio (so that eRadio can correctly identify the user as "Bob").

(b) Doesn't document sharing of personal data taking place between CoTBuy and eRadio.

(c) If Bob visits another service provider portal inside the B2B network (example: eRadio to preview music), he needs to re-authenticate with that service provider.

(d) Lacks support for measuring key performance indicators in the B2B process. For example, assume that CoTBuy wants to find out "Clickthrough rates" for each music item fetched from the different vendors that it offers to the user for purchase. The above B2B infrastructure lacks support to monitor the events required to calculate "Click through rates".

4.2 Liberty Circle of Trust Architecture

Figure 3 illustrates the Liberty Alliance Circle of Trust (CoT) architecture for which the BPEL process must be refined and extended in order to

address identity management, privacy compliance, and performance monitoring. The main components are:

- Liberty Service Providers
- Liberty Identity Provider
- Audit Trail Service
- Event Stats Service

CoT-Buy is a service provider that runs a web portal inside the CoT. eRadio, iMusic and eTunes are service providers, each of which runs a web-service (that conforms to the Liberty ID-WSF specifications).

The Identity Provider (IDP) is a trusted authority responsible for managing user's identities inside the CoT. It also offers a single-sign on service whereby a user, once logged in with IDP, can leverage his authentication state with any service provider within the CoT without having to login again. In our scenario, when Bob attempts to click on the link to search for music on CoT-Buy portal, he is redirected to the Identity Provider site for authentication. Once logged into IDP, Bob is redirected back to CoT-Buy where he obtains the results of that click (i.e. music tracks returned by the BPEL process).

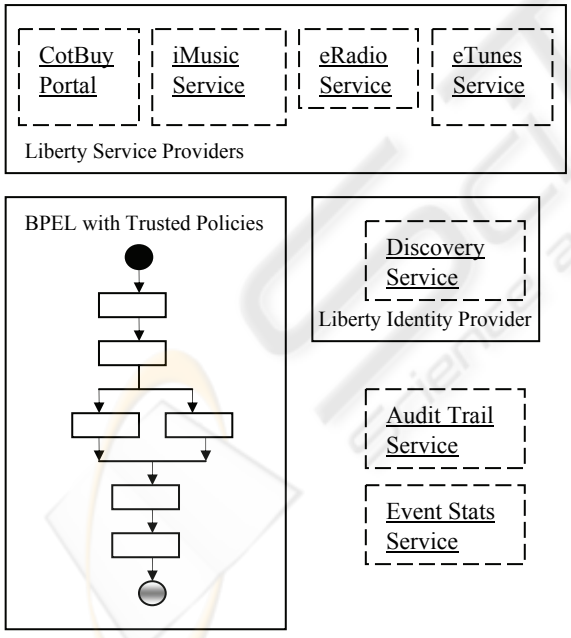

Figure 3: Liberty Alliance Circle of Trust.

The Audit Trail Service is a common service provided to log data-sharing events in the process to document privacy compliance. Every time any two service providers exchange consumer's identity or personal data, an event is logged into the Audit Trail Service. For example, when eRadio exchanges Bob's tags with CotBuy, Bob's Audit Trail records this. By providing a historical log of all such data-sharing events across the CoT, the Audit Trail Service enables the user to see how their data is shared and used. At the same time, privacy officers can verify compliance or investigate potential breaches by service providers. Businesses can make use of this audit trail service to document their compliance as required by privacy legislations. It also gives them a platform to provide transparency beyond the minimum requirements of law to build consumer's confidence.

The Event Stats Service is a common service provided to log events during the process runtime in order to measure performance. For example, in our scenario, this service can be used to log events pertaining to "views" and "clicks" on music tracks returned from eTunes. CoT-Buy can then run reports on the Event Stat Service to measure the Click-through rates for each music track from eTunes based on Total number of Views / Total Number of clicks.

Based on the Liberty Alliance federated identity standards, only the identity provider knows Bob's real identity in the Circle of Trust. Bob would then link or "federate" his local accounts at different service providers with his identity at IDP to form a network identity. The IDP assigns a unique pseudonym or an opaque identifier to each component in the CoT that uniquely identifies Bob at that particular component. For example, IDP may assign an opaque identifier "Bob_CotBuy" to CotBuy, "Bob_eRadio" to eRadio and so on. Each Service Provider only knows the user by its own pseudonym. It knows nothing about user's pseudonym at other service providers.

When any two service providers wish to exchange Bob's personal data, they cannot do so directly because of unique pseudonyms assigned to the Service Providers by IDP. Hence, there is a special Liberty Service called the Discovery Service (Hodges, J., Cahill, C., 2006) provided by the Identity provider to enable any two service providers to exchange data about the user in a privacy-preserving manner. Discovery Service works by issuing End-point references (EPRs) (Hodges, J., Cahill, C.,2006) to the data requestor (CoT-Buy) containing encrypted security and/or identity tokens that can be used by the data provider (eRadio) to dereference its opaque identifier (Bob_eRadio) for that identity (Bob).

Discovery Service, Audit Trail Service and Event Stats Service are illustrated in detail in next section.

4.3 Liberty Discovery Service

Figure 4 shows how additional steps need to be inserted into the BPEL process definition in order to integrate with the Liberty Discovery Service. In order to request Bob's music tags from eRadio, CoT-Buy first communicates with Discovery Service and then calls eRadio service. The shaded areas indicate the modified steps.

During the single-sign on process for Bob, IDP issues CoT-Buy an end-point reference (EPR) ("Bob-Discovery EPR") containing encrypted security tokens to access Discovery Service on Bob's behalf. The BPEL process passes on these tokens to Discovery Service in order to invoke it (steps 2 and 3).

The Discovery Service call returns another EPR ("Bob-eRadio EPR") containing encrypted security tokens to access eRadio service on Bob's behalf.

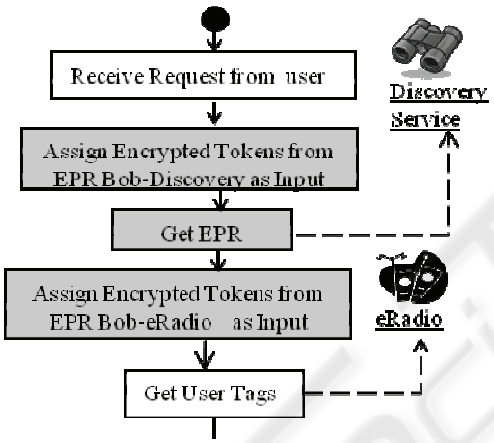

Figure 4: Liberty Discovery Service.

The BPEL process assigns these tokens as input to a call on eRadio service in order to request Bob's music tags. The encrypted tokens allow eRadio to dereference Bob's opaque identifier at eRadio (Bob_eRadio) without revealing it to CoT-Buy.

Additionally, Liberty Alliance specifications allow the Discovery Service to consult a policy decision point (PDP) to ensure Bob has granted permission to CoT-Buy to access his music tags at eRadio. The PDP may check for any access-control policies which may have been defined by Bob to control sharing of his personal data at eRadio. Thus, Discovery Service along with PDPs can enable exchange of personal data in a privacy preserving manner while giving the consumer more control over the data-sharing.

4.4 Audit Trail Service

Figure 5 shows how additional steps need to be inserted into the BPEL process definition in order to integrate with the Audit Trail service for documenting privacy compliance. Any time, one organization is calling the web service of another organization to share data about a consumer that access to data needs to be logged. Both the organization names are logged, as well as the type of data shared (but not the actual values). In our scenario, the first data access is logged after the discovery service call returns an EPR to the BPEL process enabling it to invoke eRadio service with requisite tokens. The second data access is logged after eRadio returns Bob's music tags to the process.

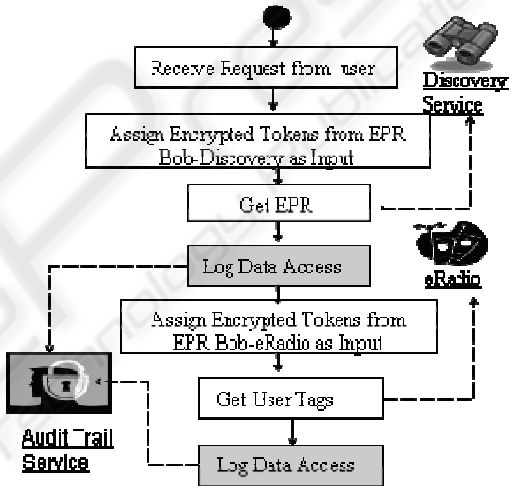

Figure 5: Audit Trail Service.

The integration of a similar audit trail service and structure of an audit event into a Liberty Alliance Circle of Trust is discussed in (Peyton, L., Doshi, C., and Seguin, P. 2007). An example audit entry is illustrated below in table 3.

Table 3: Audit Event Log.

ATS specific opaque id	Client	Attribute Provider	Attributes Shared	Time stamp
Bob_ATS	CoT_Buy BPEL	eRadio	Music tags	T1

ATS only knows the user (Bob) whose data was being shared by its opaque identifier (Bob_ATS) and not the real identity. Note that only the attribute names and not the values are stored in the audit logs.

The combination of these 2 measures ensures that Audit Trail Service itself is not a privacy breach.

4.5 Event Stats Service

Figure 6 shows how additional steps need to be inserted into the BPEL process definition in order to integrate with the Event Stats service to log events for monitoring performance. In our scenario, we look at how such an event stats service is used to help CoT-Buy measure click-through rates on music items returned by eTunes and iMusic that match user's music interests.

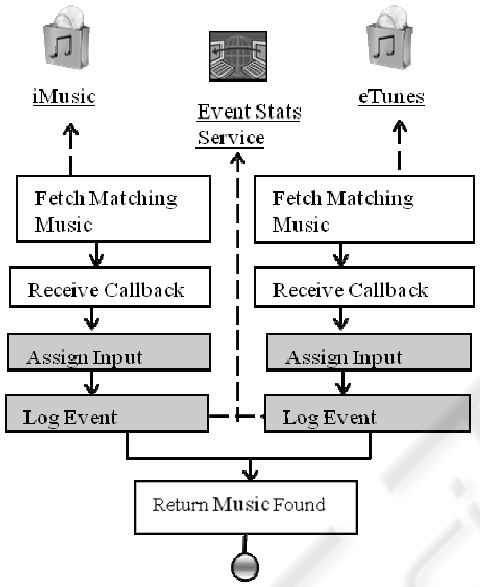

Figure 6: Event Stat Service.

In Figure 6, the event logged by the BPEL process records the music items returned and their vendor information into the Event Stat Service. This corresponds to the number of times a music item from a vendor deemed matching to some user's interests was shown to the users.

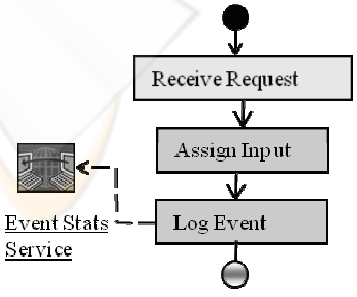

Figure 7: Event Stat Service.

Similarly, in Figure 7, when the user clicks on one of the returned music items to get more information about it, a BPEL sub-process is executed that logs the music item clicked and its vendor information. This corresponds to the number of times a music item from a vendor was clicked by the users. CoT-Buy could then run a query on the Event Stats service to generate reports on click through rates calculated as total number of times clicked / total number of times shown for any given music item from a vendor.

5 EVALUATION / CONCLUSIONS

Implementing a simple B2B scenario, we have shown how BPEL definitions of a process can be extended to integrate with a Liberty Alliance Circle of Trust architecture in order to address issues around identity management, privacy compliance and performance monitoring.

Our proposed architecture offers consumers the convenience of single-sign on and protects their identity through the use of unique pseudonyms or opaque identifiers at each service provider's site. A Discovery Service also protects user's data by allowing them to control access to their information whenever one service provider attempts to interact with another service provider. An Audit Trail Service logs events to document privacy compliance while an event stats service logs events for monitoring performance.

In the scenario, we manually extended process definitions in BPEL to accommodate the Discovery Service, Audit Trail Service and Event Stat Service. This is problematic for a number of reasons. First, it is a complex, error-prone, manual task that must be repeated for every process that is supported in the B2B network. Second, as a manual process there is no way of ensuring that all processes have been extended, so there is no way of being sure that the processes can be trusted. It would be advantageous if the required extensions to the BPEL process could be incorporated automatically, either by generating the necessary BPEL from the original definitions, or by having a special BPEL engine insert the appropriate steps dynamically. A final issue, is that the resulting BPEL defined behaviour is quite complex. Tools are needed to mediate between the original straight forward definition of the process, and the more complex version that is executing.

More work is needed to analyze these issues. It is likely that federated authentication and discovery service could be handled automatically. Audit trail

logging for privacy compliance and event logging for performance monitoring, however, require more sophisticated configuration, including the definition of a common data sharing and event model. The Liberty Alliance supports a common shared data model for defining data sharing between services. This could be leveraged to automate the process of logging access to shared data. Common events for performance monitoring could possibly be defined in a similar approach, but this would require more shared analysis between organizations to define the events for monitoring shared processes.

REFERENCES

BPEL4WS. BEA, IBM, Microsoft, SAP and Siebel, "Business Process Execution Language for Web Services ", S. Thatte, et al., May 2003. http://www.ibm.com/developerworks/library/specification/ws-bpel/ Accessed 2008/03.

Cantor, S., Kemp, I.J., Philpott, N.R. , Maler, E. (2004) "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0". *OASIS SSTC*, September 2004. http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf Accessed 2008/03.

Casati, F. and Shan, M. 2000. Process Automation as the Foundation for E-Business. In *Proceedings of the 26th international Conference on Very Large Data Bases* (September 10 - 14, 2000). A. E. Abbadi, M. L. Brodie, S. Chakravarthy, U. Dayal, N. Kamel, G. Schlageter, and K. Whang, Eds. Very Large Data Bases. Morgan Kaufmann Publishers, San Francisco, CA, 688-691.

Curbera, F., Duftler, M., Khalaf, R., Nagy, W., Mukhi, N., and Weerawarana, S., (2002). Unravelling the Web services web: an introduction to SOAP, WSDL, and UDDI. *Internet Computing, IEEE 6* (2), 86-93.

Frichman, R.G., Cronin and M.J. (2003) "Information-Rich Commerce at a Crossroads: Business and Technology Adoption Requirements", Communications of the ACM, Sept. 2003, Vol. 46, No. 9

Hammer, Michael and Champy, James (2006), Reengineering the Corporation: A Manifesto for Business Revolution (revised and updated), Collins, 2006.

Hodges, J., Cahill, C.(2006), Eds., Liberty ID-WSF Discovery Service Specification. Ver2.0, Liberty Alliance Project, New Jersey, 2006. http://www.projectliberty.org/liberty/content/download/875/6201/file/liberty-idwsf-disco-svc-v2.0.pdf, Accessed 2008/03

Huhns, M. N. and Singh, M. P. 2005. Service-Oriented Computing: Key Concepts and Principles. *IEEE Internet Computing* 9, 1 (Jan. 2005), 75-81.

Jones, S., Wilikens, M., Morris, P., and Masera, M. 2000. Trust requirements in e-business. *Commun. ACM* 43, 12 (Dec. 2000), 81-87.

Kemp, Y. (2004). Eds. "Liberty ID-WSF Web Services Framework Overview", Liberty Alliance Project, New Jersey, 2004. http://www.projectliberty.org/liberty/resource_center/papers Accessed 2008/03

Koch, M., and Möslein, K.M. (2005) "Identity Management for Ecommerce and Collaborative Applications", *International Journal of Electronic Commerce / Spring 2005*, Vol. 9, No. 3, pp. 11–29.M.E. Sharpe Inc., 2005.

Landau, S. (2003) eds., "Liberty ID-WSF Security & Privacy Overview"; version 1.0, Liberty Alliance Project, New Jersey, 2003. http://www.projectliberty.org/resource_center/specifications/liberty_alliance_id_wsf_2_0_specifications Accessed 2008/03.

Pavlou, P., Ratnasingam, P. (2003), " Technology trust in B2B electronic commerce: conceptual foundations", in Business strategies for information technology management, pp 200-215, IGI Publishing, Hershey, Pa, US. 2003.

Peyton, L., Doshi, C., and Seguin, P. 2007. An audit trail service to enhance privacy compliance in federated identity management. In *Proceedings of the 2007 Conference of the Center For Advanced Studies on Collaborative Research*,(Richmond Hill, Canada, October, 2007. ISSN:1705-7361

Ratnasingam P. (2002), "The importance of technology trust in Web services security", Information Management & Computer Security, Volume 10, Number 5, 2002 , pp. 255-260(6).

Shin D., Ahn, G-J, Shenoy, P. (2004) "Ensuring Information Assurance in Federated Identity Management", *IEEE Intl. Conference on Performance, Computing, and Communications*, 2004, p. 821-826

Venkatesh Shankar, Fareena Sultan and Glen L. Urban 2002. Online trust: a stakeholder perspective, concepts, implications, and future directions.,The Journal of Strategic Information Systems ,Volume 11, Issues 3-4, December 2002, Pages 325-344

Wason, T., eds (2003)., "Liberty ID-FF Architecture Overview"; version 1.2, Liberty Alliance Project, New Jersey, 2003. http://www.projectliberty.org/liberty/resource_center/papers Accessed 2008/03

Wilikens, M., Morris, P. and Masera, M., Eds. 1998. Defining the European.Dependability Initiative: A Strategy Document. European Communities.EUR Report, EUR 18139 EN, May 1998.