Inf Syst Front (2007) 9:469–479
DOI 10.1007/s10796-007-9055-2

An interorganizational knowledge-sharing security model with breach propagation detection

Daniel S. Soper · Haluk Demirkan · Michael Goul

Published online: 6 October 2007

Abstract The increasing adoption of Service Oriented Architecture (SOA) is allowing more and more companies to integrate themselves in interorganizational netchain environments wherein knowledge assets can be electronically shared with selected business partners. The dynamic nature of these environments implies a need for organizations to protect and monitor the flow of their valuable knowledge assets throughout the netchain if they hope to maintain their long-term competitive positions. In this paper, we propose an interorganizational knowledge-sharing security model that integrates the value chain reference model (VCOR), the federated enterprise reference architecture model (FERA), and multidimensional data warehouse technologies to allow for the proactive monitoring of shared knowledge assets across an SOA-based netchain. The proposed architecture is novel In that it supports dynamic policy revision through the automated detection of knowledge-sharing breaches within a netchain—a process whose viability is demonstrated using network flow theory and a series of simulations. Existing business intelligence infrastructures can be readily modified to support the proposed model, as multidimensional data warehousing has already been adopted in many organizations.

D. S. Soper · H. Demirkan (✉) · M. Goul
Department of Information Systems,
W.P. Carey School of Business, Arizona State University,
PO Box 874606, Tempe, AZ 85287-4606, USA
e-mail: Haluk.Demirkan@asu.edu

D. S. Soper
e-mail: Daniel.Soper@asu.edu

M. Goul
e-mail: Michael.Goul@asu.edu

Keywords Knowledge security · Netchain · Supply chain · Interdependence · Collaboration · Business intelligence

1 Introduction

Mergers and acquisitions, new regulations, rapidly changing technology, increasing competition, and heightened customer expectations all imply that organizations must find innovative ways to become more collaborative, virtual, accurate, synchronous, adaptive, and agile. Knowledge management plays a key role in this complex and rapidly evolving business environment (Sharda et al. 1999). The increasing adoption of dynamic knowledge management capabilities has endowed many organizations with the ability to share their knowledge resources and expertise with business partners in response to changing demands, and therethrough to build collective value chain-based intellectual capital. As these interorganizational knowledge-sharing relationships proliferate, however, so too does the potential for knowledge-sharing breaches (Garg et al. 2003; Majchrzak 2004). Knowledge management initiatives are commonly impeded by difficulties surrounding the establishment and maintenance of knowledge-sharing security (Grant 1996), a problem that is compounded when the initiative is extended into the realm of interorganizational knowledge-sharing. While many companies invest heavily in preventing outside information security breaches, a 2006 study found that 68% of organizations had suffered financial losses from breaches of internal information security during the preceding year (Gordon et al. 2006). Furthermore, insiders have been identified as representing the most significant threat to the security of interorganizational exchanges of information (Shih and Wen 2003). Organizations that are considering joining or establishing a

knowledge-sharing environment with their business partners must inevitably face a classic decision: do the risks outweigh the rewards? On the one hand, an interorganizational knowledge-sharing environment allows the participating organizations to collectively build valuable intellectual capital and new knowledge assets (Hardy et al. 2003)—the cornerstones of modern business. On the other hand, an organization that shares its knowledge assets outside of the boundaries of the company exposes itself to the possibility that those knowledge assets will be acquired by a non-sanctioned recipient (e.g., a direct competitor), thus eroding the originating organization's competitive position. Deciding whether or not to participate in an interorganizational knowledge-sharing environment is therefore of critical importance to many organizations, and it is hoped that the results of the simulation described later in this paper will be informative to those managers who are facing such a decision.

1.1 Netchains and interorganizational knowledge-sharing

The phrase "netchain analysis" combines the supply and value chain analysis research streams with research conducted in the area of network analysis. A netchain is a set of networks comprised of horizontal ties between firms within a particular industry or group such that these networks (or layers) are sequentially arranged based upon the vertical ties between firms in different layers (Lazzarine et al. 2001). A netchain analysis therefore differentiates between horizontal ties (i.e., transactions in the same layer) and vertical ties (i.e., transactions between layers), mapping how organizations in each layer are related to each other and to organizations in other layers. Netchains also share properties with other types of social network analyses, for example, the concepts of centrality, degrees of separation, density, etc. (Scott 2000). We posit that netchain analysis can facilitate value chain security governance for binary organizational knowledge exchanges in the short term. We also consider the question, "How can an organization govern the long term security of the knowledge assets it shares within its netchain?"

Let us consider a large technology company (C) that produces a constantly changing mix of products, and that is operating in an environment of global competition which demands ever-shorter product life cycles. Suppose this company's netchain analysis results in three distinct subsets. As a customer, this company purchases equipment from a variety of suppliers (N1). It also provides consulting services in a second subset (N2). Lastly, it designs, develops and manufactures microchips the third subset (N3).

In each of these subsets, knowledge-sharing is becoming as important as the movement of physical materials and products, as depicted in Fig. 1. To effectively govern the secure transfer of knowledge initiated from C, there is a need for a common and normalized set of business semantics to describe business processes both within the organization and with the organization's business partners; i.e., through both the vertical and horizontal connections within the netchain. In addition, as a part of the netchain agreements expressed using those semantics, each company must work closely with its relevant vertical and horizontal partners to achieve common goals through a collectively governed approach to knowledge-sharing. Also, each trusted netchain partner must extend control over key shared and un-owned knowledge assets throughout its respective netchain, such that if a breach of security in

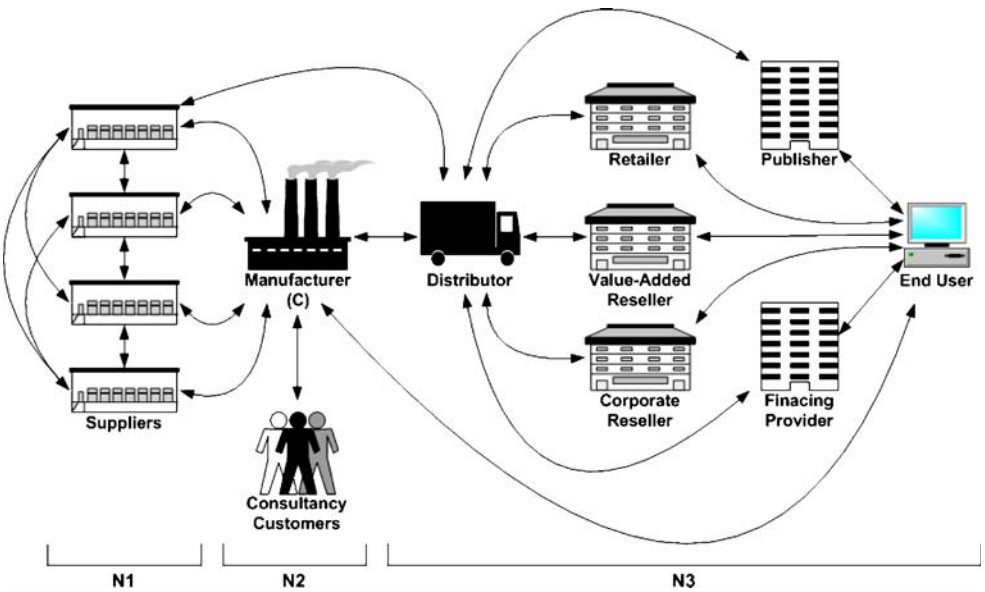

Fig. 1 Information and knowledge-sharing in netchains

any knowledge-sharing relationship occurs, containment strategies can be triggered and modifications to the governance of future knowledge-sharing relations with partners can be updated throughout relevant subsets of the netchain. All of this must take place within a context of evolving market conditions, new business requirements, and the constant acquisition of new trusted partners and mitigation of losses associated with the dismissal of untrustworthy partners. This dynamic implies the need for constant adjustments to the policies that are in place to govern knowledge-sharing security in an organization's netchain (Eisenhardt and Martin 2000; Weill et al. 2002).

2 Interorganizational knowledge-sharing security

Our approach to interorganizational knowledge-sharing security relies upon an integrated process and technology framework that is a combination of the Value Chain Group's VCOR (Value Chain Operations Reference) model (VCG 2005) and Intel Corporation's FERA (Federated Enterprise Reference Architecture) model (Semantion 2006). This framework can be used to represent and map business semantics to architectural semantics, as per the example shown in Fig. 2 (Drecun and Brown 2004). VCOR defines a unified business process framework and reference model utilizing a common language and taxonomy to facilitate effective value chain communication for information and knowledge-sharing that can be used for process modeling, gap analysis, simulation, benchmarking, and consensus building. In addition, FERA is an architectural representation that allows collaborative process models to be mapped to components of the conceptual architecture, as well as to required resources for dynamic allocations. FERA is rapidly gaining mainstream acceptance, and is currently the basis for the forthcoming ebSOA standard (OASIS 2006). These two independent but reconciled process representations facilitate the mapping of business processes to core collaboration capabilities for efficiency.

Knowledge-sharing can be considered in the context of three patterns: person-to-person, system-to-system, and person-to-system (and vice-versa). In person-to-person, knowledge is exchanged using a vocabulary and semantics. In system-to-system, information exchange is enabled by systems interpreting and processing semantics (Brown and Carpenter 2004). Finally, person-to-system depends upon a person's understanding and a system's processing. Relying upon FERA and VCOR, there are several opportunities for knowledge-sharing and exchange within a netchain, which are expressed as patterns in Table 1 along the dimensions of 'business context' and 'exchange agent type.' These patterns are described as follows (Brown and Carpenter 2004; CPDA 2004; Drecun and Brown 2004):

1. *Person-to-Person*: Personal interactions supported by collaborative software are used for the ad-hoc exchange of secure knowledge between parties with familiar contexts, e.g., dynamic replenishment, engineering changes, multi-party conceptual design, portfolio planning, etc. In this pattern, participants interact by exchanging information through direct inquiries into each other's systems through a common portal, in addition to using their own systems supported by collaborative software (Drecun and Brown 2004). For example, two planners may inquire into a sales forecast over a common portal directly with each other after each one conducts their own independent analysis. This pattern relies on manual reconciliation and shares information in an abstraction layer. It does not support detailed data moves.
2. *System-to-System*: Fixed workflow automation for industry standard content is provided by a system-to-system publish-and-subscribe pattern, e.g., warranty administration, inventory updates. In this pattern,

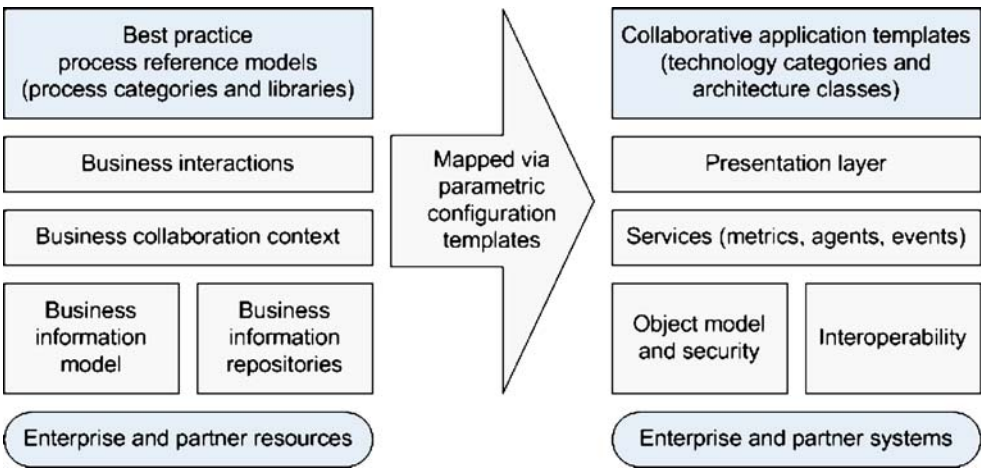

Fig. 2 Mapping business semantics to technology semantics. Adapted from Drecun and Brown (2004)

Table 1 FERA process patterns

Collaborative process patterns	People exchange business content with other people	Systems exchange business content with other systems	People exchange business content with systems
Business context is managed by a single authority			3. Bulletin boards and web meetings
Business context is distributed over several authorities	1. Personal interactions supported by collaborative software	2. System to system publish-and-subscribe	4. Collaborative business process management

Adapted from Drecun and Brown (2004)

participants generally do not need to communicate with each other; rather, the systems exchange the information (Brown and Carpenter 2004). This configuration requires that all participants agree on common semantics of the information to be exchanged. It may be used to support multiple instances of the same application, e.g., ERP, CRM. It does not support multi-threaded exchanges of information.

3. *Person-to-System*: Bulletin boards and web meetings are used for remote access to share public information and for virtual collaboration, e.g., the reporting of schedule status, changes in forecasts, inventory and delivery updates, and exchanges for conducting virtual meetings. In terms of the process, a remote participant signs in through a portal to start a session, follows the available commands, and concludes the session by leaving all of the results in the remotely administered system. Participants exchange information via the system's inbuilt upload and download functionality. Remote participants are exposed to selected information by a resource from within the control domain, whereby they consume only the sets of information presented directly by that resource (CPDA 2004). In this pattern, drilling down may expand the context that some of the parties cannot follow, which is why it is limited to those contexts that are intuitive and fully comprehensible to all parties.
4. *Person-to-System*: Collaborative business process management is used for non-deterministic complex processes that need iterative and dynamic reconciliation of multi-threaded information and knowledge-sharing, e.g., dynamic replenishment, direct and reverse allocation management, concurrent distributed security systems engineering, etc. Many processes do not have standards for the business logic that can govern distributed systems and participants. In this pattern, technically, participants map to the shared business logic representing the equivalent semantics from their internal business logic. An agent-based framework reconciles the internal logic dynamically and coordinates events between the collaborative event handlers on the federation server and the interval workflow systems (Brown and Carpenter 2004; Drecun and Brown 2004). This pattern combines the functional capabilities of other patterns in addition to full VCOR-FERA collaboration.

In light of these patterns, we propose a data warehouse-coupled knowledge security policy hub located at each netchain entity, as per Fig. 3. Policies are to be maintained and managed within each associated FERA implementation, i.e., messages passing through FERA gateways at both the sending and receiving ends of a given knowledge collaboration are governed by the hub. The hub must have the ability to manage and control the governance of knowledge-sharing security for each of the patterns discussed above in order to support binary knowledge exchanges. For example, when a manufacturer acquires equipment from a supplier, design specifications and the knowledge associated therewith must be governed by policies at both ends of the exchange.

Common FERA semantics at the two entities enable interpretation of security provisions for this type of

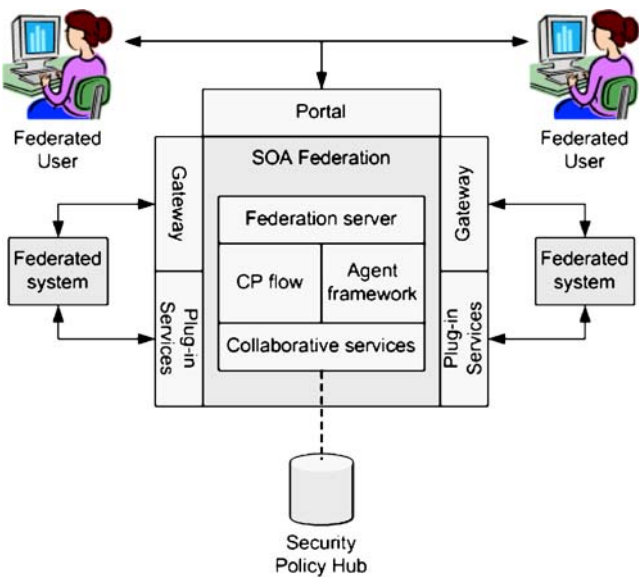

Fig. 3 A knowledge security policy hub located at each netchain entity. Adapted from Semantion (2005)

knowledge exchange, and the exchange pattern that is being used can be deployed according to those security provisions. Situated at the policy hub, then, are meta-security policy provisions for sharing a particular type of VCOR business process semantic (e.g., DESIGN) between two entities, say A and B, using a pattern, say pattern 2, from Table 1 above: personal interactions. In this way, not only can an entity C manage security for knowledge exchanges between A and B, but the process also supports recording the exchange in repositories at both ends, and C can apply, for example, non-disclosure agreements for the people involved before facilitating collaboration through the FERA gateway. Note that even in this binary exchange, C may stipulate that it is acceptable for B to share A's DESIGN exchange with some of its partners involved in marketing, i.e., indirect sharing through the MARKET semantics in VCOR (e.g., Fig. 4).

Meta-security policy provisions expressed in VCOR semantics for each FERA process pattern and deployed through FERA gateways can manage binary knowledge-sharing in a proactive way. But should those policies remain static, or can they be dynamically informed by events that take place to either reduce or increase levels of trust among netchain participants? Further, can the significance of those events be coordinated within a subset of a netchain where the impact of trust dynamics is most important? To address these questions, the policy hub discussed above must be augmented with gateway communications in order to handle breaches in security associated with specific binary knowledge-sharing arrangements. As each participating organization in the netchain possesses one of these 'security hubs', the validity of the shared knowledge received through the hub can be easily compared against the organization's set of knowledge-sharing arrangements. While the practical implementation of such a method may vary widely depending upon the context-specific needs of a given netchain, one generic approach would be to establish a 'fingerprint' for each shared knowledge 'object'—similar to a checksum in ordinary network file transfers—which would serve to uniquely identify the knowledge object throughout a netchain. As used herein, a 'knowledge object' refers to an encapsulated, structured set of knowledge that can be usefully applied within an organizational domain. When a knowledge security breach is detected (e.g., when an organization receives a knowledge object that does not correspond with any of its knowledge-sharing arrangements), it is the duty of the detecting organization and its collaborating netchain partners to broadcast the details of the breach. Each entity in the netchain must then assess how best to adjust meta-security policy provisions for the netchain linkages that ultimately led to the breach. To clarify, consider the slightly different view of a netchain shown in Fig. 5.

If the manufacturer has a new product design compromised through the release of information by a distributor that was never trusted in a binary arrangement, then that breach can be tracked through the relationships between organizations that are involved in trusted exchanges. The reduction in trust with the distributor responsible for the breach can then be relegated to more secure knowledge exchanges, governed by more restrictive non-disclosure agreements, or severed completely. Another example can be derived from the so-called "Small World" model (i.e., the social network theory which posits that all humans are connected to one another by a distance of no more than six intermediate acquaintances; Kleinberg 2000; Newman 2000). When this model is applied in the context of a fully

Fig. 4 Value Chain Operations Reference Model (VCOR) semantics and sub-processes for the execute process. Adapted from VCG (2005)

Market	Research	Develop	Source	Make	Sell	Deliver	Support	Return
K1 Analyze Market	C1 Define Opportunities	V1 Define Product Req.	S1 Identify Source	M1 Finalize Eng	L1 Target Customers	D1 Process Inquiry	U1 Register User	R1 Identify Return
K2 Analyze Performance	C2 Forecast Technology	V2 Select Technology	S2 Source Negotiation	M2 Sch. Production	L2 Quality Target	D2 Negot. Contract	U2 Manage Incidents	R2 Disposit. Return
K3 Define Needs	C3 Acquire Technology	V3 Design Product	S3 Schedule Delivery	M3 Issue Material	L3 Position Solution	D3 Enter Order	U3 Resolve Problems	R3 Request RGA
K4 Solution Ideation	C4 Validate Technology	V4 Optimize Processes	S4 Receive Product	M4 Produce & Test	L4 Inform Customer	D4 Receive Depot	U4 Educate Users	R4 Schedule Return
K5 Business Case	C5 Protect Technology	V5 Validate Product	S5 Verify Receipt	M5 Package	L5 Develop Relationship	D5 Pick Product	U5 Deliver Services	R5 Authorize Return
K6 Validate Opportunity	C6 Transfer Technology	V6 Define Life Cycle	S6 Transfer Product	M6 Stage Product	L6 Assess Needs	D6 Ship Product	U6 Monitor Experience	R6 Receive Return
K7 Create Roadmap	C7 Capitalize Technology	V7 Launch Product	S7 Authorize Payment	M7 Release Product	L7 Develop Proposal	D7 Verify Receipt		R7 Verify Condition
K8 Create Market Plan					L8 Complete Sale Cycle	D8 Install & Test		R8 Transfer Return
					L9 Close Contract	D9 Invoice		R9 Replace or Credit
					L10 Win/Lost Review			R10 Dispose or Recover

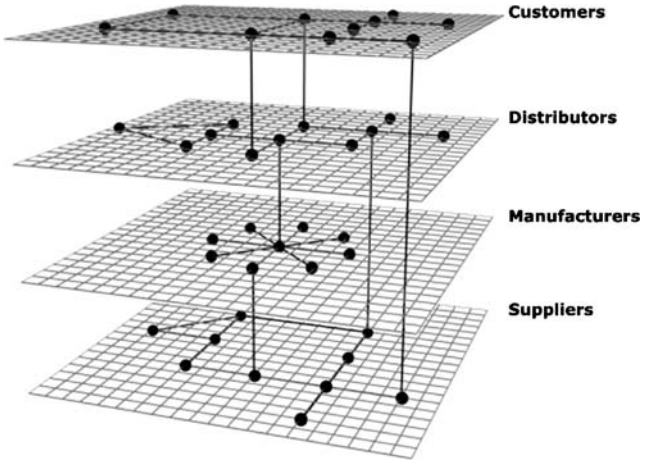

Fig. 5 A different view of a netchain

interconnected SOA business world, it becomes apparent that a manufacturer could be linked to its closest rival by a very short distance—perhaps only three or four hops on the netchain. If the manufacturer is planning to roll out a new product pricing scheme involving its distributors and customers, for example, a breach in the interorganizational knowledge-sharing netchain might lead the manufacturer's closet rival to have knowledge of the new strategy prior to its public release. Following the detection of such a breach, the manufacturer would clearly want to protect any further erosion of its competitive position by updating its security policies regarding the trustworthiness of the culprit. While the capability to dynamically manage this type of netchain-based security provisioning is dependent upon the integrity and completeness of the netchain model, the value of the model in providing for dynamic knowledge exchange is significant to sustaining a long-term competitive position.

3 Simulation

In order to evaluate the potential of the interorganizational knowledge sharing security model described above, network flow theory (Ahuja et al. 1993; Gross and Yellen 2005) was used to inform the design of a software-based simulation, the goal of which was to assess the model's breach-detection efficacy under varying conditions. Specifically, the software program was designed to simulate a dynamic interorganizational knowledge sharing environment in which the total size of the netchain, the maximum number of business partners per organization, and the maximum allowable knowledge sharing distance were allowed to vary under pseudo-random conditions. A knowledge sharing breach was then randomly introduced into the simulated environment, and the time required before the breach was detected was measured.

To facilitate understanding of the simulation process, Fig. 6 below illustrates a simple knowledge-sharing environment. Each circle or "node" in the figure represents an organization that belongs to a randomly generated netchain, with the dark lines in the figure representing the way in which the organizations are interconnected (i.e., the netchain). The organization from which the shared knowledge object originates appears in the center of the figure (labeled "Source", per network flow theory), and the concentric circles that radiate from the origin represent the number of knowledge-sharing hops or "distance" from this organization. In this particular example, the source organization has two directly connected business partners (labeled "1A" and "1B"). Knowledge transmitted from the source can therefore be directly shared with organizations 1A and 1B, but must minimally travel one additional hop through the network in order to reach organization 2B.

For the balance of this example, let us assume that the source organization has specified a maximum allowable knowledge sharing distance of two hops. This means that organizations 1A and 1B are allowed to share the knowledge object that they received from the source with their immediate business partners, with the understanding that the knowledge should not be transmitted further. In the parlance of network flow theory, this situation can be described as the source making a request of the nodes two hops away (via the intermediary nodes) that they act as "sinks", and simply absorb the shared knowledge rather than passing it on. Any organization that is within one or two hops of the source is thus allowed to possess the shared knowledge object, while all of the other organizations in the

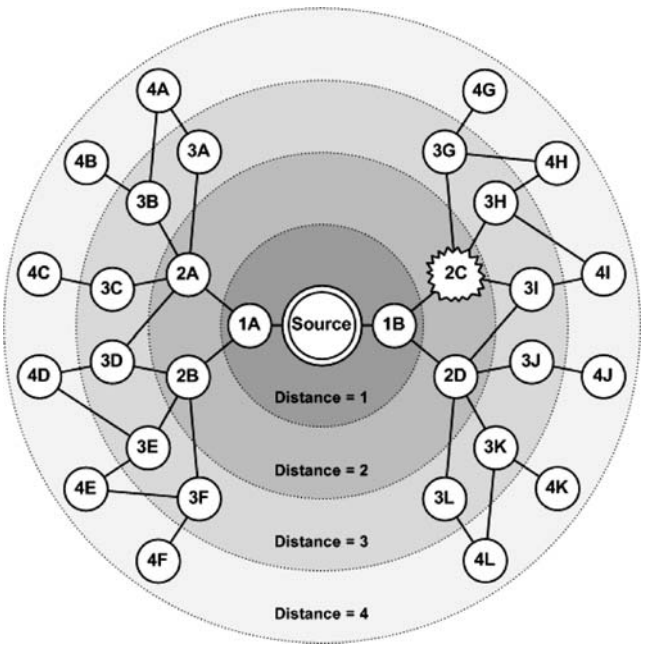

Fig. 6 A simple simulated knowledge-sharing environment

netchain are not. A knowledge-sharing breach occurs when one of the "sink" organizations on the periphery of this "circle of trust" shares the knowledge object with one or more of its business partners, thereby violating the will of the source organization. In the example above, let us assume that organization 2C breaches the agreement and shares the knowledge object with organizations 3G, 3H, and 3I (we assume that 2C does not share the knowledge with 1B, as 1B already possesses the knowledge object). If the organizations that now "illegally" possess the knowledge follow suit and share the knowledge object with their immediate business partners, how many hops along the netchain must the knowledge travel before the breach is detected? In our example, organization 3G would share the knowledge with 4G and 4H, organization 3H would share the knowledge with 4H and 4I, and organization 3I would share the knowledge with 2D and 4I. As organization 2D is within the source organization's "circle of trust" and already possesses the knowledge object, the knowledge-sharing breach would be detected as soon as 2D received the knowledge object from organization 3I. In this example, then, the breach was detected within two hops from the time that it was created, thus ending the simulated scenario. It is important to note that per network flow theory, two hops are minimally necessary for detecting any knowledge-sharing breach, regardless of the netchain size, the knowledge sharing distance, or the number of business partners.

3.1 Simulation methodology

In a similar fashion to the example described above, a total of 15,000 simulations were carried out in order to assess the breach detection potential of a dynamic interorganizational knowledge-sharing environment. The total size of each simulated netchain was varied from 10 to 500 organizations (in increments of 10 organizations), with 30 simulations being run for each size. For each simulation, the maximum knowledge sharing distance was allowed to vary randomly between one hop and three hops. The maximum number of business partners per organization was also allowed to vary randomly, with the only constraints being that the number of business partners could not be fewer than two or greater than half of the total size of the netchain. Thus, an organization participating in a 300-member netchain could have no fewer than two directly connected business partners, and no more than 150 directly connected business partners. Following the construction of each netchain, a "source" organization was chosen randomly, after which, depending upon the maximum knowledge-sharing distance, a random peripheral "sink" organization was chosen to act as the source of the knowledge-sharing breach. For each of these four random parameters, values were selected from a bounded uniform distribution so as to ensure that each value within the stated bounds had an equal chance of being chosen. With these parameters in place, the propagation of the knowledge-sharing breach throughout the netchain was simulated, and the time until the breach was detected was measured. It is this breach detection time (measured in "hops") that serves as the dependent variable in the results described below.

3.2 Simulation results

Each of the 15,000 simulations described above yielded four output values: (1) the total size of the netchain, (2) the maximum number of business partners per organization, (3) the allowable knowledge-sharing distance, and (4) the time required before the breach was detected. Descriptive statistics for each of these values are provided in Table 2 below.

As shown in the table, the average time required to detect a breach across all 15,000 simulated interorganizational knowledge-sharing environments was 2.24 hops. Given that two hops are minimally necessary for detecting any knowledge-sharing breach (per network flow theory), this result implies that knowledge-sharing breaches can on average be detected quite readily, regardless of the netchain size, the knowledge sharing distance, or the number of business partners per organization.

For predictive purposes, a linear regression analysis was undertaken to determine the extent to which variations in the simulation parameters impacted the observed breach detection time. All three of the parameters that were varied during the simulation process proved to be significant in predicting breach detection time, as detailed in Table 3 below.

As shown in the table, an increase in the size of the netchain was observed to yield a very slight—but nevertheless significant—increase in the time required to detect a knowledge-sharing breach. Conversely, an increase in the maximum number of business partners or an increase in the knowledge sharing distance were both observed to yield a decrease in the time required to detect a knowledge-sharing breach. From an interpretive perspective, these findings reveal statistically what one may expect intuitively: the time

Table 2 Simulation descriptive statistics

Parameter	*N*	Min	Max	Mean	Standard deviation
Netchain size	15,000	10	500	255.00	144.31
Maximum business partners	15,000	2	250	54.61	51.42
Knowledge sharing distance	15,000	1	3	1.46	0.63
Breach detection time	15,000	2	34	2.24	1.17

Table 3 Output of linear regression model predicting breach detection time

Parameter	Beta	Standard deviation	Sig
(Regression constant)	2.411	0.032	<0.001
Netchain size	0.001	<0.001	<0.001
Maximum business partners	−0.005	<0.001	<0.001
Knowledge sharing distance	−0.077	0.027	<0.01

required to detect a knowledge sharing breach is determined by the number of organizations that "legally" possess the shared knowledge relative to the overall size of the netchain. As the number of business partners grows or the knowledge sharing distance expands, so too does the number of organizations that can detect and report a breach. The paradox, of course, is that increasing the number of organizations trusted with the shared knowledge simultaneously increases the number of organizations that could potentially violate that trust.

Further insight into the relationship between the structural characteristics of the netchain and the knowledge-sharing breach detection time can be gained through an examination of Fig. 7 below. In the figure, a three-way median split was used to subdivide the size of the netchain into three groups (i.e., large, medium, and small), thereby allowing the impact of the netchain size on the relationship between the breach detection time and the number of partner organizations to be better understood.

As shown in the figure, the stability of the relationship between the number of partner organizations and the breach detection time decreases as the size of the netchain

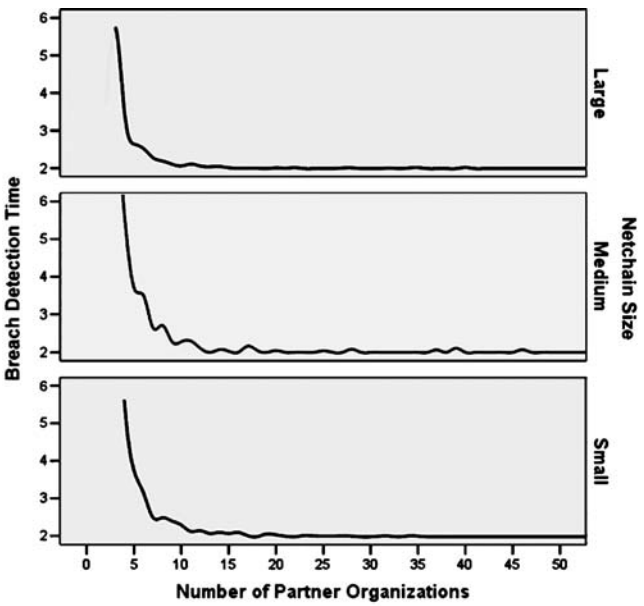

Fig. 7 Relationship between breach detection time and number of partner organizations by netchain size

increases. Despite these differences in stability, the relationship in all three groups closely approximates a negative exponential function, and becomes reasonably stable near the minimum breach detection time of two hops when approximately twelve business partners are directly connected to the source organization. With respect to mitigating risk, this finding implies that organizations with at least twelve business partners may be well-suited for participation in an interorganizational knowledge-sharing environment, regardless of the overall size of their netchain, and regardless of how far they allow their shared knowledge to travel. A closer examination of this finding is provided in Table 4 below.

As shown in the table, nearly 99% of all knowledge-sharing breaches are detected within the minimum of two hops when twelve or more business partners are present, with 100% of breaches being detected within four hops. Conversely, when six or fewer business partners are present, the breach detection efficacy drops to 39.8% and 73.6%, respectively. In light of these findings, organizations should carefully consider where they are positioned within their particular netchain before agreeing to join a knowledge-sharing alliance, especially if they have fewer than twelve directly connected business partners.

4 Schema to support a data warehouse-coupled knowledge security policy hub

All of the linkages between netchain entities need to be stored and managed in a tightly coupled data warehouse system in order to contain security breaches in a timely fashion. In this section, we extend the discussion of the exemplar infrastructure and the simulation with a multidimensional database schema design, which is depicted in Fig. 8. The schema contains tables for capturing the relationships between nodes from the perspective of the policy hub owner. For example, if an enterprise has agreed to share knowledge with a partner through one or more of the patterns discussed previously, then an entry for the link with each of those partners will be included in the NETCHAIN_LINK_STATE table. The PATTERN table

Table 4 Breach detection efficacy by number of business partners

Number of business partners	Percentage of breaches detected		
	After two hops	After three hops	After four hops
Between 2 and 6	39.8%	61.1%	73.6%
Between 7 and 11	79.5%	95.2%	99.2%
12 or more	98.7%	99.8%	100.0%

stores information about the pattern relevant to a link. The NODE and PLANE tables record the details of the nodes that exist on the various layers (planes) in the distributed netchain, e.g., suppliers, manufacturers, distributors, and customers. The POLICY table stores information regarding the internal and external security policies that exist between nodes in the netchain. The VCOR_SEMANTIC table stores semantics for each of the sub-processes and processes. The SNAPSHOT table includes a blueprint of the structure of a netchain model at a given time. The BREACH table captures historical details relevant to a breach that may be useful in containing future breaches. In short, when an organization enters into a knowledge-sharing relationship with another organization, a Start-Node and End-Node are added to the NETCHAIN_LINK_STATE table, the Policy-ID related to that link is recorded, the Plane of the organization being linked to is recorded, the relevant Pattern is entered, the VCOR semantics related to that link are recorded, and a snapshot of the current structure of the netchain is recorded. For complex knowledge-sharing arrangements such as when a company (say A) grants permission for another organization (say B) to share knowledge with B's trusted partners (say C, D, and E), then A must record in its policy hub each of these indirect linkages with B as the Start Node and C, D and E, respectively, as three distinct links. It is important to note that the schema of Fig. 8 shows a recursive relationship relevant to entries in the NETCHAIN_LINK_TABLE to reflect this type of complex knowledge-sharing arrangement. It is also important to note that each policy hub associated with an organization in a trusted federation of partners will be unique, but that subsets of recorded relationships would be common, even while planes, patterns, policies, snapshots, and breach entries would be distinct.

Now, we consider the situation of a knowledge-sharing breach in light of this security hub. When a breach is discovered by any partner in the federation, information about that breach must be broadcast to all relevant partners. Upon receipt of the breach information, an organization can query its security hub to ascertain the potential causes of the breach, to assess the policies relevant to mitigating breaches of this type in the future, and to examine the potential for the source of the breach to cause loss of security over other knowledge-sharing arrangements. It can also guide in the development of a breach impact mitigation strategy. If an organization is able to discover or find a limited subset of organizations that could have caused the breach through its search of its policy hub data, then that information would also be broadcast to the federation partners. Of course, the federation may need to develop specific policies such as when a root cause is traced to a federation member. Upon receipt of broadcast information regarding the limited set of potential culprits, another node may possess, through analysis of its policy hub data, sufficient information to further limit that subset. The ability

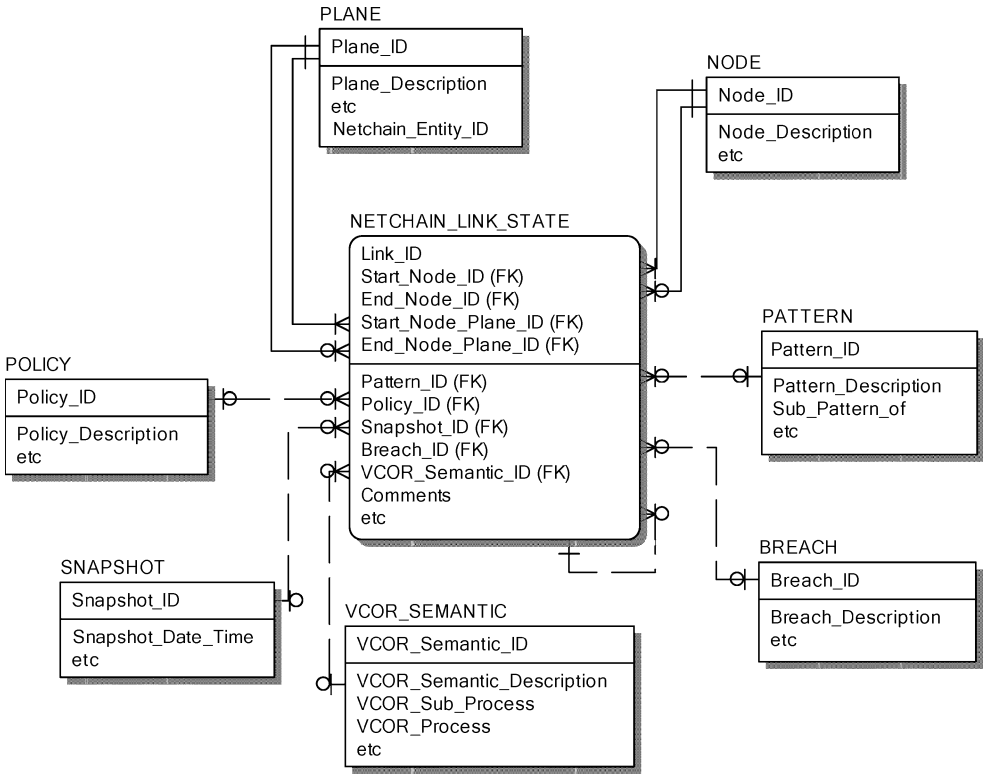

Fig. 8 A high-level entity/relationship model of a security hub

to ferret out the source of a knowledge-sharing breach in this way is a significant advantage to being a member of the federation—even while each individual organization is clearly in charge of maintaining and implementing its own, independent knowledge-sharing policies.

The approach we have described relies upon on a business intelligence infrastructure coupled with a multidimensional data warehouse to capture knowledge-sharing relationships that take into account various patterns of interorganizational collaboration. A compelling reason for adopting this approach is that many organizations already possess this technology and the expertise to manage it effectively, thereby limiting the economic cost of improving information security (Gordon and Loeb 2006). Policy hubs located at each node of a federation that maintain information about the netchains within which a node is engaged to conduct business provide a level of proactive, breach directed, knowledge-sharing policy revision capability to the members of the federation. The ability to capture states of the netchains in various snapshots could also enable more complex analyses (e.g., data mining) to discover patterns of breach combinations that may usefully inform system design improvements within context-specific netchain environments.

5 Summary, limitations, and future work

The need for new collaboration and infrastructure models relevant to today's complex and rapidly evolving business world is apparent. Novel aspects of suitable models must reflect value-chain based collaboration needs and requirements, and must carefully consider security issues. Our approach to knowledge-sharing security in netchains is targeted at proactive governance of binary knowledge exchanges through process patterns supported by the emerging FERA-based ebSOA standard. In addition, the meta-policy provisions maintained in our policy hub approach extend value-chain collaboration advantages to the realm of enhanced security. By relying on VCOR semantics, specific provisions can be applied to knowledge-sharing both vertically and horizontally within interorganizational netchains. By using dynamic adaptation mechanisms to handle breaches as facilitated through netchain analyses, the security model can be adapted as necessary by netchain partners. Given that the simulation described herein indicated that these knowledge-sharing breaches can be readily detected in most netchain environments, we are confident that much of the work necessary to identify the source of a breach can be automated, as can the exchange of revised trust profiles through closely collaborating netchain subsets. Thus, a properly designed system may serve to facilitate the establishment of long-term trust relationships between business partners—an effect that has been referred to as 'trust by design' (Keen et al. 2000).

The interorganizational knowledge-sharing security model described herein has several limitations that arise from the assumptions upon which it is built that should be acknowledged. First, as the model relies upon ebSOA-encoded binary knowledge exchanges, breaches involving knowledge that is encoded differently or is transmitted clandestinely through alternate channels would not be detectable. Unless all of the codified knowledge that is shared among netchain partners can be processed by the netchain's security hubs, this limitation may constrain the applicability of the model in its current form to knowledge-sharing breaches that are unintentional in nature. As tested in the simulation, the model also assumes that knowledge which is allowed to be shared among partner organizations at a particular distance from the originating organization will be shared among all of the organizations at that distance. In a real-world implementation, it is likely that knowledge would be made available only to those business partners who can benefit from possessing it; our future work will therefore seek to relax this assumption, thereby increasing the realism of the simulated knowledge-sharing environment. We also plan to enhance the sophistication of our simulation in the future by incorporating interorganizational trust metrics (Ziegler and Lausen 2005) and automated culprit detection. Additionally, we will seek to assess the impacts of alternative indirect breach relationships, and will endeavor to determine the best methods for supporting meta-policy updates and revisions in a dynamic setting.

Acknowledgements The authors would like to express our gratitude to the guest editors of the special issue and to the three anonymous scholars who reviewed our work; your efforts and insightful comments improved this paper significantly.

References

Ahuja, R. K., Magnanti, T. L., & Orlin, J. B. (1993). *Network flows: Theory, algorithms, and applications.* Upper Saddle River, NJ: Prentice Hall.

Brown, G., & Carpenter, R. (2004). Successful application of service-oriented architecture across the enterprise and beyond. *Intel Technology Journal, 8*(4), 343–360.

CPDA (2004). *Integrated process and technology framework.* Livermore, CA: Collaborative Research Services/Collaborative Product Development Associates.

Drecun, V., & Brown, D. H. (2004). *Closing the process/technology gap FERA.* Livermore, CA: Collaborative Product Development Associates.

Eisenhardt, K., & Martin, J. (2000). Dynamic capabilities: What are they? *Strategic Management Journal 21*, 1105–1121.

Garg, A., Curtis, J., & Halper, H. (2003). The financial impact of information technology security breaches: What do investors think? *Information Systems Security, 12*, 22–33, March–April.

Gordon, L. A., & Loeb, M. P. (2006). Economic aspects of information security: An emerging field of research. *Information Systems Frontiers, 8*, 335–337.

Gordon, L. A., Loeb, M. P., Lucyshyn, W., & Richardson, R. (2006). *Eleventh annual CSI/FBI computer crime and security survey.* San Francisco, CA: Computer Security Institute.

Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal, 17*, 109–122.

Gross, J. L., & Yellen, J. (2005). *Graph theory and its applications* (2nd ed.). Boca Raton, FL: Chapman & Hall.

Hardy, C., Phillips, N., & Lawrence, T. B. (2003). Resources, knowledge and influence: The organizational effects of interorganizational collaboration. *Journal of Management Studies, 40*(2), 321.

Keen, P., Balance, C., Chan, S., & Schrump, S. (2000). *Electronic commerce relationships: Trust by design.* Englewood Cliffs, NJ: Prentice Hall.

Kleinberg, J. M. (2000). Navigation in a small world. *Nature, 406*, 845.

Lazzarine, S. G., Chaddad, F. R., & Cook, M. L. (2001). Integrating supply chain and net-work analyses: The study of netchains. *Journal of Chain and Network Science, 1*, 7–22.

Majchrzak, A. (2004). *Human issues in secure cross-enterprise collaborative knowledge-sharing: A conceptual framework for understanding the issues and identifying critical research.* Los Angeles, CA: Center for Telecommunications Management.

Newman, M. E. J. (2000). Models of the small world. *Journal of Statistical Physics, 101*(3/4), 819–841.

OASIS (2006). *Electronic Business Service Oriented Architecture: Advancing architectural patterns for using Service Oriented Architecture in electronic business.* Billerica, MA: Organization for the Advancement of Structured Information Standards.

Scott, J. P. (2000). *Social network analysis: A handbook* (2nd ed.). Thousand Oaks, CA: Sage Publications.

Semantion (2005). *Run-time service oriented architecture.* Toronto, Ontario, Canada: Semantion, Inc.

Semantion (2006). *FERA-based SOA*. Toronto, Ontario, Canada: Semantion, Inc.

Sharda, R., Frankwick, G. L., & Turetken, O. (1999). Group knowledge networks: A framework and an Implementation. *Information Systems Frontiers, 1*(3), 221–239.

Shih, S. C., & Wen, H. J. (2003). Building e-enterprise security: A business view. *Journal of Information Systems Security, 12*(4), 41–49.

VCG (2005). *The Value Chain Operations Reference (VCOR) model.* Wexford, PA: Value Chain Group, Inc.

Weill, P., Subramani, M., & Broadbent, M. (2002). Building IT infrastructure for strategic agility. *Sloan Management Review, 44* (1), 57–65.

Ziegler, C. N., & Lausen, G. (2005). Propagation models for trust and distrust in social networks. *Information Systems Frontiers, 7*(4/5), 337–358.

Daniel S. Soper is a Ph.D. candidate in the Department of Information Systems in the W. P. Carey School of Business at Arizona State University. His research interests focus on the impact of information and communication technologies in developing and emerging societies, interorganizational knowledge-sharing security, negotiation support systems, and quantitative research methods.

Haluk Demirkan is an Assistant Professor of Information Systems in the W. P. Carey School of Business. His teaching and research interests and expertise are in services-centric computing and the management of outsourcing relationships. His research leverages his multi-disciplinary educational background and extensive professional industry experience in the fields of information logistics and strategic business engineering. Haluk holds a Ph.D. in Decision & Information Sciences; M.E. in Industrial & Systems Eng. from the University of Florida, and B.S. in Mechanical Eng. from Istanbul Technical University.

Michael Goul is a Professor of Information Systems. His research has been published in a wide range of both academic and practitioner journals, and his recent research interests are in the areas of services computing, social networking applications in value chain security models, master data management and information systems architecture. He has published in journals including Decision Support Systems, Journal of Management Information Systems, Decision Sciences, Information & Management, IEEE Expert, Data Base, Communications of the ACM, Information Systems Frontiers and other journals. He has served as journal editor, special issue editor, AIS Vice President, Conference and Program Chair, and Chair of the AIS special interest group in decision support, knowledge and data management systems.