Automating governance of open service ecosystems and conformant adaptive collaborations

Lea Kutvonen

Helsinki May 2015

Technical report C-2015-1

UNIVERSITY OF HELSINKI
Department of Computer Science

HELSINGIN YLIOPISTO — HELSINGFORS UNIVERSITET — UNIVERSITY OF HELSINKI

Tiedekunta — Fakultet — Faculty	Laitos — Institution — Department
Faculty of Science	Department of Computer Science

Tekijä — Författare — Author

Lea Kutvonen

Työn nimi — Arbetets titel — Title

Automating governance of open service ecosystems
and conformant adaptive collaborations

Oppiaine — Läroämne — Subject

Computer Science

Työn laji — Arbetets art — Level	Aika — Datum — Month and year	Sivumäärä — Sidoantal — Number of pages
Technical report C-2015-1	May 2015	15 pages

Tiivistelmä — Referat — Abstract

In the last decades, private and public organisations have started to expect agility in service interoperability from each other. However, automated governance support is still an open challenge for all societal structures that involve independent organisations. Moreover, the topical ambition for customer-centric service design arises needs for tailored collaboration management, too.

This paper proposes an automated governance support solution for ecosystem architectures. It is based on our earlier work on Pilarcos ecosystem architecture and shows that it is feasible to automate the societal sector control across independent organisations. An example study illustrates i) how healthcare sector budget boundaries and equal right for timely care can be controlled in individual patient care processes across several organisations, and ii) how collaboration case knowledge can be collated for healthcare district governors for improved regulations, budget frames and ecosystem properties.

ACM Computing Classification System (CCS):
Computing methodologies –> Artificial intelligence –> Distributed artificial intelligence,
Information systems –> Information systems applications –> Decision support systems,
Applied computing –> Enterprise computing –> Business process management

Avainsanat — Nyckelord — Keywords

inter-enterprise collaborations, case management, service ecosystem, automated regulatory governance, business-IT-al

Säilytyspaikka — Förvaringsställe — Where deposited

Department of Computer Science library, C-2015-1

Muita tietoja — övriga uppgifter — Additional information

1 Introduction

The recent decades have witnessed an increased call for networked business and cross-organisational collaboration between public and private domain partners. For that call, a variety of service ecosystem and virtual organisation breeding environments have been introduced to provide production environments for collaborative business services, as well as governmental or communal services. This development has largely been enabled by the evolution of service-oriented computing (SOC), business process management (BPM) and adaptation of Web services as a common technology.

The interest on inter-enterprise collaboration support has become evident through research (including virtual organisation breeding environments (VOBE) such as ECOLEAD [3], ADVENTURE [35], Collaborative networks [5], and CrossWork [23]; service development and execution platforms such as ATHENA [2], COIN [7], and NEFFIC [9]; supporting cloud environments like IoT [34], S-CUBE [1]; and service ecosystems like TrustCom [41], and Pilarcos [18]). A significant amount of vendor support has gained ground, such as Microsoft Azure, HP's Enterprise (Web) services, IBM WebSphere and intelligent processes, Google cloud platform and Amazon web services.

At the same time, case management (CM) [24, 11, 27] has brought the focus on an area where of individualised business process management through tools capable of accessing large amounts of data through a wide number of information sources. Knowledge workers is a new category of experts who analyse big data and innovate and design ad-hoc processes for new situations. The amount of data sources and data mining for strategy advisory and decision making have increased. Tools and information pathways are being built, for example, for cumulating information and making it accessible for healthcare, for governmental resource placement and market studies, and for understanding environmental consequences of particular industrial activities.

Analysis of VOBE (virtual organisation breeding environments), service ecosystem approaches, and case management brings up shortcomings in (automated) governance support. In this analysis, we use a generic understanding of company governance [39]: Governance is about exercising decision-making power in the system at hand. It means identifying the actors with power, their statement on intention, structure and strategy of the system, and their declaration on the norms to hold in all processes. Further, governance is involved with the delegation of subdomain governance rights to other accountable actors in the system. Open service ecosystems needs to agree, explicitly store and systematically apply its norms in order to align and automate its governance practices.

In the current digital society era, we should consider inter-enterprise collaboration an identified, unique collaboration case that has a dynamic lifecycle and need to be governed by its owners, partners, or delegated managers – and no more a final engineering product. The facilities for dynamic governance and control of conformance to relevant regulations is still falling short. The governance shortcomings include i) the lack of mechanisms that enforce regulations, such as national laws, on all commercial transactions; ii) the inability of customers, or other partners besides the key service provider, to tailor the involved business processes; iii) the lack of facilities to declare equality or fairness requirements over all service collaborations in the ecosystems. Further, the present solutions lack concepts and denotations for iv) obligations, permissions and prohibitions; and v) actions for delegation and its withdrawal, commitment, contracting, breaches and ecosystem level sanctions.

The research question of this paper seeks for feasible automated governance facilities for ecosystems comprised of independent organisations, with capability to adapt to dynamic regulations. In relation to that, we further seek ecosystem quality properties that enforce fair sharing of resources or even network reciprocity among member organisations.

This paper enhances our earlier work on eContract-governed inter-enterprise collaborations and service ecosystems [18] by facilitating the architecture with strategic governance of an ecosystem, and with conformance to the ecosystem and partner enforced requirements by all of its collaborations. In addition, the proposal facilitates the collection and collation of collaboration properties so they can be analysed at the ecosystem level against measurable goals. The collaboration partner and open service ecosystem governor activities are illustrated by a healthcare domain example that indicates their benefits. The association between ecosystem and regulatory domains is explicated.

Chapter 2 recalls the key characteristics of the Pilarcos ecosystem architecture and overlays it on our running healthcare example. Pilarcos architecture has been rigorously conseptualised and evaluated in our earlier works. Chapter 3 studies the governance metrics for services, service collaborations and service ecosystems, and proposes enhancement valuation metrics, before mapping them into the Pilarcos mechanisms. Chapter 4 elaborates the collaboration level governance policy retrieval and monitoring system enhancements. Especially, it discusses the management of an individual collaboration case. The discussion relies partially on our earlier work design on a new style of trusted business transactions [18, 16]. Further, it proposes a way for collecting collaboration case specific data for governance purposes using some CM practices. Chapter 5 introduces the new tools for ecosystem governance: collating data in to the ecosystem intelligence unit, and visualising causes of breaches in the ecosystem for change decision-management. Chapter 6 focused on the conclusions on ecosystem facility requirements and summarise the benefits and opportunities for various ecosystem members.

2 Healthcare case basics in Pilarcos ecosystem

This chapter introduces a fictitious healthcare ecosystem example, and relate it with the Pilarcos ecosystem key concepts.

The Finnish healthcare system (where a total restructuring on the political and technical architecture level is expected) is split into regions where funds are collected via taxes, and all citizens have low cost access to public healthcare services. Independent private sector services are more expensive, but at times districts need to use them to increase their care resources. Each district provides GP services or consultancy services that are available through GP control. Staff have access to stovepipe-style, in-house enterprise computing systems for medical records, digital imaging, or appointment books. A new centralised electronic prescription database is used for dispensing medications.

For simplicity of the example, we choose to model each basic or special healthcare district as an organisation in a single national ecosystem. Similar constructs can be used in a recursive manner to reach more rigor. Some major Pilarcos design choices support this kind of approach well: First, the ecosystem maturity level [17] requires that each collaboration is controlled by an eContract and that the used business processes are designed in terms of abstract roles enabling dynamic changes of service providers in those roles at operational time. Second, an extensible metamodel hierarchy [33] of ecosystem repositories is required to ensure contractual correctness and interoperability control at runtime in all collaborations. Third, there is no requirement for mutual business process execution engine, which allows freedom in computing technology selection.

The Pilarcos architecture [18, 32, 36, 31] views the *service ecosystem* as an environment – open, but strictly governed service market – where service providers and clients can meet, establish contract-governed inter-enterprise collaborations and gain experience on the business services and partners involved. A *business service* is a software-supported service with a functionality suitable for a business need on the market and thus relevant for the networked business. Each business service is an agent, in terms of being able to take initiative on some activity, being reactive to requests by other business services, and being governed by policies set by its owner.

An *inter-enterprise collaboration* is a loosely-coupled, dynamic composition of business services; it involves multiple partners through their business services and their mutual interactions. The type of the service collaboration is declared as a *business network model* (BNM), expressed in terms of the roles and interactions within the collaboration, the involved member services, and policies governing the joint behaviour. In technical terms, each BNM is a set of business processes attached to each other by requiring particular roles from adjacent processes to be fulfilled by the same service.

For healthcare, we can associate each inter-enterprise collaboration type with a patient care plan. The Current Care Guidelines (käypä hoito) [8] is an independent declaration of best practices. It provides partial advise on arranging workflows and timescales for each diagnosis. While it does not provide a full declaration, an approximate care process model can be concluded from it and stored as a BNM. In a care plan example, say for hip replacement surgery, the patient is first screened by a GP, placed in a surgery queue with a 3 or 6 mths service time associated.

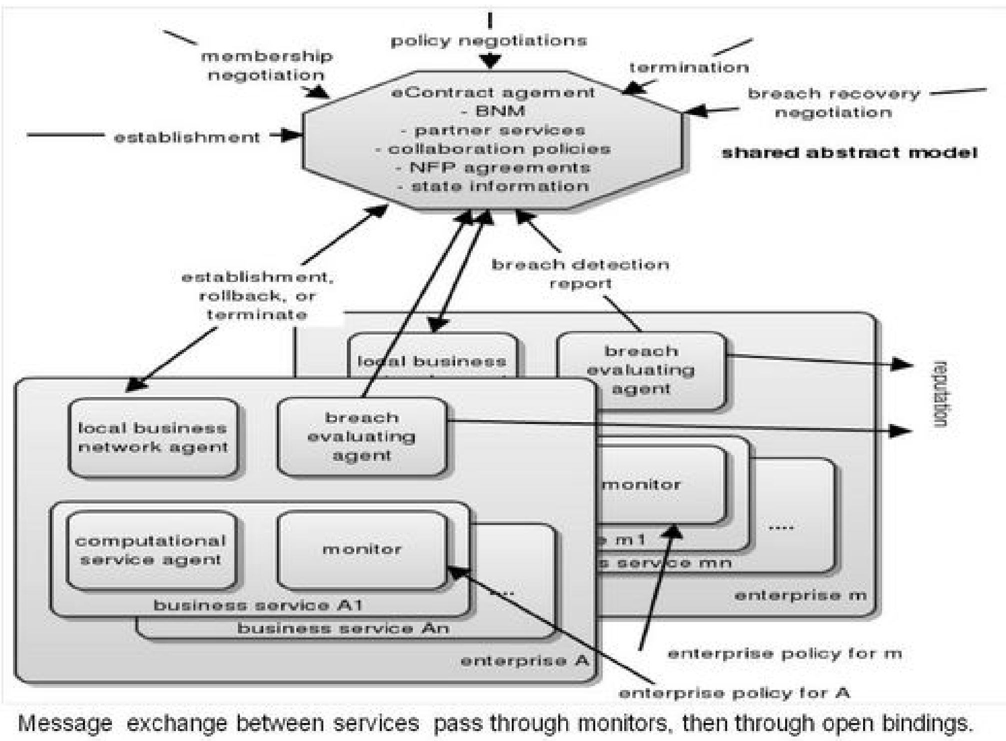

Figure 1: The eContract reflective control of inter-enterprise collaboration of business services.

An invitation letter to the operation arrives at the patient's address three weeks before the operation, to state the operation date and instructions on laboratory testing a week before. After the surgery the staff seeks a nursing home place for a few weeks, for physical therapy and daily care, or provides physical therapy advice at the patients home. The involved business services are the hospital stay for the surgery, the laboratory works, nursing home services, and therapy facilities. While most of the activities in these services are performed by humans, it is the information and control exchanges between organisations that define the BNM. Services may be organised by independent districts.

Our main interest lays in the governance and conformance of the individual *collaboration cases*. An example of such a collaboration case involves an elderly person, Anna. She is in need of hip surgery in the Helsinki area, and rehabilitation therapy in a nursing home. In addition, she been placed in custody of her daughter due to a degree of dementia.

The inter-enterprise collaboration is controlled by an *eContract agent*, both on business and technical levels, as shown in Figure 1. It is structured according to a BNM, and for each role there is knowledge of the assigned business service, with its location, communication channel requirements and expectations on nonfunctional properties. Further, the eContract carries the agreed policies that restrict the collaboration's behaviour from what is generally possible based on the BNM. The eContract includes monitorable expressions of the agreed behaviour. The eContract agent also provides interfaces for collaboration members for changing the collaboration structure, policies, membership and other properties, as well as reporting progress of collaboration milestones. The eContract also includes process descriptions for breach recovery situations and operations for changing the collaboration behaviour during its lifetime.

In the healthcare example, the eContract means a digitally accessible care contract and an interface for change management. As the contract is based on the care process model, it helps patients in orienting themselves for the procedures. which leads to better results in general. Of course the process model has alternative versions embedded, to address the mere suspect of a diagnosis, to actual treatment, and even complicated cases where multiple care processes meet. The individuality of each case is essential so the care process can be negotiated face-to-face between caretakers and patients. As the personalised care contracts are always updated at appointments, it would also become possible to consider options of care service provision with different properties, depending on the personal needs and preferences of the patient. All initial parties (Anna and her daughter, GP, surgeon) might have reasons to adapt the care process: changes in Anna's health, need to reassign service due to surgeon schedule, or adding services for the plan after the surgery, such as the therapy and nursing home. For these purposes the expressions in deontic terms (obligation, permission, prohibiition) are well suited: reassigned service is prohibited from being too far, surgery is refused in the presence of bad laboratory results, or care home placement is permitted with the condition of receiving a voucher from her home town.

From the healthcare governor point of view, the care contract provides new opportunities. The monitoring of the care cases enables the govenor to catch cases where schedules are delayed too much, for example. As the care process are explicitly stated and personalised, the governance tools can suggest alternative care services that can be offered by public and private service providers in the district (see service offer repositories). Their offer combinations may be compared to the care needs in question, and furthermore, compared also to the ecosystem governor's strategical advise as part of the decision-making process. The doctors and nurses thus act as CM knowledge workers when selecting service providers through a new interface that provides ranked suggestions from ecosystem infrastructure repositories.

An essential part of the ecosystem is its *ecosystem infrastructure* that contains shared utilities for enterprises to discover and select services available in the ecosystem, negotiate and establish collaborations, govern those collaborations through eContract agents, and utilise reputation information and collaboration type information. The infrastructure contains two groups of basic agents and knowledge repositories, local and global, for the purposes of correct behaviour of the collaborations. The local agents address the private needs of organisations, including decision-making in contract negotiations, and decisions on trust and privacy preservation. The global level includes repositories of service offers, business process definitions, and service types, as well as agents for utilising their knowledge for enforcing composite service correctness. These global services act on public information.

For healthcare, the ecosystem infrastructure repositories can be utilised as follows: i) service offer repositories as declarations of healthcare services and available resources on each GP and consultancy house in particular; ii) service type repositories as declarations of categories of information exchange plans related to each diagnostic or treatment step; and iii) business network model repositories as declarations of categories of care processes defining the roles (responsiblities in treatment) of care providers, and information and patient exchange between roles.

Although most healthcare system do not support eContracts, in the district of Forssa, a client-centric way of organising work has been developed. It involves care teams of doctors and nurses who form an approximate plan (sumea hoitopaatos [21] for each long term patient, so that the care can be better predicted and communicated between the involved parties. The care contract is also in alignment with the MyData framework [28] that encourages big data solutions where each citizen can manage the use of all public and private data on them. Many private healthcare units already now provide a web service through which clients can view their laboratory results and make appointments, also leading towards private key data and interfaces for control.

3 Governance challenges

In organisational sciences, governance, risk management and compliance there are the key concepts for assuring that organisations meet their objectives: Governance refers to the management approach the organisation leaders take for structuring the organisation into units, for selecting their work processes, for declaring organisational policies, and for using monitoring and performance indicators to optimise the results. Risk management aims at being prepared for potential obstacles and finding working methods to remove or mitigate the risks. Compliance refers to behaving in accordance to the organisations processes, policies and regulations, not forgetting that the organisation itself must comply to a legal systems, such as national laws, ethics, or best practices to be successful. The term conformance is more often used in a technical context - as a system being compliant with a technical standard - where conformance testing can be organised.

In this paper we focus on the technical solutions that at runtime try to assess that the contracted behaviour is actually met and the ecosystem targets are met. As part of that, the breach recovery behaviour must become activated in an appropriate situation – with breach frequency and type of breaches kept within agreeable limits. Thus for this study, the terms compliant and conformant become synonyms.

We start by studying the metrics of system behaviour estimation, as the question of useful and effective *system properties* is still open. The levels of service, collaboration cases, and ecosystems all need modern key performance indicators (KPI) instead of the traditional QoS metrics. Besides the indicators, suitable metrics and

service quality	suitability, completeness, accuracy of computing functionality, reliablity including failure avoidability and restorability, interoperability, security including access auditability, maintainability
collaboaration quality BP quality	financial value, productivity and efficiency, equality, reliability, informationquality, efficiency, process security and privacy interoperabilty on the technical, semantic and pragmatic levels, composability privacy, trust and willingnessto joint transactions, conflict resolution capabilities
customer satisfaction	service coverage, timeliness and responsiveness, service accessiblity, customer benefit psycology and cognition related concept of quality of experience (QoE)
organsiation quality KPI	number of business activities and amount of partners, customer loyalty, strategical partner compatibility, willingness and agility to adapt to new effective business strategies, timely resolution of differencies, collaborative planning capacity, involvement in knowledge exchange
ecosystem quality	fairness in trust evaluation (to newcomers, minorities), equity on collaboration opportunities, equity on QoS as customer, economic efficiency, fair benefit and cost distribution, environmental safety, system resiliance, ethics

Table 1: Governance metrics survey.

data flow instrumentation into the ecosystem infrastructure are necessary additions to the Pilarcos architecture. For that purpose, we need to review suggestions arising from existing research and missing aspects on the ecosystem governance area. The results of this consideration are shown in Table 1 where some constantly repeating metrics include service quality, customer satisfaction, process or collaboration quality (e.g., [6, 38, 30, 10, 22, 40]). Besides the traditional QoS metrics more psychology and cognition related concept of quality of experience (QoE) [20] is gaining ground and needs surveys and user studies for quantifying these subjective QoE factors (although physiological indicators can also be used, such as pulse changes). Another new theme to address, is to consider collaboration willingness and capability of negotiations between parties. We add to the traditional privacy and trust properties in collaboration, the availability of methods for conflict resolution. The resolution techniques should go beyond the traditional database transactions or compensation steps on business processes, especialy the techniques should involve use of new business processes at the ecosystem level. For assessing the success of the ecosystems, there are very few sources for metrics other than virtual organisation assessment metrics. Those focus around business level assessment of company assets, values, leadership and production processes, and reputation or brand value. For new metrics for ecosystem success, we have selected equality based properties that are technically more complex to collate.

For the healthcare example, some quality metrics can be picked from an open databank on Finnish and European statictis [26] that includes population and its health status, proportion of population having received certain treatments (associated with references to regulations that make it free), yearly budgets, and several efficiency and cost indicators. The statistics collators also seek for additional indicators.

However, this data focuses on the number of patients with an illness or disorder, instead of the whole care process. Therefore the statistics can only indicate serious misconducts in treatment. Breaches of the care processes are not visible. The success of the healthcare ecosystem does not have defined metrics and therefore the system goal remains unclear. The ecosystem goals and properties can partially be drawn from the Current Care Guidelines which aims for equal good service for all. However, one of the problems in Finland is that the metropolitan area has a lot of services available and everything is close by, while in the northern part of Finland, the nearest hospital could be hundreds of kilometers away, having its maternity ward open only on office hours (causing unintended home deliveries on weekends and public holidays). Computing the equality properties for maternity wards or ER units become very difficult, as distance-based delay may increase risks of complications and casualties, and thus potentially higher cost in consequent care plans than if there were a higher level of alertness. An artificial intelligence machinery could be used for producing casualty or complication risk statistics or probability charts over time.

System compliance is mostly considered to be a check that a system must go through at each of its lifecycle phases, from design, implementation and operation (e.g., [15]). The *compliance descriptor* (in [15]) is composed of compliance requirements linked with implementation-specific compliance rules and bound to entities like process activities or servers. The requirements are expressed in a generic compliance language. The compliance descriptors are created at the collaboration requirement analysis and design time. They are further refined into rule expressions that can be executed. In addition, laws and regulations are associated as structured documents located through URL-based references.

The compliance descriptor rules for Pilarcos do not form a single data unit. The corresponding conformance requirements are grouped differently: the service and collaboration type design and implementation steps produce models that are stored in the ecosystem repositories. At the collaboration establishment or renegotiation phase the eContract is associated with BPM with its policies, breach recovery processes, monitoring and breach detection rules, and policies received from each contracted partner. This way, rules become inherited from different design, declaration and governance processes as illustrated in Figure 2.

In the Pilarcos architecture, each ecosystem agent tier is permitted to use only some speech act categories. Only the ecosystem governors are allowed to declare ecosystem wide regulations, properties (norms) and policies, as well as allowed to accept the metamodel hierarchy rules for the ecosystem repositories [19]. For the governors, special interfaces are needed, with a family of aspect-related languages. The BNM design phase is separated and assumed to be an activity by consortia like standardisation committees. The ecosystem infrastructure repositories do not present plain ontologies, but are built according to MOF (meta-object facility) and MDE (model-driven engineering) principles [33]. Therefore, the repositories are able to restrict the acceptance of published models into those that are formally correct instances of the upper level modeling principles and also acceptable according to

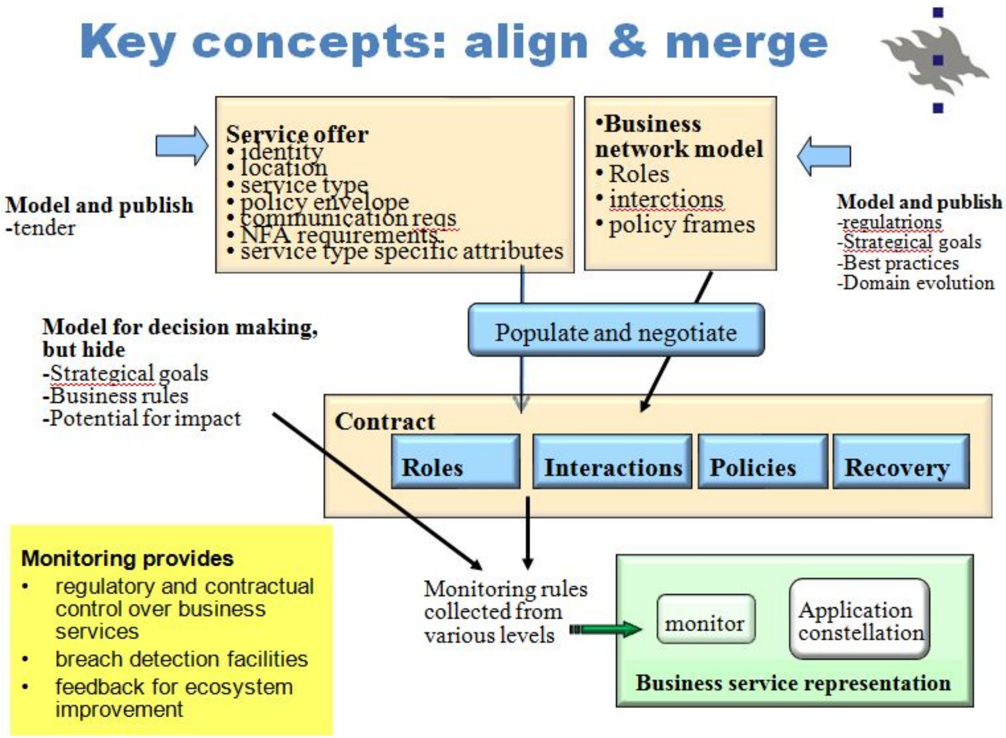

Figure 2: The eContracts inheriting conformance rules from models and policies.

the criteria set by the ecosystem governance for the repositories. Ontology based repositories would just allow reusability of units in registries. The ecosystem enacts only collaborations that are based on registered, acceptable BNMs. Similarly, the restrictions to service offer repository restricts the acceptability of business services in the ecosystem.

The policies need to be expressed in a small set of domain-specific languages, all of which have appropriate translations to effective, simple monitoring rules for run-time. There is no complete conflict checking in peer organisation policies, for several reasons. Firstly, the organisational policies are to be kept private and they are also independently constructed so there is no guarantee that conflict checking would reveal essential details. Second, the policies are subject to change and therefore focusing on catching policy discrepancies at runtime is more efficient, especially as the breach detection and resolution mechanisms fit also these situations.

For the various definitions, expressive power per language category is needed as follows:

- For deontic logic policies, we need expressions of obligations, prohibitions, permissions, delegation, withdrawal of delegation, criteria over service request by target identity, role identity, service type, and computations over message field, or computations over a sequence of messages, timeliness, and conditions on refusals, delays, or missing information. The basic ideas of the expressions required are fairly similar to [13] although the computing is arranged by first translating monitoring rules for pre- and postconditions of operations, instead of using Lisp like computing.

- Process models criteria include partial ordering, logical expressions on restrictions or requirements on the relationships of identified services in certain roles, timeliness, liveness, KPI introduction, and deontic logic expressions over roles, control flows and information flows.
- Model relationship expressions are required to express denoting similarity, replaceability, and loose matching.
- Valuation expressions for services or collaborations include KPI computations using sliding averages and occurrence frequencies, earnings and losses (monetary, reputation) according to a fairness measure between the partners, and rules for experience reports for reputation systems.
- Valuation expressions for equality measures where comparison of categories of collaboration cases should have small enough statistical variation.

The more elaborate ecosystem governance tools are especially beneficial for policy makers, politicians and their technical support groups, and governments establishing the yearly budgets. For budget decisions, or for allocation of resources to particular districts, the governing team has to study the ecosystem behaviour on the preceding period and the changes on the behavior caused by the latest modifications on the rules. The data analysis and visualisation of the ecosystem behaviour as a whole is essential for governors to understand the ecosystem properties, quality and opportunities for improvement. Especially the category of collaborations with some breach situtaions in them should be analysed, in order to be able to avoid them in future. The proposed approach opens a new opportunity to find violations against equality. For example in ambulance service times, two axis are of interest: while ambulance waiting time may stay within an acceptable variance, the cost variance may be exceeded due to high risks involved. The governance decisions need to to consider the outlier cases from both axis and have additional data available on the circumstances in which the breach occurred, such as the reasons of not using a helicopter for long distance journey for ER care.

While this chapter has studies system quality expectations, the points where the expressions are available for decision-making, and where they become dynamically computed is further studied in Chapters 4 and 5, split into collaboration case and ecosystem level events. In both, regulatory conformance is addressed.

4 Collaboration governance

This chapter shows that it is possible and feasible to perform ecosystem-wide strategic guidance of collaboration cases. For that it walks through the collaboration case

Theme	Case Management	BPM and VOBE	Pilarcos
Theory background	Not considered to have a clearly formulated theoretical background	Rely on the use of formal process description languages that allow process verification, e.g, through global state analysis or event traces.	In addition, bases its control of acceptable business process descriptions on a rigorous hierarchy of metamodels, and therefore opens up opportunities for model based engineering and model checking.
Data management	Mostly ad hoc	Rely on database thories, information flows description and checking, definition of message schema,and involvement of overarching ontologies.	Information management is grounded to information viewpoint of ODP RM, pre- and post- conditions and invariants on data schema. Uses bstract information models to allow heterogeneous data representation, partially agreed schemata for information repositories. Subjective business transaction criteria for information exchange acceptability.
Collaboration approach	Emergent information flow creation for cumulating wide range of data; ad hoc roles in deducing from data. PCM uses structured data collection processes with variation management.	Rely on predesigned processes with fixed roles, fixed interaction descriptions. Business processes mostly executed by distributed business process engines.	The eContract-agent makes each instance of BNM an unique case, where epochs and variation points in processes allow dynamic changes. The necessary adaptation to partner needs and situations is handled through subjective administration of business transactions.
Authorisation of information access or use of services	Allow changes on roles and authorisations freely.	Each business process role, associated with appropriate permissions and obligations.	Authorisation rules are defined per role in a fixed manner, but roles may change at epoch changes controlled by the rules stored into the eContract. This can further be restricted by each parther. The authorisation logic utilises deontic logic, that is, it is expressed in terms of obligations, permissions, and prohibitions.
Regulations management	Leaves this to ad hoc decisions by knowledge workers	At design time.	The first line of control is in the business process design phase where a design can only be accepted into a repository if the guarding metamodels allow it. An alternative phase for taking regulations into consideration is to embed them into the eContract. Finally, partners can subjectively refuse activities at runtime based on their private regulation related policies. Policies restrict the variation of the business processes at runtime.
Knowledge worker tools	Facilities for routing data into collections and visualising it, to support decision making.	Facilities for designing processes and providing user interfaces for breeding environments. Possible to design knowledge worker tools in a collaborative way.	Interfaces, e.g.for trust decision-making (reputation information flows are interpreted by each ecosystem member independently). Needs elaboration on interfaces for ecosystem infrastructure management, e.g. through eContract, privacy policy monitors, breach recovery agents

Table 2: Comparison of BPM and VOBE (virtual organsation), Pilarcos and CM (case management).

conformance management and the methods needed for gathering data for the collaboration behaviour for the ecosystem property evaluations. The compliance descriptor is composed by the governance properties of the ecosystem, including especially the static policies of the BNM, and dynamic requirements of the collaborating parties. The joint activities of the collaborations can be measured by extending each organisations' monitoring system with data collection modules that report to the ecosystem governors.

The eContract agent controls a novel version of business transactions [16] that is a complex interaction between multiple independent business services that

- strives to accomplish an explicitly shared business objective (either periodic or continuous);
- has a clearly defined cobehaviour leading to this objective in terms of exchanges of information and behaviour controls; that is, has i) mutually negotiated conditions of success (reached state over relevant aspects); and ii) mutually negotiated but subjectively detectable breaches:
- has clearly defined, mutually negotiated set of actors (not necessarily same as in the beginning) and their cobehaviour for breach recovery for each identified class of contract breach.

Firstly, these business transactions provide management and control capabilities of the inter-organisational collaboration to all the multiple partners, based on their subjective needs [16]. For the sake of peer independence, the agreements on goals, success criteria and breach criteria are all expressed in deontic logic terms (obligation, permission, prohibition).

Each of the collaboration peers is able to use eContract agent operations for suggesting contract changes. These operations fulfil ACID transactional properties. In the interactions between business services, there are no ACID requirements or overall data consistency requirements, but only the requirement of following the contracted BNM and policies. The eContract controls this collaborative behaviour through reports from private monitors of each organisation. In each system the local business service message exchange is overlooked for indicators of breaches either by the local service (users) or by peers. The monitoring can also collect other measurement information or user survey results for the collaboration property and ecosystem success analysis purposes. These properties may include timeliness, privacy preservation and user satisfaction. The subjective control is also supported by the reputation-based trust management system of Pilarcos. The business transaction management hooks in the trust decision making points at collaboration establishment time and further at each milestone or epoch during the collaboration life-cycle and declines collaboration if the trust level is insufficient on that particular situation [31]. Likewise, privacy valuation may lead to collaboration refusals too [36].

For governance in terms of delegation or withdrawing rights or obligations, the eContract provides operations for those partners that have permission for delegation. The operation modifies the policies recorded into the eContract, and that the change is performed as an ACID transaction. The change in the eContract is automatically distributed to the local organisational replicas of the contractual interpretation of the information, and the change is in use immediately. The permission to change delegation is a separate right, and is normally not passed on as part of the delegation of another permission. When an obligation is delegated, changes on multiple parties monitors may need to take place.

For regulatory conformance, the collaborations normally rely on the static BNM design information with the associated policies. The policies have been separated to guide choices between BNM alternative behaviours between variation points. This is especially designed for minimising the need of changing BNMs as minor changes in laws occur, and for allowing running collaboration cases to adopt new policies through the policy change operation at the eContract interface. A meaningful addition to the interface would be the change operation on collaboration properties to be collected and collated.

In collaboration modifications, it should be noted that business services are allowed to be stateful. In case of renegotiation of partners' services in the collaboration that requirement arises challenges. In Pilarcos, the partner change means an epoch change process activation where each service receives state information from the eContract the same way as in compensation activities. This state information may take an individual service backwards in its processing, or forward, into a commonly acceptable consistent point of affairs. The state information is not complete, but sufficient for a new business service to know with what contextual information it can start to act on its responsibilities. The current epoch change system needs enhancements especially on the criteria of acceptable joint continuation states. This partner change process is somewhat similar with related work by Grefen [29]. Where

Pilarcos uses eContract the related approach utilised a controller; where Pilarcos uses a contract renegoatiation, they inform all parties of the new membership. Both systems rely on their version of commitment protocols and utilisation of breach recovery facilities (compensation in both, and in addition, either rollback in their solution, and ecosystem level recovery process that is recorded into the eContract in Pilarcos). The major exception is that in Pilarcos the new partner starts working without collaboration history. For breach traceability, the eContract contains past epoch server identificators.

Statefulness of the services is relevant in healthcare services where transactions are long-lived and always leave medical records and physical patient conditions behind. Further, the choise of being ignorant on the preceding service providers is suitable, for example, for seeking a second medical opinion, or seeking care abroad.

As the eContract agent provides each partner an interface for renegotiating partnerships, policies, and changing epochs for reassociating new services to roles, all parties can dynamically govern the collaboration. At appropriate situations, each partner can check whether to continue or renegotiate the eContract, or to activate breach recovery. For human involvement at these points, user interfaces are required. Examples of eContract interface tools that support decision-making are available for Pilarcos trust decisions [14]; similar tools can be built for the patient interface to the care contract, as discussed above.

The required cumulative knowledge at the ecosystem governance system include, for each collaboration case,

- its state at completion, naming all the care and breach categories the case has belonged to; and
- when in progress, also planned care categories.

For potential feature harvesting a number of surveys were studied from business process management systems (BPM), adaptive (ACM) and production case management (PCM) systems, virtual organisation breeding environments (VOBEs), and Pilarcos ecosystem areas, e.g., [42, 11, 17, 23, 4]. Table 2 addresses the theoretical backgrounds, focus of tools and knowledge workers utilising the provided tools.

In the current Pilarcos architecture, breach detection is done by monitors that use eContracted policies and organisational policies. The eContract policies include BNM specific policies from the design time, reflecing requirements from static regulatory expectations. Any breaches can be caught and reports on the reasons collected to the eContract. Normally, organisations would make a free choise whether to inform a reputation system, but for governance that is not sufficient.

The monitoring and reporting must be enhanced in two ways. First, operational time ecosystem properties need to be retrieved from ecosystem governance policies regularly, and transformed to the local monitoring rules. This may involve signalling and transformation assistance by the ecosystem level knowledge workers. Second, the eContract agent functionality needs to be enhanced to use a specialised reporting facility that collect information from any local information systems to create the full report needed. In this, the case management (CM) practices are needed, as they enable inserting a new element into the local system although it does not fulfil the BPM practices on formal and verified processes, consistency checking based on database theories, and fixed access rights for the relevant processes. The new governance reporting element draws information from multiple local information system, and introduces a data flow to an external destination only trusted due to its role as the ecosystem level governor. This new element is designed by local knowledge workers who also need to introduce a secure communication channel that suits the local platform and represents the required data with the marshalling choise of the receiver. In some cases, such as with KPI computing across all collaborations, the solutions need to be a pervasively accessible distributed unit, comparable to a mini-architecture for a management information bus connecting all collaborations to the governors unit. The Pilarcos architecture already utilises such mini-architectures for communication channel configuration with selectable transparency properties [18, 16].

5 Ecosystem governance

This chapter discusses how the ecosystem governors interact with the ecosystem, analyse the success in reaching its goals, and design better strategic goals and properties. The ecosystem governors can express the goals for jointly produced services and their quality, and how the joint activities can be measured for comparison against the goals. Finally, regulatory conformance is discussed by associating the ecosystem with regulatory domains.

For governance, we propose four new tools. *I) Ecosystem governance modeling tool:*

The governance model is built around categories of actors and functionalities of the ecosystem, focusing on the deciding policies and controlling properties for the ecosystem governance. While the model is very rough on the ecosystem actor and functionality, it is suitable for simulation purposes and can be harvested from the ecosystem infrastructure repositories and from the goals and properties definition tool. The model needs to identify categories for normal collaboration cases, with their properties (with metrics) and ranges of acceptable values. For each normal category, there is an automatic failure category to capture breach situations as this is the fundamental mass of cases driving changes for the ecosystem. Use of this tool requires a deep understanding of the activities on the domain at the worker and client levels.

II) Periodical reporting tool: Periodic reporting allows the ecosystem governors to assess that the governance properties selected are actually guiding the ecosystem towards the intented behaviour. For enabling corrective actions, the reporting tool should keep reasons for breaches and contextual background should be in focus, but avoiding breaches on ecosystem member privacy. Normal anonymisation techniques apply to some extent, but for example geographical districts are often a unit of governance and financing, so this is particularly difficult. The use of this tool requires governance experience, knowledge on the legal responsibilities on the domain, and leadership position for being authorised to do required changes.

III) Ecosystem goals and properties definition tool: This tool is for capturing the leaders choices as conformance requirements. It needs to be able to provide a selection of practical methods for valuation and data collection. For the identified mechanisms, there should be support for organisations to apply them in their selected technological environments.

IV) Analysis tool:

The ecosystem intelligence unit constantly collects data from all collaboration cases in the ecosystem. While this sounds like too costly, it is easy to name current, popular services utilising even more costly data collection, e.g. Google maps traffic view. From the incoming data, analysis is performed according to the ecosystem goals and properties. This is not just a cost-effectiveness analysis but take societal values into consideration at large. The tool enables a what-if analysis on event history data to give insight, for example, on activities that should have been triggered earlier to save major costs in later activities. Data analysis may reveal potential early signs and simulation can be used for evaluating their validity. Data mining for long sequences of collaborations including recovery patterns may detect valuable new collaboration types.

The govenance process of an ecosystem, with these tools, starts with the creation of the ecosystem governance model, noting the categories of organisations involved based on their service roles, identifying the relevant BNM types and their breach categories, and finally the ecosystem properties and policies wanted. The organisations in this example represent healthcare districts, with two types, GP and consultancy. The relevant BNMs correspond each category of care contract block that is regularly used as evaluation basis. A suitable style of category would be a care process sequence like for Anna the initial checkup, preparation and advise before surgery, the operation itself, recovery care, and postoperation followup in a year, that is a full treatment plan. The breach categories are synthetic but necessary as they collect the essential knowledge for the redesign of the governance principles after some experience with the live systems. In addition, the same pattern with categories and synthetic breach categories is useful for queueing network simulations with multiple service classes. Repeating the same event sequencees in simulations, the governors may test what could have been a better governance strategy, or try out a few alternative governance styles for future with the assumed set(s) of patients.

Second, the collaborations start running with the set of ecosystem policies, and keep sending measurement results on successful situations and breach reports to the analysis tool. After a while, the governors want to view the situation through the periodical reporting tool. They can view how each district is doing, or how each category gets served, but cannot pry to individual patients or doctors medical records, as there are no connections to the actual systems present. Data leakage risks are however present, as the local knowledge workers at each organisation have entered probes to existing information systems through raw service interfaces, although most older systems are prepared for embedded user interface access only. Alternatively, older systems need to produce extra report databases from which required data can safely be extracted.

Third, the ecosystem intelligence service is designed to use predictive calculations based on a new method to make accurate forecasts [37] for the decision making moment. As co-operation of ecosystem members future-oriented knowledge can be created; for example, reports of surgeries being scheduled also indicate that there is a portion of them in need of rehabilitation services in a statistically observed time period. This joint forecasting and the governance system associated with it is based on a functional paradigm of futures research, which connects global responsibilities to joint short-term activities. Theories presented [37] match well with the ecosystem intelligence task as it needs to be parametrized with the ecosystem governance properties and rules, and performes multi-variable optimisation in what-if simulations.

Finally, we need to discuss compliance to regulations. While various approaches on expressing laws in rule or policy languages, or in deontic logic statements have been tried, there still is a fundamental difference in the way laws are written and how services or collaborations are managed.

In laws, we can find definition of concepts (citizen, legal entity), processes (election, applying for bankruptcy, selling property), obligations (compulsory schooling, listing to national military forces), and prohibitions (list of criminal acts, list of acts minors cannot legally perform by themselves), and rights (citizenship, minimum income, necessary healthcare). However, in different countries the definitions are not similar: for example, not all countries oblige all citizens to participate in the military forces, nor do all countries ensure that everyone in need of immediate healthcare receives that (cost free if necessary). Especially the processes may be significantly different in terms of the legal entities that holds the governing power. Differences in the organisational structures that govern and operate a community (society or system) necessarily get reflected to the processes as well.

For achieving interoperability between communities, the stability of that organisational structure is essential. Between "isomorphic" organisational structures it is possible to create an intermediating terminology. We relate the above ecosystem governance models with the rough regulatory system models. The ecosystems and regulatory domains do not need to have similar boundaries, but an ecosystem may span multiple regulatory systems, or a single regulatory system may include a number of ecosystems. In comparison to solutions (e.g., [15]) where laws are referred in natural language texts and URLs, this solution is more efficient in the sense that the monitoring criteria for services or information exchanges are readily in executable form. Naturally, there is a significant introduction cost for the initial regulatory system models.

Regulatory system model differencies can be considered as one form of pragmatic interoperability problem. When there is a close enough correspondence between the non-native partners position in the unified community model, and the applied obligations and prohibitions do not cause a discrepancy, it is possible to formulate a clearance unit that performs the intended activity in a manner that fulfils the laws of both countries. At present, certain companies have found a business for themselves in performing such transactions based on their cumulated expertise and reputation as trusted middlemen in both counties.

When arranging regulatory compliance for inter-organisational collaborations, we believe that the following foundations should be reached first:

- a unified community model should be structured to show the relationship of the essential active units;
- terminology for indicating the community functions of essence and their relationships to potential actors;
- rough process models and actors who run function for the processes in order to indicate i) the processes that are regulated, and ii) the regulation processes;
- obligation and prohibition expression language for expressing which processes are limited and how; the language should have expressions on actors (like citizens, organisations) and their power of enforcing, regulating or governing another actor or a process.

For the healthcare example, each district should be further split into actual hospitals and health centers, revealing the full organisational hierarchy. After that, the national regulations could be connected to the ecosystem level. The district level strategies would be added to the new middle layer in the hierarchy, and the knowledge on district services and agreements on private sector compensation would find a position at that level too.

For the automated governance purposes, the creation of such model in collaboration with legal system researchers would be plausible. The model enable not only the pragmatic interoperability across legal system boundaries, but also allow laws to directly govern the major societal systems, like healthcare and foreign trade. Most relevant design elements include contractual concepts, breach recovery processes at the ecosystem level, and use of deontic logic as policy language base, and for governance in particular, the dynamic collaboration conformance requirements and ecosystem level properties as main visualisable goals.

6 Conclusion

This paper addresses automation of service ecosystem governance and enhancing existing ecosystem architectures by necessary facilities. These facilities include selection of collaboration and ecosystem valuation concepts and metrics, constructing tools for declaring governance requirements, and data collation mechanisms for drawing necessary data from each involved organisation. The data collection can be organised by intrusive adaptive case management practices.

The automated inter-enterprise collaboration governance is beneficial in multiple ways. First, it improves customer experience and trustworthiness of services, due to monitoring of regulations. Second, it improves societal equality by allowing various kind of equality metrics to be declared as ecosystem level goals. Third, it allows strategic direction of business towards cost-effective markets. Overall, it enables safe societal computing support for all domains.

Building a governance environment for software-based services and their inter-organisational, dynamic collaboration requires commitment in the international level and involvement of standardisation. The awareness of service ecosystem and their dynamic governance can also be built through education and slowly arising expectations for technical support.

This adoption step involves societal changes, like unification of legal systems on digital society aspects. Governance is often focused on conforming to regulations, laws. From computing system point of view, it would be optimal if laws could be expressed as computable rules; in case of societal ecosystems, this is in a sufficient degree reachable, as the law structure reflects the system structure and the deontic logic is an appropriate tool for expressing legal clauses. Computer science and legal expert collaboration is needed for new practices.

A serious concern is that there are few laws that regulate commitments by software agents. Further there is lack of understanding whether some novel business models that innovation ecosystems get rigorous support for are actually legal. For example, law should set norms on when a business model is sufficiently visible for peers to understand the consequent threats to their privacy (e.g., web bugs or web beacons [25]). Further, there are no commonly accepted "legal" ways of technically assessing contextual facts, for example, a standard, trustworthy way of detecting where a software agent is running and storing its data (e.g, DMARC [12]).

The ecosystem governance facilities turn ecosystem architectures and their governance into a practical tool for politicians for finding deep understanding on societal systems and governing them in cost-effective and ethical way.

References

1 S-cube network of excellence final report. Technical report, 2012. FP6 NoE 215483.

2 A.-J. Berre, B Elvesater, N. Figay, C. Guglielmina, S.G. Johnsen, D. Karlsen, T. Knothe, and S. Lippe. The athena interoperability framework. In *Enterprise Interoperability II*. Springer London, 2007.

3 Luis M Camarinha-Matos and Hamideh Afsarmanesh. Virtual enterprise modeling and support infrastructures: applying multi-agent sys-

tem approaches. In *Multi-agent systems and applications*. Springer, 2001.

4 LuisM. Camarinha-Matos and Hamideh Afsarmanesh. Collaborative Systems for Smart Environments: Trends and Challenges. In LuisM. Camarinha-Matos and Hamideh Afsarmanesh, editors, *Collaborative Systems for Smart Networked Environments*, volume 434 of *IFIP Advances in Information and Communication Technology*, pages 3–15. Springer.

5 Tiago Cardoso and Luis M. Camarinha-Matos. Pro-Active Service Ecosystem Framework. *International Journal of Computer Integrated Manufacturing*, 26(11), 2013.

6 Jengchung Victor Chen, Yawen Chen, and Erik Paolo S. Capistrano. Process quality and collaboration quality on B2B e?commerce. *Industrial Management & Data Systems*, 113(6), 2013.

7 Nicolas Daclin, David Chen, and Bruno Vallespir. Developing enterprise collaboration: a methodology to implement and improve interoperability. *Enterprise Information Systems*, 0(0), 0.

8 Duodecim. Current care guidelines. Technical report, 2015.

9 Brian Elvesæter, Arne-Jørgen Berre, Henk de Man, and Man-Sze Li. Networked enterprise transformation and resource management in future internet enabled innovation clouds. In *Enterprise Interoperability IV*, pages 313–322. Springer, 2010.

10 R. Ford, M. Carvalho, L. Mayron, and M. Bishop. Antimalware software: Do we measure resilience? In *Anti-malware Testing Research (WATeR), 2013 Workshop on*, 2013.

11 M. Hauder, D. Munch, F. Michel, A. Utz, and F. Matthes. Examining adaptive case management to support processes for enterprise architecture management. In *Enterprise Distributed Object Computing Conference Workshops and Demonstrations (EDOCW), 2014 IEEE 18th International*, 2014.

12 IETF. Domain-based Message Authentication, Reporting, and Conformance (DMARC). Technical report, 2015. RFC7489.

13 Lalana Kagal, Tim Finin, and Anupam Joshi. A policy language for a pervasive computing environment. In *Policies for Distributed Systems and Networks, 2003. Proceedings. POLICY 2003. IEEE 4th International Workshop on*, pages 63–74. IEEE, 2003.

14 Puneet Kaur, Sini Ruohomaa, and Lea Kutvonen. Enabling user involvement in trust decision making for inter-enterprise collaborations. *International Journal on Advances in Intelligent Systems*, 5(3-4), 2012.

15 F. Koetter, M. Kochanowski, A. Weisbecker, C. Fehling, and F. Leymann. "integrating compliance requirements across business and it". In *Enterprise Distributed Object Computing Conference (EDOC), 2014 IEEE 18th International*, 2014.

16 L. Kutvonen, A. Norta, and S. Ruohomaa. Inter-enterprise business transaction management in open service ecosystems. In *Enterprise Distributed Object Computing Conference (EDOC), 2012 IEEE 16th International*, 2012.

17 Lea Kutvonen. Enhancing the maturity of open service ecosystems and inter-enterprise collaborations. In Marten van Sinderen, Paul Oude Luttighuis, Erwin Folmer, and Steven Bosems, editors, *Enterprise Interoperability*, volume 144 of *Lecture Notes in Business Information Processing*. Springer, 2013.

18 Lea Kutvonen. ODP RM reflections on open service ecosystems. *Computer Standards and Interfaces*, 35(3):294 – 312, 2013. Special issue on RM-ODP: Foundations, Experience and Applications.

19 Lea Kutvonen et al. Multi-tier agent architecture for open service ecosystems. In *AT*, pages 302–316, 2012.

20 K. Laghari and K. Connelly. Toward total quality of experience: A qoe model in a communication ecosystem. *Communications Magazine, IEEE*, 50(4), 2012.

21 Outi Markkanen and Markku Puro. Integraatio ratkaisuna sosiaali- ja terveydenhuollon uudistamistarpeisiin (integration as a solution for the renewal needs of social and health care in central finland). Technical report. Selvityshenkiloiden raportti Keski-Suomen sosiaali- ja terveydenhuollon palvelurakenneselvityshankkeelle 30.4.2011. Report 31.

22 A McConahy, Ben Eisenbraun, James Howison, James D Herbsleb, and P Sliz. Techniques for monitoring runtime architectures of socio-technical ecosystems. In *Workshop on Data-Intensive Collaboration in Science and Engineering (CSCW 2012)*, 2012.

23 Nikolay Mehandjiev and Paul Grefen. *Dynamic business process formation for instant virtual enterprises*. Springer, 2010.

24 H.R. Motahari-Nezhad and K.D. Swenson. Adaptive case management: Overview and research challenges. In *Business Informatics (CBI), 2013 IEEE 15th Conference on*, pages 264–269, July 2013.

25 Bhavesh Naik. Tracking user activities using web bugs. Technical report, 2013.

26 National Institute for health and welfare. Sotkanet statistics and indicator bank. Technical report, 2015.

27 Object management group. Case management model and notation (cmmn). Technical report, 2014. v 1.0.

28 Antti Poikola, Kai Kuikkaniemi, and Ossi Kuittinen. My data – johdatus ihmiskeskeiseen henkilotiedon hyodyntamiseen (my data – introduction to human-centric use of person associated information. Technical report, 2014.

29 S. Pourmirza, R. Dijkman, and P. Grefen. Switching parties in a collaboration at run-time. In *Enterprise Distributed Object Computing Conference (EDOC), 2014 IEEE 18th International*, 2014.

30 R Glenn Richey, Frank G Adams, and Vivek Dalela. Technology and flexibility: Enablers of collaboration and time-based logistics quality. *Journal of Business Logistics*, 33(1), 2012.

31 Sini Ruohomaa. *The effect of reputation on trust decisions in inter-enterprise collaborations*. PhD thesis, University of Helsinki, Department of Computer Science, May 2012.

32 T. Ruokolainen and L. Kutvonen. An architecture framework for facilitating sustainability in open service ecosystems. In *Enterprise Distributed Object Computing Conference Workshops (EDOCW), 2012 IEEE 16th International*, 2012.

33 Toni Ruokolainen. *A Model-Driven Approach to Service Ecosystem Engineering*. PhD thesis, University of Helsinki, Department of Computer Science, 2013.

34 Alexandru Serbanati, Carlo Maria Medaglia, and Ugo Biader Ceipidor. Building blocks of the internet of things: State of the art and beyond. In Cristina Turcu, editor, *Deploying RFID - Challenges, Solutions, and Open Issues*, 2011.

35 Ahm Shamsuzzoha, Filipe Ferreira, Américo Azevedo, José Faria, and Petri Helo. Business process monitoring and management in virtual enterprise through interactive user interface layer. In *Advances in Sustainable and Competitive Manufacturing Systems*, pages 451–464. Springer, 2013.

36 Y. Shen, M. Miettinen, P. Moen, and L. Kutvonen. Privacy preservation approach in service ecosystems. In *Enterprise Distributed Object*

Computing Conference Workshops (EDOCW), 2011 15th IEEE International.

37 Timo Sneck. *From hypothesis and scenarios toward forecast-initiated action platforms for decision-making in networks. Utilization of forecasting in business and governmental operations (In Finnish).* PhD thesis, University of Helsinki, Faculty of Political science., 2002. VTT Pub 468.

38 Chee-Wee Tan, Izak Benbasat, and Ronald T Cenfetelli. It-mediated customer service content and delivery in electronic governments: An empirical investigation of the antecedents of service quality. *MIS Quarterly*, 37(1), 2013.

39 B. Tricker. *Corporate governance: Principles, Policies and Practices. 2nd edition.* Oxford University Press, 2012.

40 Yan Wang, Y. Taher, and W. van den Heuvel. Towards smart service networks: An interdisciplinary service assessment metrics. In *Enterprise Distributed Object Computing Conference Workshops (EDOCW), 2012 IEEE 16th International*, 2012.

41 Michael Wilson, Alvaro Arenas, David Chadwick, Theo Dimitrakos, Jurgen Doser, Pablo Giambiagi, David Golby, Christian Geuer-Pollman, Jochen Haller, Ketil Stølen, et al. The trustcom approach to enforcing agreements between interoperating enterprises. In *Enterprise Interoperability*. Springer, 2007.

42 Zhiqiang Yan, Remco Dijkman, and Paul Grefen. Business process model repositories–framework and survey. *Information and Software Technology*, 54(4), 2012.