SDSN@RT - A middleware environment for single-instance multitenant cloud applications

Citation for published version (APA):

Kumara, I., Han, J., Colman, A., Heuvel, W.-J. V. D., Tamburri, D. A., & Kapuruge, M. (2019). SDSN@RT - A middleware environment for single-instance multitenant cloud applications. *Software : Practice and Experience*, 49(5), 813-839. <https://doi.org/10.1002/SPE.2686>

Document license:
TAVERNE

DOI:
[10.1002/SPE.2686](https://doi.org/10.1002/SPE.2686)

Document status and date:
Published: 01/05/2019

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

RESEARCH ARTICLE

WILEY

SDSN@RT: A middleware environment for single-instance multitenant cloud applications

Indika Kumara^{1,3} | Jun Han² | Alan Colman² | Willem-Jan van den Heuvel^{1,3} | Damian A. Tamburri^{1,3} | Malinda Kapuruge²

¹Tilburg School of Economics and Management, Tilburg University, Tilburg, The Netherlands

²School of Software and Electrical Engineering, Swinburne University of Technology, Hawthorn, Australia

³Universiteit Jheronimus Academy of Data Science, 's-Hertogenbosch, The Netherlands

Correspondence

Indika Kumara, Tilburg University, Tilburg, North Brabant, 5037 AB, The Netherlands.
Email: i.p.k.weerasingha.dewage@tue.nl

Funding information

Cooperative Research Centres, Australian Government Department of Industry; Swinburne University of Technology; European Commission, Grant/Award Number: 825480 (Horizon 2020, SODALITE)

Summary

With the single-instance multitenancy (SIMT) model for composite Software-as-a-Service (SaaS) applications, a single composite application instance can host multiple tenants, yielding the benefits of better service and resource utilization and reduced operational cost for the SaaS provider. An SIMT application needs to share services and their aggregation (the application) among its tenants while supporting variations in the functional and performance requirements of the tenants. The SaaS provider requires a middleware environment that can deploy, enact, and manage a designed SIMT application, to achieve the varied requirements of the different tenants in a controlled manner. This paper presents the SDSN@RT (software-defined service networks at runtime) middleware environment that can meet the aforementioned requirements. SDSN@RT represents an SIMT composite cloud application as a multitenant service network, where the same service network simultaneously hosts a set of virtual service networks, one for each tenant. A service network connects a set of services and coordinates the interactions between them. A virtual service network realizes the requirements for a specific tenant and can be deployed, configured, and logically isolated in the service network at runtime. SDSN@RT also supports the monitoring and runtime changes of the deployed multitenant service networks. We show the feasibility of SDSN@RT with a prototype implementation and demonstrate its capabilities to host SIMT applications and support their changes with a case study. The performance study of the prototype implementation shows that the runtime capabilities of our middleware incur little overhead.

KEYWORDS

cloud application, middleware, multitenancy, SaaS, SDSN, service network, virtual service network

1 | INTRODUCTION

Computing has become a utility with the advancement of cloud computing. Cloud service models offer a wide range of resources, from computing resources to business services, as utilities for individual and organizational customers.

FIGURE 1 Composite cloud applications. A, Multi-instance multitenancy model; B, Single-instance multitenancy model [Colour figure can be viewed at wileyonlinelibrary.com]

In particular, with Software-as-a-Service (SaaS) composite applications, the cloud service provider can share the use of resources, services, and their aggregations among its customers to achieve economies of scale through multitenancy.

There are two main multitenancy models for composite cloud applications^{1,2}: multi-instance multitenancy (MIMT) and single-instance multitenancy (SIMT). In the MIMT model, a tenant uses a dedicated runtime application variant and associated services (see Figure 1A). In the SIMT model, the tenants simultaneously share the same runtime application instance and services (see Figure 1B). The separation in the MIMT model can simplify the support for tenant-specific variations and changes, whereas the sharing in the SIMT model can reduce operational cost and improve utilization of the capacities of services. An important goal for multitenancy support is to achieve the best of these two multitenancy models at the same time, ie, to support tenant-specific (functional and performance) variations and changes while achieving sharing in the SIMT model.

To host and manage composite SIMT cloud applications at runtime, a specialized middleware platform is needed. First, the middleware should create and maintain the runtime models of a composite application. These models need to connect the shared services according to the relationships between the services and to coordinate the message exchanges between the services. Second, the middleware needs to be able to form and manage tenant-specific (*virtual*) composite application variants over the shared services and application (ie, slicing the shared services and application). These variants need to reflect the commonality and variability in the functional and performance requirements of the tenants. The achievement of the performance requirements of a composite service application and its tenant-specific virtual variants depends on the performance of the used service capabilities, the architectural composition of these service capabilities (referred to as *configuration design or topology*), and the regulation of the message exchanges among the relevant services (referred to as *regulation design*).³ Thus, the middleware needs to support the sharing and variations in services, configuration designs, and regulation designs among application variants. Finally, the middleware should support the monitoring of, and runtime changes to, the shared cloud application. A software engineer or a management application needs to be able to use these middleware capabilities to monitor and modify the application. Note that, in this paper, we focus on *business (IT) services*, which are surrogates for and support real-world businesses of enterprises. The underlying businesses (rather than infrastructure-level IT resources) usually constrain the capacities of such services.

While there are interesting works on middleware platforms for hosting general composite applications⁴⁻⁶ and multitenant composite applications,⁷⁻¹² they do not provide sufficient support to meet the above challenges posed by the SIMT model in achieving runtime sharing with variations. Existing middleware platforms for multitenant component compositions^{7,8} support the generation of application variants and the creation of runtime models for those variants (per a user request or per a tenant) and runtime binding of tenant-specific components to the declared locations in the component model of the application per a user request. Similarly, middleware platforms for multitenant service compositions support the creation and separate deployment of service composition variants,⁹ runtime binding of services per a user request,¹⁰ variations through conditional branches and parameters,⁹⁻¹¹ and runtime changes to service compositions.¹² In general, the works that create the runtime models per application variants can only support the MIMT model. While the runtime binding or selection of tenant-specific variations (components, services, or code branches) at the specific locations in a shared application can support the SIMT model, it only provides limited flexibility to form application variants by dynamically composing a subset of the elements in the application. For example, the fine-grained variants and variation points and the tangling of normal processing logic with variant differentiation logic can make the application complex.^{13,14} Regarding the regulation design of multitenant applications, some works support tenant-specific policy

enforcement¹⁵ (without sharing or variation of policies among tenants), quality-aware application variant generation,¹⁶ and specific^{17–20} or configurable²¹ regulation mechanisms. However, these works have not considered the runtime sharing and tenant-specific variations in the regulation design or the runtime changes to the regulation design.

In this paper, we present *SDSN@RT* (*software-defined service networks at runtime*), which is a middleware platform that can enact and manage SIMT composite cloud applications that share services among tenant-specific (virtual) composite applications while realizing the differing functional and performance requirements of the tenants. We represent such an SIMT application at runtime as a *multitenant service network* or a *software-defined service network (SDSN)*, where a set of virtual service networks (VSNs) simultaneously coexists on the same physical service network in a way that is analogous to virtual computer networks. The (physical) service network is a runtime structure that connects services and mediates the interactions between the services. It is created based on the *models@runtime* concept.²² A VSN belonging to a tenant meets the tenant's functional and performance requirements and uses a subset of the services in the underlying physical service network. The different VSNs have similarities and differences in their services, configuration designs, and regulation designs. A VSN is enacted by consuming the service capabilities used by the VSN and routing the relevant interaction messages (capability requests and responses) between relevant services in the service network. Through a separate controller, these interaction flows can be monitored and regulated according to the designs of the VSNs, and the multitenant service network can be modified and reconfigured at runtime.

The main contribution of this paper is to address the abovementioned requirements for a middleware platform that hosts SIMT composite cloud applications. To this end, we present (1) the runtime models of SIMT applications designed as multitenant service networks, (2) the support for simultaneous and controlled enactment of virtual application variants (ie, VSNs), and (3) the support for policy-based runtime management of applications and their variants. We show the feasibility of SDSN@RT with a prototype implementation, demonstrate its capabilities with a case study, and quantify its runtime overhead in supporting the SIMT model.

In an earlier paper,³ we have presented the design support for SDSN (or the programming model of the SDSN@RT middleware) that allows a software engineer to define multitenant service networks using a set of design constructs and an associated domain-specific language. The paper also included a section (around one page) on the SDSN middleware to emphasize the fact that the designs are not merely conceptual models but can be enacted/executed. However, the main focus of the paper was to describe our approach to supporting the functional and performance variability at the design specification level (design time variability). In contrast, this paper presents the SDSN middleware in detail, focusing on its support for maintaining the runtime models of the SIMT cloud application, achieving the runtime functional and performance variability via dynamic routing in service networks, and managing runtime changes to the application. In this paper, we only discuss the runtime abstractions in the application as well as the architectural and runtime characteristics and capabilities of the middleware (in contrast, the focus of our earlier paper was on design models).

Our SDSN middleware has several key middleware novelties. First, dynamic routing at the application architecture level is introduced as a flexible runtime variability mechanism. The application architecture (architectural elements and their relationships) and the conversational behaviors (message exchanges between the architectural elements) are made explicit, visible, and controllable. A set of abstractions is introduced to intercept, regulate, and route the message exchanges. Consequentially, the dynamic routing enables slicing the same runtime application at the architecture level and dynamically controlling conversational behaviors in the application slices (as in flexible network-based models^{23,24}). Thus, the dynamic routing provides higher controlled flexibility to achieve the runtime variability compared with the existing approaches, such as selecting, activating, or injecting components/services within a fixed application structure (eg, a component model). Moreover, the explicit architectural model of the composite application can improve the flexibility to dynamically manage/change the composite application at runtime (as in the applications with explicit configuration designs²⁵). Second, to the best of knowledge of the authors, our SDSN middleware is the only middleware to support the enactment of the composite service applications designed following the service network model, which is a service composition model that can provide a natural abstraction to design, enact, regulate, and manage webs of service-based businesses.^{26,27} Finally, our SDSN middleware uses the *models@runtime*, the dynamic routing, and the extended service network model to support the SIMT model for business service-based composite cloud applications.

The rest of this paper is organized as follows. Section 2 presents a business example to motivate the SIMT model for composite cloud applications and identifies the challenges for a middleware platform that hosts such applications. Section 3 presents our SDSN@RT middleware, including its system architecture, its runtime models for representing SIMT cloud applications, its capabilities to enact and manage such applications, and its prototype implementation. Section 4 presents the implementation of the two case studies with the SDSN@RT prototype and includes a performance study of the prototype. Section 5 summarizes the related work, and Section 6 concludes this paper and outlines some future work.

FIGURE 2 CampSAS, its partner organizations, and its tenants and their customers [Colour figure can be viewed at wileyonlinelibrary.com]

2 | MULTITENANT COMPOSITE CLOUD APPLICATIONS AND MIDDLEWARE CHALLENGES

In this section, we first present a business-driven case study used in this paper, highlighting the multitenancy requirements of such business scenarios. Then, we describe the key challenges pertaining to a middleware that can enact and manage multitenant cloud applications supporting such business scenarios.

2.1 | Case study

CampSAS is a company that offers camping assistance to its business clients (tenants) such as campsite owners and travel agencies (see Figure 2). Its camping assistance business is supported through a software application. This application adopts the SIMT SaaS model due to its resource utilization and operational benefits and integrates partner business services such as camping equipment providers and tour companies.

The functional and performance requirements of the different tenants exhibit commonalities and variations. Figure 2 shows three tenants of CampSAS. They use tour arrangement and case handling capabilities (commonalities) and select two more capabilities from equipment rental, grocery delivery, and taxi hire (variations). Their performance requirements also vary. For example, the expected response time is 3 hours/case for the tenant HappyTours and 4 hours/case for the tenant UniUvtClub.

The partner services also have both similar and different capabilities, even for the same service type. For example, TomGear and JackGear are both camping equipment providers. Compared with TomGear, JackGear has an internal delivery service and, thus, can deliver the rented equipment to the campers without using an external delivery company. Their rental capacities, which are the maximum number of rental orders per day, also differ (200 compared with 100). Note that these finite capacities of the business services cannot be changed by simply managing the computing resources the services use.

2.2 | Scenarios

The software engineers at CampSAS can design the camping assistance composite cloud application using a design methodology and tools for SIMT cloud applications, for example, our SDSN design support.³ Once the application is

designed, the software engineers need a middleware environment that can host the designed cloud application and enable the enactment and runtime management of the deployed application. The following scenarios illustrate the usage of the middleware environment.

- *Deployment.* First, a software engineer deploys the camping assistance cloud application on the middleware by providing the design artifacts for the application. The deployed application connects the partner services to enable and manage the interactions between them. Second, a software engineer deploys the variants of the shared camping assistance application for the three tenants. These application variants share some subset of the partner services and the application, which reflects the commonality and variability in the functional and performance requirements of the tenants.
- *Enactment and monitoring.* The campers of the tenants request assistance from CampSAS when they need some camping services. When an assistance request is received, the middleware creates and enacts an instance of the camping assistance application variant of the corresponding tenant to support the requested camping assistance case. Multiple camping assistance cases coexist and progress simultaneously. The middleware monitors the progress of each camping assistance case and measures such metrics as the time to complete a case and the available capacity of the partner services.
- *Management.* The requirements of tenants and CampSAS and the availability and capabilities of services can change over time. For example, after 1 month, the tenant SunCampsites decides that it needs 50 additional assistances per day. After 2 months, CampSAS decides to support the bike rental capability in its camping assistance business, and the tenant HappyTours would like to use the new bike rental capability instead of the currently used taxi hiring capability. To support these changes, a software engineer evolves the shared camping assistance cloud application by modifying its design at runtime without disturbing the operations of those tenants unaffected by the change (for example, the tenant UniUvtClub).

2.3 | Middleware challenges

The above scenarios represent three general key challenges for a middleware environment to support such SIMT composite cloud applications.

- *Create and maintain the runtime models of the designed composite cloud application (Req1).* To enable the enactment of and runtime incremental changes to the composite cloud application, the middleware needs to create and maintain suitable runtime models of the application.^{22,28} These models also need to facilitate the managed interactions between the services.
- *Enact and control concurrent variants of the deployed composite cloud application (Req2).* The middleware needs to create the variant instances of the composite cloud application for the instantiation requests from the customers of the tenants and to progress the created instances by executing the elements of the runtime models of the application and the service operations that are required and support the requirements of the tenant's application variant. These instances need to be able to share and selectively use the services and the runtime elements of the composite application as necessary to accommodate the similarities and differences in the functional and performance requirements of the tenants. The middleware also needs to enable fine-grained monitoring of the progress of the variant instances and the performance of the partner services.
- *Enable the runtime management of the deployed composite cloud application and its variants (Req3).* The middleware needs to support different classes of runtime changes that can potentially occur to the SIMT applications. The changes need to be realized and managed at runtime without disturbing the operations of those tenants and end users unaffected by the change. A software engineer should be able to declaratively specify changes and enact the change specification at runtime to modify the application.

3 | THE SDSN@RT MIDDLEWARE ENVIRONMENT

Software-defined service networking (SDSN)³ is an architecture-oriented approach for software engineers to design, enact, and manage SIMT service-based composite SaaS applications. SDSN represents a composite SaaS application as a *multitenant service network*, where a set of managed VSNs shares the same physical service network at runtime. A VSN represents a service composition in the service network for a particular tenant. The concept of the multitenant service

network realizes the SIMT model. In contrast, physically separated service networks realize the MIMT model. In the Introduction section, we have compared the SIMT model with the MIMT model.

This section presents SDSN@RT, which is a middleware runtime that supports the deployment, enactment, monitoring, and management of SIMT composite cloud applications designed as multitenant service networks following our SDSN approach. We first provide an overview of the SDSN@RT system architecture (Section 3.1). Upon the deployment of the design models of a multitenant service network, their runtime models are created following the *models@runtime* concept.²² The runtime model of a service network acts as the *network infrastructure* that allows the controlled interactions (exchange of messages) between services according to the design and, thus, supports the enactment of VSNs. A service network can be monitored and modified at runtime. Sections 3.2 and 3.3 describe the structure and the responsibilities of each runtime abstraction in multitenant service networks. Section 3.4 discusses the enactment of a VSN, whereas Section 3.5 presents the management capabilities of SDSN@RT. In Section 3.6, we briefly describe the prototype implementation of SDSN@RT.

3.1 | Overview of SDSN@RT system architecture

Figure 3 shows the high-level system architecture of SDSN@RT. The architecture has three main layers: services, coordination infrastructure, and management platform.

The services are Web services that support and act as proxies to the real-world business capabilities of the enterprises. Multiple tenants in a business service network share the services and their capacities. These services contractually relate to each other and collaborate at runtime by exchanging messages. These interaction messages between services need to be intercepted, coordinated, and regulated to achieve the functional and performance goals of the SaaS provider and its tenants. For this purpose, the service coordination infrastructure creates and maintains a runtime model of the service network (ie, an overlay network over the services). A service network (runtime model) has a set of regulation enforcement points (REPs) to allow the observation and regulation of the message flows throughout the service network. The coordination infrastructure also has an event manager and a state manager. The event manager records the events being generated in the service network, for example, service interaction events and performance violation events. The state manager records the states of the service network, such as the used service throughput and average response time. The events and states can trigger the regulation policies at REPs and, thus, regulate the enactment of VSNs. The event manager and the state manager support the publish-subscribe model for publishing and consuming events and state records, respectively.

The management logic of a service network resides mainly in the management platform, which acts as *the brain of the service network*. For a given multitenant service network, the management platform consists of an organizer and an operational manager that enforce configuration and regulation management policies, respectively. To implement their

FIGURE 3 High-level system architecture of software-defined service networks at runtime (with a single deployed service network)

[Colour figure can be viewed at wileyonlinelibrary.com]

capabilities, both managers use a policy manager, an event manager, and a state manager. The policy manager stores and executes the management policies. The behaviors of the event manager and the state manager are similar to those at the coordination infrastructure, except that they control the enactment of the management policies.

The management interfaces between the management platform and the coordination infrastructure define the allowed management interactions between them. There are three types of management interfaces: deployment, organizational management, and operational management. The deployment interface is at the system level and supports the deployment and undeployment of (multitenant) service networks. The latter two interfaces are at the application level (per each service network) and respectively support the management of the configuration and regulation designs of a running multitenant service network at runtime. The organizer and operational manager of the service network can use these management interfaces to dynamically change the topology of the service network as well as the routing and regulation of the messages passing over the service network.

The management platform also offers its management capabilities as Web service interfaces so that the SaaS administrator can remotely deploy the designed service networks, monitor and modify the deployed service networks by providing management policies, and undeploy the deployed service networks as necessary.

SDSN@RT can enact multiple multitenant service networks. Each deployed service network has its separate and isolated instances (Java objects) of the architectural elements, such as service network, REP, event manager, state manager, policy manager, organizational manager, and operational manager, as the runtime representations of the corresponding VSN instances. In other words, each deployed service network has a separate runtime model (ie, a Java object model). For simplicity, we only consider a single service network in this paper.

3.2 | Service networks

A composite SIMT SaaS application relies on a range of services for supporting its business. To represent, compose, and manage these services in the context of the entire application, *roles* are introduced to represent the services as well as *contracts* between services to connect the services for their interactions. The network of roles (*nodes*) and contracts (*links*) forms the *configuration* or structure of the network and is superimposed on the services for their management and interaction. The service network acts as an overlay network over the services. A role comprises a set of *tasks*, which represent the used capabilities of the service. It acts as a proxy service to the external service and can send (Web service) messages to the service and receive messages from the service. A contract consists of a set of *interaction terms* and the references to the two roles that it connects. The tasks and interaction terms define the message templates, which are used to create and validate the role-role and role-service interaction messages.

Figure 4 shows the service network for our motivating example. It consists of a number of roles (eg, CM, SC, EP1, and DC) connected by contracts (eg, CM_SC, CM_DC, CM_EP1, and SC_EP1) to form the service network supporting the coordination of the interactions between the services (eg, 24by7 support center, TomGear equipment provider, and MacDelivery delivery service) to meet the camping assistance requirements of the tenants. As shown in Figure 5, the roles include the relevant tasks, for example, *tRentEquipment* of the role EP1 (to rent camping equipment) and *tDeliver* of the role DC (to deliver the rented equipment to the campers). The contracts include the relevant interaction terms, for example, *iSendEquipmentRequirements* of the contract CM_EP1 (to represent sending the equipment requirements to the rental company) and *iOrderEquipment* of the contract SC_EP1 (to represent the approval from the support center for renting the required equipment).

The *regulation* of the message flow over the service network structure is achieved through a set of REPs at the various roles and contracts of the network. The messages flow from REP to REP, which can intercept messages and apply various operations on the intercepted messages. There are four types of REPs: synchronization (at each role), routing (at each role), pass-through (at each contract), and coordinated-pass-through (across contracts). The synchronization REP of a role synchronizes a subset of incoming interactions from the adjacent roles (and consequently their services) before invoking a task of the service represented by the role by sending a request message to the service. The routing REP of a role receives and routes a response or request message from the role's service to a subset of the adjacent roles (and consequently their services). The pass-through REP can intercept and process the interaction messages between two roles (and thus their services). The coordinated-pass-through is to regulate the interactions across different pairs of roles (and their services).

Each REP consists of a *regulation knowledgebase* and a *regulation table*. The knowledgebase contains a set of event-condition-action (ECA) rules that define the regulation logic. The regulation table maps a service network flow to one or more ECA rules, which determine the treatment of the flow. We define a service network flow broadly as a set

FIGURE 4 The camping assistance service network of CampSAS [Colour figure can be viewed at wileyonlinelibrary.com]

of messages belonging to a specific entity such as a VSN, a business process, and a VSN instance (for an end user). The messages and their interpretations (ie, events and recorded states) carry the flow information (eg, the identifiers of the relevant entities). This flow information allows the enactment engine to isolate the executions of a given composite application by the end users of different tenants (similar to the use of the tenant-context in the work of Pathirage et al⁹ for tenant execution isolation).

The condition part of a regulation rule can be a logical expression of message contents, events, and state records. Thus, the arrival of a message at an REP and a new or changed service network event or state can trigger a subset of rules. The action part of a rule uses one or more regulation functions, which tell the REP what to do with the individual interaction messages and/or the service network flow. A regulation function implements a unit of regulation decisions on messages, such as admission control, transformation, service operation execution, response time and throughput monitoring, and event publishing. A detailed discussion of the available regulation functions can be found in our previous publication.³ They collectively support the connection and interaction between services as well as the monitoring and regulation of service interactions.

There may be control flow relations between regulation functions as well as between rules. The underlying rule engine determines these control flows and the activation of the individual rules. For this purpose, we use a Rete algorithm-based²⁹ rule engine, namely, Drools (drools.org). The execution of the rules for different service network flows is independent and isolated.

Figure 5 illustrates the four types of REPs in a part of the camping assistance service network. Each role in the CampSAS service network has a routing REP and a synchronization REP, and each contract has a pass-through REP. The CampSAS service network, as a whole, also has a coordinated pass-through REP. Figure 6A shows three rules at the routing REP of the role. These rules control the admission of the assistance requests and initiate the send equipment requirements and send location role-role interactions toward roles EP1 and DC, respectively. Figure 6B shows a pass-through rule at the REP of the contract CM_DC, which evaluates the interaction message *iSendLocation* and generates the event *eSendLocToDC* to indicate the occurrence of the interaction. Figure 6C,D shows two synchronization rules: one at the role EP1 and another at the role DC. The first rule ensures that the equipment is rented based on the requirements provided by the camper, upon receiving the rental order from the 24by7 support center. The second rule ensures that MacDelivery service delivers the rented camping equipment to the location given by the camper. Figure 6E shows a regulation rule at the coordinated pass-through REP. It adjusts the used throughput for a VSN as its running instances end. The termination of a VSN

FIGURE 5 A part of the regulation design of CampSAS' service network [Colour figure can be viewed at wileyonlinelibrary.com]

instance is identified by the occurrence of the interaction *iDeliver* (response) between roles EP1 and DC (indicated by the event *eDeliveredToCM*), the interaction *iSendTourReservation* (response) between roles TC and CM (indicated by the event *eTourReservationSent*), and the interaction *iSendTaxiReservation* (response) between roles TX and CM (indicated by the event *eTaxiReservationSent*).

3.3 | Virtual service networks

To support the functional and performance requirements of a tenant, a VSN for the tenant is formed on the service network. A VSN represents one or more service compositions or business processes in the service network. Each such process has a specific service network path/topology. To enable and control the routing of the interaction messages over its service network paths, a VSN has a regulation policy, which refers to a subset of regulation rules at the REPs on the relevant paths. The virtualization of the service network is achieved by isolating the messages belonging to each VSN/process and routing them on the path of the VSN/process, while respecting the relevant regulation policy.

Multiple VSNs (of tenants) share some service network elements for their common requirements and use some other service network elements for their distinctive requirements. The service network elements include roles, contracts, tasks, interaction terms, REPs and their regulation rules, and partner services and their capabilities and throughput/capacities.

Figure 7 shows the service network topologies of the VSNs for the tenants HappyTours and UniUvtClub in our motivating scenario. The overlaps between the topologies show the sharing between the two VSNs. To support the similar requirements of taxi hiring and case handling, the two VSNs share the services 24by7Support and LocalCabs as well as the elements in the service network that correspond to the service collaborations that support these requirements. The

1	rule "rAdmissionCheckV1"	(A)
2	when	
3	\$msg : RoleServiceInteraction (opName == "requestAssist" , ...)	
4	then	
5	AdmissionControl("requestAssist" , "50", \$msg);	
6	end	
7	...	
8	rule "rRequestAssistV2"	
9	when	
10	\$msg : RoleServiceInteraction (opName == "requestAssist" , ...)	
11	then	
12	Forward ("CM_EP1.iSendEquipmentRequirements.Reg" ,	
13	Synthesize("iSendEquipmentRequirementsReqEP1.xsl" , \$msg);	
14	end	
15	rule "rRequestAssistV8"	
16	when	
17	\$msg : RoleServiceInteraction (opName == "requestAssist" , ...)	
18	then	
19	Forward ("CM_DC.iSendLocation.Reg" , Synthesize("iSendLocationReqDC.xsl" , \$msg);	
20	end	

1	rule "rSendLocation"	(B)
2	when	
3	\$msg : RoleRoleInteraction (opName == "iSendLocation" , ...)	
4	then	
5	// Examine the message content	
6	// Generate events indicating the state of the message	
7	PublishEvent ("eSendLocToDC" , \$msg);	
8	end	

1	rule "rRentEquipment"	(C)
2	when	
3	Event (id == "eOrderEquipmentFromEP1") and	
4	Event (id == "eSendEquipmentRequirementsToEP1")	
5	then	
6	RoleRoleInteraction [] in =	
7	Pull ("SC_EP1.iOrderEquipment.Reg, CM_EP1.iSendEquipmentRequirements.Reg");	
8	ExecuteTask(Synthesize (in , "tRentEquipment.xsl") , "tRentEquipment");	
9	end	

1	rule "rDeliver"	(D)
2	when	
3	Event (id == "eSendLocToDC") and Event (id == "eDelivertoCM")	
4	then	
5	RoleRoleInteraction [] in = Pull ("CM_DC.iSendLocation.Reg, EP1_DC.iDeliver.Reg");	
6	RoleServiceInteraction out = Synthesize(in , "tDeliver.xsl");	
7	ExecuteTask(out , "tDeliver");	
8	end	

1	rule "rMonitorV1"	(E)
2	when	
3	(\$e1 : Event (id == "eDeliveredToCM")) and (\$e2 : Event (id == "eTourReservationSent")	
4	and (\$e3 : Event (id == "eTaxiReservationSent"))	
5	then	
6	MonitorThroughput(\$e1.getFlowInfo());	
7	TerminateInstance (\$e1.getFlowInfo());	
8	end	

FIGURE 6 Some regulation rules at (A) the routing regulation enforcement point (REP) of the role CM, (B) the pass-through REP of the contract CM_DC, (C) the synchronization REP of the role EP1, (D) the synchronization REP of the role DC, and (E) the coordinated pass-through REP [Colour figure can be viewed at wileyonlinelibrary.com]

sharing of a service task also implies that of its throughput/capacity. To support the distinctive requirements of equipment rental and grocery delivery, the CampSAS' VSNs use different services and elements in the service network. HappyTours uses TomGear and MacDelivery, whereas UniUvtClub uses LocalGrocers. While both HappyTours and UniUvtClub need the tour arrangement capability, the capacity of the tour company ShowAround is limited. Moreover, the overall performance requirements of the two tenants (3 hours and 4 hours) are different. Thus, the two VSNs use different service collaborations for realizing the tour arrangement. A collaboration involving a vehicle rental company and a travel guide provider implements the tour arrangement for the tenant UniUvtClub.

The business processes in VSNs can be described and enacted as event-driven process chain processes.³⁰ Events are used to define the temporal dependencies between the tasks used by processes. The regulation rules that produce and/or consume events support the enactment of the processes. Figure 8 shows a part of a process in the VSN of the tenant

Kumara et al³ for more details on the designs of VSNs, their processes, and their sharing of the service network with appropriate functional and performance variations.

3.4 | Enactment of VSNs as controlled message routing in service network paths

To consume a business process of a tenant, an end user of the tenant needs to send a Web service request to the service network via its user role, which represents the end user (the role CM in our example). An instance of a business process in the tenant's VSN is created, and the created instance progresses by executing the relevant tasks (and consuming the relevant service capabilities) while respecting the control flow between the tasks. This creates interactions between services. These interaction messages are isolated from those of other process instances using the process instance identifiers included in the messages (as message headers). The isolated messages are routed in the service network path used by the business process. Below, we discuss the enactment of the business processes in VSNs in detail.

In the first stage, when the routing REP of the user role receives an instantiation Web service request, the routing REP locates the regulation rules for the VSN via the regulation table. The Web service request contains the identifier of the VSN encoded in a message header or the part of the Web service endpoint reference of the user role (similar to the use of the tenant identifier in the work of Pathirage et al⁹). The VSN identifier is the key to the regulation table. The selected rules are executed by providing the received message as inputs. The consequences of the execution of the rules depend on the regulation functions used by the rules and the execution order of the functions. In general, the admittance of an instantiation request (*AdmissionControl* function) is decided, and a business process for the admitted request is selected from the processes of the VSN to balance the load among processes (*LoadBalance* function). Next, an instance of the process is created and maintained. A process instance represents an execution of a business process for a user request. It is a state object (nonexecutable) that records the execution states of the tasks in the process. Afterward, the instantiation request is treated as a role-service interaction message and is augmented with process and process instance identifiers and reinjected into the routing REP of the user role. If an instantiation request cannot be admitted at a given moment, it can be dropped or queued by using suitable regulation functions (*Drop* and *Schedule*).

In the second stage, the created process instance progresses as an event-driven process chain process. On receiving a role-service interaction message belonging to a process instance, the routing REP of the role applies the regulation rules for the corresponding process. In general, the regulation rules may control the admittance of the interaction message, select a set of role-role interactions according to the service network path of the process instance, convert the service interaction into role-role interaction messages, and schedule the created messages to be sent to the appropriate destination roles via the pass-through REPs at the relevant contracts. When the pass-through REPs receive the role-role interaction messages, the pass-through rules belonging to the process are applied to them. These rules may examine the interaction messages, generate events indicating the states of the messages, and forward the messages to the synchronization REPs at the destination roles. When the synchronization REPs receive the role-role interaction messages, they store the messages as the execution of task may need to synchronize multiple such interaction messages. In general, the regulation rules at the synchronization REPs use a logical expression of events (generated at the pass-through REPs) as their triggering conditions. These rules retrieve the stored role-role interaction messages, create a role-service interaction message from the retrieved messages, and send the created message to the destination service to carry out the desired task and consume the service capability.

A service may send messages via its role to indicate the completion of the requested capability or to request new capabilities from other services. These messages from the services will trigger new cycles of routing, pass-through, synchronization, and task execution. These regulated message exchanges between services in the service network path of a process continue until the process instance is terminated, which is the final stage of the VSN enactment. A process instance may terminate naturally when all its tasks are completed or when a certain condition is met (eg, the completion of one or more specific tasks). In general, the regulation rules at the coordinated pass-through REP control the state of a process instance as multiple interactions between different pairs of services may need to be considered to decide the state of a process instance. Alternatively, the administrator can use the management capabilities provided by the service network to explicitly end a process instance.

Figure 9 illustrates the message flow over the service network for the business process shown in Figure 8. The annotated numbers in both figures show the correspondence between the progress of business process and the routing of messages. For example, the request assist message from the client is routed (at the role CM) to the roles DC, EP1, SC, TC, and TX (2a.1, 2b.1, 2c.1, 2d.1, and 2e.1 in Figure 9) via the relevant contracts (pass-through REPs). This generates the

TABLE 1 Organizational management application programming interface

Operation Name	Explanation
addRole (rId, rName, sBinding)	Add a new role (node) to the service network
removeRole (rId)	Remove a role from the service network
updateRole (rId, property, value)	Update a property of a role
addTask (rId, tId, inputs, outputs)	Add a new task to a role
removeTask (rId, tId)	Remove a task from a role
updateTask (rId, tId, property, value)	Update a property of a task
addContract (cId, roleA, roleB)	Add a new contract (link) between two roles
removeContract (cId)	Remove a contract from the service network
updateContract (cId, property, value)	Update a property of a contract
addTerm (cId, tmId, direction)	Add a new interaction term to a contract
removeTerm (cId, tmId)	Remove an interaction term from a contract
updateTerm (cId, tmId, property, value)	Update a property of an interaction term

Note. rId = role Id; rName = role name; sBinding = service binding; tId = task Id; cId = contract Id; tmId = term Id. The inputs and outputs of a task are references to interaction terms consumed/produced by the task. The direction of an interaction term indicates the direction of the message flow between the relevant two roles.

- **Operational management.** These management capabilities allow the changes to the regulation design of a service network, which include addition/removal/update of regulation rules, regulation table entries, and service network events/states. REPs are created or removed as part of the creation or removal of their placeholder elements (ie, roles, contracts, and service network). A process instance can also be paused, resumed, and terminated (by changing its management state). To monitor process instances, the service network events and states can be read by querying them or by subscribing to them. Table 2 shows the main operations of the operational management API. The parameter *key* in the update operations indicates what to update, for example, management state or regulation rule content. A service network flow can represent the messages belonging to a VSN, a process in a VSN, or a process instance.

To support the runtime changes to a multitenant service network in a controlled manner without compromising the consistency of the running service network, we adopted the change management scheme proposed by Kramer and Magee.²⁸ In general, an element in a service network (including a process instance) can be in three management states: *Active*, *Passive*, and *Quiescence*. The *Passive* state of an element enables the system to complete the existing process instances and to move the element to its *Quiescence* state. If a runtime change to an element can adversely affect some existing process instances, then the change must be delayed until the element reaches its *Quiescence* state. A newly added element is always in the *Passive* state and must be explicitly moved to the *Active* state so that new process instances can use the element. An element can be removed from the system when it is in *Quiescence* state. The events are generated at each management state change, allowing the organizational and operational managers to observe and make appropriate management decisions (ie, triggering ECA rules).

3.5.1 | Specification and enactment of management policies

To specify the organizational and operational management decisions, SDSN@RT provides a policy language by adopting the language of the Drools business rule engine. A policy consists of a set of ECA rules.

- **Conditions.** A policy rule needs to be able to react to service network events (eg, a management state change), service network state changes (eg, the unavailability of a service), the properties of VSNs and their instances, the time of the day or other context conditions, and the execution state of the management policies. Thus, a condition can be a logical expression of the predicates about the information in these sources.
- **Actions.** The actions can include the abovementioned organizational and operational management capabilities (eg, add role, remove contract, and add routing rule). Additionally, they can be custom management functions, eg, predicting response time violations.
- **Execution.** The correct ordering of the rules as well as that of the actions within each rule are required to achieve a desired outcome. The rules are independent and are activated based on their conditions. When multiple rules are

TABLE 2 Operational management application programming interface

Operation Name	Explanation
addSynchronizationRules (rId, ruleFile) <i>rId</i> = role Id	Add some rules to a synchronization regulation enforcement point (REP)
removeSynchronizationRules (rId, ruleIds)	Remove some rules from a synchronization REP
updateSynchronizationRule (rId, ruleId, key, value)	Update a synchronization rule
addRoutingRules (rId, ruleFile)	Add some rules to a routing REP
removeRoutingRules (rId, ruleIds)	Remove some rules from a routing REP
updateRoutingRule (rId, ruleId, key, value)	Update a routing rule
addPassthroughRules (cId, ruleFile) <i>cId</i> = contract Id	Add some rules to a pass-through REP
removePassthroughRules (cId, ruleIds)	Remove some rules from a pass-through REP
updatePassthroughRule (ruleId, key, value)	Update a pass-through rule
addCoordinatedPassthroughRules (ruleFile)	Add some rules to a coordinated pass-through REP
removeCoordinatedPassthroughRules (ruleIds)	Remove some rules from a coordinated pass-through REP
updateCoordinatedPassthroughRule (ruleId, key, value)	Update a coordinated pass-through rule
addSynchronizationTableEntries (snfId, entries) <i>snfId</i> = service network flow Id	Add some entries to the regulation tables at the synchronization REPs on a path of a service network flow
removeSynchronizationTableEntries (snfId, entries)	Remove some entries from the regulation tables at the synchronization REPs on a path of a service network flow
updateSynchronizationTableEntry (snfId, entry, key, value)	Update a synchronization regulation table entry belonging to a service network flow
addRoutingTableEntries (snfId, entries)	Add some entries to the regulation tables at the routing REPs on a path of a service network flow
removeRoutingTableEntries (snfId, entries)	Remove some entries from the regulation tables at the routing REPs on a path of a service network flow
updateRoutingTableEntry (snfId, entry, key, value)	Update a routing regulation table entry belonging to a service network flow
addPassthroughTableEntries (snfId, entries)	Add some entries to the regulation tables at the pass-through REPs on a path of a service network flow
removePassthroughTableEntries (snfId, entries)	Remove some entries from the regulation tables at the pass-through REPs on a path of a service network flow
updatePassthroughTableEntry (snfId, entry, key, value)	Update a pass-through regulation table entry belonging to a service network flow
addCoordinatedPassthroughTableEntries (snfId, entries)	Add some entries to the regulation table (belonging to a service network flow) at the coordinated pass-through REP
removeCoordinatedPassthroughTableEntries (snfId, entries)	Remove some entries from the regulation table at the coordinated pass-through REP
updateCoordinatedPassthroughTableEntry (snfId, entry, key, value)	Update a coordinated pass-through regulation table entry belonging to a service network flow
subscribe (eventPattern, subscriberEndpoint)	Subscribe to a set of service network events
unsubscribe (eventPattern, subscriberEndpoint)	Unsubscribe from a set of service network events
subscribe (stateIds, subscriberEndpoint)	Subscribe to a set of service network states
unsubscribe (stateIds, subscriberEndpoint)	Unsubscribe from a set of service network states

activated at the same time, the priorities of the rules can be used to resolve any conflict. Within a rule, if-then-else conditional constructs can be used to order the actions.

The policy manager in the management platform uses a policy engine to enact the management policies. It supports addition, removal, and update of policies. It can parse given policies, process events, state changes, and other provided information and activate the rules in the policies as their conditions are met.

Figure 10 shows a fragment of the organizational management policy that alters the camping assistance service network by incorporating the implementation of the bike rental feature. The implementation is a collaboration between the bike rental service (EuroBikes), the 24by7 support center, and the camper. To represent the bike rental service, a new role BR (a new node) is created in the service network. To represent the interconnections among the participants in the bike rental collaboration, the contracts or links, such as the contracts SC_BR and CM_BR, are added. The relevant tasks and the interaction terms are also included.

```

rule "org_mgt_rule_1"
when
  ($mgrp: ManagementPolicyState( id=="org_mgt", state=="incipient"))
then
  // role - level changes
  addRole( "BR", "EuroBikes",
    "http://localhost:8082/axis2/services/EuroBikeRentalService");
  addTask( "BR", "tRentBike");
  ...
  // Role SC
  addTask( "SC", "tSendRentInvoice");
  ...
  // contract - level changes
  addContract( "SC_BR", "SC", "BR");
  addTerm( "SC_BR", "iOrderBike");
  addTerm( "SC_BR", "iSendRentInvoice");

  addContract( "CM_BR", "CM", "BR");
  addTerm( "CM_BR", "iSendBikeRequirements");
  addTerm( "CM_BR", "iSendBikeReservation");
  ...
  $mgrp.setState( "done");
end

```

FIGURE 10 A fragment of the organizational management policies for adding the bike rental feature [Colour figure can be viewed at wileyonlinelibrary.com]

```

rule "opr_mgt_rule_1"
when
  ($mgrp1: ManagementPolicyState( id=="org_mgt", state=="done")) and
  ($mgrp2: ManagementPolicyState( id=="opr_mgt", state=="incipient"))
then
  // synchronization rule changes
  addSynchronizationRules( "BR", "BR_SYN.drl"); // all rules at the syn.REP
  addSynchronizationRules( "SC", "SC_SYN_Diff.drl");
  // routing rule changes
  addRoutingRules( "BR", "BR_Routing.drl"); // all rules at the routing REP
  addRoutingRules( "SC", "SC_Routing_Diff.drl");
  addRoutingRules( "CM", "CM_Routing_Diff.drl");
  // passthrough rule changes
  addPassthroughRules( "CM_BR", "CM_BR.drl"); // all rules at the passthrough REP
  addPassthroughRules( "SC_BR", "SC_BR.drl");
  UpdateCoordinatedPassthroughRule( "terminateV1", "replace", "CP_REP_Diff.drl");
  // configure regulation tables for HappyTours' VSN to include bike rental
  addSynchronizationTableEntries( "HappyTours", "rentBike:BR, sendRentInvoice: SC");
  addRoutingTableEntries( "HappyTours", "rentBikeResponse: BR,
    requestAssistV9: CM, analyzeResponseV8: SC, sendRentInvoiceResponse: SC");
  addPassthroughTableEntries( "HappyTours", "sendBikeRequirements: CM_BR,
    sendBikeReservation: CM_BR, OrderBike: SC_BR,
    sendRentInvoice: SC_BR, sendRentInvoiceResponse: SC_BR");
  // then, remove the table entries related to taxi hire from HappyTours VSN's path
  removeSynchronizationTableEntries( "HappyTours", "bookTaxi: TX, payTX: SC");
  ...
  $mgrp2.setState( "done");
end

```

FIGURE 11 A fragment of the operational management policies for adding the bike rental feature and configuring HappyTours' virtual service network to use it, replacing the taxi hire feature [Colour figure can be viewed at wileyonlinelibrary.com]

Figure 11 shows a fragment of the corresponding operational management policy. The regulation rules that are required to support the managed interactions between the participants in the bike rental collaboration are added to the relevant REPs. The individual rules are defined as Drools rule files (.drl). As the tenant HappyTours needs to use the new bike rental feature and to drop the taxi hire feature, the service network path of the HappyTours' VSN is modified by changing the relevant regulation tables (ie, adding and removing table entries).

3.6 | Prototype implementation

We built the prototype of the SDSN@RT middleware by adopting and further extending the ROAD/Serendip framework, which support adaptive service orchestrations.^{5,12} We used the Java programming language to develop the middleware.

The prototype follows the system architecture shown in Figure 3. We have implemented the coordination infrastructure by incorporating the support for the regulation and management capabilities into the service composition engine of ROAD. We have also extended the design and management of the service orchestrations in ROAD to fully support the configuration design of a multitenant service network. Each node (role) in a service network is an Axis2 Web service and uses the Axis2 Web service client API to send messages to and receive messages from the external service. As a Web service, a node has a Web Services Description Language (WSDL) and an endpoint reference. The external clients/systems can send the messages to the node using its WSDL and endpoint reference.

We have implemented and added the management platform as a module to Apache Axis2 Web service engine (axis.apache.org/axis2). The Drools rule engine and its language were used to define and execute the regulation rules and management policies. The SDSN implementation (including all samples) is available at <https://github.com/road-framework/SDSN>. The size of the project has 407 356 lines of code (as per Github GLOC on 3/11/2018), which include Java, XML, and Drools rule files. The base ROAD framework is also available at <https://github.com/road-framework> (ROADFactory and ROADWS sub-projects). The size of the base framework is 72 173 in lines of code.

4 | EVALUATION

Previously in evaluating our approach, we have carried out evaluations on performance differentiation, capacity utilization, and comparison to the MIMT model.³ We have implemented and used three different throughput-sharing schemes (selected from the existing works). Compared with the MIMT model, the SIMT model has been found to achieve a better utilization of the capacities or throughputs (eg, the maximum number of possible repairs at a repair shop at a given moment) of the business services with little performance overhead (see section 6.2 in the work of Kumara et al³). In the remainder of this section, we focus on the evaluation of our middleware design.

The evaluation of our middleware design consists of two case studies using our prototype and a performance study of our prototype. The former shows that our middleware design can meet the requirements for a middleware that can host SIMT composite cloud applications and facilitate the management of such applications. The latter evaluates the runtime overhead of the key capabilities of the middleware: deployment and hosting of SIMT applications, enactment of application variants, and management of SIMT applications. The case study resources are available at https://github.com/IndikaKuma/SDSN_SPE_2018.

4.1 | Case studies

As a proof of concept, we have fully developed the camping assistance SIMT application presented in this paper, using our approach and middleware. The services are Apache Axis2 Web services. We have also developed a roadside assistance case study (RoSAS), which was presented in our earlier publications on SDSN design support.^{3,31} Tables 3 and 4 show the software metrics of both case studies. We have analyzed the case studies by observing the execution of VSNs and services and by examining the logs generated by the middleware and the services (see the case study resources for the logs). All of the following exercises (for both case studies) were successfully completed with our middleware support.

- *Deployment of the service network.* First, the middleware creates the topology of the CampSAS/RoSAS service network in terms of runtime objects of roles, tasks, contracts, and interaction terms. The REPs for the service network are also created and subscribed to the event manager and the state manager. Second, the middleware populates the knowledgebases at REPs using the relevant regulation rules (Drools) that are defined as part of the regulation design of the CampSAS/RoSAS assistance service network.

TABLE 3 Software metrics for the camping assistance (CampSAS) case study

Artifact	LoC (Lines of Code)	Details
Services	945 (Java)	9 services
SIMT application	1186 (XML), 639 (Drools)	10 roles, 13 tasks, 17 contracts, 22 interaction terms, 78 regulation rules
Management policies	197 (Drools)	15 rules

Abbreviation: SIMT, single-instance multitenancy.

TABLE 4 Software metrics for the roadside assistance (RoSAS) case study

Artifact	LoC (Lines of Code)	Details
Services	1380 (Java)	8 services
SIMT application	1540 (XML), 950 (Drools)	9 roles, 21 tasks, 15 contracts, 17 interaction terms, 115 regulation rules
Management policies	280 (Drools)	18 rules

Abbreviation: SIMT, single-instance multitenancy.

- *Deployment of VSNs.* The middleware uses the VSN designs of the three tenants of CampSAS/RoSAS to populate the regulation tables at REPs with VSN-to-regulation-rule mappings.
- *Enactment of VSNs.* Upon receiving a Web service request from a user of a tenant at the client role (the role CM in the CampSAS service network), the enactment engine creates and executes an instance of the tenant's VSN. It can simultaneously enact the instances of the VSNs of multiple tenants on the same CampSAS/RoSAS service network, while ensuring the isolation among VSN instances.
- *Runtime update of the multitenant service network.* We applied the organizational and operational management policies presented in Section 3.5 to the running CampSAS service network. For the RoSAS service network, the runtime evolution scenario was adding the taxi hire feature. The online case study resources include the complete management policies for both CampSAS and RoSAS. Via the logs and the response messages of VSN executions, we validated the controlled modification of the service network. The relevant logs in the case study resources are annotated to highlight this runtime behavior. We also compared the logs and VSN responses with those of the manually created same service network to validate the changes to the service network.
- *Undeployment of the VSNs and the service network.* We first undeployed the three VSNs one by one (via management policies) and then removed the service network (via the deployment interface).
- *Coverage of runtime changes.* To validate our support for different types of runtime changes to a multitenant service network, we created and updated the same service network by incrementally adding the realizations of 10 features, all at runtime. These features are *Case Handling*, *Reimbursement*, *Towing*, *Repairing*, *Spare Parts*, *Vehicle Repair Assessment*, *Legal Assistance*, *Accommodation*, *Rental Vehicle*, and *Taxi Hire*. These features are common to both the domains of camping assistance and roadside assistance, for example, repairing/towing a car/bike/caravan and reserving special accommodation for some campers. We also considered the rollback (removal) of each feature. In general, a feature is realized by a collaboration among a subset of the services. Adding a feature requires changes to both the topology and regulation designs of the service network. For example, the addition of the feature Case Handling requires one service and one client, two roles, four tasks, one contract, one interaction term, one VSN, and 12 regulation rules (for the four types of REPs). These scenarios together cover each type of service network changes at least one time.

To realize a given feature, we first identify the differences between the current service network and the target service network after the implementation of the feature, in terms of service network changes. Then, we create the organizational and operational management policies and the regulation rules to realize the identified changes. Next, we apply the created policies at runtime on the current service network. We validated the implementation of the scenario following the same procedure used to validate the updates to the CampSAS/RoSAS service network (see above).

- *Deployment and enactment of multiple service networks.* We deployed 10 different multitenant service networks on the same middleware instance and simultaneously executed each of the VSNs in the service networks successfully. To create these service networks, we have incrementally used the abovementioned 10 features. The k th multitenant service network realizes the first to k th features ($1 \leq k \leq 10$). Each service network has a single VSN whose enactment covers the whole service network (ie, execute each task and rule).

4.2 | Performance evaluation

In these experiments, we quantify the absolute and relative performance overhead induced by our middleware design and its prototype. We consider the runtime overhead of deploying service networks and VSNs, enacting VSNs, and modifying service networks and VSNs. Our experiment setup consists of a workload generator (Apache AB), middleware (SDSN@RT), and Web services (Apache Axis2). Each system is deployed in a virtual machine with four 3400-MHz CPUs (Intel Broadwell), 8-GB memory, and Ubuntu 16.04.4 LTS. The Java version is Oracle JDK 1.8.0-161. To measure the memory usage, we used the top utility in Ubuntu.

FIGURE 12 Overhead of service network deployment. A, Deployment time vs size for service network; B, Physical memory usage vs size for service network [Colour figure can be viewed at wileyonlinelibrary.com]

4.2.1 | Deployment overhead

We used the 10 different SIMT applications mentioned in Section 4.1. For each SIMT application, we deployed 50 copies of the service network of the SIMT application and took the average deployment time and physical memory consumption. We used 100 copies of the VSN in the SIMT application to quantify the average deployment time and physical memory usage for VSNs.

Figure 12 shows the deployment time and memory consumption for service networks with respect to their sizes in LoC (Lines of Code). Each of two measures increases approximately linearly. The size of a service network (approximately) proportionally corresponds to the number of runtime objects created during the deployment. Thus, as the size of the service network increases, its deployment time and memory usage also increase.

In general, the deployment of a VSN simply adds a set of regulation table entries (the mappings between VSN identifiers and regulation rule identifiers). Thus, the deployment overheads for the VSNs are approximately similar. The deployment time of a VSN has a mean of 3.374 ms and a standard deviation of 0.23 ms. The memory usage has a mean of 12.76 KB and a standard deviation of 2.09 KB.

4.2.2 | Enactment overhead

To quantify the execution overhead for VSNs, we used the same 10 SIMT applications mentioned in Section 4.1. In these scenarios, the enactment of the VSN covers the entire service network for the SIMT application. As we only quantify the overhead included by our middleware, we exclude the service execution time from the enactment time of the VSN. We first issued 100 (with the concurrency level of 10) warm-up calls. The concurrency level is the number of multiple requests to perform at a time (see Apache AB documentation). Next, we issued 2000 calls (with the concurrency level of 100) and measured the average response time at the client side. Figure 13 shows the enactment time (across 100 concurrent requests) for VSNs against the size of the service network. The enactment overheads increase approximately linearly with the size of the service network path that the VSN covers.

4.2.3 | Management overhead

To quantify the runtime management overhead, we used the 10 change scenarios mentioned in Section 4.1, where each scenario adds a new feature. For each scenario, we have measured the runtime change enactment time (RCET), which is the time difference between the middleware receiving the management policies and its being ready for use after applying policies. As active VSN instances can delay the execution of the management policies, we enacted the policies when there are no active VSN instances.

Figures 14 and 15 show the RCET and size of management policies for the realization and rollback of changes. The size of the management policy also includes the regulation rules added or updated. For the 10 scenarios, the RCET for the realization of changes is within 730–939 ms, and that for the rollback of the realizations is within 73–773 ms. We believe that this is reasonable. In general, RCET depends on the number of atomic change operations (which approximately

FIGURE 13 Overhead of virtual service network (VSN) enactment (across 100 concurrent requests) [Colour figure can be viewed at wileyonlinelibrary.com]

FIGURE 14 Runtime change enactment time (RCET) and management policy size for change scenarios: feature addition or change realization [Colour figure can be viewed at wileyonlinelibrary.com]

FIGURE 15 Runtime change enactment time (RCET) and management policy size for change scenarios: feature removal or change rollback [Colour figure can be viewed at wileyonlinelibrary.com]

corresponds to the size of the management policies) as well as the types of atomic change operations (addition, update, or removal). The addition of an object (eg, role or a regulation rule) is expensive compared to the removal of the corresponding object. The rollback of a change realization can involve the addition of removed objects, the removal of added objects, and the update of the changed objects.

4.2.4 | Relative enactment overhead

In our earlier publications on our SDSN design support,^{3,31} we have compared the SIMT model in our SDSN approach and the SIMT model in BPEL (Business Process Execution Language)³² with Apache ODE (ode.apache.org) in terms of runtime enactment overhead.* We measured the end-to-end application enactment time (including service execution times) at the client side. We have used the roadside assistance case study in this experiment. Our SDSN middleware is shown to have similar enactment overhead to Apache ODE.

5 | RELATED WORK

The functional and performance differentiation in an SIMT composite application requires support for the runtime sharing and variation of services and the configuration design (architectural composition of services) and regulation design (regulation of conversational behaviors between services) of the application. A middleware for hosting SIMT applications needs to be able to create and maintain the runtime models of the applications, simultaneously enact the application variants in a controlled manner, and enable the runtime management of the applications and their variants. This section compares our middleware with related works in middleware for component/service compositions and multitenant applications.

5.1 | Middleware for composite applications

Middleware has been proposed for two main types of composite applications: component-based applications and service-based applications. The former is a precursor to the latter. For a comparison of these two types of applications, interested readers are referred to the relevant studies.³³⁻³⁵ This research focuses on service-based applications. With service orientation, the implementations of the services are hidden from service consumers (the services are black boxes). They have discoverable interfaces (eg, WSDL) and are consumed by sending messages through Web technologies. The composition of services is done through coordinating interactions (message exchanges) between services. The main coordination model types are orchestration and choreography. The orchestration model requires a special middleware, and this paper has presented a service orchestration middleware for service-based single-instance multitenant cloud applications.

Most middleware systems for component-based applications^{4-6,36,37} can deploy a specification of a component assembly. The deployment generally creates runtime models. These middleware systems also provide management capabilities for these runtime models such as intercepting and monitoring message exchanges between components, injecting extensions/plugin-ins during invocation at predefined hooks (extension points), and modifying runtime models. The component models supported by these middleware systems include a subset of the common software architectural abstractions such as components, connectors, composites, and (provided and required) interfaces.

Similarly, most existing middleware systems for composite service applications³⁸⁻⁴⁰ can enact a specification of the composite service by creating runtime models for representing the composite service. The runtime management capabilities such as discovering services, monitoring message exchanges between services, replacing services, modifying the definition and instances of the composite applications, and providing quality of service for services by managing computing resources used by the middleware. Common service composite models supported by these middleware systems include process-centric models (eg, BPEL or Business Process Model and Notation⁴¹), component-based models,⁴ and organization-based models.⁵ The process-centric models consider how business activities or interactions are organized and ordered to achieve business goals. Component-based models define composite applications as assemblies of components that provide and request services. Organization-based models represent service compositions as goal-oriented

*This exercise was for comparing the middleware performance only. The SIMT model in BPEL does not have the dynamic sharing and adaptation capabilities of the SIMT model in our SDSN approach.³

organizations using roles, players realizing roles, interaction and normative relationships between roles, and controller(s) managing the organizations.

The service network model can also be considered as a service composition model. It can provide a natural abstraction to design, enact, regulate, and manage webs of real-world business service networks.^{26,27} Most existing works on service networks consider the modeling and analysis of service networks from specific aspects²⁷ such as business value flow⁴² and service relationships.²⁶ Their realizations have relied on process-centric models, which fail to represent service networks naturally. That is, the domain concepts and their representations are mismatched, and the domain concepts (eg, services, service capabilities, service relationships, service interactions, interaction routing and regulation, service network paths, and virtualization) are not directly represented or managed in the realization, limiting their utility. To the best of our knowledge, SDSN@RT is the first middleware to support the deployment, enactment, and management of service-based applications designed following the (extended) service network model. Note that, recently, the concept of service mesh^{43,44} was used to describe the interservice communication infrastructure for microservices. However, compared with the service network model, the service mesh model does not have the abstractions required to represent the abovementioned domain concepts of multitenant business service networks.

5.2 | Middleware for multitenant applications

There exist middleware systems for multitenant service-based applications as well as multitenant component-based applications. They can be further categorized based on two multitenant models: MIMT and SIMT. Their capabilities and approaches to support functional and performance differentiation for tenants also differ. A few of these works support runtime changes to multitenant applications.

Some process engines⁹ and enterprise service buses⁴⁵ implement the MIMT model for composite service applications. The service provider hosts the middleware, and tenants create and manage their own services and processes. The tenant-specific metadata are used to customize software elements (services or enterprise service bus components), and each created variant is separately deployed and managed (ie, the MIMT model). Tenants can define their applications (processes or message mediation flows) that use their own customized software elements. To achieve the performance differentiation among tenants, the regulation functions such as global admission controllers or schedulers are used to fairly allocate middleware-level computing resources among tenants.¹⁷

We identify three approaches to enacting the variant instances of an SIMT composite application for different tenants and their users: generative approach, dynamic-binding approach, and virtualization (slicing) approach.

The generative approach^{16,46,47} derives application variants by customizing the composite application for each user request and deploys and enacts the derived variant. This approach can utilize the existing works on the customization support for composite applications. In our earlier works, we also supported customizing service compositions⁴⁸ and multitenant applications.⁴⁹ The composite service application is represented as a customizable application, which includes all allowed variations. The model-driven engineering techniques are used to resolve variability and to generate executable application variants. An alternative approach is to select a subset of application fragments (eg, process fragments⁵⁰ or Java code fragments⁵¹) and merge them to form the application variant. The generative approach can potentially incur a considerable runtime overhead (variant generation time + deployment and compilation time + enactment time) as the variability and size of the composite cloud application increase. Changing the design of the application variant while it is being enacted requires the capability to modify the running application instance or regeneration of a new application variant and migrating the running instance to the new variant.

In the dynamic-binding approach,^{8,10,15} specific elements are selected, attached, and executed at the extension points of the application, for example, dynamically selecting/binding a service, executing ECA rules, selecting a control branch in a business process, and injecting a component via dependency injection. At the design time, a software engineer needs to annotate the extension points in the application specification and codify the mappings between tenant requirements/features, extension points, and choices at each extension point. The middleware needs to deploy such application specifications and enact application variants at runtime by resolving these mappings (runtime variability resolution). A tenant selects a set of extension points and the choices for them. Using tenant information, the composite application identifies the extension points for the tenant and binds/selects the relevant choices at them during the execution of the application.

There are several limitations with the existing works that use the dynamic-binding approach. First, a tenant requirement or feature is generally not mapped to a single component or service but to a collaboration among a set of components or

services.^{3,13,52} Thus, activating/deactivating a feature requires activating/deactivating a collaboration, which requires the ability to dynamically change the architectural topology of the application variant and the conversational behaviors within it. This capability is not sufficiently supported by the existing works. Second, the performance differentiation requires the regulation of interactions (message or object exchanges) between services/components to dictate progression of the instances of the application variants. The application should be able to execute tenant-specific regulation policies (with appropriate sharing of policies between tenants) to enforce the tenant-specific regulations on each service interaction. These capabilities are also not sufficiently available in the related studies. Third, most of these approaches annotate the variation points and/or their options within the composite application at design time, which can increase the design complexity and reduce the maintainability of the application.¹³ For example, it is difficult to represent modular abstractions such as collaborations, and it involves the tangling of normal processing logic with variant differentiation logic.

In the virtualization approach,¹² the complete virtual variants of the shared composite application are formed per tenant and/or per user request. The middleware maintains the runtime model of the application (models@runtime). The enactment of an application instance executes a subset of the elements of the shared runtime model of the application (no recompilation or changes to the runtime model). In one of our earlier works,¹² the virtualization approach was applied to realize the SIMT model for the composite cloud applications designed following the organization-based service composition model. The support for runtime sharing with variations was limited to the configuration design of the application. Moreover, an assumption was made that changing service providers (role-players) can be done by changing role-player/service bindings, which is not always valid for business services.

To the best of our knowledge, we are the first to support the SIMT model for composite service applications with the virtualization approach and the service network model. We extend the service network model with virtualization and regulatory control of the conversational behaviors in the service network. SDSN@RT creates and maintains the runtime model of a multitenant service network. This runtime model acts as an overlay network over the external services. The REPs in the overlay network make the message routing over it explicit, visible, and controllable at the appropriate granularity. With managed dynamic routing, VSNs can be dynamically formed per tenant and user request. The organizational and operational managers of the service network can change its topology and the routing and regulation logic at REPs at runtime. This virtualization of a composite application is analogous to the virtualization of a computer network. Virtual networks share the same physical computer network, each virtual network has a subtopology of the physical network, and the network packets are routed and regulated in the subtopology according to the policies defined for the virtual network. Some approaches to software-defined networking (SDN)-based network virtualization apply the concept of slicing a physical network via controlled and programmable routing of network data packets.⁵³

Intercepting, regulating, and routing message/object/data exchanges between applications have been used by middleware systems for component-based applications^{6,54} and service-based applications.^{4,55} In general, these works focus on dynamically selecting a specific service/component, intercepting method calls to add crosscutting functionalities, or regulating messages to a specific service/component. However, in SDSN@RT, dynamic routing is used to form complete managed virtual application variants in an SIMT composite service application. It creates the virtual slices of services and the configuration and regulation designs of the composite service application to support the functional and performance requirements of different tenants. Moreover, in the traditional reflective middleware for distributed object systems, the interceptors are generally registered at the middleware level to intercept method calls of any application. In SDSN@RT, each composite application has its own intercepting and routing capabilities, which can be specified in the application design by a software engineer. The sharing and variations of these capabilities among multiple tenants/application variants are defined during the application design and enforced during execution by the middleware.

Moreover, as mentioned above, in general, a collaboration among a set of services realizes a high-level tenant requirement or a feature. Thus, to form an application variant for a tenant, a set of collaborations needs to be composed/wired, and the interactions in the collaborations need to be regulated (locally and globally). Furthermore, the business services are heterogeneous. Even the services in the same category (see our motivating example) can have different capabilities, capacities, and relationships with other services. Thus, replacing a service cannot be done simply by changing the service binding. The message routing in the application needs to be altered. If an existing service needs to be replaced with a new service, the runtime model of the application needs to be modified structurally. SDSN@RT uses the models@runtime, service network model, and managed dynamic routing (via the REPs) in the service network to support the creation of VSNs (application variants) by composing regulated and managed collaborations (between business services) dynamically. The domain-specific language of SDSN@RT supports the design of an SIMT application following a compositional approach, where the configuration and regulation designs of the service network are decomposed into modular units by

TABLE 5 A summary of the comparative analysis of the related works (including our SDSN@RT approach)

Criteria\Approach		8	12	16	10	15	43	9	SDSN@RT
Service/component orientation		CO	SO	SO	SO	SO	SO	SO	SO
Service composition model		–	OR	PC	PC	PC	SM	PC	Service network model extended with collaborations, processes, virtualization, regulation of conversational behaviors, and management/control
Multitenancy model		SIMT	SIMT	MIMT	SIMT	SIMT	MIMT	MIMT	SIMT
Runtime variation mechanism		DB	VR	CU	DB	DB	DB	–	Virtualization (VR) via dynamic routing
Services	Sharing with variations	–	~	~	+	–	~	~	+
	Business services	–	~	–	~	~	~	~	+
Configuration design	Models@runtime	~	+	–	~	~	~	–	+
	Sharing with variations	~	+	–	~	–	–	–	+
Regulation design	Models@runtime	–	–	–	–	~	~	–	+
	Granularity	–	–	–	–	FG	FG	–	FG
	Sharing with variations	–	–	–	–	–	–	–	+
	Integration with configuration structure	–	–	–	–	~	~	–	+
Runtime change support (SIMT model)		–	~	–	–	–	–	–	+
Middleware DSL (SIMT model)		–	~CV	–	+ AV	~AV	–	–	+ CV

Abbreviations: +, supported; –, not supported; ~, partially supported; AV, annotative variability; CB, component-based; CU, customization; CV, compositional variability; DB, dynamic-binding; DSL, domain-specific language; FG, fine-grained regulation; GB, global regulation; MIMT, multi-instance multitenancy; OR, organization-based; PC, process-centric; SIMT, single-instance multitenancy; SM, service mesh; VR, virtualization/slicing.

using the collaboration over services as the abstraction (see the work of Kumara et al³ for more information). Through the improved modularity, the compositional approach can potentially alleviate the drawbacks of the annotative approach.¹³

To support the runtime evolution of a multitenant composite cloud application, the middleware needs to support the potential classes of runtime changes to each key abstraction in the application. Moreover, the runtime models (models@runtime) of the application need to be maintained to enable changes to the application without redeploying, recompiling, or halting it. The existing works^{56–58} are limited to the changes in configuration designs and do not support the fine-grained changes to the runtime models of regulation designs. They support changes at the level of the choices for the variation points such as the addition and removal of choices, versioning of services or process definitions, the addition and removal of alternative services of the same type, and the modification of service dependencies.

Table 5 summarizes the comparison of SDSN@RT with the related works. SDSN@RT realizes the SIMT model for business service compositions. It supports the deployment, enactment, and runtime change management of SIMT composite applications that are modeled using the extended service network model. An SIMT application is represented as a multi-tenant service network, where a set of VSNs (per each tenant) coexists on the same service network at runtime. The REPs at the service network make the message passing over the service network observable and controllable at the appropriate granularity. At runtime, the application variants with different configuration and regulation designs are dynamically formed via routing, slicing the runtime model of the service network. The network view or controlled message passing between entities can enhance the capability to control conversational behaviors (interactions and message exchanges among participants) and dynamically form application variants in composite applications, thus supporting such management goals as virtualization and dynamic policy enforcement (as in flexible network-based models^{23,24} and dynamic routing in computer networks). SDSN@RT covers the potential classes of runtime changes to each of the key architectural abstractions in a multitenant service network. The architectural abstractions are represented explicitly using the models@runtime concept to enable the runtime changes to them. A software engineer can specify the desired runtime changes as policies consist of ECA rules and enact the created policies at runtime without halting the running service network (or the SIMT application).

6 | CONCLUSION AND FUTURE WORK

SDSN@RT is a middleware system that can deploy, enact, and manage SIMT composite SaaS applications. SDSN@RT represents an SIMT application as a multitenant service network, where a set of VSNs simultaneously coexists on the same physical service network at runtime. Each VSN realizes the functional and performance requirements of a particular tenant while sharing the underlying service network with other VSNs. The service network architecturally connects a set of services and enables disciplined and fine-grained regulation of the interactions between services. VSNs may overlap as well as differ in terms of services and their capabilities, service network topologies, and regulation policies being applied to the interaction messages passing over the topologies. SDSN@RT can create and maintain the runtime models of the service network and enact VSNs for multiple tenants over the shared service network simultaneously and in a controlled manner. It also facilitates the runtime changes to the service network and its VSNs. We have evaluated SDSN@RT and its support for enacting and managing SIMT composite cloud applications with a prototype, a case study, and a performance study. The experiment results have shown that SDSN@RT can deploy and execute SIMT applications and support the runtime modifications to the deployed applications, all with little performance overhead.

We have been exploring further such issues as regulation functions and strategies for enacting VSNs with the optimal service sharing while meeting tenant expectations, VSN instance migration, runtime evolution of the multitenant service networks for performance-related changes (eg, performance violation and abnormal termination of VSN instances), and high-level management of multitenant service networks (eg, high-level abstractions and policies). Another future work is to improve SDSN@RT to support the deployment, enactment, and management of multitenant service networks in data centers with greater scalability and efficiency. We also plan to explore further aspects of SDSN@RT as our approach evolves over time, especially through the research insights concerning the dynamic management of communication networks from the SDN²³ and conventional computer network communities. We also plan to apply SDSN (or SDN concepts) to compose and manage microservices in a way that a logically centralized control is realized over choreography-style decentralized service interactions.

ACKNOWLEDGEMENTS

This research was supported in part by the Smart Services Cooperative Research Center (CRC) through the Australian Government's CRC Program, by Swinburne University of Technology through research training support, and by the European Commission under grant 825480 (Horizon 2020, SODALITE).

ORCID

Indika Kumara <https://orcid.org/0000-0003-4355-0494>

REFERENCES

- Guo CJ, Sun W, Huang Y, Wang ZH, Gao B. A framework for native multi-tenancy application development and management. Paper presented at: 9th IEEE International Conference on E-Commerce Technology and 4th IEEE International Conference on Enterprise Computing, E-Commerce and E-Services (CEC-EEE); 2007; Tokyo, Japan.
- Chong F, Carraro G. *Architecture Strategies for Catching the Long Tail*. Redmond, WA: MSDN Library Microsoft Corporation; 2006.
- Kumara I, Han J, Colman A, Kapuruge M. Software-defined service networking: performance differentiation in shared multi-tenant cloud applications. *IEEE Trans Serv Comput*. 2017;10(1):9-22.
- Seinturier L, Merle P, Rouvoy R, Romero D, Schiavoni V, Stefani JB. A component-based middleware platform for reconfigurable service-oriented architectures. *Softw Pract Exper*. 2011;42(5):559-583.
- Colman A, Han J. Using role-based coordination to achieve software adaptability. *Sci Comput Program*. 2007;64(2):223-245.
- Truyen E, Vanhaute B, Jørgensen BN, Joosen W, Verbaeten P. Dynamic and selective combination of extensions in component-based applications. In: *Proceedings of the 23rd International Conference on Software Engineering (ICSE)*; 2001; Toronto, Canada.
- Truyen E, Cardozo N, Walraven S, et al. Context-oriented programming for customizable SaaS applications. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing*; 2012; Trento, Italy.
- Walraven S, Truyen E, Joosen W. A middleware layer for flexible and cost-efficient multi-tenant applications. In: *Middleware 2011: ACM/IFIP/USENIX 12th International Middleware Conference, Lisbon, Portugal, December 12-16, 2011. Proceedings*. Berlin, Germany: Springer; 2011:370-389.
- Pathirage M, Perera S, Kumara I, Weerasiri D, Weerawarana S. A scalable multi-tenant architecture for business process executions. *Int J Web Serv Res*. 2012;9(2):21-41.

10. Gey F, Walraven S, Van Landuyt D, Joosen W. Building a customizable business-process-as-a-service application with current state-of-practice. In: *Software Composition: 12th International Conference, SC 2013, Budapest, Hungary, June 19, 2013. Proceedings*. Berlin, Germany: Springer; 2013:113-127.
11. Mietzner R, Leymann F, Papazoglou MP. Defining composite configurable SaaS application packages using SCA, variability descriptors and multi-tenancy patterns. Paper presented at: 2008 Third International Conference on Internet and Web Applications and Services; 2008; Athens, Greece.
12. Kapuruge M, Han J, Colman A. *Service Orchestration as Organization: Building Multi-Tenant Service Applications in the Cloud*. Waltham, MA: Morgan Kaufmann; 2014.
13. Kästner C, Apel S, Kuhlemann M. Granularity in software product lines. In: *Proceedings of the 30th International Conference on Software Engineering*; 2008; Leipzig, Germany.
14. Reichert M, Hallerbach A, Bauer T. Lifecycle management of business process variants. In: *Handbook on Business Process Management*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2015:251-278.
15. Wang MX, Bandara KY, Pahl C. Process as a service distributed multi-tenant policy-based process runtime governance. Paper presented: 2010 IEEE International Conference on Services Computing; 2010; Miami, FL.
16. Geebelen K, Walraven S, Truyen E, et al. An open middleware for proactive QoS-aware service composition in a multi-tenant SaaS environment. In: *Proceedings of the International Conference on Internet Computing*; 2012; Las Vegas, NV.
17. Patikirikorala T, Kumara I, Colman A, et al. Dynamic performance management in multi-tenanted business process servers using non-linear control. In: *Service-Oriented Computing: 10th International Conference, ICSOC 2012, Shanghai, China, November 12-15, 2012. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2012:206-221.
18. Lin H, Sun K, Zhao S, Han Y. Feedback-control-based performance regulation for multi-tenant applications. In: *Proceedings of 15th International Conference on Parallel and Distributed Systems*; 2009; Shenzhen, China.
19. Krebs R, Momm C, Kounev S. Metrics and techniques for quantifying performance isolation in cloud environments. *Sci Comput Program*. 2014;90:116-134.
20. Li XH, Liu TC, Li Y, Chen Y. SPIN: service performance isolation infrastructure in multi-tenancy environment. In: *Service-Oriented Computing - ICSOC 2008: 6th International Conference, Sydney, Australia, December 1-5, 2008. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2008:649-663.
21. Krebs R, Loesch M, Kounev S. Platform-as-a-service architecture for performance isolated multi-tenant applications. In: *Proceedings of IEEE 7th International Conference on Cloud Computing (CLOUD)*; 2014; Anchorage, AK.
22. Blair G, Bencomo N, France RB. Models@run.time. *Computer*. 2009;42(10):22-27.
23. Kreutz D, Ramos FM, Verissimo PE, Rothenberg CE, Azodolmolky S, Uhlig S. Software-defined networking: a comprehensive survey. *Proc IEEE*. 2015;103(1):14-76.
24. Kagal L, Finin T. Modeling conversation policies using permissions and obligations. *Auton Agents Multi Agent Syst*. 2007;14(2):187-206.
25. Kramer J, Magee J. Dynamic configuration for distributed systems. *IEEE Trans Softw Eng*. 1985;SE-11(4):424-436.
26. Danylevych O, Karastoyanova D, Leymann F. Service networks modeling: an SOA & BPM standpoint. *J Univers Comput Sci*. 2010;16(13):1668-1693.
27. Razo-Zapata IS, De Leenheer P, Gordijn J, Akkermans H. Service network approaches. In: *Handbook of Service Description: USDL and Its Methods*. Berlin, Germany: Springer Science+Business Media; 2012:45-74.
28. Kramer J, Magee J. The evolving philosophers problem: dynamic change management. *IEEE Trans Softw Eng*. 1990;16(11):1293-1306.
29. Forgy CL. Rete: a fast algorithm for the many pattern/many object pattern match problem. In: *Readings in Artificial Intelligence and Databases*. Burlington, MA: Morgan Kaufmann; 1988:547-559.
30. Scheer A-W, Thomas O, Adam O. Process modeling using event-driven process chains. In: *Process-Aware Information Systems: Bridging People and Software Through Process Technology*. Hoboken, NJ: John Wiley & Sons; 2005:119-145.
31. Kumara I, Han J, Colman A, Kapuruge M. Software-defined service networking: runtime sharing with performance differentiation in multi-tenant SaaS applications. Paper presented at: 2015 IEEE International Conference on Services Computing; 2015; New York, NY.
32. Alves A, Arkin A, Askary S, et al. Web services business process execution language version 2.0. Burlington, MA: OASIS; 2006.
33. Papazoglou MP, Van Den Heuvel WJ. Service oriented architectures: approaches, technologies and research issues. *VLDB J*. 2007;16(3):389-415.
34. Vogels W. Web services are not distributed objects. *IEEE Internet Comput*. 2003;7(6):59-66.
35. Baker S, Dobson S. Comparing service-oriented and distributed object architectures. In: *On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE: OTM Confederated International Conferences, CoopIS, DOA, and ODBASE 2005, Agia Napa, Cyprus, October 31 - November 4, 2005, Proceedings, Part I*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2005:631-645.
36. Morin B, Barais O, Jezequel JM, Fleurey F, Solberg A. Models@run.time to support dynamic adaptation. *Computer*. 2009;42(10):44-51.
37. Boyer F, Gruber O, Pous D. Robust reconfiguration protocol for the dynamic update of component-based software systems. *Softw Pract Exper*. 2017;47(11):1729-1753.
38. Al-Jaroodi J, Mohamed N. Service-oriented middleware: a survey. *J Netw Comput Appl*. 2012;35(1):211-220.
39. Loyall JP, Gillen M, Paulos A, et al. Dynamic policy-driven quality of service in service-oriented information management systems. *Softw Pract Exper*. 2011;41(12):1459-1489.
40. Lemos AL, Daniel F, Benatallah B. Web service composition: a survey of techniques and tools. *ACM Comput Surv*. 2016;48(3). Article No. 33.

41. Object Management Group (OMG). Business Process Model Notation (BPMN) version 2.0. OMG Specification. 2011.
42. Allee V. Value network analysis and value conversion of tangible and intangible assets. *J Intellect Cap.* 2008;9(1):5-24.
43. Istio Framework. <https://istio.io/>. Accessed January 10, 2019.
44. Jamshidi P, Pahl C, Mendonça NC, Lewis J, Tilkov S. Microservices: the journey so far and challenges ahead. *IEEE Softw.* 2018;35(3):24-35.
45. Strauch S, Andrikopoulos V, Sáez SG, Leymann F. Implementation and evaluation of a multi-tenant open-source ESB. In: *Service-Oriented and Cloud Computing: Second European Conference, ESOC 2013, Málaga, Spain, September 11-13, 2013. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2013:79-93.
46. Blake M, Tan W, Rosenberg F. Composition as a service [web-scale workflow]. *IEEE Internet Comput.* 2010;14(1):78-82.
47. Pasquale L, Guinea S, Baresi L. Service-oriented dynamic software product lines with DyBPEL. *Computer.* 2012;45(10):42-48.
48. Nguyen T, Colman A, Han J. Modeling and managing variability in process-based service compositions. In: *Service-Oriented Computing: 9th International Conference, ICSOC 2011, Paphos, Cyprus, December 5-8, 2011 Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2011:404-420.
49. Kumara I, Han J, Colman A, Nguyen T, Kapuruge M. Sharing with a difference: realizing service-based SaaS applications with runtime sharing and variation in dynamic software product lines. Paper presented at: 2013 IEEE International Conference on Services Computing; 2013; Santa Clara, CA.
50. Eberle H, Unger T, Leymann F. Process fragments. In: *On the Move to Meaningful Internet Systems: OTM 2009: Confederated International Conferences, CoopIS, DOA, IS, and ODBASE 2009, Vilamoura, Portugal, November 1-6, 2009, Proceedings, Part I*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2009:398-405.
51. Walraven S, Lagaisse B, Truyen E, Joosen W. Policy-driven customization of cross-organizational features in distributed service systems. *Softw Pract Exper.* 2013;43(10):1145-1163.
52. Haesevoets R, Weyns D, Holvoet T. Architecture-centric support for adaptive service collaborations. *ACM Trans Softw Eng Methodol.* 2014;23(1):1-40.
53. Sherwood R, Gibb G, Yap KK, et al. Can the production network be the testbed? Paper presented at: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI); 2016; Savannah, GA.
54. Schmidt DC. Middleware for real-time and embedded systems. *Commun ACM.* 2002;45(6):43-48.
55. Canfora G, Di Penta M, Esposito R, Perfetto F, Villani ML. Service composition (re) binding driven by application-specific QoS. In: *Service-Oriented Computing - ICSOC 2006: 4th International Conference, Chicago, IL, USA, December 4-7, 2006. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2006:141-152.
56. Van Landuyt D, Gey F, Truyen E, Joosen W. Middleware for dynamic upgrade activation and compensations in multi-tenant SaaS. In: *Service-Oriented Computing: 15th International Conference, ICSOC 2017, Málaga, Spain, November 13-16, 2017, Proceedings*. Cham, Switzerland: Springer International Publishing; 2017:340-348.
57. Kapuruge M, Han J, Colman A, Kumara I. ROAD4Saas: scalable business service-based SaaS applications. In: *Advanced Information Systems Engineering: 25th International Conference, CAiSE 2013, Valencia, Spain, June 17-21, 2013. Proceedings*. Berlin, Germany: Springer-Verlag Berlin Heidelberg; 2013:338-352.
58. Zúñiga-Prieto M, González-Huerta J, Insfran E, Abrahão S. Dynamic reconfiguration of cloud application architectures. *Softw Pract Exper.* 2018;48(2):327-344.

How to cite this article: Kumara I, Han J, Colman A, van den Heuvel W-J, Tamburri DA, Kapuruge M. SDSN@RT: A middleware environment for single-instance multitenant cloud applications. *Softw: Pract Exper.* 2019;49:813-839. <https://doi.org/10.1002/spe.2686>