

Publish–Subscribe Grows Up

Support for Management, Visibility Control, and Heterogeneity

Message-oriented middleware is used to decouple the operation of cooperating applications. Existing approaches have concentrated mainly on scalability issues, but dynamic business processes and the integration of a wide range of data sources and applications require a middleware that is customizable. The Rebeca publish–subscribe service uses scoping to structure both middleware and applications. It thus offers advanced routing mechanisms to subsystems that need high scalability and it allows for heterogeneous message models that are transparently mapped onto each other.

**Ludger Fiege,
Mariano Cilia, Gero Mühl,
and Alejandro Buchmann**
Darmstadt University of Technology

Messaging systems have proven to be an important building block in modern computing systems – they facilitate the paradigm shift from centralized applications and data stores toward data-driven systems that comprise autonomously operating components and services. Data management and scalability in terms of sustainable message throughput and reliability have been major factors driving the development of messaging services.

With the advent of radio frequency identification (RFID) tags, the automation of business processes continues to include an ever-broader range of data sources.¹ From low-level sensors to high-level business objects, communication must scale to span various levels of detail to facilitate process automation. But when boundaries are crossed, problems arise:

administrative boundaries raise management and security issues; geographic boundaries in large systems require scalable systems; and application boundaries add heterogeneity of data models.

Researches have investigated scalability in detail,² whereas mostly ad hoc solutions are available for the rest of these problems.

In this article, we concentrate on publish–subscribe services as a special case of message-oriented middleware (MOM)³ and on the resulting system architectures (see Figure 1). With a publish–subscribe service, producers publish notifications to inform about events they have observed. The classic examples are stock quotations, weather news, and load-monitoring events. Consumers subscribe to notifications in which they're interested. For instance, a user

might be interested in weather news about Berlin or might ask for notifications if CPU load increases above 70 percent. The underlying publish-subscribe service delivers notifications to consumers with matching subscriptions. The service itself consists of a network of brokers that convey and filter the notifications. Sometimes, producers send advertisements to announce the kinds of notifications they're going to publish. This simplifies the routing of notifications and subscriptions.

On higher levels of abstraction, producers and consumers, which are arbitrary software components, are composed into applications that in turn form a system of collaborating applications. Figure 1 sketches two applications that are composed of producers and consumers (red lines), which communicate with each other by sending notifications on network lines (green lines). The resulting event-driven architecture is loosely coupled, easily adaptable, and facilitates scalable implementations of networked services.

These benefits can be exploited in large settings, but the indirect communication also hides system structure. A weakness of many existing publish-subscribe systems is their inability to structure and control communication without impeding the desired loose coupling. Deployment, orchestration, and management of these systems is then mostly considered as an afterthought.

In this article, we introduce a scoping concept for structuring publish-subscribe systems. It provides primitives to control the disseminated information's visibility, and it allows for the integration of different messaging and routing implementations (there is no one-size-fits-all solution). Based on an analysis of application requirements, we introduce *scopes* as a system engineering tool that provides modularization and customizability in a publish-subscribe middleware. Scopes are the basis for handling security and management issues as well as for improving scalability and dealing with heterogeneity.

Application Requirements

Application requirements are changing; as the speed at which business is conducted increases, service-orientation and automation require faster and deeper integration of business processes than were possible in the past.

For instance, smart-item technology, such as SAP's Auto-ID infrastructure,¹ links physical goods, their movements, and their transport con-

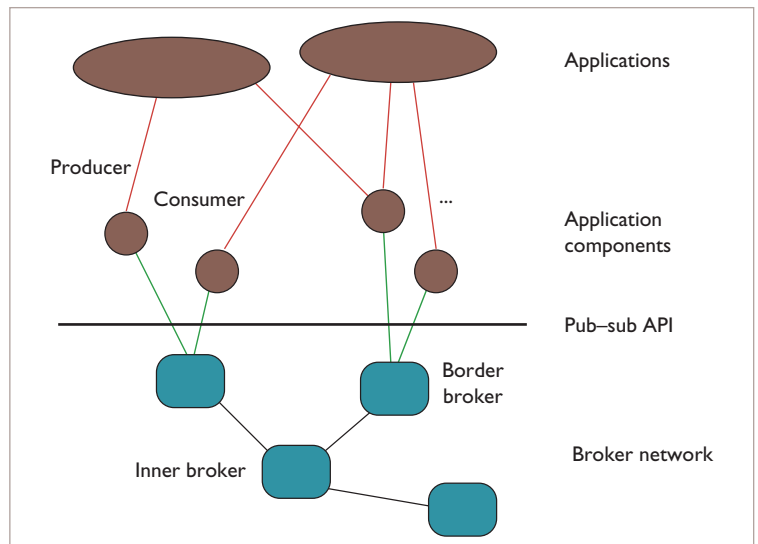


Figure 1. Participants of a distributed event-based system. Producers and consumers (red lines) communicate with each other by sending notifications on network lines (green lines).

ditions with business processes, all without human interaction. The lessons learned from these pilot projects are that cross-organizational cooperation is necessary; different applications running on the same infrastructure require different quality of service (QoS); and current solutions still lack the full support they need for distribution of functionality and data.

Smart-item technology is just one example. Other application domains have similar requirements.⁴ But the common denominator is the requirement for a networked IT infrastructure that integrates data sources and drives functionality on a global scale. Key requirements for such systems typically go far beyond mere scalability issues.

Open Issues

Frequent changes of processes, collaboration partners, and, thus, communication requirements have an important impact on networked infrastructure. A one-size-fits-all solution won't sustain the full range of processes that we outlined earlier. The necessary federation of notification services and applications is still a field of ongoing research. Namely, we need support for management and customization, heterogeneity, and security.

The core problem here is the lack of support for managing federated notification services; thus, the customization of the QoS provided to applications is insufficient. If a component is used in multiple applications, its notifications could be subject to different QoS constraints. To maintain




Figure 2. An exemplary scope graph that structures a publish-subscribe application. The scope graph hierarchically structures publish-subscribe applications and imposes visibility constraints.

loose coupling, we don't want producers to distinguish any destinations themselves – this would impede deployment time and any future runtime changes. Instead, the mediating infrastructure must be aware of the participating applications – that is, of the system's structure. Once the communication subsystem is aware of this structure, we can provide differentiated services and reconfigure the system without modifying producers and consumers.

In loosely coupled systems, the exchanged data must be able to interact and understand the data beyond the closed confines of a single component or application. That includes applications that interact across traditional borders regardless of economic, cultural, or linguistic differences (in the simplest form, for example, systems of units, currency, or date/time format). Most existing MOMs expose the messages' data structure but not the explicit semantics. They require that each producer and consumer application share the same homogeneous namespace. This reflects the low-level support of current infrastructures for the integration of heterogeneous data.

Finally, security measures seem to contradict the loose coupling of publish-subscribe functionality, but they are essential in any but toy settings, such as sandbox games. The rather limited support for security is one reason that MOMs can't fully offer their benefits as open-integration platforms.

Scoping

To support system management tasks, we must identify system parts – that is, administrative domains on which the tasks operate.

Administrative domains must be delimited to control and limit the effects of customization. On the other hand, interaction between these domains must be possible in a controlled way to prevent a complete partitioning of the system.

Taken together, the fundamental problem is

visibility control, and to face it, we've introduced the notion of scoping in event-based systems.

Scoping Model

The main idea of the scoping concept is to control the visibility of notifications outside of application components and orthogonal to their subscriptions.⁵ A scope bundles a set of application components and can contain other scopes. A directed acyclic graph of simple and complex components (scopes) gives the system's resulting structure (see Figure 2). This scope graph hierarchically structures publish-subscribe applications and imposes visibility constraints.

Notifications' visibility is initially limited to the scope in which they're published. The transition of notifications between scopes is governed by *scope interfaces* – that is, a scope issues subscriptions and advertisements in order to act as a regular producer and consumer in its superscopes. In Figure 2, scope R is the superscope of T and U, which look like simple components within R. A scope's interface selects the internal notifications that are forwarded to its superscopes; the external notifications are then relayed toward the scope's subcomponents. For instance, in Figure 2, the publish-subscribe service delivers a notification that Z publishes to Y and to any other consumers in T and U if their subscriptions match. This notification is also visible in R if it matches T or U's output interface, but it's not visible in S.

Scoping isn't an alternative to using established organization schemes such as topics, types, and filters. It controls visibility in addition and orthogonal to these schemes. Individual components still need to express their interests with subscriptions of any kind.

Using Scopes

Obviously, we can use scopes to compose applications. Scopes offer a structuring mechanism, on both the application and the infrastructure level (as we will see later). They govern the communication without intrusive changes to application components.

Using scopes is about creating and maintaining the scope graph. We have created a scope graph specification language that we use to model scopes as well as deploy preconfigured scopes and update them at runtime. The following example defines a scope named `temp` containing components of the existing scope `world`, plus components A and B, in which the attribute `has-temp-sensor` is set to one:

```

DEFINE SCOPE temp AS
ALL FROM MEMBERS(world), A, B
WHERE has-temp-sensor = 1

```

We now extend the definition, allowing the components of `temp` to send only temperature notifications, while `temp` itself forwards only alarm notifications to its superscopes

```

DEFINE SCOPE temp AS
ALL FROM MEMBERS(world)
WHERE has-temp-sensor = 1 : {
INTERFACES OUTPUT(TempNotification),
INPUT(0)
}
INTERFACES OUTPUT(Alarm
Notification)

```

We can deploy such a scope in several superscopes `S1, S2,...` by

```

DEPLOY temp
SUPERSCOPE ALL FROM S1, S2,...

```

The deployment command might carry additional implementation specific parameters.

Alternatively, programmers can access scoping through an API at runtime. In addition to the plain API functions (`pub`, `sub`), four new functions are necessary for maintaining a scope graph: creating and destroying scopes with `cscope` and `dscope`, and joining and leaving an existing scope with `jscope` and `lscope`.

Implementing Scopes

We now look to the distributed implementation of scopes. This approach opens the black box of a notification service and determines groups of brokers that implement a specific scope, thus correlating groups on the application and system levels.

Integrated routing reconciles distributed notification routing with the visibility constraints that the scope graph defines. Each broker's original routing table is divided into multiple tables, one for each locally available scope. Thus, for each scope, a connected subset of brokers constitutes an overlay within the broker network that conveys scope-internal traffic. Another routing table, the *scope routing table*, records scope-link pairs in each broker to signify in which directions we can find brokers of the respective scope.

Upon scope creation, an initially empty routing table is created at a broker, together with any

management information regarding this scope, such as interface definitions. A notification announces the creation and distributes it throughout the network to update the scope routing tables. The overlay can either be extended manually by administrative commands to preset a certain extent of the overlay, or dynamically when other components join the scope. Either way, a scope *join request* is always issued at a broker that's currently not part of the overlay. A request travels in the direction stored in the scope's routing table, leaving a temporary trail of references to the request source. The first broker encountered that's part of the requested scope processes the request and sends a reply back along the trail. If affirmative, the reply contains management information needed to set up the routing tables in the involved brokers; they become part of the scope's overlay.

The transition of notifications between two scopes requires the two scope overlays to share at least one broker. Consider scopes `T` and `R` of Figure 2; `T` is a component of and has joined `R`. For each subscription of `T`, a respective entry gets added to `R`'s routing table that points to `T`'s table. For each advertisement, an entry is added in `T`'s table that points to `R`. Mechanisms are in place to prevent multiple transitions at different brokers, but they are discussed here.

With this implementation, scopes not only group clients of the publish-subscribe service on the application level, but they're also an important tool for grouping brokers, thereby extending their structuring capabilities to the infrastructure. They determine which subset of brokers belongs to the same grouping and even allow for different routing algorithms in separate overlays, as long as the transition between the scopes adheres to the scope graph's constraints.

Data Heterogeneity

The exchanged notifications encapsulate data about a given event of interest, which users and developers can properly interpret and use only when they have sufficient context information. In traditional systems, this context information is typically left implicit and is normally lost when data crosses component or institutional boundaries.

As we mentioned earlier, today's MOM (publish-subscribe) infrastructures don't support data integration aspects such as an explicit descriptions of the intended meaning of notification content (for example, the implicit assumptions made by event/data producers). Without this kind

of information, data producers and consumers are expected to comply with implicit assumptions made by participating software components or applications. Even in the case of a very small set of applications within an enterprise, this approach is questionable.

When trying to tackle the problem of data/event heterogeneity, we must consider two main issues. The first concentrates on the vocabulary that all participants share. The second relates to the contextual information of applications that produce and consume data.

Processing exchanged data in a semantically meaningful way requires explicit information about the semantics of events and data. Because a scope groups participants who share commonalities, we decided to enrich it with the association of a vocabulary and the corresponding assumptions (metadata) about the data that's produced and consumed within the scope in question. Vocabularies (ontologies) represent the semantics, structure, and relationships of the terms (concepts) in a given business domain. Data assumptions (also known as contextual information) are represented by a set of properties and their values. Those properties refer to concepts in a vocabulary, making this self-containing. Additionally, vocabularies incorporate conversion functions that make possible an automatic data transformation between different contexts (for instance, date/time formats, systems of units, and currency conversions).

Scopes explicitly support the association of metadata about the notifications flowing within. To be in the same scope implies a shared context. For example, notifications sent by RFID readers in a warehouse need not specify their origin as long as they're processed within the warehouse. Additional information is necessary only when notifications leave the shared context – in this case, the warehouse. Generally, when messages cross scope boundaries, we might need to add information or map the terms used within one scope into the terms used in the other. If they share the same ontology but refer to different contexts (such as metric and American units), the conversion can occur automatically. If they use different ontologies, the *scope administrator* must manually specify the mapping of terms because integrating ontologies is a very difficult task. Once we solve the vocabulary problems, the notifications exchanged across scopes can be automatically transformed to the context of the target scope with the conversion functions. This

approach greatly simplifies data integration (even if using a single vocabulary) by moving code related to integration from participating applications into the infrastructure that orchestrates their integration.

Security

The preceding discussion introduced scopes as a means to group application and infrastructure components. They are therefore an apparent place to implement groups of trust – that is, those whose members, belong to the same authentication domain.⁶ To establish trust relationships in scopes, we have to deal with client access control and with securing the networked infrastructure itself.

In many scenarios, access to the publish–subscribe service must be controlled on the subscriptions and publications level. Only authorized clients should have access to the network of brokers to publish and subscribe to authorized notifications.

Our prototype implementation, Rebeca, uses rather simple policies because the main focus lies on how security is integrated – more sophisticated policies would be available if role-based access control schemes are bound to scopes.⁷

We use attribute certificates (AC; as specified in RFC 3281) to encode privileges and a public key infrastructure to bind the certificates to the clients. An AC is a credential with a digitally signed identity and a set of attributes. It carries the commands a component is allowed to issue, such as authorized filter expressions or scope creation commands. ACs are issued either by the provider of the broker network, by some other previously authorized (scope) administrator, or by the notifications' producers. In this way, we can build an authorization hierarchy.

When a client subscribes to some information at a border broker, it also gives its credentials in the form of an AC. The border broker checks the signature of the certificate with the network provider key and the key associated with the client's scope. Only if the certificate covers the subscription is it processed further, such as in a standard publish–subscribe case.

To secure the communication within a scope, the administrator can mandate that all traffic within the scope be encrypted. In the presented scope architecture, scopes connect with subgraphs in the broker network – that is, each communication link is between scope participants that can share a session key for encryption. If new components join, the session key is propagated to them.

If new brokers are added to the overlay in order to reach new components, these brokers can either be included in the trusted overlay, or the infrastructure can create tunnels toward the new components. If the authorization hierarchy includes network providers that host brokers in a globally operating publish–subscribe network, we can use attribute certificates to establish a trusted overlay network of brokers of a trusted provider, whereas tunneling can be used to bridge intermediate untrusted brokers.

Our Platform for Experiments: Rebeca

Our distributed notification service prototype, Rebeca, implements the scoping concept, serves as a testbed for routing algorithms, and supports data integration aspects to transparently map between different data models. Several PhD and masters' theses contributed to the prototype during the past years.

Event brokers constitute the routing network and also connect individual components to the network – that is, the fan-out of the network. They are implemented as separate processes either on designated nodes or mixed with other processes. The default broker maintains a routing table for unscoped traffic and TCP connections to other brokers and to clients. We have left the use of IP multicast as an option for future optimizing, for example, to improve the implementation of scope overlays or the network fan-out.

Brokers are customizable software containers⁸ and, as such, we can configure the routing engine, connection pooling, and transmission protocols at deployment time. All software building blocks act as message handlers that register themselves for internal events, so that incoming messages pass through deserialization handlers, protocol implementations, scope-specific routing tables, and so on. A node-internal network of connected handlers constitutes the broker implementation.

We are currently working on a new architecture to investigate the construction of the notification service.⁹ Modules of major MOM building blocks, such as matching, network topology maintenance, and so on, are composed using a service-oriented approach, such as the Open Services Gateway Initiative (OSGi; www.osgi.org). Thus, exactly the desired functions are packed into a specific notification service instance, and even runtime modifications of service implementation become possible.

Scopes are a runtime mechanism that delimit different implementations of the publish–subscribe API. Starting from a model of the system's structure as given in the scope graph, the administrator deploys and manages the systems at scope granularity. An administrator responsible for some sub-graph chooses (and deploys) a MOM instance for each scope and customizes its deployment parameters. In our prototype, scope-graph design and deployment is supported with a plug-in to the Eclipse development environment (www.eclipse.org) that offers a graphical interface for graph layout, as well as a text editor for editing the XML representation of the graph directly. Furthermore, the prototype is implemented in Java (an alternative version running on .NET is available), and we can use the Eclipse development environment to modify functionality just before deployment. We access the configuration of running scopes through a remote management interface using Java Management Extensions (JMX).

The desired functions are packed into a specific notification service instance, and even runtime modifications of service implementation become possible.

The current Rebeca architecture doesn't allow for the easy inclusion of security policies because several core classes are involved and would have to be reimplemented. To achieve greater flexibility, we employ aspect-oriented programming (AOP) techniques¹⁰ to implement the security aspects of scopes. Briefly, let's look at two security extensions.

First, access control on the API level is required for the authorization to invoke management functions. Provider keys are stored on all brokers and clients send certificates with each call to the broker. The broker then checks the certificates before granting access to management functionality.

Second, when we extend an overlay in integrated routing, we ensure that the new next-hop broker belongs to a trusted network provider by checking its certificates. This test checks chains of certificates in the authorization hierarchy and is evaluated prior to calling the handler that processes the actual scope extension. An encrypted fan-

out to consumers uses point-to-point connections; in the case of performance problems, administrators might employ caching schemes like the one described in Opyrchal and Prakash's work.¹¹

Rebeca Extensions

We used Rebeca to investigate several extensions.¹² *Mobile systems* are a natural application domain of messaging. We investigated physical and logical mobility as extensions to a static notification infrastructure.¹³ Physical – that is, client – mobility brings location transparency, whereas logical mobility supports navigating in a location/state space that's orthogonal to the notification infrastructure's layout.

We recently extended this work into the area of multipurpose *wireless sensor networks* (WSNs), in which scoping could serve as a generic node selection scheme that subsumes existing approaches.¹⁴

Message dissemination's instantaneous nature leads to isochronous communication. Caching

appropriate way to define rules. This conceptual representation enables the use of a "generic" reactive functionality service for different domains, making the underlying service independent from the rule specification.

In this article, we've shown how scopes, as first-class structuring mechanisms, help manage and control publish-subscribe infrastructures. Scopes also provide the necessary primitives for visibility and access control, as well as the containers to which context information and mapping functions can be attached.

We implemented these concepts in the Rebeca notification service, and we're extending them to deal with new sources of data with limited resources, such as RFID readers, heterogeneous sensor nodes, and mobile components.

Although scopes are our response to the need for management, security, and heterogeneity in publish-subscribe infrastructures, we're convinced that similar abstractions and primitives are essential for mature message-oriented middleware. □

References

1. C. Bornhövd et al., "Integrating Automatic Data Acquisition with Business Processes – Experiences with SAP's Auto-ID Infrastructure," *Proc. 30th Int'l Conf. Very Large Data Bases (VLDB 04)*, M.A. Nascimento et al., eds., Morgan Kaufmann, 2004, pp. 1182–1188.
2. G. Mühl et al., "Evaluating Advanced Routing Algorithms for Content-Based Publish-Subscribe Systems," *Proc. 10th IEEE/ACM Int'l Symp. Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 02)*, pp. 167–176.
3. P.T. Eugster et al., "The Many Faces of Publish-Subscribe," *ACM Computing Surveys*, vol. 35, no. 2, 2003, pp. 114–131.
4. G. Banavar et al., "A Case for Message-Oriented Middleware," *Proc. 13th Int'l Symp. Distributed Comp. (DISC 99)*, LNCS 1693, P. Jayanti, ed., Springer-Verlag, 1999, pp. 1–17.
5. L. Fiege et al., "Engineering Event-Based Systems with Scopes," *Proc. European Conf. Object-Oriented Programming (ECOOP)*, LNCS 2374, B. Magnusson, ed., Springer-Verlag, 2002, pp. 309–333.
6. L. Fiege et al., "Security Aspects in Publish-Subscribe Systems," *Proc. 3rd Int'l Workshop Distributed Event-Based Sys. (DEBS 04)*, A. Carzaniga and P. Fenkam, eds., IEE, 2004, pp. 44–49.
7. A. Belokosztolszki et al., "Role-Based Access Control for Publish-Subscribe Middleware Architectures," *Proc. 2nd Int'l Workshop Distributed Event-Based Sys. (DEBS 03)*, H.-

A data-driven system depends on the appropriate reaction to notifications.

mechanisms are needed to decouple components in time. Message queues on top of database systems offer flexible and powerful solutions, but we can also incorporate handling of historic data into the distributed notification service.^{15,16} In particular, when considering mobility, mobile devices might require a bootstrap phase to get in sync with the notifications flow. For this particular purpose, we developed a best-effort in-network caching approach to deliver recently published notifications, thereby reducing bootstrapping time.

A data-driven system depends on the appropriate reaction to notifications. We have developed a reactive functionality service that processes event-condition-action rules (ECA-rules). This allows a smooth integration of notifications into business process automation. This service provides several benefits: rule definitions could include contextual information to more easily facilitate the integration, and we could tailor rule-definition languages for different domains using a conceptual representation, providing end users the most

- Arno Jacobsen, ed., ACM Press, 2003, pp. 1-8.
8. M. Fowler, "Inversion of Control Containers and the Dependency Injection Pattern," Jan. 2004; <http://martinfowler.com/articles/injection.html#InversionOfControl>.
 9. C. Fiorentino et al., "Building a Configurable Publish-Subscribe Notification Service," *IFIP Int'l Conf. Distributed Applications and Interoperable Sys.* (DAIS 05), LNCS 3543, Springer-Verlag, 2005, pp. 136-147.
 10. T. Elrad, R.E. Filman, and A. Bader, "Aspect-Oriented Programming: Introduction," special issue on aspect-oriented programming, *Comm. ACM*, vol. 44, no. 10, 2001, pp. 29-32.
 11. L. Opyrchal and A. Prakash, "Secure Distribution of Events in Content-Based Publish-Subscribe Systems," *Proc. 10th Usenix Security Symp.*, Usenix Assoc., 2001, pp. 281-295.
 12. A. Buchmann et al., "Dream: Distributed Reliable Event-Based Application Management," *Web Dynamics-Adapting to Change in Content, Size, Topology and Use*, M. Levene and A. Pouloussis, eds., Springer-Verlag, 2004, pp. 319-349.
 13. L. Fiege et al., "Supporting Mobility in Content-Based Publish-Subscribe Middleware," *ACM/IFIP/Usenix Int'l Middleware Conf.* (Middleware 03), M. Endler and D.C. Schmidt, eds., LNCS 2672, Springer-Verlag, 2003, pp. 103-122.
 14. J. Steffan et al., "Towards Multi-Purpose Wireless Sensor Networks," *Int'l Conf. Sensor Networks* (SENET 05), P. Dini et al., eds., IEEE CS Press, 2005, pp. 336-341.
 15. A. Ulbrich et al., "Programming Abstractions for Content-Based Publish-Subscribe in Object-Oriented Languages," *Confederated Int'l Conf. CoopIS, DOA, and ODBASE 2004*, LNCS 3291, Springer-Verlag, 2004, pp. 1538-1557.
 16. J. Bacon et al., "Event Storage and Federation using ODMG," *Proc. 9th Int'l Workshop on Persistent Object Sys.* (POS 9), LNCS 2135, G. Kirby, A. Dearle, and D. Sjøberg, eds., Springer-Verlag, pp. 265-281.

Ludger Fiege is an engineer at the corporate technology department of Siemens AG in Munich, Germany. His research interests include software architecture of enterprise systems, event-based systems, and configurable middleware and programming paradigms. Fiege has a PhD in computer science from the Technische Universität Darmstadt, Germany. Contact him at fiege@acm.org.

Mariano Cilia is a postdoctoral researcher at the Department of Computer Science at Technische Universität Darmstadt. He is also visiting professor at the Faculty of Sciences, Universidad Nacional del Centro de la Provincia de Buenos Aires (UNICEN), Argentina. His research interests include data dissemination, semantic data integration, pervasive computing, event-driven systems, and middleware. Cilia has an MSc in computer science from University of Campinas (UNICAMP), Brazil, and a PhD in

computer science from the Technische Universität Darmstadt. He is a member of the IEEE and the ACM. Contact him at mcilia@acm.org.

Gero Mühl is a postdoctoral researcher at the Berlin University of Technology. His research interests include middleware, event-based systems, self-organization, and mobile computing. Mühl has a PhD in computer science from the Technische Universität Darmstadt. Contact him at g_muehl@acm.org.

Alejandro Buchmann is a professor of databases and distributed systems and chairman of the Department of Computer Science of the Technische Universität Darmstadt. His research interests include event-based systems, ambient intelligence, data dissemination, peer-to-peer and sensor networks, and performance of middleware-based systems. Buchmann has a BS degree from Universidad Nacional Autónoma de México (UNAM) and an MS and a PhD in chemical engineering from the University of Texas, Austin. Contact him at buchmann@dvs1.informatik.tu-darmstadt.de.

DON'T RUN THE RISK.

BE SECURE.

IEEE SECURITY & PRIVACY

Ensure that your networks operate safely and provide critical services even in the face of attacks. Develop lasting security solutions, with this peer-reviewed publication.

Top security professionals in the field share information you can rely on:

Wireless Security • Securing the Enterprise • Designing for Security Infrastructure • Privacy Issues • Legal Issues • Cybercrime • Digital Rights Management • Intellectual Property Protection and Piracy • The Security Profession • Education

Order your subscription today.
www.computer.org/security/

IEEE
 computer society
 60th anniversary