

IF 2002
INFORMATION FUSION
Annapolis, Maryland


THE FIFTH INTERNATIONAL CONFERENCE ON INFORMATION FUSION

Riding the Information Wave

Sponsored by ISIF and IEEE

"Complimentary copy of Fusion 2002 paper from Tim Bass ..."

The Federation of Critical Infrastructure Information via Publish-Subscribe Enabled Multisensor Data Fusion



7-11 July 2002 + Loews Annapolis Hotel + Annapolis, Maryland, U.S.A.

Invited Session: Information Fusion for Critical Infrastructure Protection
Copyright © Fusion 2002 and Silk Road, Updated: April 28, 2002
Paper Author: Tim Bass

The Federation of Critical Infrastructure Information via Publish-Subscribe Enabled Multisensor Data Fusion

Tim Bass
Silk Road
bass@silkroad.com

Abstract - *The art and science of multisensor data fusion is the emerging foundation for the development of next generation network-centric decision support systems, including critical infrastructure protection. These challenging technical objectives require the cooperative signal processing of a federation of critical infrastructures. Publish-subscribe architectures provide process-to-process messaging infrastructures that enable a communications framework for the distribution and delivery of information between sensor fusion processes. In this paper we discuss high level service-oriented architectural issues for critical infrastructure multisensor data fusion including event notification services, wide-area network topology, and the publish-subscribe subscription language.*

Keywords: multisensor data fusion, publish-subscribe, event notification service, intrusion detection, critical infrastructure, homeland defense and cooperative federation

1 Introduction

In [1,2] we suggested that the art and science of multisensor data fusion is directly applicable to detection theory in cyberspace situational awareness, network management and network intrusion detection systems. Recent political events have increased the visibility of the computational challenges in the design, analysis and survivability of critical infrastructures such as computer and communication networks, electric power grids, and similar distributed computing systems. Correlating the health and real-time security of interconnected distributed systems are also socially challenging because event notifications and other information objects must be shared across political, organizational and administrative boundaries.

In this paper we continue the development of the models presented in [1,2] by examining how a subscription-based data processing architecture might be applied to address many of the challenges associated with cooperative distributive processing between federated information

sources. The key concepts of this architecture are loosely coupled service-oriented information fusion processes administered by a federation of organizational services including:

- federated sensor, processing, assessment and storage.
- publication, subscription and other fusion services.
- attribute-based data and information publishing.
- subscription-based data and information availability.

Applying these characteristics to cyberspace data fusion models [2] forms a federation of distributed services that publish, and subscribe to, relevant information and data sets. This architecture appears to provide a scaleable foundation for developing new fusion and decision support applications with minimal impact to established information systems and services. In addition, the architectural model is well suited for federated interdomain information sharing.

Homeland defense critical infrastructures such as computer and communication networks, electric power grids, intelligence networks, immigration systems, air traffic control and transportation systems are controlled and administered by numerous autonomous organizations. A subscription-based information infrastructure enables cooperative multisensor, service-oriented fusion in a federated processing model. The next section discusses the subscription-based service-oriented architectural model in the context of sharing and fusing federated information.

2 Publish-Subscribe Architecture

The architectural approach suggested for homeland defense and critical infrastructure protection in this paper is a service-oriented multisensor fusion model enabled by a publish-subscribe event notification communications network. Figure 1 illustrates the relationships between federated sensors, event notification, storage and fusion services. The key concepts are twofold: (a) the processing abstraction is a federated system processing architecture and (b) information moves between systems based on publish-subscribe communication models. Distributed services in the architecture performs local or regional data

and information processing. Each service may be a subscriber and/or a publisher. Sensor systems may receive sensor information (S_i) from many sources including, but not limited to, subscription services.

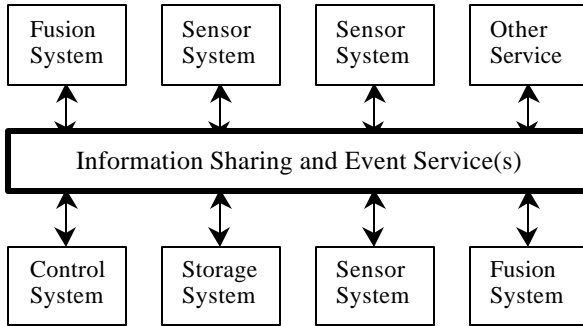


Fig. 1. Federated Information Processing

We expand this concept in Figure 2, illustrating multisensor data fusion data and object refinement process relationships [3] in context to receiving input from subscription-based services. The object database publishes objects to information sharing or event notification services. Subscribers to these information objects include situation refinement, threat assessment and situation knowledge-base fusion services [3]. Figure 3 continues the extension of these concepts to publish-subscribe services between multisensor fusion object bases (O_b) and associated fusion services. The interested reader is kindly referred to the referenced literature [1-3] for a more detailed background discussion of the cyberspace multisensor data fusion model.

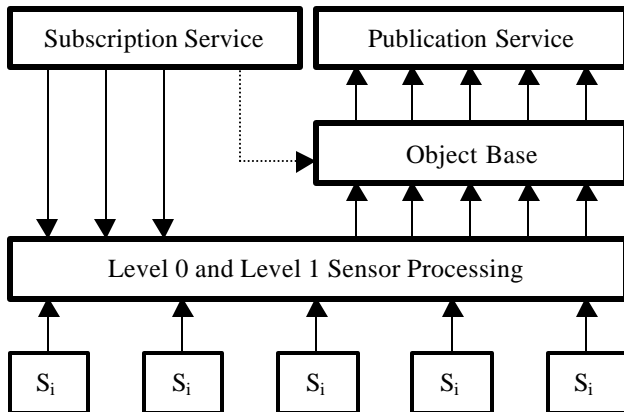


Fig. 2. Publish-Subscribe Data/Object Subsystem

In addition to fusion services, complimentary service-oriented architectures for data mining and knowledge management may be logically constructed from the same publish-subscribe infrastructure. The communications component in all of these conceptual models is an event notification capability that is politically, economically, and technically scalable.

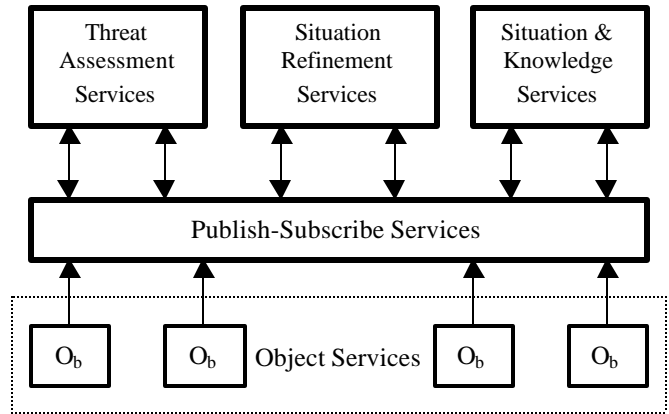


Fig. 3. Fusion Situation/Threat Subsystem

2.1 Critical Infrastructure Event Federation

A properly designed publish-subscribe architecture is a scalable distributed computing model that shows considerable promise to enable the federation of fusion processing in a service-oriented, complex environment. This distributed computing architecture could enable scalable collaborative information processing between critical infrastructures as a federation of administrative domains. The suggested approach is complimentary to emerging U.S. homeland defense situation assessment and command and control activities.

Homeland defense requires a federation of numerous organizations including FEMA, over forty federation agencies, state and local authorities, the CIA, NSA, DIA and other intelligence agencies, Department of Defense and other organizations. The cross-organizational correlation of threats, attack profiling and course-of-action dissemination requires the federation of autonomous event-driven decision support systems. A July 2001 Defense Department briefing articulated three critical capabilities for homeland defense information fusion and command and control activities [4]:

- assured connectivity,
- attribution (situational awareness), and
- crisis coordination.

The operational capabilities have identified the publish-subscribe communications model as one core competency required for future homeland defense. According to numerous experts, this capability would enable autonomous decision support systems to access and track threats across multiple fronts, providing high confidence and timely alerts.

2.2 Example Publish-Subscribe Models

Many publish-subscribe architectures are based on an event-driven communications model. Members of the

information federation who wish to participate in the event notification architecture may join the publish-subscribe network. The remainder of this section summarizes a few representative system architectures.

The InfoBus Repeater [5] is an interesting architecture where members register with the publish-subscribe service as information producers, consumers or both, depending on the service or role of the federated member. When a producer generates an event the producer notifies the publish-subscribe bus, that in turn, notifies appropriate event consumers. One of the interesting characteristics of the InfoBus architecture is that subscribers (consumers) may trigger events in publishers (producers).

Another interesting architecture that uses publish-subscribe concepts is Rio [6]. Rio claims to provide the basis to federate a loose coupling of event producers that advertise event attributes. This architecture allows consumers of event information to discover unknown event publishers. This capability is well suited for a service-oriented multisensor data fusion architecture that has many federates and the requirement to adapt to change rapidly.

SIENA [7] is a research-oriented publish-subscribe architecture that provides a wide-area event notification service in a scaleable and flexible Internet framework. Examples of other publish-subscribe research projects worthy of review are Gryphon [8] and Elvin [9]. The Java Messaging Service (JMS) [10] is noteworthy; however the current JMS specification does not articulate a complete event notification service or publish-subscribe infrastructure. There are many other examples in the literature of publish-subscribe communications models being used for federated information processing and event notification services. Publish-subscribe architectures will enable numerous next generation distributed network applications in the immediate future.

2.3 Security Considerations

Scaleable engineering solutions are achievable when processing components are minimally coupled in an architectural model that adheres to complexity management principles. It logically follows that multisensor fusion models are scaleable when security services are loosely coupled with the underlying publish-subscribe infrastructure. This implies that a subset of confidentiality, integrity and non-repudiation services for multisensor fusion architecture should be provided by security infrastructure services such as virtual private networks or end-to-end cryptographic systems germane to the application.

Furthermore, designers of infrastructure services must consider security services for publish-subscribe architectures in the context of a qualitative risk

management model. Interested readers are referred to [11] for an applicable discussion on defense-in-depth risk management topics. In the next section we review publish-subscribe architectures in the technical context of the event notification service and wide-area networking.

3 Event Notification Service

The asynchronous, heterogeneous loose coupling that characterizes software applications in wide-area networks points to event interaction as the abstraction for multisensor fusion systems design. An emerging building block for these services is an infrastructure service known as an event notification service [7]. An event notification service accessible from organizations in the network federation is required to support cooperative information fusion processing.

An event notification service is typically implemented as a network of servers that provide service access points to client processes. Client sensor and fusion processes use the service access points to advertise information about events and to publish event notifications (per notification type previously advertised). The access points are also used by interested parties to subscribe to notifications of interest [7]. This type of interaction permits selection, filtering and pattern matching subscriptions that are key enablers in distributed multisensor data fusion models.




Fig. 4. Distributed Event Notification Service

The event notification service implements two important actions; notification selection and notification delivery. The service must be configured as a distributed system that exploits the benefits of locality or community. In the most general form an event notification service is a network of servers as shown in Figure 4. The clients in the event notification service are either objects of interest that generate events or interested parties who are event notification consumers. Network clients in the federation may be both event producers and consumers.

Conceptually the service corresponds to a wide-area network of routers with pattern matching engines transposed on top of an underlying communications

facility such as the Internet [7]. In context of the architecture suggested in this paper, sensors correspond to objects of interest in the event notification model. Multisensor data fusion architectural processes such as data refinement, object refinement, object and situation bases, situation refinement and threat assessment act as both objects of interest and interested parties.

Following well established wide-area network design principles, it is logical to place at least one critical infrastructure event notification server within each administrative domain in the federation. This motivates a discussion of the topology of a network of servers, with issues generally centered on three design issues [7]:

- interconnection topology,
- routing algorithm, and
- processing strategy.

Similar design issues have been extensively studied for many years in many different networking scenarios. For example, there are generally three classes of interconnection topologies; (a) a hierarchical organization, (b) a generalized graph of peers, and (c) hybrid clusters of both topologies. Cargzaniga *et al* do an fine job of discussing these issues relative to a generalized event notification service including *a priori* knowledge of locality, the event notification model, notification semantics, attributes, filters, patterns and timing. Interested readers are strongly encouraged to review [7] for an excellent comprehensive discussion in this area.

The event notification service model is, in essence, *the glue* that binds together a distributed network of federated critical infrastructure sensor processes in space and time. Euster, P. *et al* [12] summarized the space-time relationship between different communication models, presented in Table 1.

	Time	Space
Request/Reply	Coupled	Coupled
Asynchronous Send	Decoupled	Coupled
Shared Memory	Coupled	Decoupled
Publish/Subscribe	Decoupled	Decoupled

Tab. 1. Summary of Communications Models

Time decoupling implies that the objects of interest and interested parties in the event notification model (Figure 4) do not need to be up and available at the same time. Likewise, space decoupling implies that the clients of the event notification service are not required to have *a priori* knowledge of each other.

As stated, the event notification service model appears to be well suited for emerging critical infrastructure fusion applications. In the next section we turn our attention to server network topologies and protocols. These

discussions will be helpful to readers working on large distributed enterprise and/or localized architectures.

4 Server Topologies & Protocols

In the previous section we discussed the generalized event notification service and how servers might communicate in a wide-area network topology to cooperatively distribute event selection and delivery tasks to interested parties in the federation. The underlying network service must be arranged in a distributed network topology that makes use of a server-to-server communications protocol. This implies an underlying communications protocol between servers such as TCP/IP (and the Internet). The actual messaging protocol between the servers may be a wide range of network application protocols that are not the focus of this paper.

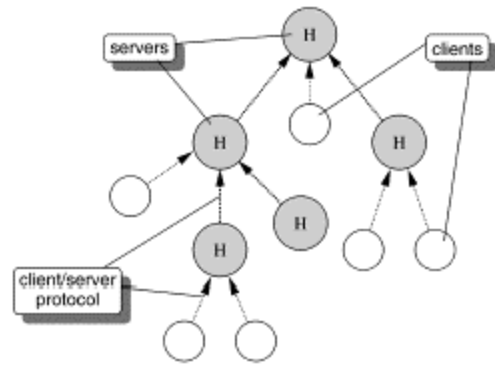


Fig. 5. Hierarchical Client/Server Architecture

In this section we summarize applicable prior work [7] by examining three basic architectures: (a) hierarchical client/server, (b) acyclic peer-to-peer and (c) general peer-to-peer. In addition, we briefly discuss a few hybrid network topologies. The degenerate case of a centralized architecture having a single server is not discussed because it is generally not scaleable to large distributed applications [7].

The hierarchical architecture of Figure 5 is a logical extension of a centralized architecture. One obvious problem with this architecture is the fact that the processing capability of the servers at the top of the hierarchy tends to limit the performance and scaleability of the system. Another problem is that each server in the network is a critical node for servers lower in the hierarchy [7]. Therefore, this topology does not appear to have the necessary reliability for the general case of a critical infrastructure event notification service. However, hierarchical topologies in hybrid architectures are useful for event processing in localized fusion communities.

The acyclic peer-to-peer architecture provides symmetric bi-directional flow of event subscriptions, advertisements

and notifications. However, as in the hierarchical architecture, the acyclic peer-to-peer architecture (Figure 6) does not have the optimal redundancy to guarantee highly robust event notification delivery. The solution [7] is to adopt a generalized peer-to-peer architecture, illustrated in Figure 7.

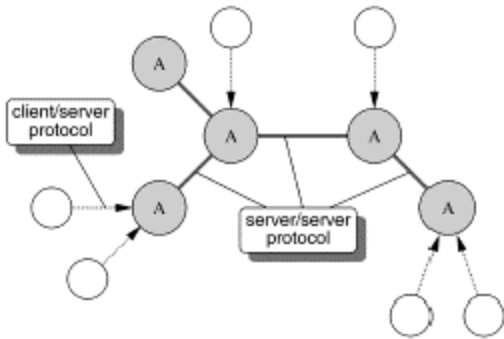


Fig. 6. Acyclic Peer-to-Peer Architecture

An advantage of the general peer-to-peer architecture is that it offers more flexibility and reliability over the other architectures, including path redundancy. However, a disadvantage is that complex routing algorithms must be used that avoid cycles and choose the best path. Routing in the Internet uses similar forwarding algorithms; therefore, designers of highly reliable and scaleable event notification services should carefully consider the benefits and disadvantages of peer-to-peer networking.

Networking is an organic process and the history of routing protocol developing in the Internet teaches us that the issues that determine the routing topologies tend to be dominated by political, acceptable usage and governance issues [13]. On the other hand, networks under a single administrative control may easily take advantage of the benefits of robust peer-to-peer architectures with less concern about complex routing trade-offs.




Fig. 7. General Peer-to-Peer Server Architecture

In practice, federated architectures of many autonomous systems tend to be acyclic and directive because of political and other administrative policies. The social-

economic influences on interconnected wide-area networks tend to give rise to large federated hybrid general/acyclic architectures, illustrated in Figure 8. Complexities and challenges in developing policies for sharing information between numerous critical infrastructure domains tend to favor a hybrid architectural approach composed of local and regional (community of interest) networking and acyclic or directed acyclic gateways between large federated organizations.

Political and cultural influences on homeland defense and critical infrastructure domains will drive networks toward hybrid architectures composed of hierarchical, acyclic and peer-to-peer topologies. These cooperative models for information sharing follow the federation of autonomous networks that characterize the Internet today. The primary difference is a higher abstraction at the information and event layer compared to the data packet and network address level in the TCP/IP model.

Due to size constraints of this paper, the brief summary of topology in this section was constructed to introduce the key concepts, not to be complete or conclusive. In the next section we direct the readers attention to another important design issue unique to this level of abstraction, the publish-subscribe subscription language.

5 Subscription Language

In summary, publish-subscribe is a communications architecture where information flow is directed by the interest of the consumers of information rather than by specific addresses determined by the sender. Carzaniga, Rosenblum and Wolf summarized four event notification subscription languages in [7,14] which we briefly review in this section:

- channel-based subscriptions,
- subject-based subscriptions,
- context-based subscriptions, and
- context-based subscriptions with patterns.

The concept of the subscription channel is similar to tuning in to a television or radio channel. The channel could broadcast information that is not of interest to all potential channel subscribers. It is not unreasonable to envision that interested parties might subscribe and unsubscribe to various channels depending on their changing information requirements. The channel-based architecture is the least complex of the subscription languages above. In this architecture, subscription services are based on interested parties listening to, or subscribing to, a single channel. Event notifications posted to the channel are delivered by the event service to all the interested parties that listen to that channel [14]. Subscribing to web-based streaming media is a form of channel subscription. Another example of a channel subscription is a when subscribers listen for activity on a particular queue.

Event-driven subject-based subscription is one of the current commercial trends in publish-subscribe networking. The primary difference between channel-based and subject-based subscription languages is that subject-based subscriptions are more expressive [14]. Event notification is generally based on the topic or subject. Attributes other than subject are not used; however, pattern matching within the subject attribute may be supported. For example, if the event notification service was capable of subject-based subscriptions, all interested parties subscribing to {electric.generation.*} would receive the notification (Table 2). Subject-based subscriptions extend the basic concept of channel-based subscriptions by offering a more flexible addressing mechanism. A subject-based subscription language is currently used in state-of-the-art commercial publish-subscribe offerings, for example, the TIBCO/Rendezvous suite of products [15].

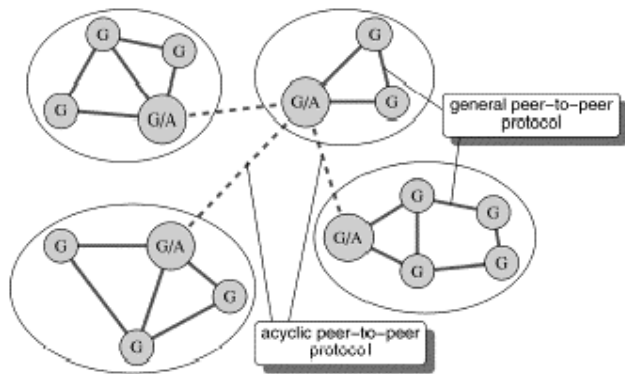


Fig. 8. General/Acyclic Hybrid Architecture

It is generally accepted that the success of subject-based networking will drive industry toward context-based subscriptions. Context-based subscription extends the capability of event notification with more expressive subscription filters compared to subject-based or channel-based architectures. Information is forwarded in a context-based architecture with software algorithms that match message context to context-subscribers.

string	subject=	electric.generation.output
time	date=	Mar 24 12:00:01 EST 2002
string	location=	Canefield Generating Station
string	region=	Southern District
float	change=	40.2
string	unit=	percent

Tab. 2 Example Critical Infrastructure Notification

Event notification in a context-based network is performed on a set of typed attributes and each individual attribute has a type, a name and a value [7]. Event notification delivery is then based on a structural value derived from

the attributes. The forwarding algorithm provides the motivation for context-based routing protocols that permit nodes to exchange context-subscriber information and formulate information routing tables. Turning again to Table 2, a context-based subscriber would could refine their subscription to receive only {electric.generation.*} messages from the Canefield Generating Station.

The most advanced and expressive of all of these subscription languages is the context-based subscription language with pattern matching. For example, from our example message of Table 2, context-based subscriptions with expressive patterns could be used to permit interested parties to subscribe only to messages where {electric.generation.output} in the Southern District dropped greater than 20 percent. The interested reader is referred to [7,14] for additional discussions on the expressiveness of subscription languages in event notification architectures.

6 Conclusions

Critical infrastructure monitoring and protection requires the federation of a vast amount of internetworked resources. A service-oriented multisensor data fusion architectural model with publish-subscribe event notification services appears to provide the distributed computing infrastructure required for critical infrastructure protection. The underlying infrastructure is strikingly similar to the cooperative federation of autonomous wide-area and local-area networks in the global Internet today.

Channel, subject and context-based subscription languages also appear to have important roles in the future design of event notification services. Furthermore, federated multisensor data fusion concepts, enabled by a service-oriented messaging infrastructure, appear to be a leading architectural model for enabling critical infrastructure protection and emerging homeland defense initiatives.

Emerging *Web Services (WS)* for peer-to-peer information sharing is an interesting technology that we did not have space to discuss in this paper. Current *WS* architectures tends to resemble the traditional point-to-point request-reply networking model. Web-based request-reply architectures also appear to be less complex than evolving publish-subscribe event notification systems. Having stated that, the publish-subscribe model is significantly more mature and robust than *WS* at this point in time. Significant evaluation is required as myriad competing *Web Services* implementations mature.

Perhaps the most pragmatic conclusion is the necessity for critical infrastructure sensor fusion architects to carefully examine their requirements and to consider the utility of federated service-oriented architectures. Critical infrastructure protection and homeland defense requires

the federation of information from many global autonomous infrastructures. Publish-subscribe services provide an interesting and promising way to federate event information from loosely coupled networking domains.

Acknowledgments

The author would like to thank Mr. Ed Waltz for his pioneering work in the field of multisensor data fusion, the SIENA team at the University of Colorado Boulder for their comprehensive research in publish-subscribe event notification services and the men and women of the United States Air Force for day to day inspiration and motivation. A special acknowledgment and word of appreciation to Dr. Alex Wolf and Dr. Antonio Carzaniga of the University of Colorado Boulder for permission to reproduce Figures 4 – 8. This work was partially sponsored by Silk Road purchase order DSD-01-302-6801.33, Project HAF/CIO.

References

- [1] Bass, T., “*Multisensor Data Fusion for Next Generation Distributed Intrusion Detection Systems*,” 1999 IRIS National Symposium on Sensor and Data Fusion, The Johns Hopkins University Applied Physics Laboratory, 24-27 May 1999.
- [2] Bass, T., “*Intrusion Detection Systems & Multisensor Data Fusion*,” *Communications of the ACM*, Vol. 43, No. 4, April 2000, pp. 99-105.
- [3] Waltz, E. and Llinas, J., “*Multisensor Data Fusion*,” 1990, Artech House, Boston, MA.
- [4] Eddington, D. Mr. and Malone, M. Maj., “*Homeland Defense Command & Control (HLD C2) FY02 ACTD Proposal*,” July 2001.
- [5] Uramoto, N., and Maruyama, H., “*InfoBus Repeater: A Secure and Distributed Publish/Subscribe Middleware*,” Proceedings, 1999 International Workshops on Parallel Processing, Sept. 21-24, 1999, pp. 260-265.
- [6] “*Rio Architecture Overview*,” Sun Microsystems, Version 1.0, 2001.
- [7] Carzaniga, A., Rosenblum, D., and Wolf, A., “*Design and Evaluation of a Wide Area Event Notification Service*,” *ACM Transactions on Computer Systems*, Vol. 19, No. 3, August 2001, pp. 332-383.
- [8] Banavar, G. *et al*, “*An Efficient Multicast Protocol for Content-Based Publish-Subscribe Systems*,” The 19th IEEE International Conference on Distributed Computing Systems (ICDSC '99), May 1999, pp. 262-272.

[9] Segall, B., and Arnold, D., “*Elvin Has Left the Building: A Publish/Subscribe Notification Service with Quenching*,” Proceedings of AUUG97, Sept. 3-5, 1997, pp. 243-255.

[10] Sun Microsystems Inc., “*Java Messaging Service*,” Mountain View, CA, 1999.

[11] Bass, T. and Robichaux, R., *Defense-In-Depth Revisited: Qualitative Risk Analysis Methodology for Complex Network-Centric Operations*, IEEE MILCOM 2001, October 28-31, 2001.

[12] Euster, P. *et al*, “*Distributed Asynchronous Collections: Abstractions for Publish/Subscribe Interaction*,” ECOOP 2000, LNCS 1850, 2000, pp. 252-276.

[13] Bass, T., “*Internet Exterior Routing Protocol Development: Problems, Issues, and Misconceptions*,” IEEE Network Magazine, July/Aug 1997, pp. 50-55.

[14] Carzaniga, A., “*Architectures for an Event Notification Service Scaleable to Wide-area Networks*,” PhD Thesis, Politecnico di Milano, December 1998.

[15] Tanenbaum, A. and van Steen, M., *Distributed Systems Principles and Paradigms*, Prentice-Hall, Inc., Upper Saddle River, New Jersey, 2002, pp.702-716.

Additional Reading

Moro, G. and Virile, M., “*Enabling Business Cooperation Using a Publish-Subscribe Architecture Aware of Transactions*,” Proceedings of the 34th Annual Hawaii International Conference on System Sciences, Jan. 3-6, 2001, pp. 4116–4125.

Biography

Tim Bass (bass@silkroad.com) is the president of Silk Road, a consulting firm specializing in operational concepts, design, architecture and security of IP networks and Internet applications. Mr. Bass provides senior subject matter expertise to the USAF, Office of the CIO. He graduated B.S.E., Tulane University, School of Engineering, 1987 *Magna Cum Laude*, Electrical Engineering and has been providing Internet domain expertise as an independent trusted advisor to the US military and commercial industry for over 10 years. His list of publications, detailed biography and company information is available at the Silk Road web site (www.silkroad.com).