Service-Oriented Horizontal Fusion in Distributed Coordination-Based Systems

Tim Bass
Silk Road
bass@silkroad.com

Abstract - *The art and science of multisensor data fusion has emerged as the foundation for the development of next generation net-centric decision support systems, including horizontal fusion systems. These decision support systems require the coordination of service-oriented sensors and fusion components. Distributed coordination-based architectures provide a process-to-process communications infrastructure that supports horizontal fusion services. In this paper we discuss architectural considerations for distributed service-oriented horizontal fusion including distributed coordination-based architectures, service access, data transformation, adaption, and end-to-end visualization.*

Keywords: adaption, coordination-based systems, distributed systems, end-to-end visualization, horizontal fusion, information dissemination management, information transformation, multisensor data fusion, service-oriented architectures.

1 Introduction

We suggested in [1, 2] that the art and science of multisensor data fusion [3 - 5] and, by inference, the Joint Directors of Laboratories (JDL) data fusion model is directly applicable to detection theory in cyberspace situational awareness, network management, and network intrusion detection systems. These concepts were expanded in [6] to suggest a high level service-oriented architecture for service-oriented federated critical infrastructure protection.

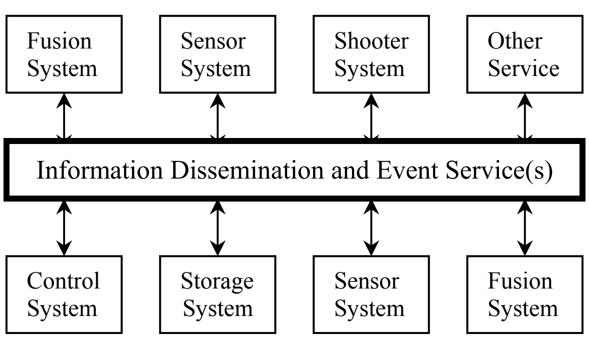

Fig. 1. Functional Service-Oriented Architecture

A generalized functional architecture for service-oriented horizontal fusion is illustrated in Fig. 1. This paper augments earlier concepts to support service-oriented horizontal fusion in distributed coordination-based systems [7].

The JDL data fusion model [5], illustrated in Fig. 2, is the dominant functional fusion model [5, 8], and has been the basis for numerous multisensor data fusion architectures. It is logical to assert that the JDL data fusion model is a key foundation for horizontal fusion architectures. The focus of this paper is a discussion on how distributed coordination-based systems support horizontal fusion and the JDL data fusion model. The interested reader is kindly referred to the references for very detailed discussions on the terms and abbreviations used in Fig. 2 [3 - 5].

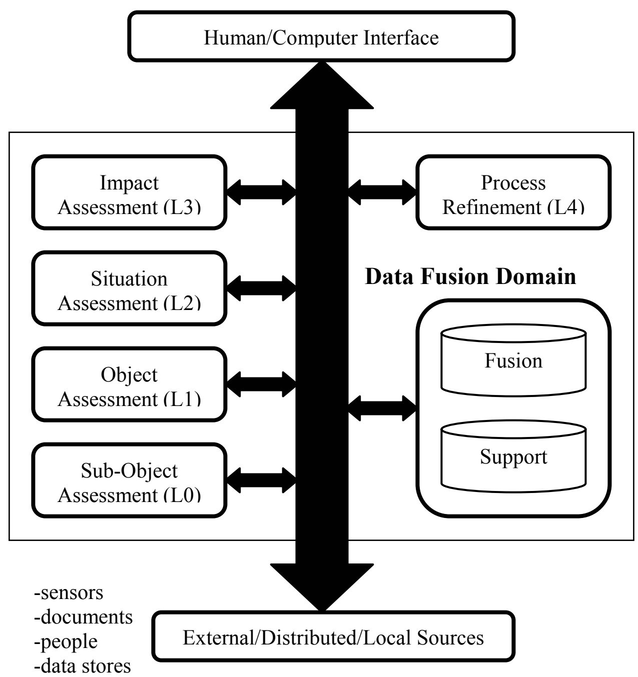

Fig. 2. The JDL Data Fusion Model (Revised 1998)

Distributed data fusion can provide a degree of scalability and robustness that cannot be achieved with centralized architectures. Decentralization can also permit a network of nodes to exchange information and coordinate activities in a flexible and scaleable architecture that would be impractical or impossible to achieve with a single, monolithic systems platform [5].

The notion that a node on the communications grid should have the capability of establishing a communications link with any other node to obtain the information needed to perform its mission is well established; however, achieving this robustness has proven very difficult [5]. Following the network concepts of a distributed service-oriented architecture suggested in [1, 2 and 6], we apply these characteristics to distributed coordination-based systems [7] to develop a high level architecture for horizontal fusion [9] in the context of the JDL fusion model. We suggest that inexpensive commercially available software can support distributed coordination-based systems, and in turn, support emerging distributed service-oriented fusion processing requirements.

2 Horizontal Fusion

Horizontal fusion initiatives in the Department of Defense are envisioned to enhance decision support capabilities to enable time sensitive, net-centric collaborative operations among distributed organization components. A stated goal of horizontal fusion is to take advantage of battlefield intelligence and information sources such as advanced sensors to enable end-to-end coordination and traceability of organizational missions and supporting infrastructures, illustrated in Fig 3. These capabilities are expected to increase the speed-of-command of distributed and dispersed force components emphasizing improved situational awareness and more rapid and cost effective integration of operational and intelligence planning [9].

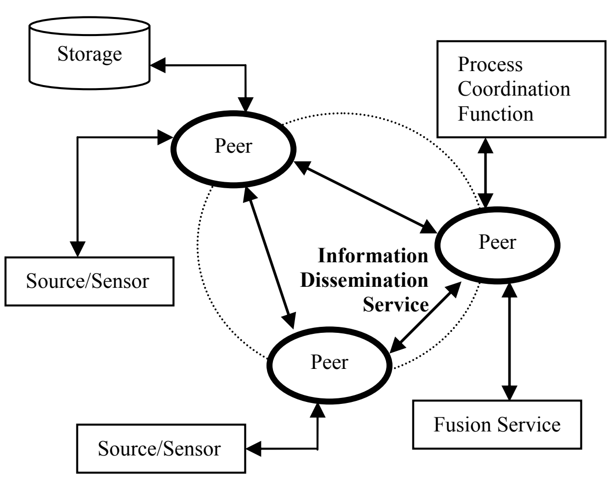

Fig. 3. System Architecture for Horizontal Fusion

Horizontal fusion is envisioned to enable the interoperability and collaboration needed for agile warfighting, combat support and business system transformation. The underlying infrastructure is divided into three large classes of assets [10]: an end-to-end Internet Protocol-based network, a set of shared data resources, and a new set of services currently referred to as Net-centric Enterprise Service (NCES) in the Department of Defense. NCES is defined as a common set of interoperable information capabilities for accessing, processing, storing and sharing information securely across dynamically established communities of interest. NCES services include messaging, discovery (tools to find relevant information), mediation (brokering, aggregation and fusion), collaboration/coordination, storage, and other services [11]. These net-centric services facilitate horizontal fusion.

One key architectural design pattern suggested for horizontal fusion is a service-oriented fusion network enabled by distributed coordination-based network services. Figs. 1, 2 and 3 illustrate the relationships between sensors, fusion services and the underlying communications infrastructure [5, 6]. The key concepts are twofold: (a) the processing abstraction is a distributed system architecture of semi-autonomous nodes and (b) information moves between systems based on a coordination-based publish-subscribe communication model. These concepts were discussed in [7].

In the following sections, we elaborate on the suggested coordination-based architectural design pattern [7, 12] in the context of horizontal fusion [5, 9, and 11] to develop a more detailed functional model that elaborates on three key aspects of horizontal fusion:

- Distributed coordination-based processing;
- Adaption and transformation; and,
- End-to-end process visualization.

For purposes of discussion, we assume the starting point for the development of a horizontal fusion architecture is a hybrid peer-to-peer architecture distributed across a wide-area network [13] with coordination-based interprocess communications [7].

3 Distributed Coordination Systems

This paper suggests that an important architecture in support of horizontal fusion is based on a hybrid distributed peer-to-peer coordination-based system, illustrated in Fig. 4 [7]. Tanenbaum and van Steen discuss coordination-based systems as a collection of autonomous distributed computational processes. The process coordination functionality manages interprocess communications and cooperation. We focus on a particular subclass of publish/subscribe based coordination systems where processes can subscribe to messages containing information on specific subjects.

The central design principles are that the core communications system is highly application independent; the interprocess messages are self describing; and the processes are not required to be directly referenced. These requirements allow fusion information to flow in a fault-tolerant service-oriented architecture that compensates for network and process outages. This architecture also permits the dynamic introduction of fusion services, sensors and clients as required.

Tanenbaum, van Steen *et al* suggest that commercial-off-the-shelf software can be utilized to build state-of-the-art distributed subject-based messaging architectures that do not require direct reference to service processes [7]. In this approach, a fusion process that is sending a message is not required to know the physical destination of the message. Instead, the sending process tags the message with a subject name and passes the message to the communications system for network transmission. In turn, receiving processes are not required to know the physical location of sending fusion processes. Receiving processes inform the communications system about the subjects they are interested in receiving.

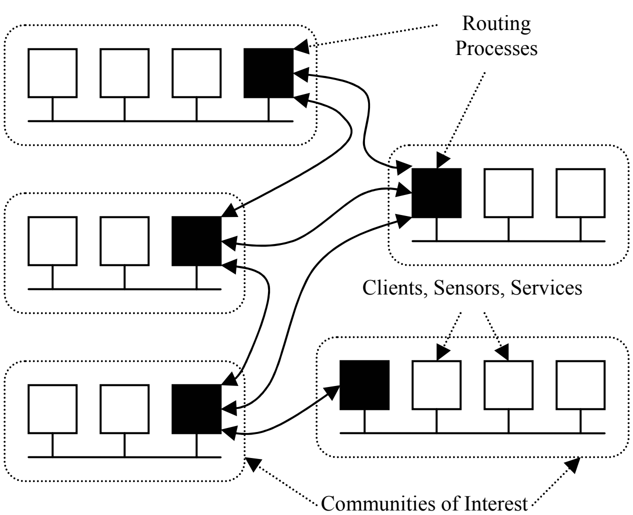

Fig. 4. Generalized Hybrid Peer-to-Peer Architecture

The communications system matches the interests of receivers with the information products of sending processes. This model creates a wide-area network of subject-based policy routing. For example, sensors and information sources would send a message to the network using subject-based addressing. This activity is referred to as publishing. All interested fusion application processes would receive the message because the process had previously subscribed to subjects published by sending processes. In addition, the suggested publish/subscribe communications network provides software quality-of-service options such as priority and guaranteed delivery, delivery receipts and other services that applications may require.

Process coordination for service-oriented horizontal fusion is a workflow service that manages and coordinates distributed communications between distributed network-centric processes. The architecture uses efficient distributed subject-based addressing to build a fault-tolerant fusion support capability. Fault-tolerant functionality may be established by heartbeat messages within process groups. The interested reader is kindly referred to Tanenbaum and van Steen for additional features and functionality of distributed coordination-based systems [7].

In the next section we describe adaptation and information transformation as software services for horizontal fusion clients and fusion services. Commercial-off-the-shelf adaption software facilitates interoperability with minimal shared design criteria that results in significantly lower life-cycle integration costs.

4 Adaption & Transformation

As previously discussed, publish-subscribe architectures are generally built on message-oriented communications networks. Component members of the service-oriented architecture who participate as coordinating processes create logically distributed virtual computing machines as depicted in Fig. 4. The peers in the network are software processes that support one or more communities-of-interest by routing information objects to interested information consumers [7, 10].

Information discovery, mediation, access, brokerage, quality-of-service are net-centric enterprise services provided by the enterprise information dissemination infrastructure. Clients may be information producers, consumers, fusion processes, storage and warehousing services, visualization, coordination, or other services that provide or utilize the infrastructure service provided by the network. Information access and transformation is provided at multiple levels in the architecture based on functionality. The peer-to-client and peer-to-service interface is a key interface that must be accurately profiled and securely managed. Software adaption and information transformation functionality is managed across this interface, functionally illustrated in Fig. 5.

Access to enterprise services can be facilitated by inexpensive commercial-off-the-shelf software adapters. These component adapters provide peer-to-client or peer-to-service software services that permit distributed systems in one community-of-interest to share information with other communities-of-interest. For example, a relational database client that uses structured query language (SQL) for access would require an SQL adapter. Information that resides in a directory that uses lightweight directory access protocol (LDAP) would require an LDAP adapter to connect to the service. Simple network management protocol (SNMP) alerts are another example of sensor information that could utilize a COTS adapter to connect to the wide-area enterprise information dissemination service.

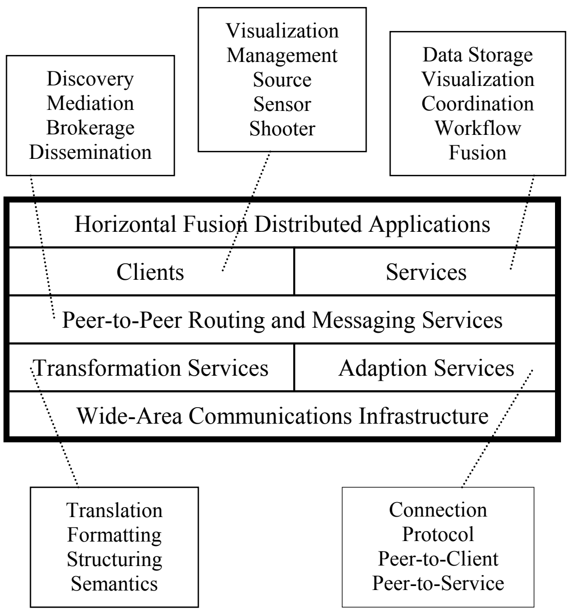

Fig. 5. Peer-to-Client and Peer-to-Service Architecture

The supporting client and service adaptation could be categorized as information dissemination management services. For purposes of discussion, we define transformation services as software services that transform information from one format/structure to another format/structure. For example, source information in an SQL database might be of interest to a net-centric fusion service that only accepts extensible markup language (XML) structured data. The transformation service would provide the structural translation that would facilitate interoperability. Unstructured text could be also transformed into XML if required by a fusion application or service.

The purpose of the adaption layer is to provide software connectors to facilitate cost-effective communications between distributed fusion services at the communications protocol design layer. Likewise, the transformation layer facilitates information interoperability at the application process layer. Transformation functionality may occur as various stages of the overall coordination-based system architecture.

5 End-to-End Visualization

In this section we turn our attention to an activity that is of interest to system managers and administrators, the visualization and network management of coordination-based systems. Horizontal fusion event management activity should monitor every stage in the decision support process from sensor output to a final situational assessment at the operational site, including alerts when things go wrong. Horizontal fusion performance management functionality monitors and reports on key indicators and objectives of the decision support cycle, including speed-of-command and information reliability across the fusion network.

Visualization also fosters collaborative planning and enables decision makers to collaborate on dynamic sensor activation and synchronize events based on a dynamic information exchange. This architectural design facilitates aligning decision support infrastructures to changing mission objectives and conditions, such as new military campaigns, strategic objectives and new intelligence segments, reducing command cycle time.

Collaborative horizontal fusion integrates net-centric processes, including rule-based activities, automated sensor updates, and multiple sensor support. Collaborative fusion manages decision support throughout all stages of the decision support process, even across communities-of-interest and enterprise boundaries. The visualization of the network enables operators to quickly determine where and when an intelligence product can be obtained and facilitates sensor management, availability checks, and network management.

The discussion in this section reminds us that visualization tools for coordination-based systems used for horizontal fusion are very similar to visualization tools for distributed supply chain management. Our observation aligns with the desire to use commercial-off-the-shelf software in horizontal fusion architectures. For this reason, we observe that system architects can rapidly insert COTS distributed supply chain management visualization software to build and manage horizontal fusion infrastructures. Recent prototyping activity appears to verify this observation.

Another important visualization functionality in horizontal fusion is the requirement to create mission profile templates and policies that can be used to configure coordination-based access control across the network. Configuration information created by visualization templates could be applied to distributed routing processes (refer to Fig. 4) to implement information dissemination policies, both static and dynamic. Mission profiles and policies for many strategic and tactical mission scenarios could be created and managed with commercial-off-the-shelf visualization tools.

Before leaving this topic it is important to point out that the human-computer interface (HCI) is considered to be the bottleneck in the performance of fusion systems. Ultimately the results of fusion processes are displayed for humans to visualize. The speed-of-command is significantly increased when appropriate supporting processing is performed without humans in the loop.

Further elaboration on horizontal fusion and visualization will be the subject of future papers. The reader is referred to [1-6] for earlier work on fusion and data mining visualization in the context of multisensor data fusion, military systems, intrusion detection systems and critical infrastructure protection.

6 A Caveat on Peer-to-Peer Services

Today's market-driven Web Services (WS) architectural design patterns are built on client-server request-reply networking models with directory-based discovery services. This architectural pattern is well suited for open commercial electronic business models where discovering the business service increases the commercial value of the service. However, for decision support models that require a greater degree of access control and accountability, the distributed coordination-based system design pattern becomes an attractive design choice for system architects and should not be overlooked. There is recent activity to develop open standards for WS process-coordination and orchestration.

There is also a trend toward peer-to-peer services that have familiar names like Napster, Grokster, Morpheus, and Kazaa. These services were designed to obfuscate the transfer of information between peers. Because of that basic design premise, these peer-to-peer services do not have robust software features for access control, accounting and auditability. Furthermore, these Napster-style peer-to-peer services do not support coordination between distributed processes and are primarily designed for ad hoc file sharing without accountability.

When considering horizontal fusion architectures, distributed coordination-based systems have numerous advantages over ad hoc peer-to-peer networks and directory-based discovery services. Conceptually the coordination-based architectural model is a wide-area network of information routers with pattern matching engines transposed on top of an underlying communications facility such as the Internet [13]. Each client or service that uses the communications infrastructure is only required to establish a service-level-agreement with the enterprise service, illustrated in Figures 3 and 4. The users of the enterprise service are not required to establish service-level-agreements with all other users of a coordination-based system. In other words, the coordination-based, or broker/mediation architectures tend to scale from an operational and economic perspective.

When service relationships are established with an enterprise service, information flows across the access nodes (peers), denied by default, and permitted by explicit policies that are applied to each peer node in the network. This architectural pattern creates a network of peers that route information based on security policies that are normally part of larger information sharing policies that comprise mission profiles and templates. Also, because many supply chain architectures are based on publish/subscribe architectural design patterns, horizontal fusion applications benefit from the similarity.

Therefore, we recommend that system architects consider appropriate architectural design patterns based on the transaction model best suited to accomplish the mission. One size does not fit all and caution is advised when following market-based trends. Web-services are an architectural pattern that has utility but should not be considered as a complete solution. In fact, for distributed coordination-based systems, Web-services are very useful for client-to-peer and service-to-peer adaption services, illustrated in Fig. 5. On the other hand, the vast majority of systems that must be a part of a horizontal fusion network have interfaces that pre-date Web-services. It is prudent for horizontal fusion architects to build an adaption and transformation strategy that leverages existing infrastructure investments.

7 Security Considerations

Scaleable engineering solutions are achievable when processing components are minimally coupled in an architectural model that adheres to complexity management principles. Coordination-based systems should be decoupled from confidentiality, integrity and non-repudiation services whenever possible. These security services should be provided by the enterprise, utilizing enterprise-class virtual private networks or end-to-end cryptographic systems in a logical defense-in-depth strategy [14].

Many messaging systems, including coordination-based systems, provide security services. Architects should carefully consider the benefits of avoiding unique and proprietary cryptographic services that sacrifice interoperability. The need for security should be balanced against established operational risk criteria and care should be exercised to avoid non-standard or vendor unique security services when building horizontal fusion architectures.

8 Conclusions

It is very important for horizontal fusion architects to consider numerous design patterns and, in particular, service-oriented architectures built on distributed coordination-based models. These system architectures appear to provide an economically scaleable distributed computing infrastructure, required for fault tolerant horizontal fusion applications. Architects should be careful not to follow a single architectural pattern and should take care that proposed solutions fit the operational environment, combining good ideas with organizational, economic and geographic constraints.

Coordination-based systems are very useful in supply chain and horizontal fusion applications because they are designed to leverage existing information technology investments with a heterogeneous adaption and transformation architecture. Emerging open standards in Web Services process orchestration and collaboration are encouraging and the fusion community could benefit from active participation in the development of these emerging standards.

Coordination-based systems can provide quality-of-service capabilities and can be designed to be fault-tolerant. Subject-based messaging used in conjunction with coordination-based architectures permit fusion sources to publish information without being aware of the physical location of the receiver. Likewise, receivers of fusion information may operate without specific knowledge of the physical location of the sender. The author encourages system architects to review Tanenbaum and van Steen [7] and also Hall and Llinas, *editors* [5], when considering communications architectures for horizontal fusion and, in particular, distributed coordination-based fusion architectures.

Acknowledgments

This paper was motivated by numerous excellent discussions on net-centric enterprise services and information dissemination management with Mr. Jack Eller, Program Manager for the DISA CS/IDM program and many friends and colleagues in the United States Air Force and across the federal government.

In addition, we thank Mr. Don Adams, Principal Architect, Office of the CTO, Tibco Software, Inc., for his valuable insight into real-time distributed coordination-based architectures.

Also, words of appreciation are in order to Dr. Mark T. Maybury, Executive Director, Information Technology Division of The MITRE Corporation for his reminder to include a discussion on the JDL and Mr. Nino Bonura. Senior Solutions Consultant, Tibco Software, Inc. for his review and comments.

Special words of admiration are offered for the many dedicated scientists who work to foster the emerging art and science of multisensor data fusion, including, David L. Hall, Ph.D., Associate Dean of Research and Graduate Studies for the Pennsylvania State University School of Information Sciences and Technology; and James Llinas, Ph.D., Adjunct Research Professor at the State University of New York at Buffalo.

This work was partially supported by Silk Road contract GS-35F-0290K, Order Number DCA100-02-F-4522; DISA Information Dissemination Management (IDM) Architecture Analysis and Definition Task.

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,"* Artech House, Boston, MA, 1990.

[4] Hall, D., and Llinas, J., *"An Introduction to Multisensor Data Fusion,"* Proceedings of the IEEE, Vol. 85, No. 1, IEEE Press, 1997.

[5] Hall, D. and Llinas, J. *editors,* Handbook of Multisensor Data Fusion, CRC Press, Boca Raton, Florida, 2001.

[6] Bass, T., "*The Federation of Critical Infrastructure Information via Publish and Subscribe Enabled Multisensor Data Fusion,"* Proceedings of the Fifth International Conference on Information Fusion: Fusion 2002, Annapolis, MD, 8-11 July 2002, pp. 1076-1083.

[7] Tanenbaum, A. and van Steen, M., Distributed Systems Principles and Paradigms, Prentice-Hall, Inc., Upper Saddle River, New Jersey, 2002, pp.700-716.

[8] Lambert, Dr. D. A., "*Assessing Situations*," IEEE Proceedings of Information Decision and Control, February 1999, pp. 503-508.

[9] RDT&E Project Justification, Project Name: Horizontal Fusion, 0305190D/C3I, August 2002.

[10] Bond, A., *"Open Distributed Services ODSI White Paper,"* Distributed Systems Technology Centre (DSTC), 2001.
[11] Hutton, R., *"Transforming Joint Command and Control Capability,"* Military Information Technology, Volume: 7 Issue: 3, March 30, 2003.

[12] Gamma, E., Vlissides, J., Johnson, R., and Helm, R., Design Patterns, Addison-Wesley Professional; Boston, MA, 1st edition, January 15, 1995.

[13] 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.

[14] 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.

Additional Reading

Ko, I., Yao, K., and Neches, R., *"Dynamic Coordination of Information Management Services for Processing Dynamic Web Content,"* Information Sciences Institute, University of Southern California, *WWW 2002*, May 7-11, 2002.

Eugster, P., Felber, P., Guerraoui, R., and. Kerrmarec, A.-M., "*The Many Faces of Publish/Subscribe*," Submitted for Publication, Tech. Report DSC ID:200104, 2001.

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

Public Release Approval

Review of this conference paper by the Department of Defense for public release does not imply Department of Defense endorsement of factual accuracy or opinion.

This conference paper was cleared for public release by the Department of Defense on June 15, 2004. DD Form 1910 tracking number 04-S-1637.

Biography

Tim Bass (bass@silkroad.com) is a Senior Net-Centric Subject Matter Expert with Silk Road, specializing in operational concepts, systems design, architecture and security of distributed Internet applications. He provides senior subject matter expertise to the USAF, Office of the CIO, DoD Office of the CIO, ASD/NII, the Defense Information Systems Agency (DISA) and other large enterprises. 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 12 years. Mr. Bass' work on Internet security and countermeasures for the USAF has been featured in Popular Science Magazine and Federal Computer Week. His list of publications, detailed biography and company information is available at the Silk Road web site (www.silkroad.com).

