



Boundaryless Information Flow Reference Architecture

DRAFT VERSION FOR REVIEW AND COMMENT

Prepared by:

Eliot Solomon
Principal, Eliot M. Solomon Consulting, Inc.
eliot@eliotsoolomon.com

Terence J. Blevins
Chief Information Officer, The Open Group
t.blevins@opengroup.org

January 24th, 2003

***Boundaryless Information Flow
Reference Architecture –
DRAFT VERSION***

Copyright © 2003, The Open Group

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the copyright owners.

Boundaryless Information Flow Reference Architecture

Issued by The Open Group, February 2003.

Any comments relating to the material contained in this document may be submitted to:

The Open Group
44 Montgomery St. #960
San Francisco, CA 94104

or by Electronic Mail to:

ogpubs@opengroup.org

Table of Contents

Executive Summary	1
Introduction	2
Background	2
Need for Architecture Continuum	2
Business Models Demand IT Response	3
Identification of Generic Business Models	4
Map of IT Architecture Styles to Business Models	4
Focusing on Boundaryless Business Models	5
Boundaryless Information Flow Reference Architecture	6
Overall Model Description	6
Sub-architecture Descriptions	7
Workflow Management Architecture	8
Messaging Architecture	10
Security Architecture	11
Directory Architecture	13
System Management Architecture	13
Information Architecture	16
User Interface and Ontology Architecture	18
Transaction Management Architecture	19
Mobility Architecture	21
Summary	21
About the Authors	24
About The Open Group	25

Table of Figures

Figure 1 The Business Continuum of IT Architecture	3
Figure 2 Mapping Business Strategy to IT Styles	5
Figure 3 Common System Architectures of Boundaryless Architecture	6
Figure 4 Security and its related services	11



*Boundaryless Information Flow™
achieved through global interoperability
in a secure, reliable and timely manner*

Executive Summary

In 2002 The Open Group updated its vision to be “Boundaryless Information Flow achieved through global interoperability in a secure, reliable and timely manner.” To help the membership of The Open Group engage in efforts to achieve this vision, this document is presented as a “straw man” of the high level strategic recommendations for Reference Architecture. The customers and members of The Open Group are encouraged to critique the document, ideally offering comments regarding specifics of strategy and the implementation of these strategies to one of the authors. These comments will then be incorporated into a White Paper published by The Open Group.

This document explains a framework called the Boundaryless Information Flow Reference Architecture. The Reference Architecture will guide the creation of specific architectures, solutions, and systems for use by companies whose key IT objective is the reduction of the friction of internal and/or external boundaries. It includes descriptions of business strategies, scenarios and profiles to help an organization identify the significance of Boundaryless Information Flow architectures to its business strategy and goals.

This document provides an initial explanation of the objectives of the Boundaryless Information Flow Reference Architecture. It relates this architecture to The Open Group’s TOGAF, and especially to the concept of the Architecture Continuum. In this context, this document identifies the parts of the generic IT architecture — The Common System Architectures — that require particular attention and have the most influence when designing IT architectures that accomplish boundarylessness.

This document on the Boundaryless Information Flow Reference Architecture provides a framework to elicit, evaluate, and position further architectural contributions, from within and outside The Open Group, for Boundaryless Information Flow.

Introduction

Background

Organizations that choose to move toward the Boundaryless Organization¹ to improve their operational effectiveness are finding Information Technology resistant. The Open Group is seeking to help organizations address that resistance, thereby achieving Boundaryless Information Flow in support of their movement toward the Boundaryless Organization.

Boundaryless Information Flow represents the vision that The Open Group is pursuing. This document takes a step toward that pursuit by providing a framework to elicit, evaluate, and position necessary architectural contributions for Boundaryless Information Flow.

Need for Architecture Continuum

An Enterprise's IT Architecture needs cannot be addressed with a single architecture. It requires the support of a continuum of architecture. The concept in the continuum is that there is a range of related areas that require separate consideration and attention. Each area addressed through architectural building blocks organized in specific architectural styles representing architecture models that contribute to the enterprise architecture.

In the paper "Family of Architectures" an application of TOGAF's Architecture Continuum sets the stage for this need and the specific areas that demand attention.

The scope of the problem an architecture framework is intended to address determines the depth and breadth of subjects it must explicitly address – that scope determines how complex the framework's Architecture Continuum must be. In this white paper we discuss the minimum necessary breadth of an IT architecture that will guide the creation of IT systems that support a Boundaryless Organization. Based on the approach presented in the "Family of Architectures" paper, we present a set of "Common System Architectures" that will guide the various choices that most influence and are most constrained by the objective of Boundaryless Information Flow.

By understanding this specific interpretation of Architecture Continuum we better understand the requirements for addressing the issues surrounding this problem.

¹ "The Boundaryless Organization: Breaking the Chains of Organizational Structure" by Ron Ashkenas et al.

Business Models Demand IT Response

An essential part of the Architecture continuum is its mapping to business space. One dimension of the continuum, which we call the “business continuum of IT architecture,” addresses the creation (or evolution) of IT architectures to meet the needs of business domains. The business domains of the business continuum extend in two dimensions.

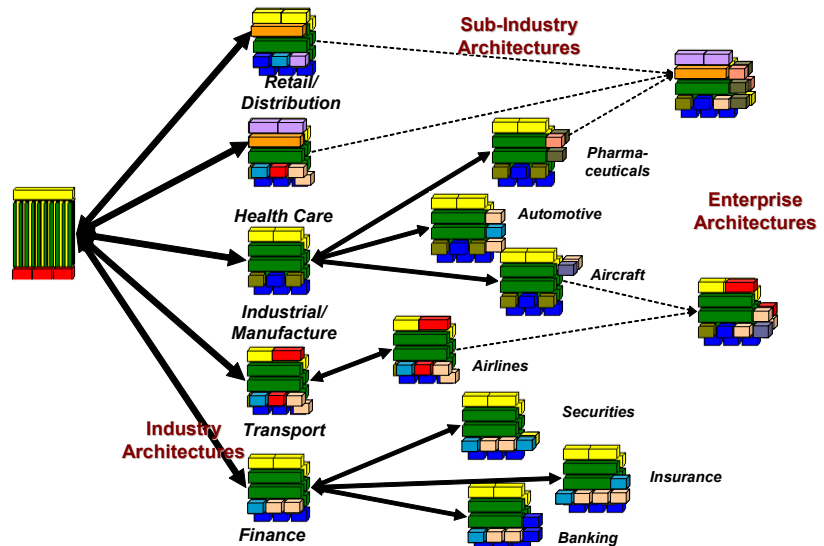


Figure 1 The Business Continuum of IT Architecture

In one dimension, architectures differentiate by the line of business they address. In the other dimension, architectures become increasingly specific, addressing the needs of smaller and smaller entities. The first branching, as shown in Figure 1, differentiates architectures for different industries. What constitutes an industry is, to some extent, subjective. Certainly, though, any grouping of enterprises that has common business processes, competitive objectives, and that develops or adopts IT standards or conventional solutions should be considered to be one that has, implicitly or explicitly, an IT architecture.

The progression from less specific to more specific architectures is accomplished by successive differentiations, dividing industries into finer and finer sub-industries, and, eventually, to individual enterprises. Beyond the enterprise, it may be appropriate to further refine an IT architecture to meet the needs of specific lines of business, business units, or departments of the enterprise.

The process of creating increasingly specific architectures is not just one of differentiation. In many cases, a more specific architecture may be created from two more general architectures. In Figure 1, each enterprise shown is a participant in more than one industry or sub-industry. The refinement of two architectures into a single, more-specific architecture may call for the inclusion of variants of the same element from each of the contributing architectures. In some cases, it may be appropriate to select the variant

from one of the architectures over the forms of the element in the other contributing architectures. And, in some cases, the consolidation of two or more architectures may require the rejection of all the predecessor architectures' variants of an element, and the creation of a wholly new solution that addresses the requirements of the specific enterprise.

Identification of Generic Business Models

The process of transforming broadly applicable architectures, such as industry architectures, into directly applicable architectures such as enterprise architectures involves, at each step, the substitution of specific solutions for architectural possibilities. Ideally, users of IT will find appropriate solutions for their specific objectives available on a commercial basis, saving the cost and uncertainty of custom development.

To achieve the cost efficiencies and improvement of quality that standard elements can provide, there must be relatively few solutions available to meet many businesses' different needs. Some of these solutions will be specific to an industry or other vertical market, but most IT solutions will be "horizontally" positioned. Specific solutions for increasingly specialized business architectures will often be chosen based on its addressing a business objective that, while specific, is generic. That is, it is well suited to a business objective or strategy that is not unique to a particular business or industry.

Map of IT Architecture Styles to Business Models

Generic business models or strategies describe ways enterprises approach doing their business and, more specifically, how they do business in ways that get them an appropriate competitive advantage. There are many ways to identify and classify business models. Porter's three generic competitive strategies is a frequently used model. Boston Consulting Group's quadrant model of stars, cows, and dogs is another approach to describing corporate business strategy. IT choices made by a company should be appropriate to the company's strategy, no matter how it is analyzed. We recommend an approach to analyzing and classifying corporate strategy whose perspective (or viewpoint) emphasizes aspects that strongly influence IT architectural options. Figure 2 illustrates the approach. The first differentiation is between enterprises who compete on information and those that compete with processes or capabilities. Some businesses base their competitive position on having more or better information than their competitors, or on being better able to manage the information. Other companies base their competitive edge on having better processes to work with information (or whatever resources the systems process). A business of this sort might, in fact, have no information of its own, but process information that belongs to and is managed by its clients.

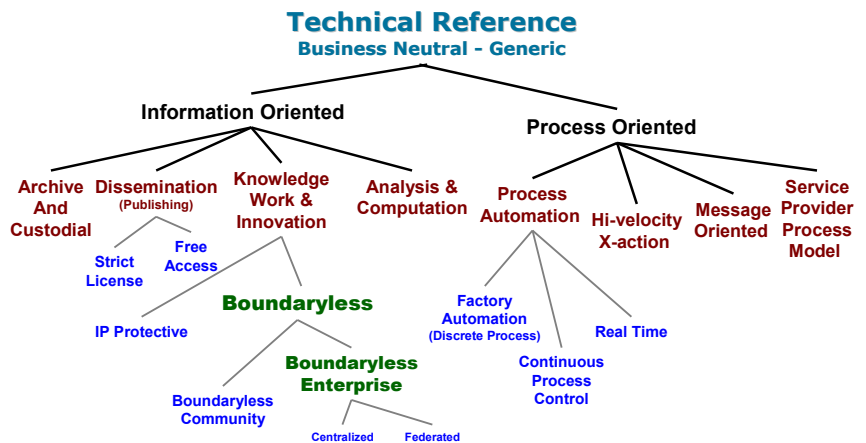


Figure 2 Mapping Business Strategy to IT Styles

Info-centric companies are further differentiated by the uses to which they (or their clients) put the information. We identify in the figure four major ways in which information is used which might call for meaningfully different IT architectures. One is Information Dissemination. (Digital publishing is an example of an industry that has this model.) Another is “Knowledge work,” often in the form of research and development, but also represented as professional services, education, and similar businesses.

Information dissemination and knowledge work each are shown with two major subdivisions. In both cases, one of the two subdivisions represents a business model that (generally) emphasizes strict control over the information, and one that emphasizes easy access and broad use of the information. The latter are the “Boundaryless” business models that strongly benefit from a “boundaryless” IT architecture.

Focusing on Boundaryless Business Models

This document describes in more detail the elements of an IT architecture that is particularly well suited to the needs of a company that chooses boundarylessness as a key element of their business and competitive strategy.

Boundaryless Information Flow Reference Architecture

Overall Model Description

Figure 3 represents the overall model of the Common Systems Architectures required to address the needs of Boundaryless Information Flow. The box at the top of the figure marked “Boundaryless Information Flow” represents this document. The boxes attached to Boundaryless Information flow each represent a specific subject area that requires its own sub-architecture.

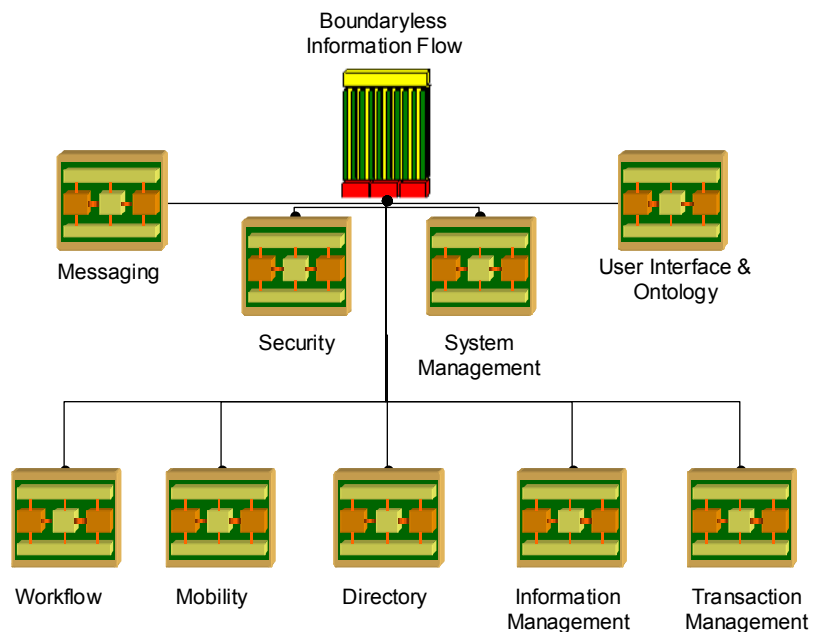


Figure 3 Common System Architectures of Boundaryless Architecture

This document will describe each of the sub-architectures from a functional perspective.

The model above was created from a viewpoint of understanding the functionality required to address the Boundaryless Information Flow problems in context to separable technical domains. At the high level to provide Boundaryless Information Flow one needs to have:

- A messaging infrastructure that delivers information to where it is needed, all the way to the end user at “the edge.”
- A security infrastructure to ensure that only the people or applications that are permitted to see the information are provided the information.
- A management infrastructure to ensure that the resources are

available to deliver and process the information any and every time there is a need.

- An information management infrastructure to ensure that the information resources are available any and every time there is a need.
- A transaction management infrastructure to assure that the information is consistent and accurate within the context of its usage.
- A directory infrastructure to make the location and access mechanisms of resources known to all trusted users and applications.
- A workflow infrastructure to manage complex flows of information and information processing involving multiple usages.
- A user interface and ontology driven infrastructure to provide the information to the information consumer in the context of their particular use – usually driven by the process in play.
- A mobility infrastructure to support the movement of users and resources.

Each of the above infrastructures requires a separate sub-architecture. The needs of those sub-architectures are elaborated upon in the following sections.

Sub-architecture Descriptions

The Common Elements of Boundarylessness

Boundaryless Information Flow architectures, like all IT architectures, are built of “sub-architectures.” Sub-architectures address specific aspects of the overall IT architecture, like security or messaging. Each type of sub-architecture may, in practice, be represented by several different architectures or architecture styles. For example, data management may be done using centralized architectures or distributed architectures; messaging services may be implemented in peer-to-peer or client/server architecture models or styles.

A comprehensive architecture, like the boundaryless information flow architecture, is made of selected forms of these sub-architectures. For many classes of sub-architecture, there are examples that would be better suited for inclusion in an overall architecture, based on that architecture’s objectives, than are others.

In creating the Boundaryless Information Flow Reference Architecture, we need to identify the specific forms of each sub-architecture that are best suited to achieve the goals of boundarylessness. Other forms of these sub-

architectures that are better suited for other objectives may also be identified, for purposes of comparison, or to show how other objectives may be accommodated within an architecture primarily seeking boundarylessness.

Workflow Management Architecture

The workflow management architecture describes the tools, services and techniques that support the automation of a business process. Process definition, workflow interoperability services, application communication services, user interface services, and services that provide system monitoring are typically included in workflow management. Workflow systems are typically based on an asynchronous messaging model.

*Boundarylessness emphasizes
the flow of work over the
arrangement of systems*

The Workflow Management Coalition defines workflow as “The computerized facilitation or automation of a business process, in whole or part.” Additionally the WfMC defines a Workflow Management System as “A system that completely defines, manages and executes “workflows” through the execution of software whose order of execution is driven by a computer representation of the workflow logic.”

Important Viewpoints to Use

One can view workflow from the perspective of business process managers. These managers are primarily concerned with effective and efficient execution of the business processes sometime measured in terms of output quality, execution time, amount of scrap, costs, and reduction in iterations of the processes of parts of the processes. Of course there are other views to consider, this view was presented as one example.

The above definitions and views are described from the perspective of business processes.

Typically Included Services

The WfMC further describes the areas that need coverage:

- specifications for process definition data and its interchange
- interfaces to support interoperability between different workflow systems
- interfaces to support interaction with a variety of IT application types
- interfaces to support interaction with user interface desktop functions
- interfaces to provide system monitoring and metric functions to facilitate the management of composite workflow application environments

The generic model of most workflow management systems has the following types of component according to WfMC:

- software components which provide support for various functions within the workflow system
- various types of system definition and control data which are used by one or more software components
- applications and application databases which are not part of the workflow product, but which may be invoked by it as part of the total workflow system

The reference architecture and the standards of the WfMC should be considered a starting point for this area.

A Workflow Management Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Transparently supports geographic distribution of work flow activities
- Transparently supports multiple organizations involved in work flows
- Operates over an open ubiquitous infrastructure
- Operates within the constraints of an open security infrastructure
- Utilizes open interfaces, including formats and protocols, for information exchange

For consideration

At this time there are a number of competing and/or distracting movements that relate to workflow:

- <http://www.biztalk.org>
- <http://www.ebxml.org/>
- <http://www.rosettanet.org/>
- <http://www.wfmc.org/>
- <http://www.gotdotnet.com>
- <http://www-4.ibm.com/software/solutions/webservices/>

For future research: A good reference paper “Business Transactions in Workflow and Business Process Management” from OASIS.

*Moving information in
discrete packages*

Messaging Architecture

The messaging architecture describes the services and techniques that support the transfer of information from one process to another in the form of discrete messages. Message assembly, delivery and disassembly capabilities are always included in messaging. Reliability, archiving and audit, translation, and delivery fairness guarantees may also be included. Messages may be synchronous or asynchronous, explicitly addressed, broadcast, or multicast.

Important Viewpoints to Use

A viewpoint that can be used when considering the messaging infrastructure is that of the users of the system. The users of the system are primarily concerned with ensuring that messages can be sent and/or received at anytime given the global nature of today's operations. As users are moving toward conducting business with electronic messages security is a major consideration, as well as, reliable delivery and integrity of messages. Again other views would be considered in the detailed architecture work.

Typically Included Services

The messaging infrastructure is not only about communications but includes the applications that deliver the messages end to end. Therefore messaging includes applications and services such as:

- E-mail
- Telephony – Phone Messaging
- Telephony – Pager Messaging
- Instant Messenger
- Message unification
- Message queuing
- Publish and subscribe messaging
- Synchronous/reliable messaging
- Asynchronous/reliable messaging
- Message brokering

A Messaging Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Transparently supports messaging among multiple organizations

- Transparently supports reliable and secure messaging
- Provides asynchronous communications models
- Supports publish and subscribe to messages of interest
- Operates over an open ubiquitous infrastructure
- Operates within the constraints of an open security infrastructure

*Giving access when
appropriate and only when
appropriate*

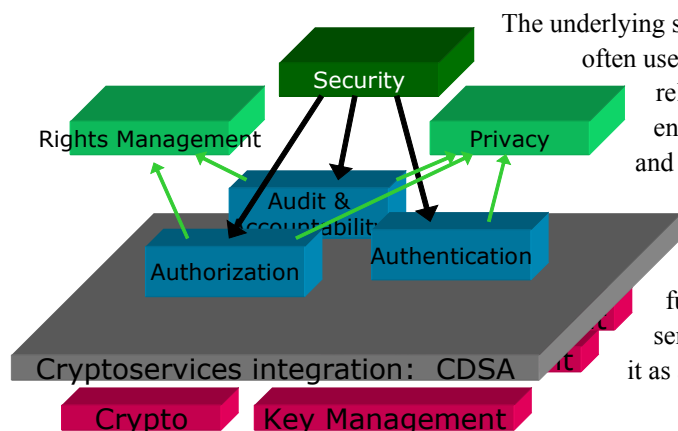
Security Architecture

Security is often described as the services or characteristics of a system that ensure that only authorized use of resources is permitted. This definition is certainly correct, but in practice it is too broad to provide specific guidance in IT architecture.

Security subsystem architectures

Security practice conventionally breaks security down into a number of sub systems or services, each of which may be described with an architecture. Subsystems commonly identified include authentication, authorization, audit, and data integrity. Technical services that underlie these functions are often treated as architected subsystems. These include cryptographic services and key management services, for example. PKI may be treated as a subsystem architecture for key management and distribution. CDSA represents both a solution and an architecture for cryptographic and key service integration.

Security and related solutions architectures



The underlying services (subsystems) of a security architecture are often used to create additional overlying solutions. These often relate in some way to the broad definition of security as ensuring that resources are only used in authorized ways and for authorized purposes. Some authorities use the term security narrowly, to represent only protective functions; others use it broadly, to include accountability and affirmative policy enforcement functions. For example, some authorities treat privacy services as part of security services, other authorities treat it as a distinct system.

Figure 4 Security and its related services

In the Boundaryless Information Flow reference architecture we consider all services that rely directly on the underlying services of the security architecture to be part of the security family of common subsystems or solutions architectures.

Viewpoints to be Considered

Security services must address the needs and constraints of a wide variety of stakeholders and interested parties. Because of its broad applicability and the broad consequences of its use, there has been much attention given to formal methods of security analysis and solutions documentation. Because of this, security architectures must be considered from several viewpoints, including user viewpoints, and formally defined technical viewpoints.

The key end (business) user viewpoints to be considered are:

- The owner of the resource(s) being secured
- The user of the resources being secured
- The custodian of the resources being secured

Operational and technical viewpoints to be considered are:

- Systems operations and management staff
- Audit, accounting, and compliance functions
- Regulatory and legal

Typically Included Services and Solutions

Services and solutions commonly included in or described by security architectures include:

- Password and similar system protective mechanisms
- Firewalls
- Secured networks
- Intrusion detection services
- Encryption and key management services
- Digital signature and document integrity services
- Identification and authentication services
- Authorization services
- Digital rights management, including licensing, labeling, etc.
- Privacy management

A security architecture that is well suited to a Boundaryless enterprise would emphasize mechanisms that facilitate the acquisition of privileges and permissions; provide means to ensure the authenticity of services that may be accessed by users without direct knowledge of the service or its

provider; facilitate federated or distributed administration of access control and authorization policies; facilitate the introduction of participants in ad hoc interactions in ways that enhance the trustworthiness of those interactions.

Directory Architecture

The directory architecture provides services for specialized support to locate required resources and for mediation between service consumers and service providers, and between and among users and other participants in interactions facilitated by the system.

The directory architecture includes tools and services that support the naming of participants and resources in a system, allow information to be published concerning the attributes of named entities, and facilitate the use of that published information either to identify a named entity based on its attributes, or to obtain the characteristics of an identified entity. Directory services are commonly used to provide the address of a resource, information about how that resource can be accessed, and information that allows the identify of the resource to be authenticated.

Important Viewpoints to Use

End user viewpoints

- Users seeking information
- Users publishing information
- Entities about whom information is published

Typically Included Services

- Directory Access Services
- Special Purpose Naming Services
- Service Location Services
- Registration Services
- Filtering Services
- Accounting Services

System Management Architecture

The system management architecture describes the utilities required to manage resources, services to handle events that require attention, services required to implement changes guided by management utilities, and information models that describe the data structures of the information that is used in the system. Utilities that predict and report on upcoming events

based upon current events are sometime included in a system management architecture.

Important Viewpoints to Use

System Management can be viewed through the eyes of the Operations, Administration, and Management (OA&M) personnel. The personnel of the OA&M organization are primarily concerned with making sure that the system performs as expected and is predictable in terms of behavior, response time, and availability.

The Open Group 's work in making sure that each and every component in a system can be manageable is an essential ingredient for System Management. In context to Boundaryless Information Flow this includes information assets as well as the resources that capture, update, access, deliver, and use information assets.

Typically Included Services

System Management constitutes a system of its own. It includes components and services that provide:

- User management services
- Configuration management (CM) services
- Performance management services
- Availability and fault management services
- Accounting management services
- Security management services
- Print management services
- Network management services
- Backup and Restore services
- On-line Disk Management services
- License Management services
- Capacity Management services
- Software Installation services
- Trouble Ticketing functions

Covering assets such as:

- Information
- Networks and networking components

- Software, system and user applications
- Computer hardware and peripherals
- User interface devices and peripherals

In context to Boundaryless Information Flow information is the key resource we wish to manage. The most important services are:

- User management services to manage user preferences that dictate how information is to be supplied to them.
- Configuration management (CM) services to ensure that the infrastructure can deliver information on demand.
- Availability and fault management services to ensure that the infrastructure can deliver information on demand.
- Network management services to ensure that the infrastructure can deliver information on demand.
- Performance management services to ensure that performance levels are maintained in context to the supported business processes.
- Accounting management services to track usage of information assets.
- Security management services (this should be covered in the management architecture)
- Print management services (this is a specialization of information delivery – the end point is a device, much like any other device. I'm not sure it belongs in our scope.)
- Backup and Restore services to ensure that information is recoverable.
- On-line Disk Management services to ensure that information is accessible on demand.
- License Management services to ensure that applications are usable when they are needed (as licenses expire and software may become unusable).
- Capacity Management services to ensure that the infrastructure can deliver information on demand.
- Software Installation services to ensure that necessary applications can be made available on demand.
- Trouble Ticketing functions to maintain the health of the system to ensure that the infrastructure can deliver information on demand.

A System Management Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Transparently supports geographic distribution of managed resources
- Transparently supports multiple organizations
- Transparently support semantically aware replication and synchronization of related information sources
- Operates over an open ubiquitous infrastructure
- Operates within the constraints of an open security infrastructure
- Utilizes open interfaces, including formats and protocols, for management information exchanges
- Address utilities required to manage the resources
- Address event handling of resources
- Address information models related to managing information
- Address instrumentation services that effect controls and commends from utilities

*How Can We All Talk the
Same Language?*

Information Architecture

The information architecture describes the services that provides ad-hoc and transparent information access including create, retrieve, update, and delete. Transparency is provided from geographic distribution, storage, sources, semantics, and different database management system models. Information architectures also typically include transformation and data cleansing services and support for relational, object oriented, and flat file systems.

Important Viewpoints to Use

A useful viewpoint when considering the information infrastructure is that of the users of the information. The users of information, human users and applications, are primarily concerned with understanding the semantics of the data, and being able to create, retrieve, update, and delete data. The users of information want the information in their context of use, and want only the information they need, rather than drowning in data much of which may be irrelevant. Consider situations where there are demands to make mission critical decisions in seconds. Along these lines users of information need the information at any time.

Typically Included Services

An information infrastructure constitutes a system of its own. It includes

components and services that provide:

- Information Storage and Access Services providing on-line services to access information including:
 - Create services
 - Retrieve services
 - Update services
 - Delete services
- Semantic Services
 - Metadata (data dictionary) Create services
 - Metadata (data dictionary) Retrieve services
 - Metadata (data dictionary) Update services
 - Metadata (data dictionary) Delete services
 - Application level transformation and mapping services
- Information Delivery Services
 - Interface generation services
 - Report generation services

An Information Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Support ad hoc access to information
- Transparently supports geographic distribution of information
- Transparently supports a hierarchy of storage
- Transparently supports information for multiple source organizations
- Transparently supports semantically aware federation of information operations
 - Provides on-line services to create, retrieve, update, and delete information
- Transparently supports information obtained from multiple database management system models, e.g. flat file, relational, object oriented, etc...
- Operates over an open ubiquitous infrastructure

- Operates within the constraints of an open security infrastructure
- Utilizes open interfaces, including formats and protocols, for information exchanges

User Interface and Ontology Architecture

User Interface and Ontology Services provide the means to consistently present data to the end user in the appropriate format. The interaction and ontology components provide services that assist in the formulation of customer data requests, and enable visualization and presentation of the data accessed within the context of its use.

Important Viewpoints to Use

The User Interface and Ontology infrastructure is viewed from the perspective of those that execute the business processes. Those that execute the business processes are primarily concerned with getting the right information at the right time to execute their part of a process efficiently and effectively. Those that execute a business process need information in the context of the business process, not in context to how information is stored in the system. This means that information is filtered and translated to fit the context of its use.

The notion of fitting information into the context of its use is supported by the concept of ontologies. Ontologies support domain specific languages and accepted relationships between concepts within a language. When ontologies are documented it enables context to be shared, understood and operated. As a simple example lets take two vary different ontologies; the marketing ontology and the battlefield ontology. Within these two domains there is different meaning to the word “target.” There is also two different meanings to the word “message” given a specific relationship to “target.” In a battlefield ontology “send the target a message” does not mean “send the enemy some brochures” as when this phrase is used in a marketing ontology.

Typically Included Services

User Interface and Ontology Services typically include:

- Ontology definition facility
 - Including an ontology map to business processes
- Ontology access services
- User interface presentation and transformation services based on the pertinent ontologies
- Meta indices of information categorized by ontologies

- Portal and personalization services based on the pertinent ontologies

A User Interface and Ontology Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Address ontologies from different domains
- Operates over an open ubiquitous infrastructure
- Operates within the constraints of an open security infrastructure
- Supports ad hoc access to information
- Transparently supports multiple organizations involved in work flows
- Transparently supports semantically aware federation of information operations
- Utilizes open interfaces, including formats and protocols, for information exchanges

Transaction Management Architecture

The transaction management architecture describes the services and techniques that support a complex global business transaction. Fundamentally supporting a complex global business transaction means that at the end of the business transaction everyone engaged in the transaction has fulfilled their portion of the transaction – there are no surprises in context of computing records. To achieve this a transaction management architecture must be able to describe a global transaction, demark a global transaction and subsequent sub-transactions, co-ordinate compensating sub-transactions, communicate with underlying distributed transaction processing managers, handle timeout events, and monitor the overall status of a global transaction.

Important Viewpoints to Use

Transaction Management is being viewed from the perspective of the business process. It is the business process that actually drove the need for the earlier versions of transaction management where the major concern was getting a high integrity record of a given business event, such as a deposit, withdrawal, sale, etc. This led to the two-phase commit protocol that ensured that there was high integrity of the database in distributed situations.

Supporting business processes today require the same integrity, yet more due to the nature of the business processes requiring support. In today's environment where business processes are creating more complex

heterogeneous business transactions the concerns are shifting. Where two-phase commit is assumed, there are additional concerns that sit on top of it.

For example in heterogeneous business transactions the timing of a business transaction is not short as a bank deposit or withdrawal. The duration of a business transaction can span days or more. So the concern is to be able to complete the business transaction where there are connected sets of alternatives and sub-transactions. When alternatives are executed other alternatives may need to be backed-out. This mimics what is done today, but done manually by buyers and suppliers communicating over the phone and using faxes.

So the concerns today are supporting complex business transactions that include connected alternatives. This requires some support for things such as compensating transactions, or an appropriate alternative.

Typically Included Services

A Transaction Management System that supports Boundaryless Information Flow would include service to:

- Describe/encode a complex global transaction
- Start a complex global transaction
- Associate sub-transaction to the global transaction
- Co-ordinate compensating alternatives
- Accessing underlying sub-transaction controls
 - Start, commit, or roll back
- Controlling timeouts on transactions
- Chaining transactions together as necessary
- Monitoring overall complex global transaction status

A Transaction Management Architecture, and subsequent system, that would be appropriate for Boundaryless Information Flow would have the following characteristics:

- Transparently supports geographic distribution of business process
- Transparently supports multiple organizations
- Transparently support compensating transactions (or alternatives)
- Operates over an open ubiquitous infrastructure
- Operates within the constraints of an open security infrastructure
- Utilizes open interfaces, including formats and protocols, for

transaction management information exchanges

- Communicates with open DTP transaction managers

Mobility Architecture

Emerging Technology that Transcends Boundaries

Architectural approaches to mobility are becoming increasingly important as a wider variety of mobility-enabling technologies need to be consistently incorporated into Enterprise IT planning. At this writing, the authors do not profess to have a concise or consistent description of what a Mobility Architecture would include. In particular, its relation to other architected systems, such as directory services and networking services need further clarification.

There is no doubt that Mobility and its evil twin Ubiquity are increasingly important in addressing business needs and competitive opportunities, and that they have a strong affinity to the Boundaryless Enterprise.

Typically Included Services

- Session management
- Global location services
- User preference and profile management
- Desktop restoration and process resumption

Summary

Within The Open Group are working groups that have the knowledge and experience to advance the definition of the Boundaryless Information Flow Reference Architecture, and to our understanding of the Family of Architectures in general. This White Paper was written to engage those elements within The Open Group in the process. With this as our goal, the authors humbly offer these suggestions to the members of The Open Group.

The **Customer Council** has helped establish this problem definition through its work in requirements capture and understanding – the primary document that captured and expanded on this requirement being the Interoperable Enterprise Business Scenario. The customer Council is still capturing new requirements that need to be positioned with this problem space, or positioned in the context of a new problem space. We suggest that the Customer Council

The **Messaging Forum** can take a more detailed look into alternative architectures and solutions for the messaging infrastructure.

The **Security Forum** can take a more detailed look into alternative architectures and solutions for the security infrastructure.

The **Management Forum** can take a more detailed look into alternative architectures and solutions for the management infrastructure.

The **Directory Forum** can take a more detailed look into alternative architectures and solutions for the directory infrastructure.

The **Mobility Management Forum** can take a more detailed look into alternative architectures and solutions for mobility.

The **Real-Time and Embedded System Forum** should consider how boundarylessness as a driving strategy serves or is served by the objectives and capabilities of typical real-time and embedded systems. The authors believe that real-time and embedded systems may be best served by another architecture in the Family of Architectures. The forum should consider what such an architecture would emphasize, and offer it to be used, at first, as a counterexample to the Boundaryless Information Flow Reference Architecture and, possibly, as the next objective of a company-wide architecture initiative.

The **Base Platform Forum** should consider operating system solutions and specifications that support this problem space. They should work to develop solutions guidance for companies seeking to develop an operating systems procurement and deployment strategy that will support the likely path of boundaryless services and architectures over the near- mid- and long term.

The **Architecture Forum** can provide assistance to other groups in creating architectures. They should help develop The Open Group's specific methodologies (processes, responsibilities, documentation, etc.) for developing architectures in the TOGAF way. The experience gained in creating the Boundaryless Information Flow Reference Architecture should be captured and published within TOGAF 9.

There are areas that are currently uncovered in our existing groups. We as an organization need to work together and determine how we fill these areas, whether through special interest groups or working with outside consortia that are addressing these areas.

- information infrastructure
- transaction management infrastructure
- workflow infrastructure
- user interface and ontology driven infrastructure

In the summary we can present a table of work areas that can contribute, inspire contributions, and inform how contributions can be submitted.

***Boundaryless Information Flow
Reference Architecture –
DRAFT VERSION***

About the Authors



Eliot M. Solomon has worked on the leading edge of information technology for more than thirty years. He gained experience in such diverse fields as electronic warfare, military C³, international telecommunications, medical electronics, and office automation equipment. Common to all this work was real time operation, mission- or life-critical significance, distributed and networked computing, and a need for security, privacy, and assured integrity for the information being processed.

For the last seventeen years Solomon has brought his expertise and creativity to the Securities Industry. In fifteen years at the Securities Industry Automation Corporation (SIAC) and its subsidiary SECTOR, Solomon made significant contributions to the architectures of systems and networks the entire market relies on. He was appointed SIAC's first Distinguished Technologist and Vice President, in recognition not only of his contributions to SIAC, but also to the entire industry.

Solomon is founder and chair of The Securities Industry Middleware Council, Inc. (SIMC), an industry organization that works to improve the infrastructure of the Securities Industry. Solomon has guided SIMC since its founding in 1996, and helped it become a significant influence on the software vendors deliver to the industry, and the way the Securities Industry uses infrastructure technology. At The Open Group Solomon chaired the DCE program and is now a member of the Security Forum steering committee. He is frequently invited to speak at major conferences on the subject of linking information security and business policy, and the management of risk and trust. He holds an A.B. from Columbia University, and an M.S. from Polytechnic University.



Terence (Terry) Blevins is Vice President and CIO of The Open Group LLC. Terry has two major roles within The Open Group. First, Terry is working with the customer community to better understand key business oriented issues and driving programs to best address those issues, most recently helping to formulate Boundaryless Information Flow. Secondly, Terry looks over the internal information technology strategy applying the same principles and techniques from The Open Group's Architecture Framework that he uses in the first major role.

Previously, Terry was with the NCR corporation (21 years), last holding the position of Director of Strategic Architecture where he had responsibility for the management and oversight of NCR's Corporate Architecture, Architecture Development Process, and NCR's Architecture Capability Maturity. Terry chaired NCR's Architecture Council and was Co-chair of The Open Group's Architecture Program Group in 1999 and 2000.

Terry has BA and MS degrees in Mathematics from the Youngstown State University

About The Open Group

The Open Group is a vendor-neutral and technology-neutral consortium, committed to a vision of **Boundaryless Information Flow** achieved through global interoperability in a secure, reliable and timely manner.

The Open Group's mission is to drive the creation of **Boundaryless Information Flow** by:

- Working with customers to capture, understand and address current and emerging requirements, establish policies, and share best practices
- Working with suppliers, consortia and standards bodies to develop consensus and facilitate interoperability, to evolve and integrate specifications and open source technologies
- Offering a comprehensive set of services to enhance the operational efficiency of consortia
- Developing and operating the industry's premier certification service and encouraging procurement of certified products.

The interoperability that characterizes **Boundaryless Information Flow** results in gaining operational efficiencies and competitive advantages. Through access to integrated information, across the extended enterprise and beyond, employees, trading partners, and customers are enabled and empowered.