


Technical forum



The role of IBM's Open Blueprint approach in distributed computing

What is the Open Blueprint and why is it important? Knowing the answer to this question has far-reaching implications for IBM's customers, software vendors, consultants, systems integrators, and service providers. The answer affects the way in which these people plan their information technology investments or deliver solutions in a distributed or client/server environment, or both.

What it is. IBM's Open Blueprint is a structure for distributed systems and provides the base on which to build, run, and manage distributed applications in a multivendor, heterogeneous environment. It is built using industry standards and enables IBM to respond to customer requirements with products and solutions that connect seamlessly, interoperate with many industry hardware and software platforms, and integrate additional products and solutions from IBM and other suppliers.

Open distributed computing, or client/server computing, is the underlying model that is driving the evolution of information technology today. Any strategy for this rapidly changing environment must be flexible and dynamic to accommodate technological advances, while at the same time it must protect existing investments. So, any description of the Open Blueprint infrastructure is current at a specific point in time as the infrastructure evolves in accordance with the introduction of new technologies, an example of which is object technology.

Business environment. For businesses to be profitable, they have to be better than the competition in pleasing customers, providing value, and producing products or services at lower costs. Often it has meant that the businesses required re-engineering and their investments needed to be re-examined, especially in the area of information technology.

In their book *Re-engineering the Corporation*, Michael Hammer and James Champy write that re-engineering is "the fundamental re-thinking and radical re-design of business processes to achieve dramatic improvements in critical contemporary measures of performance, such as cost, quality, service and speed."¹ Don Tapscott and Art Caston in their book *Paradigm Shift* suggest that creating a new enterprise is a process of "generative learning. Rather than fixing the old, we are on a course to create the new. This involves considerable change in virtually everything we have come to know in our work lives."²

Changing information technology (I/T) to support re-engineered business processes often means connecting any user at a workstation to data and applications throughout the company, as well as extending that connection to companies where there is a business relationship.

©Copyright 1995 by International Business Machines Corporation.

In the 1970s, information technology was typically characterized by large centralized computers managed by highly skilled and specialized staffs. Users were connected by display terminals and had little clout in obtaining new or modified applications. A change in this setting began in the 1980s with the emergence of personal computers (PCs) and local area networks (LANs), the client/server revolution, the growing importance of the UNIX** operating system, and the emphasis on open systems. The worldwide recession also played a role in forcing organizations to rethink the way they do business.

The I/T challenge was how to integrate PCs, LANs, and emerging new technologies while protecting past investments and maintaining the integrity of company data and other I/T assets. It led to requirements for transparent global access, interoperability, desktop usability, platform flexibility, application development tools, manageability, and investment protection.

These customer requirements motivated the development of the Open Blueprint. The Open Blueprint enables the development, execution, and management of distributed, client/server applications that work together in an integrated fashion in today's heterogeneous, multivendor environment.


Open Blueprint details. The Open Blueprint addresses the challenges of the open environment by viewing a system as part of a distributed network and viewing the network as if it were a single system.

The Open Blueprint serves four major roles:

1. It helps customers to develop their own architectures and organize products and applications in an open distributed environment.
2. It describes IBM's directions for products and solutions in the open distributed environment.
3. It guides developers as they meet customers' needs by supplying products and solutions that include the appropriate functions and that can be integrated and can interoperate with other installed products.
4. It provides a context for the incorporation of new technologies into a distributed environment.

The key building blocks of the Open Blueprint are resource managers that can be functionally grouped into sets of services. The services are

Figure 1 Building blocks of the Open Blueprint



described below, where the letters and numbers in parentheses correspond to the letters and numbers shown on the structure depicted in Figure 1.

Network services (A) provide for the transport of data from an end point in one system to an end point in another. This group includes common transport semantics; the physical network, which may encompass local area networks (1), wide area networks (2), channels (3), and emerging forms (4); and protocols for transporting the data such as Systems Network Architecture (5), Advanced Peer-to-Peer Networking* (6), Transmission Control Protocol/Internet Protocol (7), Open Systems Interconnection (8), NetBIOS (9), and IPX** (Internetwork Packet Exchange) (10).

Distributed systems services (B) comprise three services. One is *communication services* (B1). They are mechanisms that allow parts of a distributed application or resource managers to talk with one another. These mechanisms are the conversational method (11), remote procedure call (12), or messaging and queuing (13). Another is *object management services* (B2). They provide transparent access to local and remote objects and include the object manager (14). The third one of this group is *distribution services* (B3), which provide common mechanisms to assist the parts of a distributed application or resource manager to communicate. Included in this group are the directory (15), security (16), time (17), and the transaction manager (18).

Applications and application enabling services (C) encompass three types of services. One is *presentation services* (C1) that define the interaction between applications and the user via various devices. Among them are user interfaces (19), print/view (20), and multimedia (21). A second is *data access services* (C2) that allow applications to access various types of data that may be in files (22) or in databases (23). A third is *application services* (C3) that consist of common functions and have been developed once and in a standard way. They include the transaction monitor (24), workflow manager (25), and mail (26).

Underlying all of these services are *systems management services* (D1) that provide the systems administrator with the facilities to manage a distributed computing environment and *local operating system services* (D2) that operate within the confines of a single system in a network.

Resource managers do not necessarily correspond to any specific product. The Open Blueprint is implemented by different products on different system platforms. It describes the technical attributes and characteristics of supporting software, reflects desirable functional modularity, provides software principles and guidelines, and specifies important boundaries and interfaces.

The background. The Open Blueprint has been evolving since the initial task force was formed to address questions regarding how IBM's systems interoperate in distributed environments. But it became readily apparent that this work effort had

to have a broader focus to include all platforms belonging to IBM as well as other vendors.

Over the years, the Open Blueprint has evolved from previous internal versions known as Open Distributed Systems Structure and Open Distributed Computing Structure. The basis for the lower layers in the structure is the Networking Blueprint announced by IBM in 1992. As the original structure evolved, the work group validated their efforts by soliciting feedback from customers such as American Airlines, Inc. and the Royal Bank of Canada. Those customers indicated they found the structure valuable, both as a tool for understanding how IBM products fit into a distributed environment and as an approach to develop their own infrastructure for distributed or client/server computing. As a result, the Open Blueprint was made available to the industry in April 1994. A cross-functional team representing IBM's divisions is responsible for ensuring that all the components or resource managers are kept current.

Standards. Standards are included from industry consortia, e.g., X/Open, Object Management Group (OMG), and the Open Software Foundation (OSF), as well as from standards bodies, e.g., the American National Standards Institute (ANSI), the Institute of Electrical and Electronics Engineers (IEEE), and the International Organization for Standardization (ISO). *De facto* standards are included when there are popular product implementations to provide investment protection and satisfy functional requirements.

Summary. IBM has come a long way since the days of Systems Application Architecture* (SAA*), which had as its objectives consistency and interoperability across IBM systems. Today, the Open Blueprint goes much further in meeting current and future customer needs by describing industry standard interfaces and protocols that will enable interoperability and integration across IBM's and other vendors' platforms.

Over time, the Open Blueprint will be expanded with additional function, and additional product implementations will be provided. For example, object technology is a key factor to the evolution of the Open Blueprint, and it is expected to become more prevalent in applications and in system and network components. The cross-divisional functional team mentioned earlier is the

same one responsible for the evolution of the Open Blueprint.

In answer to the question posed at the beginning of this forum, the Open Blueprint is a structure for distributed computing that will help IBM and others deliver integrated, interoperable products and solutions. It has value across the industry. For end users, it hides the complexities of the network and makes it appear as a single system. For application developers, standard interfaces enable a single system view of the network and allow for the development of interoperable applications that can run on many platforms. For system administrators, the Open Blueprint defines a consistent way to manage the network to hide the complexities from application developers and end users.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of X/Open, Ltd. or Novell Corp., Inc.

Cited references

1. M. Hammer and J. Champy, *Re-engineering the Corporation*, Harper Collins, New York (1993).
2. D. Tapscott and A. Caston, *Paradigm Shift*, McGraw-Hill Inc., New York (1993).

General references

Introduction to the Open Blueprint, G326-0395, IBM Corporation (1994); available through IBM branch offices.

Open Blueprint Technical Overview, GC23-3808, IBM Corporation (1994); available through IBM branch offices.

J. A. Colosimo
IBM Corporate Headquarters
Somers
New York