

A Brief Overview of the Digital Object Architecture

Robert E. Kahn and Jay Allen Sears
Corporation for National Research Initiatives
October 2003

The Digital Object Architecture was developed by CNRI with funding from DARPA as a means of identifying, managing and tracking information in the Internet environment. A key component of the architecture is the “Digital Object” (or DO). The term Digital Object is used to denote structured data in the form of a set of bit sequences that can be interpreted by a computer or other computational facility. Each DO has a unique identifier that is part of the DO and which may be used with a resolution system to locate the DO within a network (the DO may exist in multiple locations). Each DO consists of multiple elements each of which consists of a <type, value> pair. The resolution system provides for administration of identifiers over time.

Each DO may have associated with it a properties record that contains relevant metadata about the DO, such as the terms and conditions for its use. Some of the metadata may also be contained in the DO itself, such as its identifier, its length and all the type fields. This is known as key metadata. It may also have a transaction record associated with it to track usage, but the information in this record is normally intended only for the owner of the DO.

The DO is a data structure that is machine or device independent, location independent and it may be easily ported from one platform to another. Device and location independence will be key attributes when implementing a distributed system managed by cooperating countries for passport identity management. Thus, an investment in creating a DO will have value long after the initial system that housed it is retired. DOs may be stored in “Repositories” which are themselves Digital Objects containing other DOs. Since DOs can be mobile in the network, it follows that Repositories can also be mobile as well. Furthermore, DOs need not be contained in Repositories, but may be separately identifiable data structures incorporated in other digital objects. A composite DO can consist, in part, of other DOs, which may themselves exist in other Repositories, all wired together with handles. In most cases, however, Repositories will be associated with specific network-based storage systems from which DOs may be accessed and into which they may be deposited.

The Repository is an interface specification. It is a portal to a storage system, but does not itself specify either the hardware or software to be used. The Repository interface may be arbitrarily sophisticated, but at its core it must support two key functions. One is the ability to accept a deposit of a DO from a user or another Repository. Two is the ability to support access to a DO by a user that supplies the identifier for the DO and an appropriate service request. The service request generally consists of running a program against the DO to produce a “dissemination” over the net. The service request would involve a means of identifying the appropriate program to run against the DO.

By use of this standard interface, each repository is guaranteed the ability to interoperate with every other repository at the level of moving digital objects between them. While these can be the basic passport objects, it is more likely that these will be subsets of passport objects or objects that reflect country-specific decisions about travel based on real-time information shared between countries at or near the point of departure and time of travel. This will be especially useful in

PassPort (PP) processing since each government can manage their own repository and yet interoperate with other countries using individual country policy constraints. The simplest service requests would be 1) running an identity program that produces the original DO in its entirety, 2) running a program that extracts a particular element or subset of elements of the DO, 3) running a program that executes the DO, if it is executable and 4) running a program that extracts the first N bytes of the DO or plays the first N minutes of the DO (if it is a time based object).

Access control is used to limit access to the Repository. This may be done on a DO by DO basis or some other means. Thus, knowing the identity of the Repository and the unique identifier of a DO in that Repository does not guarantee access to the object.

Resolution of identifiers is accomplished by a Resolution System, which is called the Handle System. Each identifier is known as a handle. The result of a handle resolution is a “handle record” which looks like a DO and provides state information about the designated digital object. For example, it might identify the IP addresses of some or all the locations that contain that DO. If the identifier were given to a DO that reflected an access page for something else, like an individual, the handle record might contain access information for the individual (e.g. telephone number, cell number, fax, or even public key).

The Handle System was designed to be flat, non-nodal and scaleable. It is a distributed system consisting of at least as many servers as needed to provide the then current handle service. Only a subset of the servers needs to contain each handle record. One would suffice, if each were perfectly reliable and accessible. Multiple servers guarantee reliability in the face of outages and can be used by clients to improve resolution performance through load balancing and network proximity. If the system consists of N servers, each providing a separate handle service, the user software needs to know which of the N servers will resolve a handle resolution request. This is determined precisely by a hashing algorithm that maps each handle into a specific handle server guaranteed to contain the handle record. This initial version is called the Global Handle Registry (GHR).

For many reasons, organizations may wish to maintain control of their handle records and resolution systems. Local Handle Servers (LHSs) may be used for that purpose. In general, each local handle server is made known to the GHR when it is brought up. A handle consists of a prefix and a suffix separated by a slash “/”. Each prefix is unique to an organization or individual and can be derived in any way the organization wishes. The suffix may be any bit string. The ability to manage handle records with these prefixes resides in the particular LHS. If the handle record is not in the GHR, the prefix will resolve to the location of the LHS. Like the GHR, any given LHS can also consist of multiple servers, which adds to the scalability of the overall system.

Metadata Registries are used for discovery of DOs and for certain types of access control to DOs. Each country in a PP experiment would manage a PP registry used to authenticate and control the access to PP DOs. Authentication of DOs may be done by use of separate authentication servers. However, a self-authenticating design for verification exists and can be implemented if the identifier suffix provides a cryptographically generated “fingerprint” of the DO in question. The style used in creating metadata registries may vary according to subject matter, organization or use, to name a few possibilities. Access to and use of such registries will be subject to group and organizational control for the PP initiative.

The Digital Object Architecture and processing components provide a provable and scalable approach to PassPort management. This, in turn, will provide an important first step to on-line identity management. In the passport processing context, this approach will improve the security of current and future passport processing (using associated biometric data), while concurrently improving privacy for individuals. This ability to address both security and privacy within an identity management framework is what makes this approach so powerful.