

DEMONSTRATION FILE FOR ADOBE ACROBAT HANDLE PLUG-IN
(Full document available at <http://hdl.handle.net/4263537/5022>)

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.