ITU - Telecommunication Standardization Sector
Focus Group on Identity Management: Geneva 13-16 February 2007

Source: Tertius Ltd: Dr. Norman Paskin
(Handle System Advisory Committee)

Title: The Handle System®

Purpose: For information

Summary

This document accompanies a presentation of the same title to be made at the FG IdM meeting in Geneva 13-16 February 2007.

The Handle System is a general-purpose distributed information system used to assign, manage, and resolve persistent identifiers, known as "handles", for digital objects and other resources on the Internet. Some applications of this are in content, and others in a variety of identity management applications. The Corporation for National Research Initiatives manages this through its Handle System Advisory Committee (composed of external interested parties).

References

Key reference: The Handle System web site: http://www.handle.net

Supplementary references:

- "A framework for distributed digital object services": Robert Kahn & Robert Wilensky http://www.cnri.reston.va.us/k-w.html
- "Naming And Meaning: Key To The Management Of Intellectual Property In Digital Media": N. Paskin *(applications in content management)* http://www.doi.org/topics/060922IPDM_China_Paskin_preprint.pdf

The Handle System®

The Handle System is highly relevant to the ITU FG scope of "management of...attributes of an entity". It is a non-commercial, openly available protocol and reference implementation of a general-purpose distributed information system used to assign, manage, and resolve persistent identifiers, developed at the Corporation for National Research Initiatives (US) by Robert Kahn, one of the co-inventors of TCP/IP and a pioneer of internet technologies. The Handle System can utilise existing or new numbering schemes and protocols, adding value to them.

Digital information needs to be a first class citizen (one that has an identity independent of any other item) in the networked environment. The original Internet design conflated addresses to serve two purposes: an indication of the location of the end point, and an indication of its identity, now recognised as a limitation: see e.g.

- Future generation Internet architecture http://www.isi.edu/newarch/
- Future internet network design http://find.isi.edu/

The fundamental characteristic of digital information is that it is processable data, enabling re-use and hence new forms of electronic commerce, creativity and social benefit. Managing these units of digital information, the "citizens" in the network, requires that they have unique names (or "identifiers") denoting a specific referent, and the ability to manage their attributes. The objects ("citizens") may be representations of content, people, parties, resources, licences, avatars, sensors, etc.

The Handle system is a basic Internet resolution system that identifies digital objects, not servers:

- Optimized for speed, reliability, scaling; logically centralized, physically distributed
- Open defined protocol and data model (IETF RFC 3650,1,2)
- Free protocol; service at low cost (non-profit);
- Freely available to be used as engine underneath other named identifiers.
- Separates control of the handle and who runs the servers
- Distributed administration, granularity at the handle level
- Supports any Unicode character set and hence internationalisation
- All transactions can be secure and certified, both registration and resolution
- Not all data is public: individual values within a handle can be private.
- Carries no semantics in the identifier
- Does not need DNS, but can work with DNS: may be deployed via tools e.g http proxies, client plug-ins, server software, etc

A Handle consists of a prefix and suffix, e.g. 123/4567. The prefix and suffix may be any length. The suffix may incorporate another identifier numbering scheme. Shorter prefixes (1-3 digits) are reserved for major projects, countries, etc.

A handle resolves to a a set of values. Each value's <type> field defines the syntax and semantics of a value's data (e.g. URL = resolving to current web location). A pre-defined set of handle data types are defined for administrative use; extensible registered or non-registered handle data types are for non-administrative use. Types may include public or secret encryption keys, descriptions, etc.

- See http://www.handle.net/overviews/types.html

The Handle System provides infrastructure for many application domains, e.g. digital libraries & publishing, network management, identity management (http://www.handle.net/apps.html). Current applications include:

- *International DOI Foundation*: a federation of several independent applications with an added layer of social infrastructure (and specific rules) including CrossRef (scholarly journal consortium); Office of Publications of the European Community (EC documents); MEDRA (Multilingual European DOI Registration Agency); Nielsen BookData, R.R. Bowker, et al (bibliographic data - ISBN); German Nat. Lib. Science and Technology (science data)
- *Defense Virtual Information Architecture* (Defense Technical Information Center DARPA, CNRI: context sensitive distribution of data and metadata)
- *GRID computing* (Handle System - Globus Toolkit Integration Project)
- *DSpace* - MIT Libraries/Hewlett-Packard (Digital Repository System)
- *National Digital Library Program* (NDLP)
- *Los Alamos National Laboratory*
- Several Digital Library projects

There are currently over 3500 assigned namespaces to users, and an estimated several hundred million individual "Handles" (identifiers within each namespace); the total per namespace is known only to each namespace manager.

Future projected applications of particular relevance to identity management include:

- *Transient Network Architecture* (Pervasive transient mobile network in which all communications occur between persistently identified entities. Under NSF's FIND (Future Internet Network Design)
- *Using PKI for persistent trustworthy identity*.
- *Representing Value as Digital Objects* ("Transferable records" structured as digital objects; transferability and anonymity as attributes of digital representations of deeds of trust, mortgages, bills of lading, digital cash etc.)
- *Application of Handles to licences and parties*

Security is a major feature of the Global Handle Registry service:

- Protected service information and public key pair used to sign global service information.
- Handle protocol allows handle servers to authenticate their clients and to provide data integrity service on client request.

- Handle servers can be set to explicitly asked to generate or return a digital signature for every service response
- Public key and/or secret key cryptography may be used.
- Server authentication may be used to prevent eavesdroppers from forging client requests or tampering with server responses.
- Client applications can (if wished) only accept information from the authoritative Global Handle Registry (not any mirrors) and check its integrity on each update.

Compared to the Domain Name System, there are similarities and differences in both the design and intended use (see http://www.handle.net/overviews/dns.html). These include:

- Naming. DNS naming hierarchy reflects a control hierarchy, Handle system need not. Handle separates control of handle (id) from control of server (location)
- Distributed Administration. Handle administrators can add/delete identifier and identifier values securely over the public Internet.
- Proxies. Technical advantages regarding resolution work behind SOCKS or HTTP proxies, both supported in Handle client library.
- Unicode. Handle has full native Unicode support. There are hacks to make DNS support 8-bit character sets, but they are not widely implemented.
- Replication. In DNS, if a single record is updated all records must be copied to mirror servers. The Handle System has finer granularity: if a single record is updated, the server will copy only that record to the mirror servers.
- Certification. DNS has to be fast, especially at the root, hence not very good for alternative uses, e.g. certificates. Handle System has more flexible and robust certification support.
- Access Control. Handle System has support for access control and authentication; DNS does not.
- Record Size. Technical advantages regarding UDP and TCP handling: more efficient request handling; much larger storage in a record (DNS 64KB, Handle System 4GB).
- Handles avoid broken URLs when control changes
- Ownership: In DNS, the system administrator is considered the owner of the data, in the Handle System the prefix administrator is considered the owner: each Handle identifier and prefix can have its own set of administrators independent from the system administrator.

The Handle System provides a highly effective general-purpose distributed information system to assign, manage, and resolve persistent identifiers for digital objects and other resources on the Internet. Notable key functionalities widely considered to be essential for identity management include:

- Functional Granularity: “it should be possible to identify an entity whenever it needs to be distinguished”
- First class naming: “Digital objects should have first class names”