A modest proposal for e-health identifiers based on DNS.

My colleague Phil Johnson and I have provisional patents filed for an identifier system known as ***DIAD** (StarDIAD). Here's an example of the way that *part* of the *DIAD naming tree might look to implement individual identifiers (iDIADs).

Root domain-name for e-health **`oz-ehealth.org.`**

This domain-name is controlled by the Australian e-health naming authority. The authority only authorises the names of top-level e-health domains, which are delegated from the DNS server of this domain.

This domain-name is to some extent arbitrary. However, the optimal domain root exists under a DNSSEC signed domain. In practice, that is .ORG. The public key for the root servers is to distributed by 1st July 2010. .ORG is now signed, and owners of .ORG domains can now participate in operational testing of the procedures.

General Practice Queensland **`gpq.`**`oz-ehealth.org.`

A top level naming authority, authorised by the root level naming authority. This authority only authorises the names of practices within its ambit. Those names will occur within the domain's DNS server.

Sunshine Coast Division of General Practice **`scdgp.`**`gpq.oz-ehealth.org.`

authorised by General Practice Queensland. This naming authority only authorises the names of practices within the division. Those names occur within the domain's DNS server.

The fictional dh-clinic **`dh-clinic.`**`scdgp.oz-ehealth.org.`

The clinic is on the Sunshine Coast. This naming authority only authorises patients of the clinic. Those names occur within the domain's DNS server.

A patient **`thegecko.`**`dh-clinic.scdgp.oz-ehealth.org.`

The patient accepts an optional identifier from the clinic. This is a terminal in the tree, and is not a naming authority. That is, there is no DNS server for this domain. Note the use of a *handle*, as is common on social networking sites, where people have learned to protect their identities.

It's a simple enough structure, but it has some notable differences from a normal DNS structure. For a start, all of these names are just that - *names*. None of them denote any device that is connected to the Internet.

None of them have an associated IP address.

What do they have? Naming authorities will have NS records with the names of their name servers. Others have, at the least, a TXT record with an empty string.

Opaque Identifiers

At each level then, there is a set of names which is available, by default, only to the controlling name servers. Discovering the set of names without the co-operation of the naming authority is akin to cracking passwords by brute force methods over the Net.

Each name uniquely denotes an individual entity, if the naming authority does its job correctly. There is no restriction on the number of names that an entity may possess. Each naming authority creates names unique to the authority, which guarantees global uniqueness for all names.

Resolving Identities

The association between name and identity is maintained by the naming authority. The identifier can be passed around safely, and resolution of the identity requires some interaction with the naming authority, for example, dh-clinic.

Another way of resolving identity is to put a sticky label with the identifier on a driver's licence or Medicare card or credit card. The presence of this sticker indicates that the person referred to on the card is also the owner of the identifier. No electronic connection here. But if an unconscious person is attended by para-medics, that association will be available.

Knowledge of the identifier set is restricted. Knowledge of the identity associated with the identifier is restricted. Yet information associated with the identifier is public.

Payloads

Let's say we associate a TXT record of the following form with the identifier.

```
TXT "allergies=penicillin,latex medication=warfarin:2mg"
```

Now anyone with an internet connection can issue a command like

```
    $ dig thegecko.dh-clinic.scdgp.oz-ehealth.org txt
```

to read the content of that TXT record. Anyone. Anywhere.

The identifier of the issuing clinic is part of the patient identifier, giving access to additional information about the patient.

The identifiers are also a gateway to external systems. NAPTR records provide a means of constructing a variety of strings from a given input. The input might be, for example, be the patient identifier. An enquiry against the issuing clinic's identifier returns a series of recipes for manitpulating the patient identifier in order to determine more information about the patient. At the simplest level, that might be the telephone numbers of the clinic and treating doctor. At the other extreme, a web service might be defined by which qualifier enquirers could determine demographic and contact details about the patient, or extract a summary medical history.

Note that these are externally defined and managed systems. All the identifier can provide is an address to attempt to access them.