NICC ND 1512 V1.1.1 (2010-01)

NICC Document

Report into the implications of usage of alphanumeric (i.e. non-E.164-based) naming in NGNs

NICC Standards Limited

Michael Faraday House,
Six Dials Way,

Stevenage
SG1 2AY

Tel.: +44(0) 20 7036 3636

Registered in England and Wales under number 6613589

If you have any comments concerning the accuracy of the contents of this document, please write to:
The Technical Secretary, NICC Standards Ltd.,
Michael Faraday House,
Six Dials Way,
Stevenage
SG1 2AY

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to NICC.

Pursuant to the NICC IPR Policy, no investigation, including IPR searches, has been carried out by NICC. No guarantee can be given as to the existence of other IPRs which are, or may be, or may become, essential to the present document.

Foreword

This NICC Document (ND) has been produced by NICC TSG NNA.

0. Executive Summary

This report examines the issues associated with the support of alphanumeric names, and explores potential architectures. The format of names to be supported are set out in Section 4.2, and generically are of the form <local_part>@<domain>, where <domain> could be owned by a Communications Provider but equally could be a customer-owned domain.

It is concluded that the alphanumeric name as used by the calling customer should be converted to a destination group for the purposes of network routeing, which identifies the terminating network and potentially the location within it of the called customer. This is accomplished in two stages;

- The Name Resolution Locator Function identifies the Name Resolution Function(s) which contain information for a given alphanumeric domain

- The Name Resolution Function then provides information about the mapping from the full alphanumeric name to destination group.

The report examines various implementations of the overall architecture. It concludes that the Name Resolution Locator Function should be provided within public DNS, and allow for multiple Name Resolution Function federations to be identified. Access to Name Resolution Functions should be restricted to CPs.

Techniques are considered for sharing information within federations of Name Resolution Functions. The report doesn’t draw conclusions about the best internal architecture for Name Resolution Function federations, but presents a series of implementation options.

A number of issues related to maintaining end-to-end service are identified for further consideration.

The report concludes that to move towards implementation, standardisation at an international level will be necessary.

1 Scope

The present document examines the implication of usage of naming schemes in Next Generation Networks which are not based upon E.164 numbering. For the purposes of this report, these are termed “alphanumeric names”.

2 References

For the particular version of a document applicable to this release see ND1610 [1].

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For non-specific references, the latest edition of the referenced document (including any amendments) applies.

[1] ND1610 Multi-service Interconnect of UK Next Generation Networks

[2] ITU E.164 Public Telecommunication Numbering Plan

[3] ND1631 NGN; PSTN/ISDN Service Interconnect; Architecture for usage of Common Numbering Database

[4] ND1415 NGN; PSTN/ISDN Service Interconnect; Guide to Common Numbering Database standards

[5] IETF RFC5322 Internet Message Format

[6] IETF STD013 Domain Names – Concepts and Facilities

3 Definitions, symbols and abbreviations

3.1 Definitions

For the purposes of the present document, the following termapplies:

Alphanumeric name : A name that is not taken from the E.164 numbering plan.

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AXFR	Authoritative Transfer
ANRF	Authoritative Name Resolution Function
CDB	Common numbering DataBase
CLI	Calling Line Identity
CNAME	Canonical Name
CP	Communications Provider
DNS	Domain Name System
IANA	Internet Assigned Numbers Authority
IDN	Internationalised Domain Name
IETF	Internet Engineering Task Force
INRF	Initiating Name Resolution Function
IP	Internet Protocol
IXFR	Incremental Transfer
LDAP	Lightweight Directory Access Protocol
NAPTR	Name Authority PoinTer Record
NGN	Next Generation Network
NRF	Name Resolution Funtion
NRLF	Name Resolution Locator Function
RR	Resource Record
SIP	Session Initiation Protocol
URI	Uniform Resource Identifier

4 Requirements

4.1 Overview

Within individual NGNs, it is common for callers to be able to indicate the destination of calls using not just telephone numbers, but names as well. In particular, it is frequently possible to use e-mail style addresses to indicate the desired destination of calls. In comparison, for calls between NGNs, the only agreed addressing format is to use E.164 telephone numbers [2]. This means that there is an asymmetry between the name formats that can be used “on” and “off” net : non-numeric naming is only possible for “on-net” calls.

In the public internet, calling using e-mail style addresses is possible, because usage of public DNS facilitates resolving these to IP addresses for routeing purposes. However because NGNs are characterised by a walled-garden approach to maintain network integrity and quality of service, typically the internal IP addressing is private, making such an approach difficult.

The purpose of this study is to examine what would be necessary to make usage of alphanumeric naming possible between NGNs, rather than just within them.

4.2 Name formats to be supported

In examining the implications for NGN interconnect, it has been assumed that the following name formats will be supported;

1. jo.bloggs@commsprovider.com

   In this format the user identity “jo.bloggs” is assigned by the communications provider to the end user.

2. jo@bloggs.commsprovider.com

   In this format the communications provider assigns the identity group “bloggs” to the end user, who is then free to assign individual user identities such as “jo” within their group.

3. jo.bloggs@namespaceprovider.com

   In this format an independent name provider assigns the user identity “jo.bloggs” to the end user from their higher-level domain, who is then at liberty to make their own selection of NGN communications provider.

4. jo@bloggs.namespaceprovider.com

   In this format an independent name provider assigns an identity group “bloggs” to the end user from their higher-level domain, who is then free to assign individual end user identities such as “jo” within their group, and to make their own selection of NGN communications provider for each.

5. jo@bloggs.com

   In this format the end user secures their own domain, assigns individual end user identities such as “jo” and makes their own selection of NGN communications provider for each.

Additionally, IDN formatted derivatives would need to be supported, for example to encode the full Welsh language set.

The general form of all of the above is <local_part>@<domain>. It is required that any domain name used be properly registered in the public DNS under the IANA maintained root and is registered to the same entity that is using it for NGN addressing. Where third parties provision names on behalf of the registrant, procedures will need to be in place to validate that they have permission.

4.3 Service requirements

NICC has co-ordinated with NGNuk to determine likely service requirements.

It is a service requirement that any names be reachable globally. For the enterprise market, it is expected that multinational companies will require different <local_parts> within the same domain to terminate in different countries, probably utilising different communications providers (both within a country and on a CP-per-country basis). Further, it is possible that enterprises may wish individual services utilising the same namespace to use different providers.

It is highly desireable that CPs be able to download the contents of any resolution database into their own network, so that they are not dependent upon real-time queries to third parties to determine how to route calls, indeed one CP has stated that for service reasons they require this. The technical sections of this report will address the extent to which it is possible to achieve this. Further discussion in the appropriate commercial body will then be required..

Assuming that some form of resolution database is required, a commercial framework will need to be established for who has access to it, both to populate it and read the contents. From a read perspective, due to data sensitivity (both commercial and data privacy), is is a requirement that only CPs be able to access the mapping of name to CP, and then only for the purposes of network operation. It is assumed that population of the database will be possible by both CPs and accredited agencies. It is not intended that individual end-users would generally populate their own data, but this does not preclude e.g. enterprises becoming suitably accredited in order to manage their own namespace.

As far as is technically feasible at reasonable cost, it is a commercial requirement that portability of names between CPs is not precluded. However, it must be stressed that this should not be taken as asserting a requirement that names will be portable : this is a commercial and regulatory matter which is beyond the scope of this report.

It will be necessary to understand how the usage of alphanumeric names would interact with provision of end-to-end services : this subject is discussed in Section 8.

5 Architecture for support of alphanumeric names

5.1 Basis of routeing

Names of the form described in Section 4 are not directly suitable for routeing of communications. When used in the public internet, they are generally translated into IP addresses (frequently via some form of intermediate record). In NGNs, although IP addresses will ultimately be used for routeing of messages, some form of abstraction is essential because the IP addresses used within an individual NGN are not typically published/routeable from other networks.

For call routeing using E.164 numbers, NICC has concluded that the best approach for the future is the usage of Destination Groups, as described in ND1631 [3] and ND1415 [4]; in essence an E.164 number is translated to a Destinaton Group which is then used to route the call. Foralphanumeric names, NICC considered two options;

- Convert the alphanumeric name to its equivalent in the E.164 space, which would then use the ND1631 call routeing constructs to convert to a suitable Destination Group, or
- Convert the alphanumeric name directly to a Destination Group.

These approaches are depicted in Figure 5.1.a.

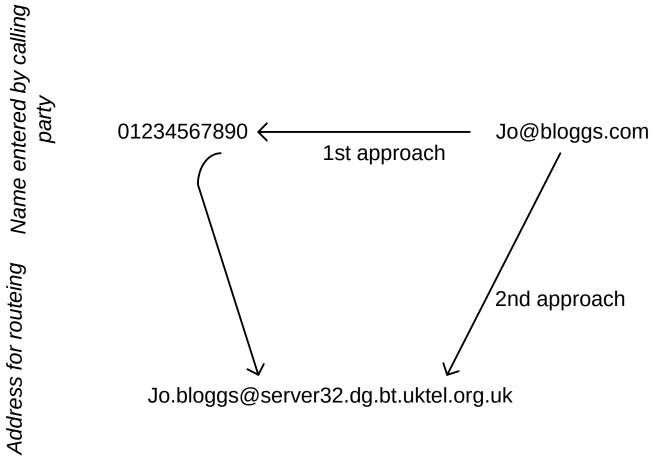

Figure 5.1.a : Alternate approaches to routeing on names

The first approach adopts the paradigm that names should interwork to names rather than addresses (as an analogy, in ENUM telephone numbers are mapped to NAPTR records rather than direct to IP addresses). It also has the advantage that portability of alphanumeric names between CPs can be accommodated at the E.164 layer. Set against this, arguably E.164 numbers are a backward looking resource, whereas NGN call routeing should be forward looking. The approach requires an equivalent E.164 number for every name, which could in the long term be incredibly resource inefficient. For these reasons, NICC has concluded that a direct mapping of non-numeric names to Destination Groups, i.e. the second approach, is the best solution.

5.2 Generic Architecture

NICC has established that the most appropriate approach would be to have one or more name-resolution databases, which map the name formats described in Section 4.2 into a form that could more readily be used for routeing in NGNs. As described in Section 5.1, this would draw upon the Destination Group concept as described in ND1631 [3] and ND1415 [4]. An overview architecture is shown in Figure 5.2.a.

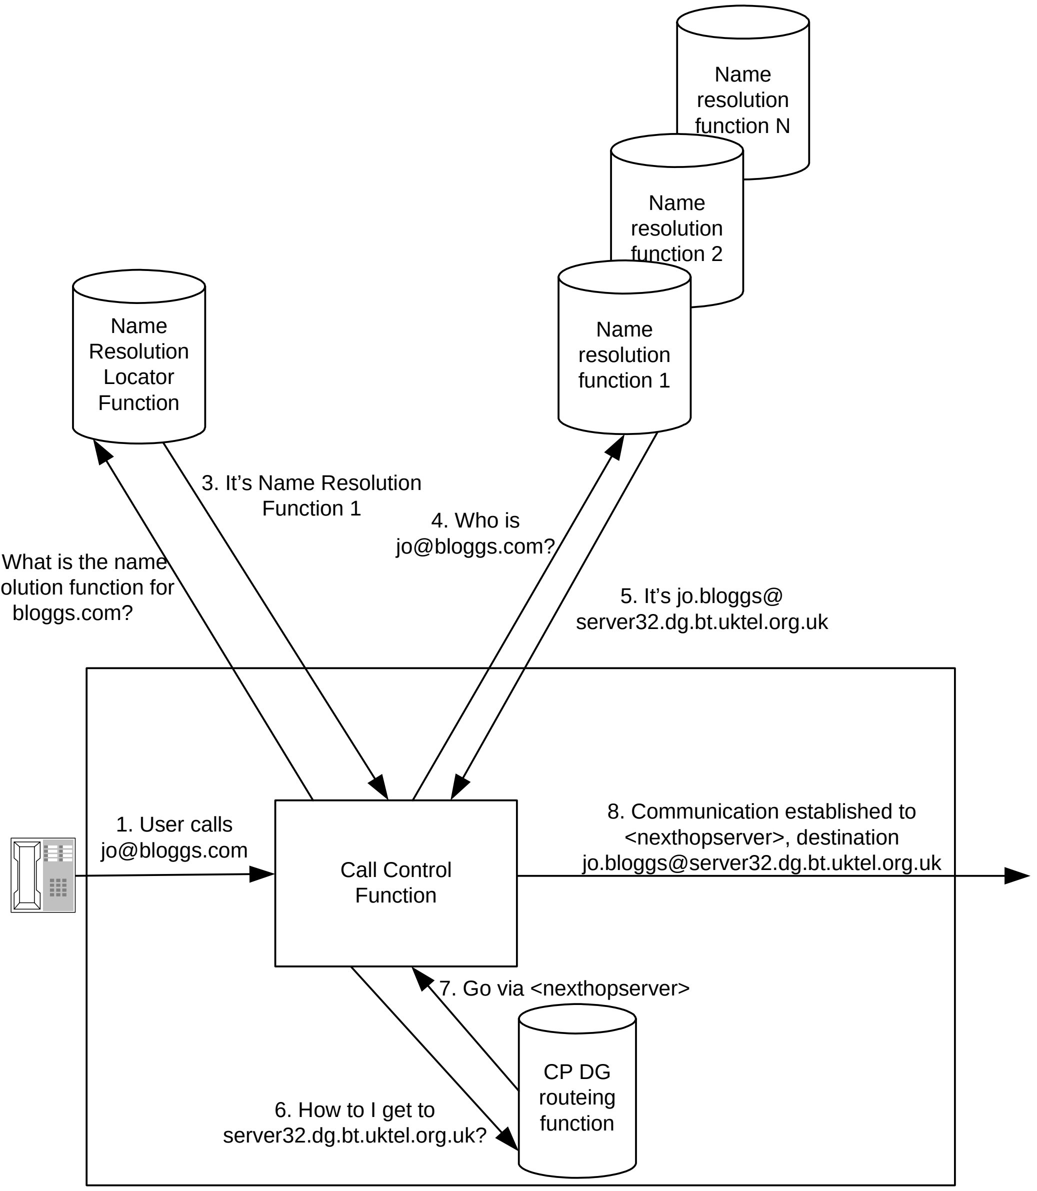
Name resolution function N
Name resolution function 2
Name Resolution Locator Function
Name resolution function 1
3. It's Name Resolution Function 1
4. Who is jo@bloggs.com?
2. What is the name resolution function for bloggs.com?
5. It's jo.bloggs@ server32.dg.bt.uktel.org.uk
1. User calls jo@bloggs.com
Call Control Function
8. Communication established to <nexthopserver>, destination jo.bloggs@server32.dg.bt.uktel.org.uk
7. Go via <nexthopserver>
CP DG routeing function
6. How to I get to server32.dg.bt.uktel.org.uk?

Figure 5.2.a : Generic Architecture

Core to the architecture is the concept of the Name Resolution Function. This would hold the identification of the termination point for each name. The CP would then use this information to consult their local data to determine the best way to route calls to that location. In this generic architecture, there can be competition in the Name Resolution Function market, from the standpoint that domain holders could choose which Name Resolution Function would be used to host the information about their domain. In order for originating CPs to know which Name Resolution Function to consult, they would first query a Name Resolution Locator Function that would provide this information : in practise this could be a simple DNS query.

The Call Control Function is shown as a single entity in the diagrams in this section. It is not intended that it will be required to implement all of the look-ups in a single instance of hardware or software. The Call Control Function may be distributed in any manner which aligns with implementer strategies, whilst satisfying the overall functional goals described in this report.

When the call arrives at the chosen terminating network, it may carry out a further lookup. This may be particularly beneficial if different levels of information are available to the terminating network’s resolution function. For example an originating network resolution function may know that a domain is served by a particular terminating network, but only the terminating network may be aware of the location of the names within it. It should be noted, however, that this would make far-end-handover impossible and any implementation of name portability would necessarily have to be on an onward routeing basis.

Although Figure 5.2.a provides a generic architecture, there are certain specific approaches that could be taken, as described in the following sections.

5.3 Implementations of the generic architecture

5.3.1 Architecture 1 : monopoly Name Resolution Function

In this architecture, a single Name Resolution Function would be selected for a set of originating networks, in practise probably for the UK. This means that the Name Resolution Locator Function could probably be dispensed with, as depicted in Figure 5.3.1.a.

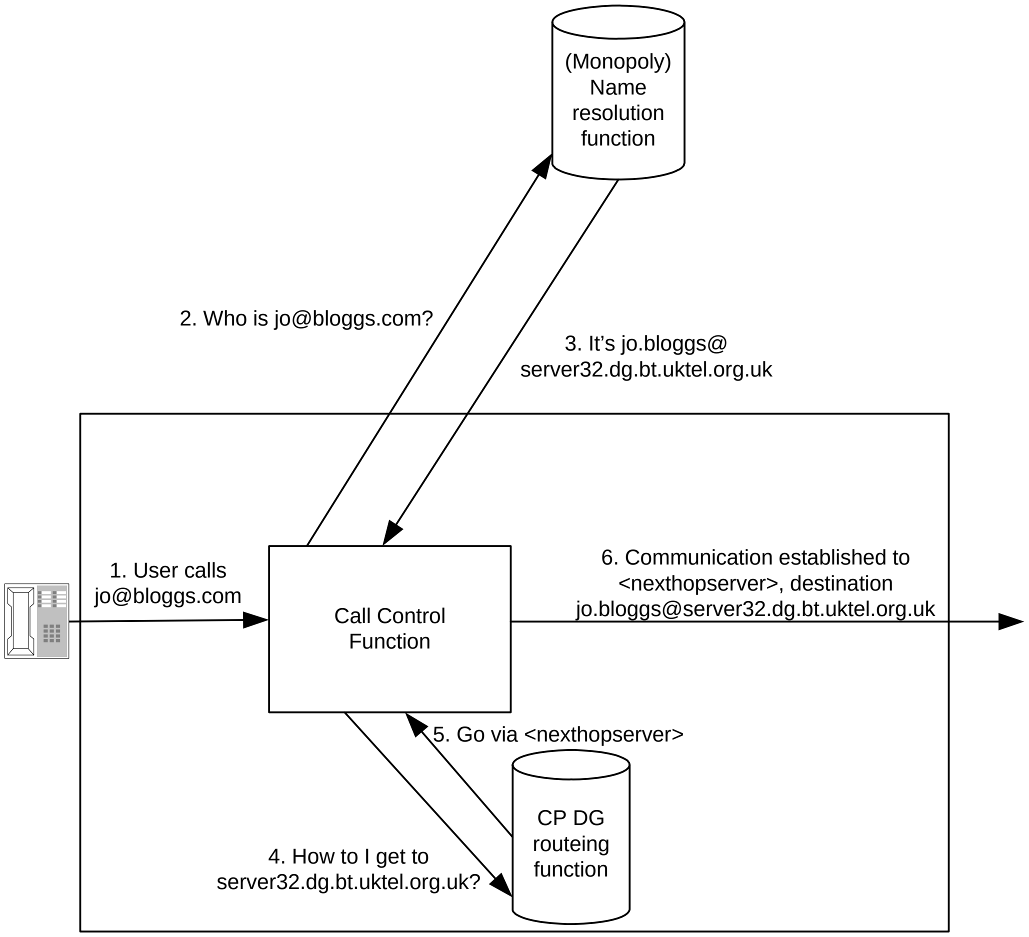
(Monopoly) Name resolution function
2. Who is jo@bloggs.com?
3. It's jo.bloggs@ server32.dg.bt.uktel.org.uk
1. User calls jo@bloggs.com
Call Control Function
6. Communication established to <nexthopserver>, destination jo.bloggs@server32.dg.bt.uktel.org.uk
5. Go via <nexthopserver>
CP DG routeing function
4. How to I get to server32.dg.bt.uktel.org.uk?

Figure 5.3.1.a : Architecture 1, monopoly Name Resolution Function

The approach has the advantage that an originating CP need only have a relationship with a single Name Resolution Function provider rather than potentially with multiple ones. With a single Name Resolutuion Function, establishing standards for its design and operation would be a simpler proposition. Set against this, inherently the Name Resolution Function provider is a monopoly, which could be considered undesirable.

Support of enterprise domains which span terminations in multiple countries would be particularly problematic. For example globalmonolith.com could have employees in multiple countries served by multiple CPs, wishing to have calls routed to all of them. Since they may not even be based in the UK, a UK-centric solution is inherently limited. There are potential solutions to this, for example that globalmonolith.com would place a contract with a registrar that would replicate their name-resolution data in each national Name Resolution Function, or that national Name Resolution Functions would federate data in a similar manner to that described in Section 5.3.3.

5.3.2 Architecture 2 : public DNS as the Name Resolution Function

This architecture represents a particular implementation of the first architecture, where the single Name Resolution Function is agreed to be the public DNS. This means that the Name Resolution Locator Function would be dispensed with, as depicted in Figure 5.3.2.a.

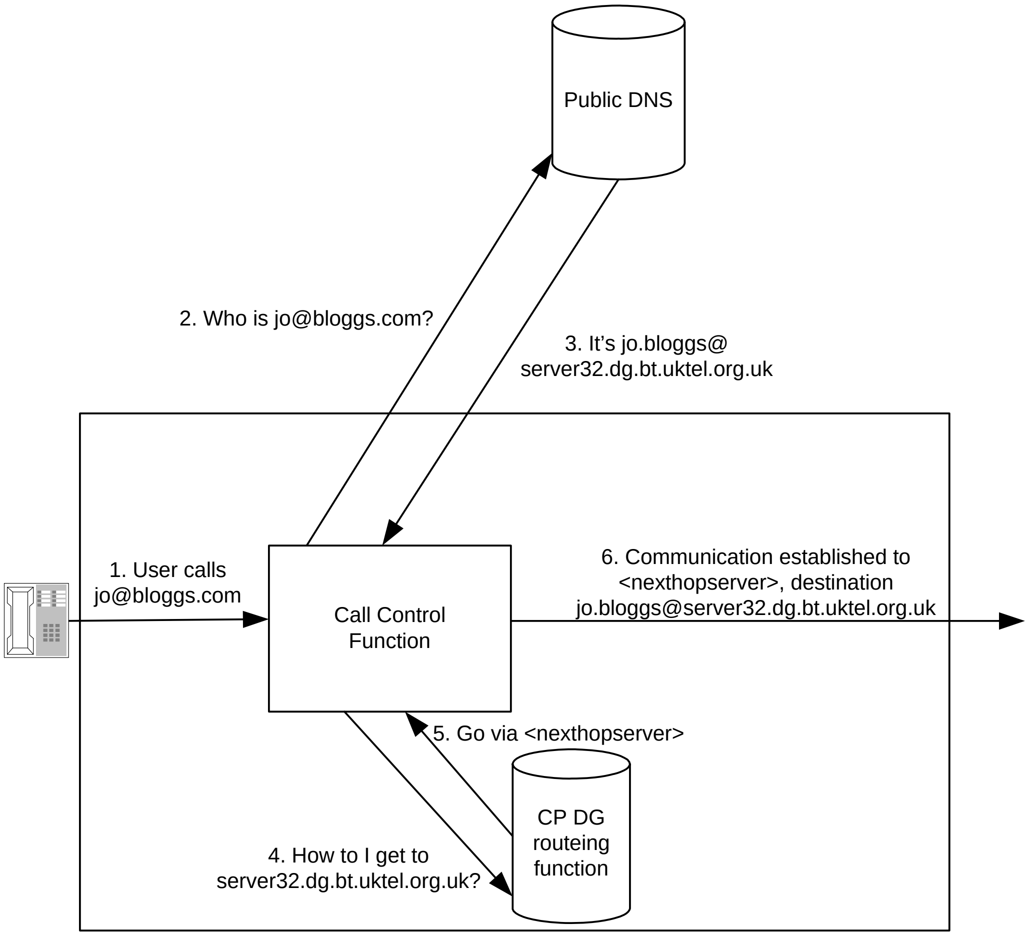
Public DNS
2. Who is jo@bloggs.com?
3. It's jo.bloggs@
server32.dg.bt.uktel.org.uk
1. User calls
jo@bloggs.com
Call Control
Function
6. Communication established to
<nexthopserver>, destination
jo.bloggs@server32.dg.bt.uktel.org.uk
5. Go via <nexthopserver>
CP DG
routeing
function
4. How to I get to
server32.dg.bt.uktel.org.uk?

Figure 5.3.2.a : Architecture 2, public DNS as Name Resolution Function

The approach has the advantage that no special infrastructure is required for the Name Resolution Function. Set against this, it would mean the the data would be publically available (see Section 7.1.2) and that performance would be dependent upon that of the public DNS.

Support of enterprise domains which span terminations in multiple countries would be readily possible.

5.3.3 Architecture 3 : federated Name Resolution Functions

In this architecture, an originating network would have a chosen Name Resolution Function provider. Since this provider might not have been populated with the relevant name resolution information for all domains directly by each domain holder (or their registrar), arrangements would be in place between Name Resolution Function providers to act as a federation to share data. Under this architecture it is possible that the Name Resolution Locator Function could be dispensed with, as depicted in Figure 5.3.3.a. However, it could also be possible to have multiple federations of Name Resolution Functions, with a given domain served in one or more of these federations (i.e. if the domain is in more than one federation, the individual name data would be replicated in each of those federations). In that scenario, shown in Figure 5.3.3.b, the Name Resolution Locator Function would be present and would return details of multiple federations that contain the resolution data.

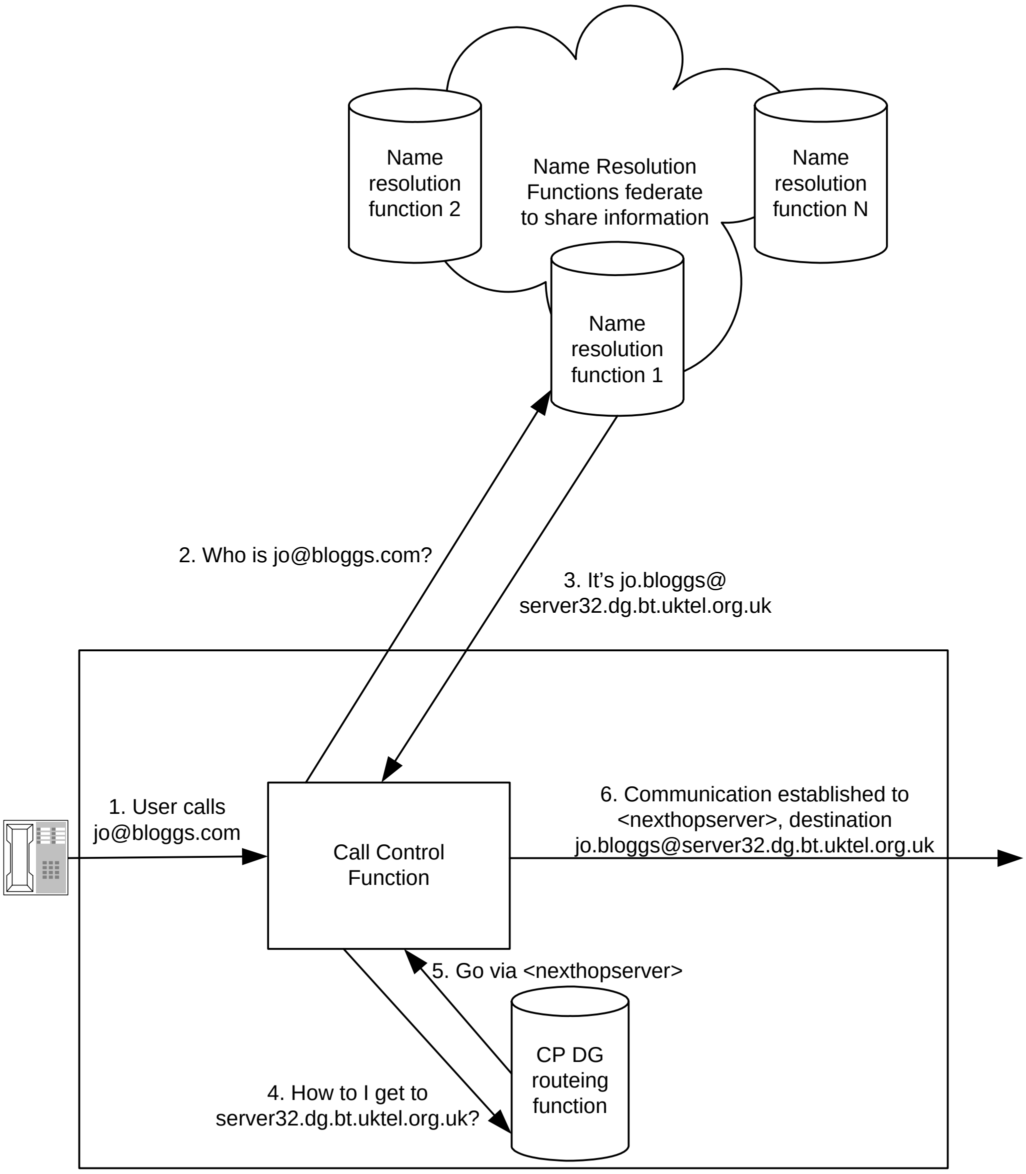
Name resolution function 2
Name Resolution Functions federate to share information
Name resolution function N
Name resolution function 1
2. Who is jo@bloggs.com?
3. It's jo.bloggs@ server32.dg.bt.uktel.org.uk
1. User calls jo@bloggs.com
Call Control Function
6. Communication established to <nexthopserver>, destination jo.bloggs@server32.dg.bt.uktel.org.uk
5. Go via <nexthopserver>
CP DG routeing function
4. How to I get to server32.dg.bt.uktel.org.uk?

Figure 5.3.3.a : Architecture 3a, federated Name Resolution Function with no Name Resolution Location Function

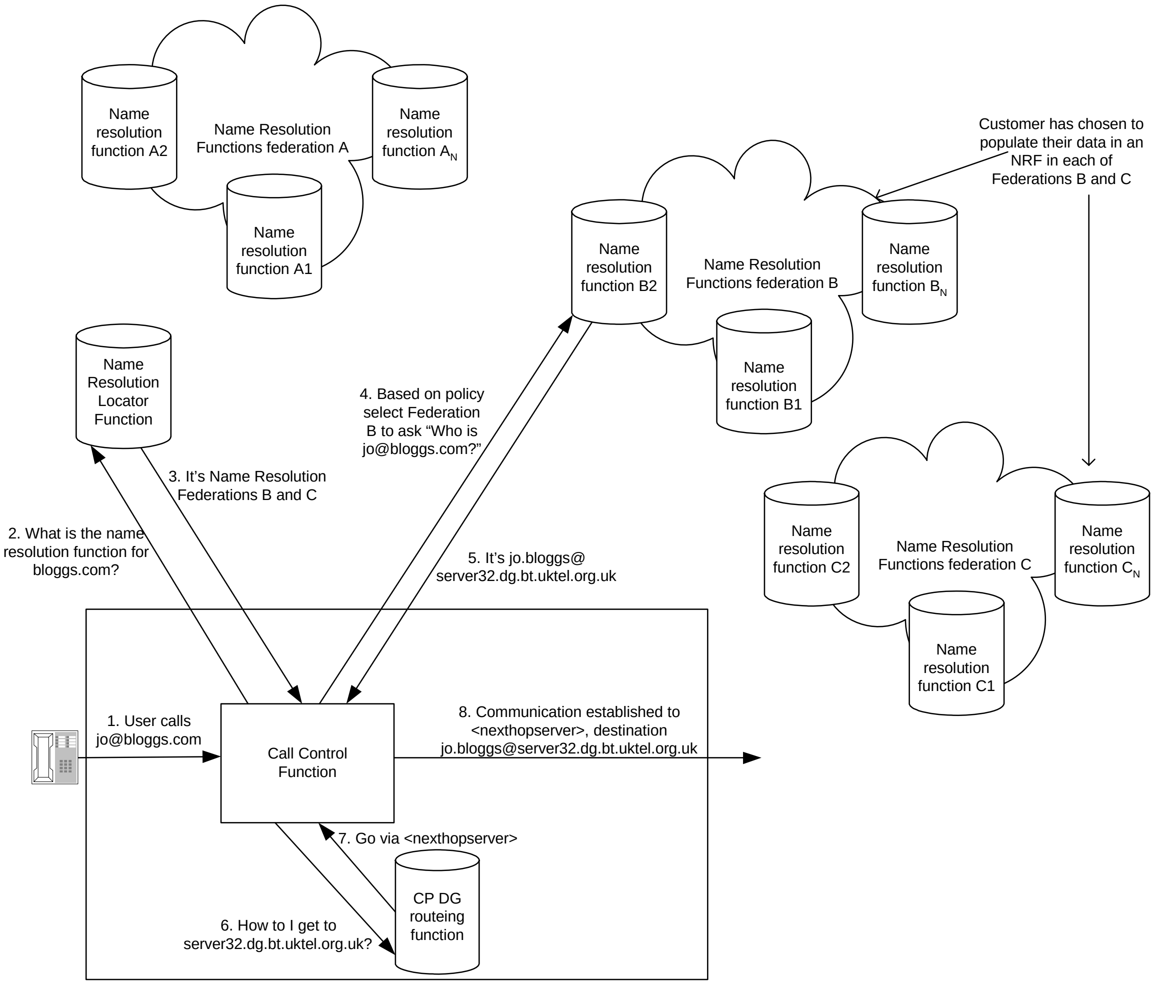
Name resolution function A2
Name Resolution Functions federation A
Name resolution function A$_N$
Name resolution function A1
Customer has chosen to populate their data in an NRF in each of Federations B and C
Name resolution function B2
Name Resolution Functions federation B
Name resolution function B$_N$
Name resolution function B1
Name Resolution Locator Function
4. Based on policy select Federation B to ask "Who is jo@bloggs.com?"
3. It's Name Resolution Federations B and C
2. What is the name resolution function for bloggs.com?
5. It's jo.bloggs@ server32.dg.bt.uktel.org.uk
Name resolution function C2
Name Resolution Functions federation C
Name resolution function C$_N$
Name resolution function C1
1. User calls jo@bloggs.com
Call Control Function
8. Communication established to <nexthopserver>, destination jo.bloggs@server32.dg.bt.uktel.org.uk
7. Go via <nexthopserver>
CP DG routeing function
6. How to I get to server32.dg.bt.uktel.org.uk?

Figure 5.3.3.b : Architecture 3b, federated Name Resolution Function with multiple federations

The mechanism to federate the information could take many forms. For example, individual Name Resolution Function providers could propagate changes to their data to other Resolution providers in the federation. Alternatively, a given Name Resolution Function, on receiving a query to a domain for which they had no information, could propagate this query through to a third party Name Resolution Function. Section 7.3 elaborates on this aspect.

The approach has the advantage that an originating CP could only need to have a relationship with a single Name Resolution Function provider rather than potentially with multiple ones. Set against this, arrangements would need to be put in place for information sharing within the federation.

It should be noted that although the architecture depicts an originating CP having a relationship with a single Name Resolution Function provider, this need not be the case. In practical terms, they may generate multiple simultaneous queries to a series of Name Resolution Function Providers, each a member of a different federation that holds partial information. Clearly, this raises questions of what the CP would do should they receive differing responses from these queries (or indeed no useful response).

Support of enterprise domains which span terminations in multiple countries would be possible, so long as the domain holder (or their registrar) ensured that their information was populated in at least one federation that had a presence in each originating country. For example globalmonolith.com could have employees in multiple countries served by multiple CPs, and wish to have calls routed to them. Globalmonolith.com would place a contract with a registrar that would replicate their name-resolution data in a series of Name Resolution Functions that spanned multiple federations, ensuring that the federations spanned all countries.

5.3.4 Architecture 4 : competitive Name Resolution Functions

This architecture is in essence the generic architecture described in Section 5.2, see Figure 5.3.4.a. The originating CP would have to query a Name Resolution Locator Function in order to determine the correct Name Resolution Function to use. Necessarily this will imply that the CP will have a relationship with all Name Resolution Function providers, potentially globally. Since information privacy requirements may dictate that the information in the Name Resolution Function is not publicly available (see Section 7.2), this could imply an extremely large web of private agreements.

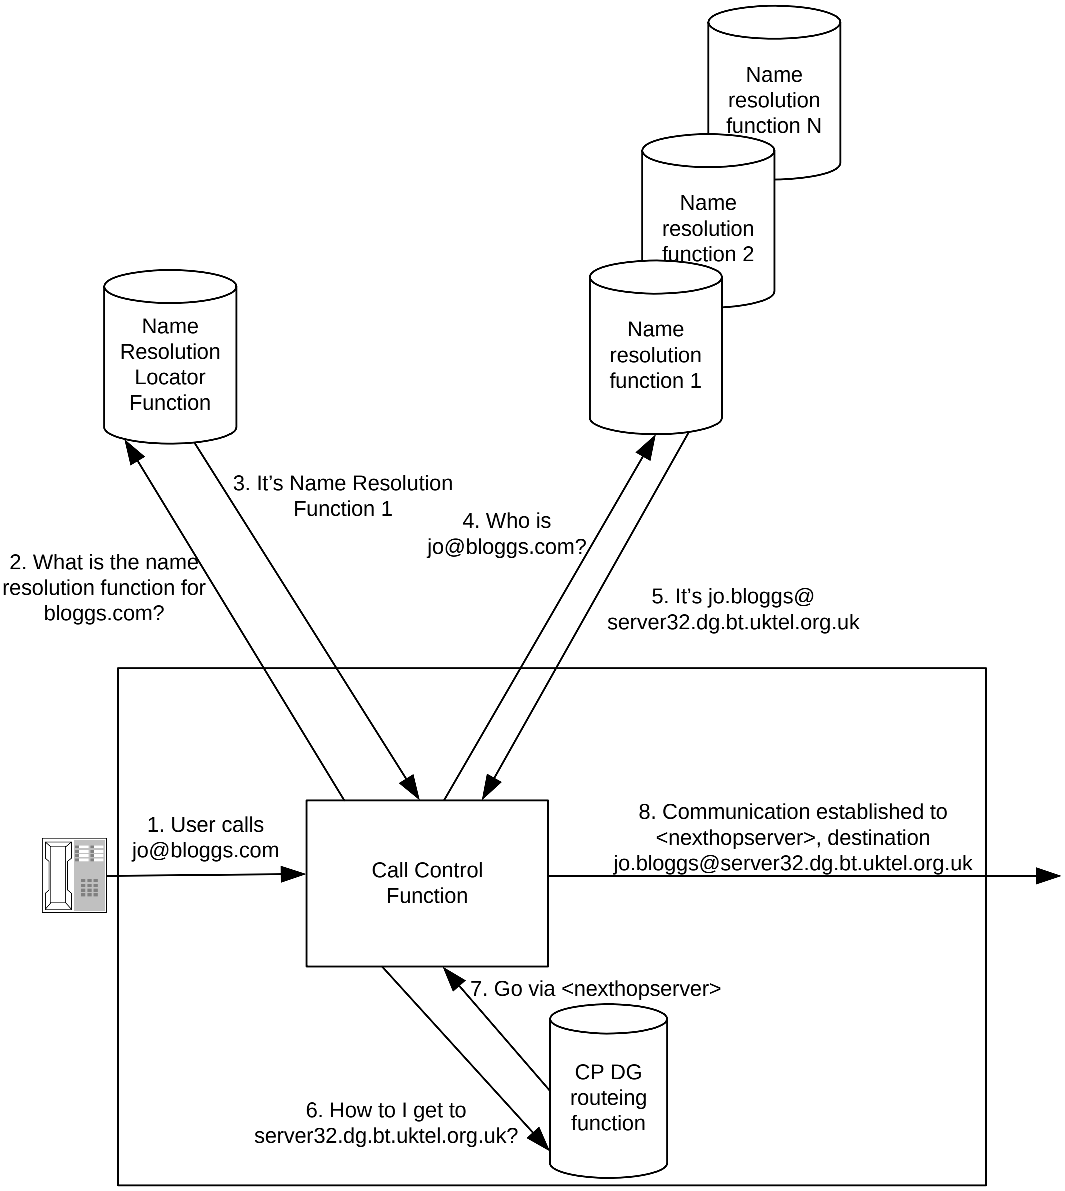
Name resolution function N
Name resolution function 2
Name Resolution Locator Function
Name resolution function 1
3. It's Name Resolution Function 1
4. Who is jo@bloggs.com?
2. What is the name resolution function for bloggs.com?
5. It's jo.bloggs@ server32.dg.bt.uktel.org.uk
1. User calls jo@bloggs.com
Call Control Function
8. Communication established to <nexthopserver>, destination jo.bloggs@server32.dg.bt.uktel.org.uk
7. Go via <nexthopserver>
CP DG routeing function
6. How to I get to server32.dg.bt.uktel.org.uk?

Figure 5.3.4.a : Architecture 4, competitive Name Resolution Functions

6 Name Resolution Locator Function

6.1 Requirements of Name Resolution Locator Function

6.1.1 Inputs and outputs

The Name Resolution Locator Function takes the domain part of the user-friendly name and identifies the Name Resolution Function(s) or federations of Name Resolution Functions that are able to transform the individual names within that namespace to something more useful for routeing.

Input

The input to the Name Resolution Locator Function is the domain name part of the non-numeric name.

Output

The output from the Name Resolution Locator Function is a list of identities of Name Resolution Function**s** or Name Resolution Function federation**s**, probably in the form of a domain name or URI, each tagged with protocols that can be used for access

6.1.2 Access controls : public or private?

NICC has concluded that there is no reason why there should be restrictions on access to this function. This is because;

- The input is a domain, rather than individual name. As such inherently it isn't private data.
- The output only represents an intermediate step in the resolution from called name to serving CP. As such, it reveals nothing of the choice of communications provider by the end user.

6.2 Potential implementations

In principle it would be possible to create an implementation of new technology to provide the Name Resolution Locator Function. However, NICC can see no compelling case to do so and can see merit in the usage of public DNS for the following reasons;

- it automatically enforces the need for domains to be used in NGNs to be properly registered in the IANA structure,
- it puts the control of choice of NRF under the owner of the domain and
- it is readily available without prior negotiation on global basis

There are several alternate methods by which the NRLF could be implemented using DNS:

1. A new Resource Record type at the domain apex defined specifically for this purpose

2. Re-use of an appropriate existing RR type (e.g. NAPTR)

3. Using CNAME or DNAME RRs to “alias” a part of the end user’s name space to a domain controlled by the relevant NRF.

Both approaches (1) and (2) would require standardisation via IETF. Whilst it’s now relatively easy to register a new Resource Record type with IANA, the existing NAPTR record format appears to satisfy all requirements. In particular the “Service Type” field could be used to indicate the protocol supported by the target NRF (i.e. DNS, LDAP, SIP Redirect).

The requirement for a domain to be able to return a list of NRFs would in many circumstances rule out use of CNAME or DNAME since they can only create a one-to-one mapping from source domain to destination domain. However if an enterprise only used the services of one NRF then a CNAME might be appropriate.

Such use would also require a “standard” prefix label (e.g. _nrlf.example.com) so that the use of these domain aliasing records would not interfere with existing Resource Records in the domain (and vice-versa).

7 Name Resolution function

7.1 Requirements of Name Resolution function

7.1.1 Inputs and outputs

Input

The Name Resolution Function transforms a user-friendly name into a format that is suitable for routeing in CP networks. It is proposed that the Name Resolution Function is based upon DNS technology. Since DNS is specified to be queried using domain names, and the name forms described in Section 4.1 are RFC5322 [5] addresses, incorporating both a user and domain component, the name “dialled” by the calling party needs to be transformed into a suitable format for usage in a DNS query. Per STD013 [6], this is carried out by encapsulating the “user” component (i.e. the local part) into a label, and replacing the “@” with a label-separator, as used to separate labels in DNS queries. So, using the examples from Section 4.1, the query would be to:

jo.bloggs@commsprovider.com	→ [jo.bloggs].[commsprovider].[com]
jo@bloggs.commsprovider.com	→ [jo].[bloggs].[commsprovider].[com]
jo.bloggs@namespaceprovider.com	→ [jo.bloggs].[namespaceprovider].[com]
jo@bloggs.namespaceprovider.com	→ [jo].[bloggs].[namespaceprovider].[com]
jo@bloggs.com	→ [jo].[bloggs].[com]

(where “.” between the square brackets is a label separator).

The above inputs assume a single Name Resolution Function, and that queries are directed to it using a specific IP address. In practical terms, particularly for Architectures three and four where there are multiple Name Resolution Functions, the identity of the desired one could be appended to the input query. For example rather than

jo.bloggs@commsprovider.com

resulting in a query to

[jo.bloggs].[commsprovider].[com]

it could result in a query to

[jo.bloggs].[commsprovider].[com].[resolutionprovider1].[com]

Output

The response to the query needs to be something that networks can use as a key for routeing. NICC has determined that the IMS form of Destination Group entry as specified in ND1631 Annex A.3.2 would be appropriate, albeit modified so that the "user" field need no longer be an E.164 number, for example;

df9asdsa@312.dg.commsprovider.uktel.org.uk

where

- df8asdsa is the user name on the terminating call control function
- 312 is the identifier for a destination call control function (or group of them) on the terminating communication provider's network.

It should be noted that the originally "dialled" name entered by the calling party must be carried through the call control signalling. This being the case, there is no need for the user name to bear any resemblance to that entered by the calling party - indeed it may be desirable for there not to be such a linkage.

7.1.2 Access controls : public or private?

The question of whether the mapping of user identity to communications provider represents personal data is a vexed one. For E.164 telephone numbers, the Office of the Information Commissioner has indicated that whilst in of itself a telephone number isn't personal data, in reality it needs to be treated as such due to the concept of "linked" databases. In brief, although the telephone number doesn't identify an individual so publishing a mapping of telephone number to communications provider discloses nothing, by linking that information with a publically available telephone directory, it is possible to determine the communications provider used by a given individual. This could be considered personal data. NICC has not engaged with the Office of the Information Commissioner to determine whether the same principle would be applicable foralphanumeric names (for example via usage of whois databases), but it is possible that this could be the case.

Notwithstanding the legal position, there could be a case to say that from the perspective of protection from SPIT, it would be unwise for the Name Resolution Function to be publically available.

Arguments around the data being commercially sensitive, as it is an enabler for marketing to entice customers to switch communications provider, are weak because the very entities that could engage in such activities (communications providers) are the ones who would be allowed access to the Name Resolution Function whether it was a public database, or had more closed access controls. However, there is an argument that the data as a whole is commercially sensitive, as it would allow third parties to glean business intelligence of the relative market shares of communications providers.

Whilst it is not the role of NICC to make commercial decisions such as whether the Name Resolution Function should be private to participating communications providers, it does seem that there are significant drivers for access to be restricted.

Notwithstanding this, insisting on data privacy does impose costs and constraints on the implementation. In particular, for Architecture Four, because each originating CP needs to have a relationship with each Name Resolution Function provider (potentially globally), this suggests a complex web of private connectivity should a public approach not be taken.

For Architecture Two which relies upon public DNS, clearly the data is inherently made public. Alternative approaches could be possible to making the data itself private rather than access to it. For example, the data could be made publically available, but encrypted in such a way that only those CPs with the appropriate access keys could decrypt and utilise it.

7.1.3 Local Storage Requirements

Some originating communications providers express a desire to maintain a copy of Name Resolution Function data within their own networks and systems, rather than having to query an external function. The advantage of this approach is that the communications provider is in full control of the performance of their network. Set against this, it can be problematic to download the data if it is held in a distributed architecture, especially if the full scope of the naming domain isn't known. The complexity this implies varies according to the architecture adopted for the Name Resolution Function; this issue is addressed in Section 7.2.

7.2 Architecture of Name Resolution Function

A series of internal architectural models can be envisaged for the Name Resolution Function. Each has implications for the size of the components of the Name Resolution Function, the speed of query response, the relative independence of operation of CP networks and the ability to download data into CP networks. The models and associated issues are described in Annex A.

7.3 Federation techniques

7.3.1 Background

In a federation approach, Name Resolution Function providers, who may be competitors, agree to peer to share their name resolution data. The sharing of data could be done off-line in advance of a query, in real-time when a query for a particular name is made, or using a hybrid approach (for example advance sharing of the most frequently called names, real-time for less frequently called names).

Sharing all data in advance removes the dependency on inter-domain communications links between NRFs for any real-time session initiation. This decoupling of real-time set-up from inter-NRF links means that the availability of these links becomes an issue only for updates to the data. This will ensure that there is no greater risk to service in alphanumeric addressing than with existing E.164 addressing today, and may help ensure that customer expectations regarding service availability and reliability are not unnecessarily compromised.

For the purposes of this document, the Name Resolution Function which is queried by the Call Control Function is termed the Initiating Name Resolution Function (INRF). The Name Resolution Function that has been provided with the name resolution data by the customer is termed the Authoritative Name Resolution Function (ANRF). Although this clause uses DNS terminology, there is no reason why other technologies couldn't be used, for example LDAP.

7.3.2 Real-time data sharing techniques

Typically within a federation, a namespace such as bloggs.com could be split across multiple ANRFs, each one being authoritative for a given set of individual names. The challenge for an INRF, therefore, is to determine to whom it should route a query.

7.3.2.1 Blanket query

In this approach, as illustrated in Figure 7.3.2.1.a, on receiving a query about a given name, an INRF would generate queries to all of its peered Name Resolution Functions, either in parallel or series. One of these should be the ANRF, and provide the required information (which could be that the individual name doesn't exist anywhere). The remainder would respond indicating that they possess no relevant information. If more than one Name Resolution Function responded with data, the INRF would need to implement some process to determine which "ANRF" to trust : this would not be a trivial exercise.

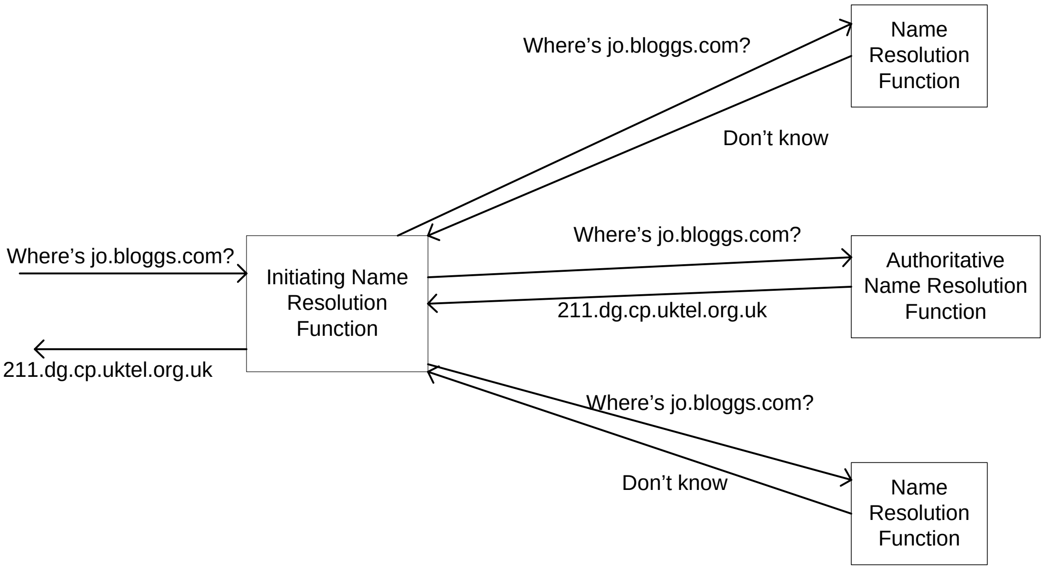

Figure 7.3.2.1.a : blanket query

The advantage of this approach is that the Name Resolution Function providers are truly peers, and no advance configuration is required for a given name. Set against this, clearly a single inbound query results in a series of queries within the federation, so there are issues of scaleability, particularly where a large number of Name Resolution Functions are federated.

7.3.2.2 Intelligent blanket query

In this approach, the INRF once again generates multiple queries to peered Name Resolution Functions, but in a more intelligent manner to those peers likely to possess the data. This is illustrated in Figure 7.3.2.2.a.

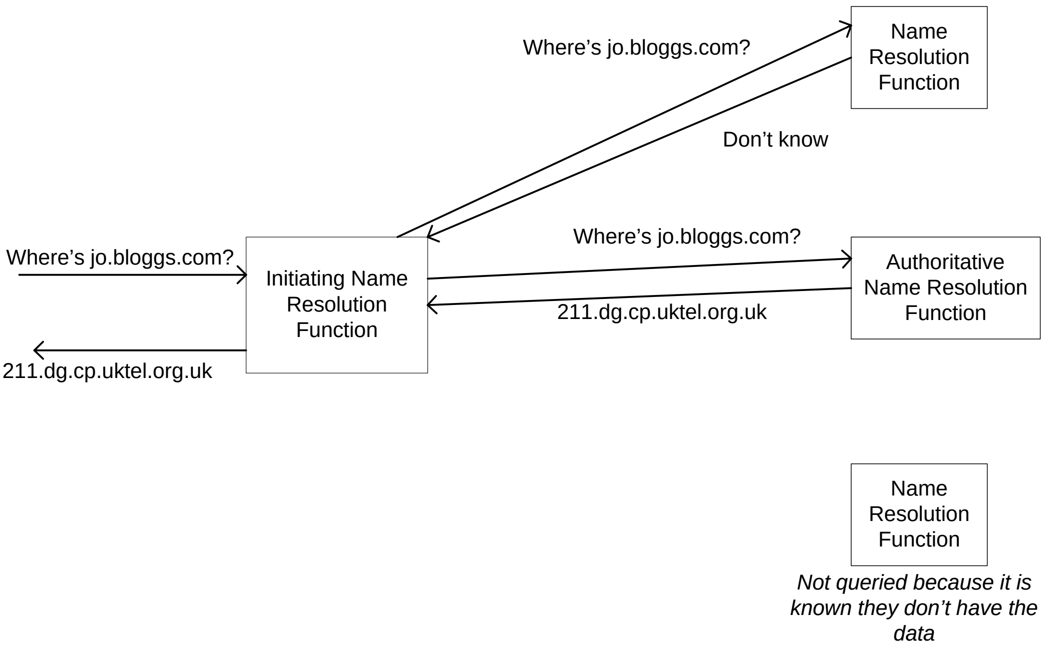

Figure 7.3.2.2.a : Intelligent blanket query

The mechanism for including/omitting a Name Resolution Function from a blanket query could be;

- That the Name Resolution Function has declared in advance whether they possess information within the namespace concerned, or

- That the Name Resolution Function has provided responses to queries on the namespace concerned within a given recent period.

If all of the Name Resolution Functions included within the intelligent blanket query yield no result, then the INRF could then resort to a blanket query to the remaining Name Resolution Functions.

This approach has the advantages of the blanket query approach, but generates fewer queries.

7.3.2.3 Lead ANRF

In this approach, those Name Resolution Functions which contain data for a given namespace would co-ordinate to agree a Lead ANRF for that namespace.

When the INRF received a query to a given name, they would initiate a single query to the Lead ANRF. This Lead ANRF could then either respond with the ANRF for that particular name (see Figure 7.3.2.3.a), or generate a subsequent query to the ANRF and respond to the INRF with the resolution data (see Figure 7.3.2.3.b). In both cases, if the Lead ANRF for the domain was actually the ANRF for the particular name queried, it would obviously respond with the resolution data.

The choice of Lead ANRF would be based on a commercially agreed approach, for example on the basis of the volumes of names that they contain, the anticipated volume of queries to the names they contain, or on a round-robin basis. The Lead ANRF for a given domain could additionally be split, for example names with local parts commencing a-m with one Lead ANRF, commencing n-z with another Lead ANRF.

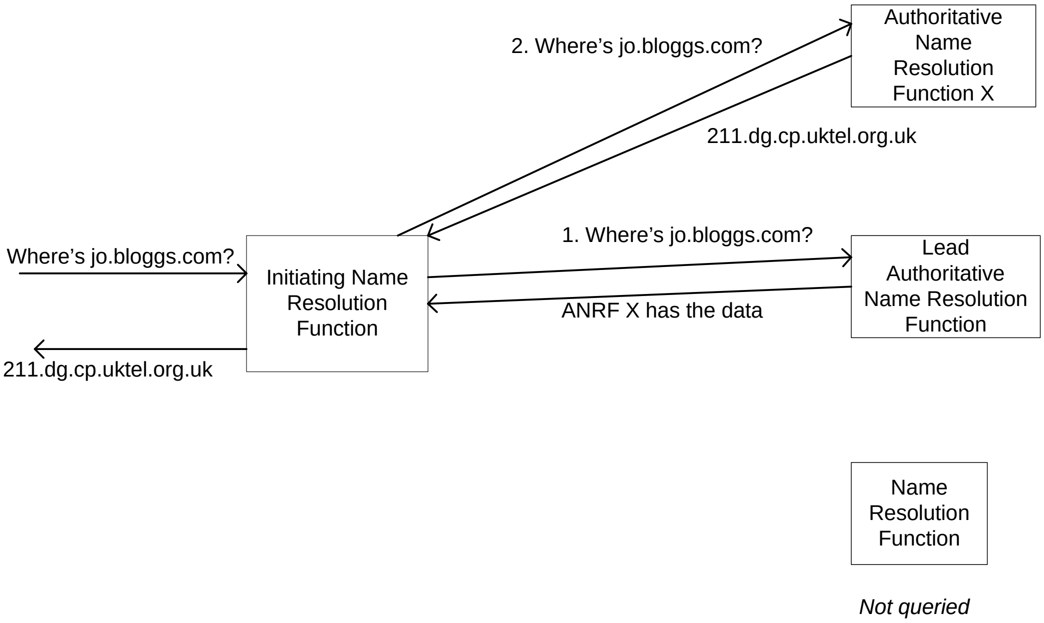
Authoritative Name Resolution Function X
2. Where’s jo.bloggs.com?
211.dg.cp.uktel.org.uk
1. Where’s jo.bloggs.com?
Where’s jo.bloggs.com?
Initiating Name Resolution Function
Lead Authoritative Name Resolution Function
ANRF X has the data
211.dg.cp.uktel.org.uk
Name Resolution Function
Not queried

Figure 7.3.2.3.a : Lead ANRF for domain provides identity of ANRF

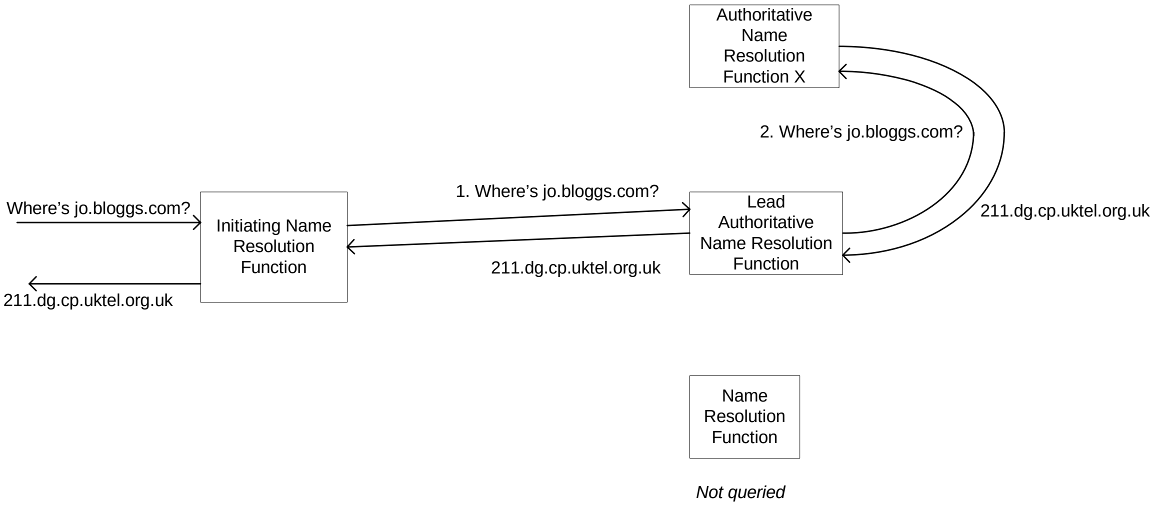
Authoritative Name Resolution Function X
2. Where’s jo.bloggs.com?
1. Where’s jo.bloggs.com?
Where’s jo.bloggs.com?
Initiating Name Resolution Function
Lead Authoritative Name Resolution Function
211.dg.cp.uktel.org.uk
211.dg.cp.uktel.org.uk
211.dg.cp.uktel.org.uk
Name Resolution Function
Not queried

Figure 7.3.2.3.b : Lead ANRF for domain generates query to ANRF to provide response

This approach has the advantage of limiting the number of queries required within the federation. However, agreement is required between the peering Name Resolution Functions about who should be the Lead ANRF for a particular domain.

7.3.3 Techniques for data sharing in advance

The issue of data sharing in advance can be subdivided into two approaches, namely sharing the actual data in advance, or simply sharing which Name Resolution Function is the ANRF in advance. For the former, the INRF would contain the actual data, which arguably is more resilient but comes at the expense of potentially large amounts of storage required. In the latter case, the INRF would just contain pointers to the relevant ANRF.

The information could be shared from ANRFs to INRFs via transmission of files of the whole Name Resolution Function contents, followed by periodic updates of any changes that have occurred. In DNS terms, this could in theory be accomplished using AXFR, NOTIFY and IXFR, but this would be dependent upon the INRF actually knowing the domains for which the ANRF contains authoritative name data : mechanisms would need to be put in place to enable this, which is beyond the scope of this document. Equivalent techniques could be utilised if the technology was not DNS-based. Further, on receiving the data from various ANRFs (whether via DNS technology or in some other form), the INRF would have to merge this together : given multiple ANRFs would potentially be authoritative for various names within a given domain, this would not be a trivial exercise. For this reason, it is considered only realistic to share the data in advance where it is for a known domain and there is a single ANRF for that domain. In practical terms this could only apply where the domain relates to a CP, and it may not be feasible to share the data in advance e.g. for a multi-national enterprise with their own domain.

A mechanism for ensuring the integrity of the shared data should be considered. One possible solution would be to implement a system based on checksums as was envisaged for use between the proposed CDB and CPs' own copies of this.

Consideration should be given to whether a domain should be sub-divided such that only a part of the total domain namespace need to be downloadedFor example, a domain could be divided on the basis of an agreed algorithm, for example the first character in the name. The number of subdivisions would logically be related to the total population of names. Further work in this area could assess what the optimum size of a partition would be and sub-division on that basis employed. As domains would contain different volumes of names (with some probably at least an order of magnitude greater than others), a hard-and-fast rule for where sub-division should take place would lead to an unnecessary fragmentation of smaller domains.

8 Interaction with end-to-end services

As alphanumeric names would essentially replace E.164 numbers for networks supporting this capability, there is a need to consider how such networks would interact with networks only supporting E.164 numbers, and the implication for services such as Calling Line Identity (CLI). This report does not seek to provide answers to the questions raised, instead the intention is to provide a list of issues to be addressed.

8.1 Retention of original called party name

This report does not specify any signalling mechanisms but the appropriate standards, if developed, must provide for the originally requested URI to be conveyed across the network to the final called party. This must be accomplished in such a way that appropriately standardised end-to-end identification systems are not compromised.

8.2 Interaction with networks supporting only E.164-based numbers

Two scenarios require consideration, namely where the called customer is on a network that supports alphanumeric names but the calling customer is not (case A), and where the calling customer is on a network that supports alphanumeric names but the called customer is not (case B).

Case A : Called customer has alphanumeric name but calling customer does not support this

For Case A, the calling customer will have no mechanism of addressing the called customer (probably both from a perspective of ability of their terminal equipment to provide anything other than numeric addresses, and from a perspective of their CP's network to interpret this). Clearly the only option available is for them to enter an equivalent E.164 number for the intended customer and for an interworking to apply at some point in the call path. However, as highlighted in Section 5.1, if the implication is that every end-user with an alphanumeric name also requires an equivalent E.164 number, this could be very resource hungry. Alternatives will need to be examined, for example usage of a common number with the final destination customer information entered in-band.

Case B : Calling customer has alphanumeric name but called customer does not support this

For Case B, the calling customer should, of course, be able to address the called customer using a standard E.164 telephone number. The issue, though, is what CLI (see Section 8.2) will be presented to the called customer. To support Case A, the CLI would have to be E.164-based. However, if the caller was to be placing a call to a customer whose network did support alphanumeric names, then clearly it would be preferable to have a CLI that reflected their alphanumeric name. The implication of this is either that the originating network would need to examine the called name/number and associate an appropriate CLI, or that both the E.164 and alphanumeric CLI be conveyed with the call. For usage of E.164 CLIs, there are particular issues if the measures to preserve E.164 numbers raised in the previous paragraph are implemented.

Even where both the called and calling customer support alphanumeric names, the originating CP must plan their routeing in order that the call only traverses transit networks that are able to convey all the required signalling. In isolate situations, it is possible that the originating and terminating networks may represent islands so there is no route between them that is able to convey the alphanumeric names. In this situation the considerations of Case A and B would apply, but would be exacerbated by it not being known at the outset that interworking is required.

International calls represent a particularly complex example of these issues, undoubtedly necessitating global agreement of the approach to be adopted.

8.3 CLI

As described in Section 8.2, in some cases it will be desirable to use a CLI based upon the alphanumeric name. To do this, in essence all systems that currently assume a numeric identifier would need to be upgraded to support alphanumeric identifiers. However, where this occurs, it will be necessary to consider the implications for:

- Presentation services (see Section 8.2)
- Anonymous Call Reject
- Malicious Call Identification
- Diversion services
- Voicemail

- Emergency services – would it be necessary to upgrade emergency call handling centres to support alphanumeric CLIs? Additionally, alphanumeric names would need to be associated with a location.
- Directory enquiry databases
- Billing systems
- Network usage such as validation of the calling customer for carrier selection / pre-selection
- Location-based routeing for intelligent network services.

9 Recommendations

9.1 Architecture

Having examined various architectures and the issues associated with each, NICC concludes that Architecture 3b appears at this stage to best fulfil the current requirements as currently known in a scalable, cost effective manner. This means that the Name Resolution Locator Function would be accomplished via DNS, and would provide the identities of federations of Name Resolution Functions that can provide termination information for the domain in question. Information about a given domain could be provided by one or more federations of Name Resolution Functions (indeed nothing precludes a federation consisting of a single Name Resolution Function instance).

NICC does not make any recommendation about the internal architecture of Name Resolution Functions, as this is a matter for the providers of such federations. However, the material in Section 7.3 of this report provides guidance as to potential solutions.

9.2 Standards requirements

Further work in this area should look to the industry standards bodies to ensure that any solution for the UK will be compliant to more universal standards. This should track work not only in IETF, but also in TISPAN & 3GPP to ensure that carrier-class standards and methodologies are adopted across the industry. This is particularly salient in relation to the presumed requirement of end-users identified by name wishing to contact end-users located in CP domains in other countries.

In this respect, one use of this report would be as the basis for input into the above bodies to extend the dialogue beyond just the UK.

Annex A : Internal Name Resolution Function Architectural Models

A series of internal architectural Models can be envisaged for the Name Resolution Function, as described in this Annex.

For Architectures One and Two, as a monopoly provider of Name Resolution Functions, the data contained within the Model will be the entire accessible namespace. For Architecture Four, the data held within the each instance of the Model will be solely that for which the particular Name Resolution Function is authoritative. For Architecture Three, the data contained within each instance of the Model will be that held locally within locally the Name Resolution Function and will either be just that for which it is authoritative (in the case of dynamic query – see Section 7.3) or all data, including that obtained from federated Name Resolution Functions (in the case of *a priori* sharing of data – see Section 7.3).

A.1 Name Resolution Function Model A : monolithic database

In this model, within a given Name Resolution Function the data would be held in a monolithic database holding all of the mappings from names to termination locations, as depicted in Figure A.1.a.

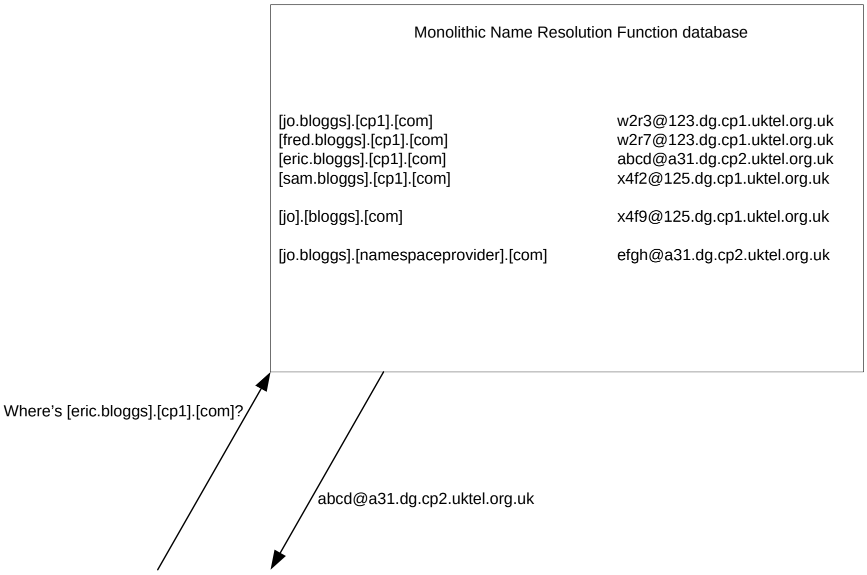
Monolithic Name Resolution Function database
Where's [eric.bloggs].[cp1].[com]?
abcd@a31.dg.cp2.uktel.org.uk

[jo.bloggs].[cp1].[com]	w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com]	w2r7@123.dg.cp1.uktel.org.uk
[eric.bloggs].[cp1].[com]	abcd@a31.dg.cp2.uktel.org.uk
[sam.bloggs].[cp1].[com]	x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com]	x4f9@125.dg.cp1.uktel.org.uk
[jo.bloggs].[namespaceprovider].[com]	efgh@a31.dg.cp2.uktel.org.uk

Figure A.1.a : Model A, monolithic database

This model would be best suited to a requirement to be able to store all of the data locally in originating CPs, because as a single database it should be more readily downloadable. Set against this, the database would be very large.

For existing names, changes to the termination location by the current terminating CP would require a change to this monolithic database. If there was a need to port a name from one CP to another, this would require a change to the monolithic database. Whenever new names needed to be added, this would also require a change to the monolithic database.

For overall Architecture Three, exactly what the monolithic Name Resolution Function would comprise would vary according to the federation model adopted : see Section 7.3. For example, if the model was to proactively share data, then the monolith could contain data both populated directly into that Name Resolution Function, and also that acquired via sharing within a federation. If, however, the model was to share data only on demand, then clearly the Name Resolution Function could not be monolithic and must to a degree be dependent on subsequent external queries.

A.2 Name Resolution Function Model B : Tiered with thick Tier One

In this model, within a given Name Resolution Function there would be a Tier One database that pointed on an individual name basis to Tier Two databases containing the actual records, as depicted in Figure A.2.a.

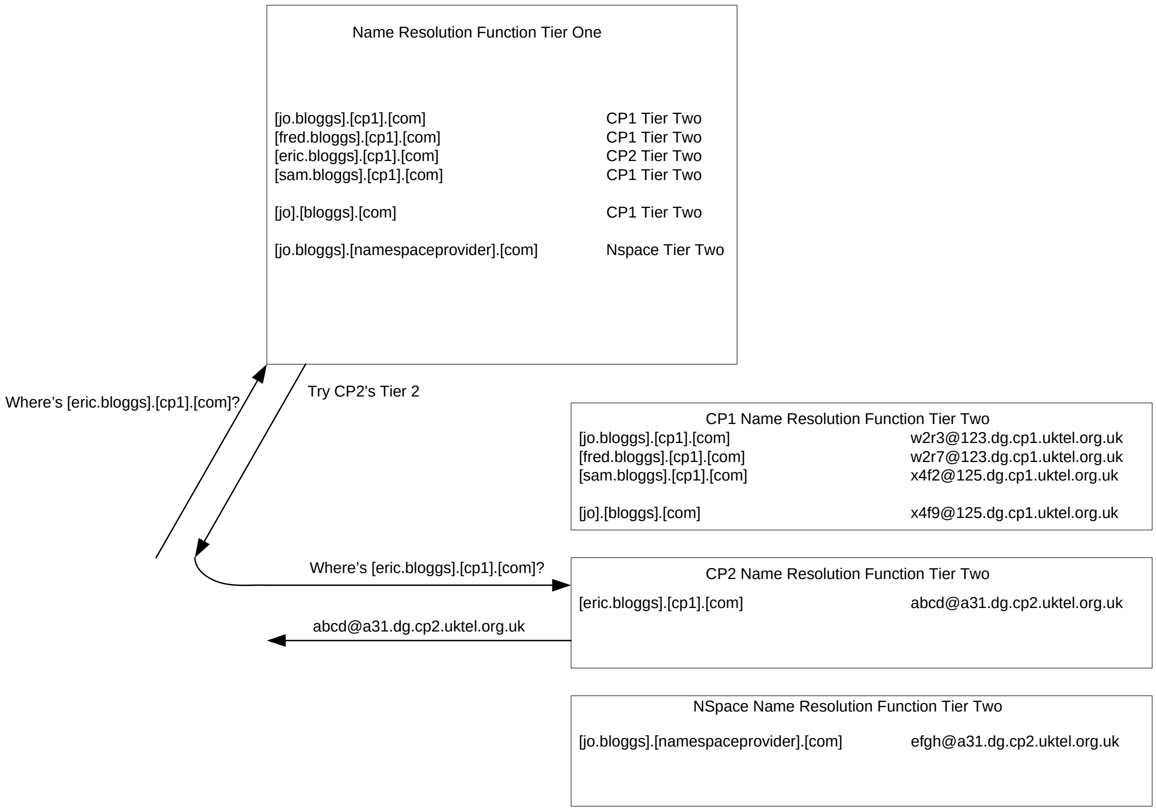
Name Resolution Function Tier One
[jo.bloggs].[cp1].[com] CP1 Tier Two
[fred.bloggs].[cp1].[com] CP1 Tier Two
[eric.bloggs].[cp1].[com] CP2 Tier Two
[sam.bloggs].[cp1].[com] CP1 Tier Two
[jo].[bloggs].[com] CP1 Tier Two
[jo.bloggs].[namespaceprovider].[com] Nspace Tier Two
Where's [eric.bloggs].[cp1].[com]?
Try CP2's Tier 2
CP1 Name Resolution Function Tier Two
[jo.bloggs].[cp1].[com] w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com] w2r7@123.dg.cp1.uktel.org.uk
[sam.bloggs].[cp1].[com] x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com] x4f9@125.dg.cp1.uktel.org.uk
Where's [eric.bloggs].[cp1].[com]?
CP2 Name Resolution Function Tier Two
[eric.bloggs].[cp1].[com] abcd@a31.dg.cp2.uktel.org.uk
abcd@a31.dg.cp2.uktel.org.uk
NSpace Name Resolution Function Tier Two
[jo.bloggs].[namespaceprovider].[com] efgh@a31.dg.cp2.uktel.org.uk

Figure A.2.a : Model B, Tiered with thick Tier One

This model treads a middle line in the context of the size of the Tier One database. Each individual name needs to be listed, but only with a pointer to an authoritative Tier Two database, rather than the full data in Model A. The database would nevertheless be significant in size.

For existing names, changes to the termination location by terminating CPs would require only a change to the Tier Two database. If there was a need to port a name from one CP to another, this would however require a change to the Tier One database. Similarly, whenever new names needed to be added, this would require a change to the Tier One database.

Downloading the contents of the Name Resolution Function to individual originating CPs would be problematic. In principle, the downloading CP would need to obtain copies of the contents of each individual Tier Two database. In order to do this, they would need *a priori* knowledge of the Tier Two identities. A compromise approach could be adopted whereby well known domains were downloaded (e.g. assuming CPs would operate Tier Twos for their own terminations, the main CPs), combining this with analysis to determine the most queried Tier Twos from their originating call patterns. However, this is unavoidably complex.

With this approach, a single CP query to the Name Resolution Function involves two database lookups, to the Tier One then Tier Two : this could adversely impact call setup time.

For overall Architecture Three, exactly what the Tier One of Name Resolution Function would consist of would vary according to the federation model adopted : see Section 7.3.

A.3 Name Resolution Function Model C : Tiered with lean Tier One

In this model, within a given Name Resolution Function there would be a Tier One database that pointed to Tier Two databases containing the actual records. Within the Tier One database there would be a mixture of individual and wildcard pointers. So, if the situation was that communications provider CP1 operated their own Tier Two database for all of the names they served within the domain cp1.com, but a customer had ported their individual name eric.bloggs@cp1.com to communications provider CP2, then there would be a wildcard entry which pointed all queries to CP1's Tier Two database, other than the specific [fred.bloggs].[cp1].[com] entry which would point to the recipient CP2's Tier Two database. This arrangement is depicted in Figure A.3.a. and A.3.b for non-ported and ported names respectively.

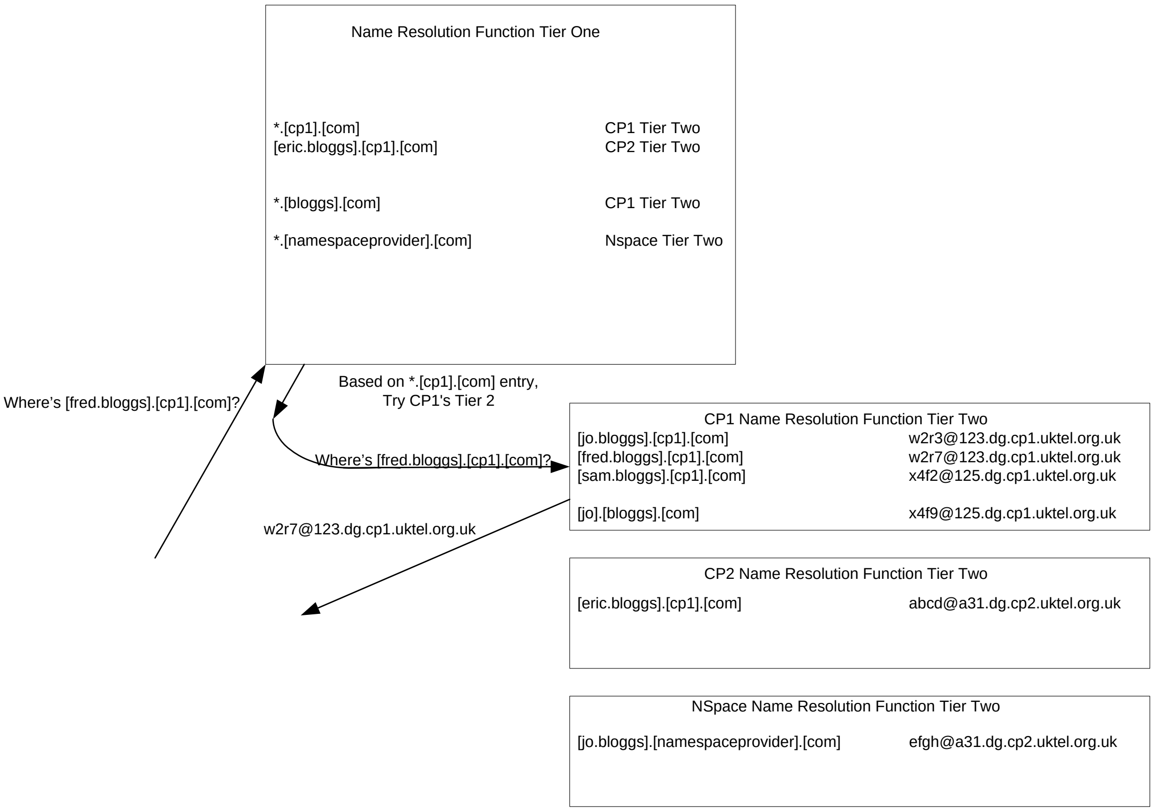
Name Resolution Function Tier One
*.[cp1].[com]
CP1 Tier Two
[eric.bloggs].[cp1].[com]
CP2 Tier Two
*.[bloggs].[com]
CP1 Tier Two
*.[namespaceprovider].[com]
Nspace Tier Two
Where's [fred.bloggs].[cp1].[com]?
Based on *.[cp1].[com] entry,
Try CP1's Tier 2
Where's [fred.bloggs].[cp1].[com]?
w2r7@123.dg.cp1.uktel.org.uk
CP1 Name Resolution Function Tier Two
[jo.bloggs].[cp1].[com]
w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com]
w2r7@123.dg.cp1.uktel.org.uk
[sam.bloggs].[cp1].[com]
x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com]
x4f9@125.dg.cp1.uktel.org.uk
CP2 Name Resolution Function Tier Two
[eric.bloggs].[cp1].[com]
abcd@a31.dg.cp2.uktel.org.uk
NSpace Name Resolution Function Tier Two
[jo.bloggs].[namespaceprovider].[com]
efgh@a31.dg.cp2.uktel.org.uk

Figure A.3.a : Model C, Tiered with lean Tier One, non-ported name

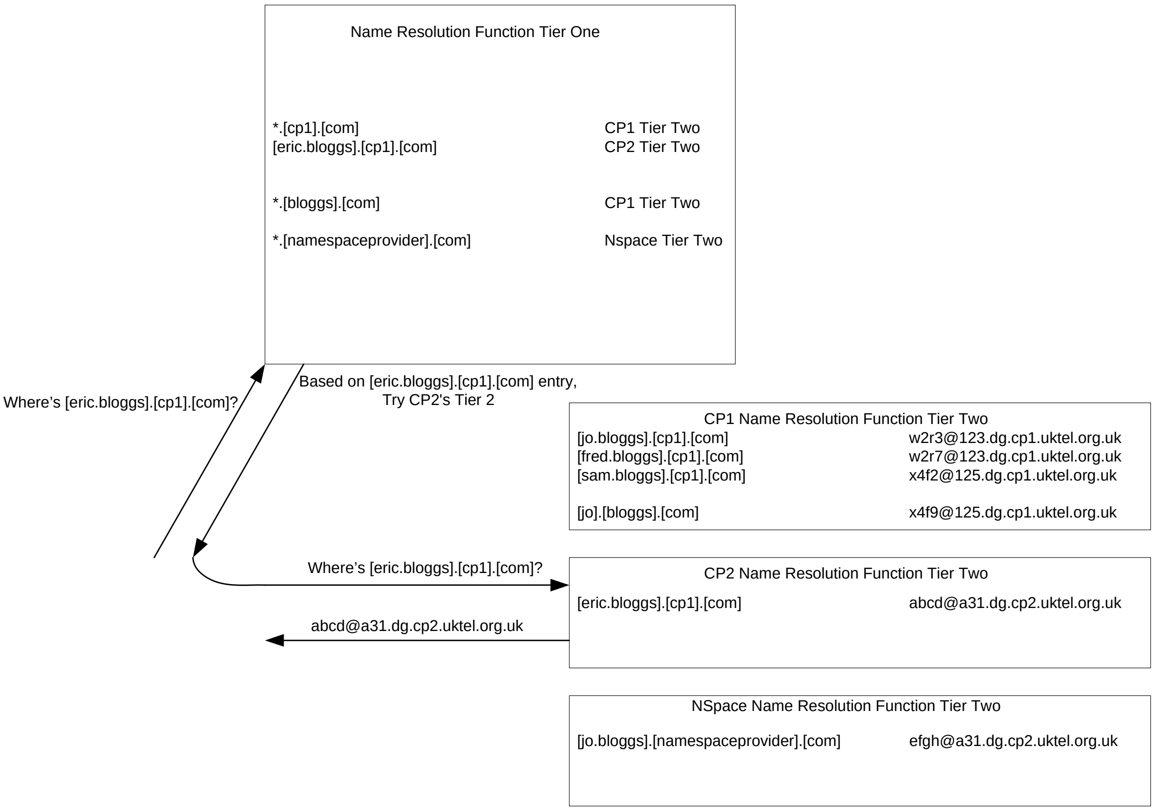
Name Resolution Function Tier One
*.[cp1].[com] CP1 Tier Two
[eric.bloggs].[cp1].[com] CP2 Tier Two
*.[bloggs].[com] CP1 Tier Two
*.[namespaceprovider].[com] Nspace Tier Two
Where's [eric.bloggs].[cp1].[com]?
Based on [eric.bloggs].[cp1].[com] entry, Try CP2's Tier 2
CP1 Name Resolution Function Tier Two
[jo.bloggs].[cp1].[com] w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com] w2r7@123.dg.cp1.uktel.org.uk
[sam.bloggs].[cp1].[com] x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com] x4f9@125.dg.cp1.uktel.org.uk
Where's [eric.bloggs].[cp1].[com]?
CP2 Name Resolution Function Tier Two
[eric.bloggs].[cp1].[com] abcd@a31.dg.cp2.uktel.org.uk
abcd@a31.dg.cp2.uktel.org.uk
NSpace Name Resolution Function Tier Two
[jo.bloggs].[namespaceprovider].[com] efgh@a31.dg.cp2.uktel.org.uk

Figure A.3.b : Model C, Tiered with lean Tier One, ported name

This model leads to a smaller Tier One database than Models A and B. In essence, the only individual names that would need to be broken out in the Tier One are those that have been ported, and those where the customer uses their own domain and chooses to split its management across multiple Tier Two providers.

For existing names, changes to the termination location by terminating CPs would require only a change to the Tier Two database. If there was a need to port a name from one CP to another, this would require a change to the Tier One database. However, whenever new names needed to be added, this would once again require a change only to the Tier Two database.

Downloading the contents of the Name Resolution Function to individual originating CPs would cause similar issues to that described for Model B.

Similarly, for overall Architecture Three, exactly what the Tier One of Name Resolution Function would consist of would vary according to the federation model adopted : see Section 7.3.

As with Model B, with this approach, a single CP query to the Name Resolution Function involves two database lookups, to the Tier One then Tier Two : this could adversely impact call setup time.

A.4 Name Resolution Function Model D : Tiered with skinny Tier One

In this model, within a given Name Resolution Function there would be a Tier One database that pointed to Tier Two databases containing the actual records. Within the Tier One database there would only be entries at a domain level basis, pointing to a relevant Tier Two database. Tier Two would contain the detailed data for each name : where a name was ported from one communications provider to another, the data in the Tier Two (which probably would be maintained by the donor communications provider) would reflect the termination information of the recipient communications provider. This arrangement is depicted in Figure A.4.a.

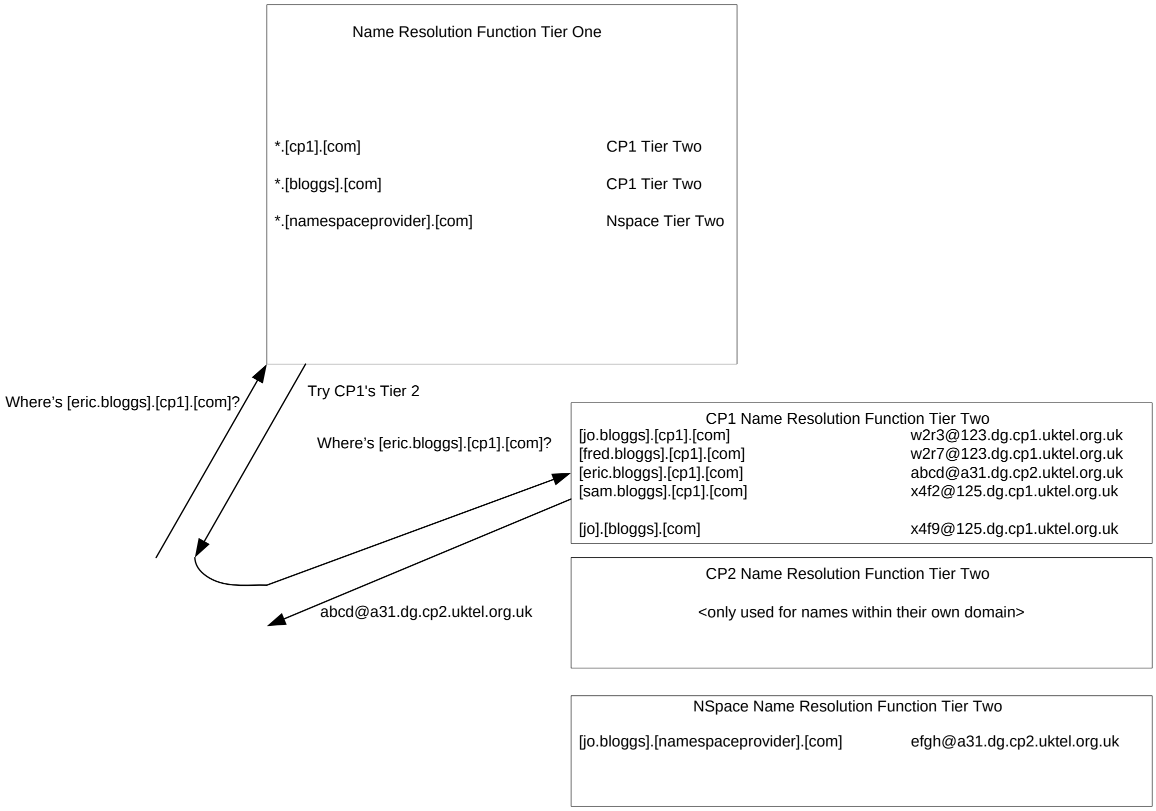
Name Resolution Function Tier One
*.[cp1].[com]
CP1 Tier Two
*.[bloggs].[com]
CP1 Tier Two
*.[namespaceprovider].[com]
Nspace Tier Two
Where's [eric.bloggs].[cp1].[com]?
Try CP1's Tier 2
Where's [eric.bloggs].[cp1].[com]?
abcd@a31.dg.cp2.uktel.org.uk
CP1 Name Resolution Function Tier Two
[jo.bloggs].[cp1].[com]
w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com]
w2r7@123.dg.cp1.uktel.org.uk
[eric.bloggs].[cp1].[com]
abcd@a31.dg.cp2.uktel.org.uk
[sam.bloggs].[cp1].[com]
x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com]
x4f9@125.dg.cp1.uktel.org.uk
CP2 Name Resolution Function Tier Two
<only used for names within their own domain>
NSpace Name Resolution Function Tier Two
[jo.bloggs].[namespaceprovider].[com]
efgh@a31.dg.cp2.uktel.org.uk

Figure A.4.a : Model D, Tiered with skinny Tier One

This model leads to the smallest Tier One database as it would only contain domain names. Indeed, for overall Architecture Three, the contents of the Name Resolution Function would be pretty much the same as the Name Resolution Locator Function, so to a large degree this Model can be discounted for that architecture.

Changes to names, whether introduction of new ones, rehosting within the terminating network (i.e. changing the Destination Group) or porting between communications providers would impact only the Tier Two. However, it does imply that in the case of a ported name, the recipient provider is dependent upon the performance of the donor's Tier Two. For new names that involved new domains, clearly this would involve the Tier One of the Name Resolution Function as well.

Downloading the contents of the Name Resolution Function to individual originating CPs would cause similar issues to that described for Model B.

Similarly, for overall Architecture Two, exactly what the Tier One of Name Resolution Function would consist of would vary according to the federation model adopted : see Section 7.3.

With this approach, a single CP query to the Name Resolution Function involves two database lookups, to the Tier One then Tier Two : this could adversely impact call setup time.

A.5 Name Resolution Function Model E : Tiered with skinny Tier One and Tier Two redirection

In this model, within a given Name Resolution Function there would be a Tier One database that pointed to Tier Two databases containing the actual records. Like Model D, within the Tier One database there would only be entries at a domain level basis, pointing to a relevant Tier Two database. Tier Two would contain the detailed data for each name : however where a name was ported from one communications provider to another, the data in the Tier Two (which probably would be maintained by the donor communications provider) would redirect the query to the correct Tier Two database. This arrangement is depicted in Figure A.5.a.

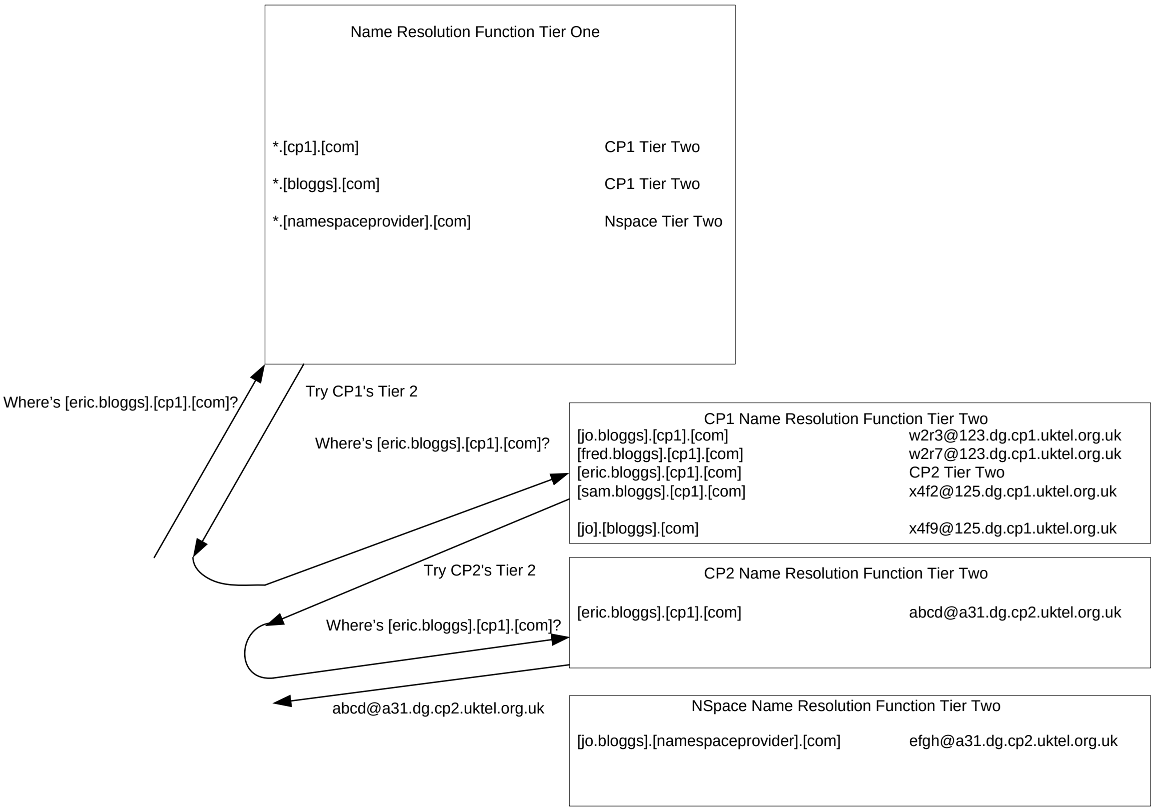
Name Resolution Function Tier One
*.[cp1].[com] CP1 Tier Two
*.[bloggs].[com] CP1 Tier Two
*.[namespaceprovider].[com] Nspace Tier Two
Where’s [eric.bloggs].[cp1].[com]?
Try CP1's Tier 2
Where’s [eric.bloggs].[cp1].[com]?
CP1 Name Resolution Function Tier Two
[jo.bloggs].[cp1].[com] w2r3@123.dg.cp1.uktel.org.uk
[fred.bloggs].[cp1].[com] w2r7@123.dg.cp1.uktel.org.uk
[eric.bloggs].[cp1].[com] CP2 Tier Two
[sam.bloggs].[cp1].[com] x4f2@125.dg.cp1.uktel.org.uk
[jo].[bloggs].[com] x4f9@125.dg.cp1.uktel.org.uk
Try CP2's Tier 2
Where’s [eric.bloggs].[cp1].[com]?
CP2 Name Resolution Function Tier Two
[eric.bloggs].[cp1].[com] abcd@a31.dg.cp2.uktel.org.uk
abcd@a31.dg.cp2.uktel.org.uk
NSpace Name Resolution Function Tier Two
[jo.bloggs].[namespaceprovider].[com] efgh@a31.dg.cp2.uktel.org.uk

Figure A.5.a : Model E, Tiered with skinny Tier One and Tier Two redirection

This model leads to the smallest Tier One database as it would only contain domain names. For overall Architecture Four, the contents of the Name Resolution Function would be pretty much the same as the Name Resolution Locator Function, so to a large degree this Model can be discounted for that architecture.

Changes to names, whether introduction of new ones, rehosting within the terminating network or porting between communications providers would impact only the Tier Two. However, it does imply that in the case of a ported name, the recipient provider is dependent upon the performance of the donor, albeit unlike Model D, they can make changes to the termination data without the direct involvement of the donor. For new names that involved new domains, clearly this would involve the Tier One of the Name Resolution Function as well.

Downloading the contents of the Name Resolution Function to individual originating CPs would cause similar issues to that described for Model B.

Similarly, for overall Architecture Three, exactly what the Tier One of Name Resolution Function would consist of would vary according to the federation model adopted : see Section 7.3.

With this approach, a single CP query to the Name Resolution Function involves two database lookups (to the Tier One then Tier Two) for a non-ported name, and three database lookups (to the Tier One, then donor Tier Two then recipient Tier Two) for a ported name : this could adversely impact call setup time.

History

Document history		
<Version>	<Date>	<Milestone>
1.1.1	Jan 10	Initial issue