NICC ND 1415 V1.1.2 (2008-08)

NICC Document

NGN; PSTN/ISDN Service Interconnect; Guide to Common Numbering Database standards

Network Interoperability Consultative Committee,
Ofcom,
2a Southwark Bridge Road,
London,
SE1 9HA

If you have any comments concerning the accuracy of the contents of this document, please write to:

The Technical Secretary,
Network Interoperability Consultative Committee,
Ofcom,
2a Southwark Bridge Road,
London SE1 9HA.

Contents

Intellectual Property Rights 5
Foreword 5
Introduction 5
1 Scope 6
2 References 6
2.1 Informative references 6
3 Definitions and abbreviations 6
3.1 Definitions 6
3.3 Abbreviations 7
4 Executive Summary 8
4.1 Database Concept 8
4.2 Routeing Concept 10
4.3 Security of Data 10
4.4 Porting Numbers 11
4.5 Numbering Management 11
5 Requirements and Architecture Overview 12
5.1 Overall Requirement 12
5.2 Functional Entities and Interfaces 12
5.3.1 Functions 14
5.3.1.1 CP Functions 14
5.3.1.1.1 Call Control Function 14
5.3.1.1.2 Routeing Function 14
5.3.1.2 Common Numbering Database 14
5.3.1.2.1 Central Numbering Database 14
5.3.1.2.2 CP Copy of Central Numbering Database 15
6 What the Architecture means in practice 16
6.1 Destination Groups 16
7 Database considerations 19
7.1 What numbers should be in the Database 19
7.2 Contents – What should be in the records 19
7.3 Should the database be accessed in real-time or only administrative 21
7.4 Potential physical implementations of the database 21
7.5 How the database is structured, managed and queried in practice 21
7.5.1 Access 21
7.5.2 Sections 21
7.5.3 Querying a copy of the data hosted locally on CP systems 22
7.5.3.1 Format and compression of bulk download data 23
7.5.4 Querying the central database as and when required. 24
7.6 Provisioning the database 24
7.6.1 The management Interface 24
7.6.2 Number Groups 24
7.6.3 Access Keys 25
7.6.4 Typical Processes 25
7.6.4.1 Numbers are allocated by Ofcom 26
7.6.4.2 CP provisions numbers 26
7.6.4.3 CP reconfigures their network 26
7.6.4.4. Number is ceased 26
7.6.4.5 Number is ported between networks 26
7.6.4.6 Ported number is ceased 27
7.6.4.7 Numbers are reclaimed by Ofcom 28

8 Routeing Rules 28

Annex A : Studies that determined that NICC Standards should incorporate individual number routeing, i.e. All Call Query 29

A1 Introduction 29

A2 Background 29

A3 Number Portability Techniques 30
A3.1 Block routeing techniques 30
A3.1.1 Onward Routeing : continued usage of prefixes 30
A3.1.1.1 Solution overview 30
A3.1.1.2 Advantages & Disadvantages.... 31
A3.1.1.3 Conclusions.... 31
A3.1.2 Onward Routeing : Enhanced Onward Routeing via Modification of Destination SIP header by Donor Network.... 32
A3.1.2.1 Solution overview 32
A3.1.2.2 Advantages and disadvantages.... 33
A3.1.2.3 Conclusions.... 33
A3.1.3 Block Routeing enhanced with SIP-redirect techniques.... 34
A3.1.3.1 Solution overview 34
A3.1.3.2 Advantages & Disadvantages.... 35
A3.1.3.3 Conclusions.... 35
A3.2 Shared Database Techniques 36
A3.2.1 Single Common Database.... 36
A3.2.1.1 Solution overview 36
A3.2.1.2 Advantages & Disadvantages.... 37
A3.2.1.3 Conclusions.... 37
A3.2.2 Multiple Shared Databases 38
A3.2.2.1 Solution Overview 38
A3.2.2.2 Advantages & Disadvantages.... 40
A3.2.2.3 Conclusions.... 40

A4 Conclusions....40

Annex B : Use Cases – from ND1631....41

Annex C : Issues considered in determining which numbers to include in the database 45
C1 Defining the Question 45
C2 Administrative Information. 45
C3 Destination Information.... 46
C4 Conclusions 47

Annex D : Examination of possible physical database architectures....48

Annex E : Partitioning of the database....51

Annex F : Impacts of compression on the bulk XML download of a fully populated number section 54

History 60

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 This NICC Document (ND) has been produced by NICC Naming Numbering and Addressing Working Group

Introduction

The present document provides an overview of the material considered by the NNA group with respect to fulfilment of Requirement L of Green Release.

This work concluded that an "All Call Query" solution using data from a Central Numbering Database was the appropriate solution.

Prior to establishing the requirement to support an All Call Query approach within Green Release, NNA examined various other architectures in an NGN context, to determine the optimal approach. A summary of this work is provided in Annex A of this document.

The documents within Green Release which set out the standards for the common numbering database are as follows;

ND1610	Green Release specification and documentation structure
ND1631	Architecture associated with the common numbering database
ND1022	Specification of the DNS implementation of the real-time and bulk download interfaces
ND1023	Framework for webservice implementation of interfaces to the common numbering database
ND1024	Specification of the webservice implementation of the real-time and bulk download interfaces
ND1025	Specification of the webservice implementation of the management interface
ND1415	This narrative guide to the conclusions reached by NICC.

The current version of these documents is specified in ND1610 [2].

The requirements of the modified General Condition 18 are supported by the All Call Query approach adopted.

1 Scope

The present document provides an overview of the material considered by the NNA group in determining the architecture for use of a Common Numbering Database with respect to fulfilment of Requirement L of Green Release, and the requirements of General Condition 18.

This document is informative only.

2 References

2.1 Informative references

[1] NICC ND1631: "NGN; PSTN/ISDN Service Interconnect; Architecture for usage of Common Numbering Database"

[2] ND1610 "Multi-Service Interconnect of UK Next Generation Networks"

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

Service : For the purposes of Green Release, this means IMS or PSTN (whether provided by TDM technologies or using a PSTN Emulation System)

C-D Digit boundary : the boundary between the fifth and sixth digits in a UK telephone number including the leading zero.

.

NOTE: The standard format of a National Significant Number ('NSN') within the UK National Telephone Numbering Plan is notated as follows: 'SABCDEFGHI'. 'S' equates to the first digit after the '0' (for example S is 1 or 2 for Geographic Numbers). The C-D digit boundary is chosen as the dividing point for Sections of the Common Numbering Database so that all numbers with the same 0SABC prefix will be in the same Section e.g. for the number 01632 960123, C is '2' and D is '9': all numbers beginning 01632 will be in the same Section. Most geographic area codes comprise at least 3 digits after the S digit (i.e. ABC), hence this means that in many cases all numbers from a particular geographic area code will be contained in the same database Section.

Common Numbering Database : the generic function which holds reference mapping of telephone numbers to Destination Group. The Common Numbering Database is the general term for the superset of the Central Numbering Database and all copies of it held locally by individual CPs.

Central Numbering Database : The central reference database holding the mapping of telephone numbers to Destination Group.

CP numbering database : A CP copy of the Central Numbering Database.

Routeing function : takes the Destination Group from the Central/CP numbering database, and based on CP routeing policy, determines routeing of session/messages.

Destination Group : A Destination Group represents a set of numbers (which may or may not be contiguous) served by a communications provider for which another provider would make the same routeing decisions for all numbers within it

Owning Communications Provider (Owning CP): The Communications Provider which currently has rights to populate Administrative and Destination Group information associated with a number into the Common Numbering Database.

Pending Communications Provider (Pending CP): The Communications Provider which, following granting of permission by the Owning Communications Provider, has the right to take over control of the Administrative and Destination Group information associated with a number in the Common Numbering Database.

Recipient Communications Provider (Recipient CP): In process terms, the Communications Provider which will be the Owning Communications Provider for a number which is to be subject to Number Portability

3.3 Abbreviations

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

CDB	Central numbering DataBase
CP	Communications Provider
DG	Destination Group
DNS	Domain Name System
ebXML	Electronic Business using XML
IN	Intelligent Network
IRN	Intermediate Routing Number, defined in ND 1208
MAP	Mobile Application Part (of C7 signalling)
SRF	Signalling Relay Function
XML	eXtensible Markup Language

4 Executive Summary

NICC has concluded that the proposed revisions to General Condition 18 would mean that UK realisation of number portability would be best served by a common numbering database – that is a central database and multiple authoritative copies held by CPs. NICC recommends that this should contain all UK numbers in service (whether ported or not, covering all parts of the numbering space with the exception of network-internal numbers), with records for each containing information about where calls should be terminated.

Note : The solution also supports real-time query by a CP, provided either directly by the central database or as a third party provided service.

Note : In this context "calls" refers to communications whether voice, SMS or MMS.

However, the NICC standards are compatible with a system whereby destination information for a lesser set of numbers, e.g. ported numbers only, is stored using the common numbering database. It should be noted that at least information about which CP has been assigned each and every number would still need to be held in the database even if a decision is taken that destination information is to be stored only for some subset of numbers.

If the decision is taken that the common numbering database contains destination information for some sub-set of numbers, then traditional block routing will continue to be necessary for the remainder.

The Central Numbering Database has **not** been scoped to provide:

- Numbering Administration
- Number Portability "Clearing House" / "Hub" functionality.

4.1 Database Concept

Access to the database will be available to CPs and other authorised parties (e.g. law enforcement agencies and commercial service providers)

Note : The rules for determining what constitutes a CP will be a matter for UKPorting.

The read interface will facilitate a periodic bulk download model and may also offer a real-time query service.

The provisioning (B2B) interface may be defined to use ebXML or an appropriate alternative, and will allow a CP with the appropriate access key for a number to, amongst other things;

- Specify a high priority change to the record contents (e.g. to port a number or change the node on which it is hosted). Such a change could be specified to occur immediately, or at a given time in the future. The system has been specified such that the data change should propagate to all CP networks within 20 minutes.
- Specify a lower priority change to the record (e.g. to pre-provision a range of numbers that has been assigned by Ofcom). Such changes will be propagated over a defined (longer) time to smooth the level of information being propagated.
- Cancel changes that have been scheduled but not yet executed.

The read interface will facilitate both real-time and periodic bulk download models;

- For the real-time approach, the originating CP will query the central database and receive records, together with a “time to live” flag which indicates for how long it is safe to cache the data locally. This flag will be set uniformly at 12 minutes.
- For the periodic download approach, the database will be structured into multiple smaller partitions at the C-D digit boundary (e.g. 0121 5) level, termed “Sections”. This approach has been chosen to balance the desire to have sections of manageable size, with the desire to have a manageable volume of sections. CPs can choose to subscribe to some or all Sections. The CP will initially download all of the data for each subscribed Section, then be notified by the database of any subsequent changes; receipt of the notification will trigger the CP to download any incremental changes to the relevant Section of the database. For the avoidance of doubt, the mechanisms proposed are such that full downloads will be an exception (for example following catastrophic loss of local systems or data corruption), and the norm will be to download the databases once then rely on incremental change information thereafter.

Reference DNS and XML implementations of the read interface have been incorporated into the standards, but this does not preclude database providers proposing alternatives that meet the functional requirements.

Depending upon the nature of the destination, the database will contain either one or two records for each number. The record format will be different according to whether the number is served by a PSTN call model (whether traditional TDM network or PES on an NGN), or within an IMS network;

- For numbers served by PSTN, a single record will be provided.
- For numbers served by an IMS network, up to two records will be provided, one format for use where the database is being queried in the knowledge that the call must be routed by a PSTN network, and another for use where the database is being queried by an IMS network hence potentially allowing routeing of the call end-to-end across IMS networks.

Note : A technical detail is that there will actually be 2 or 4 records, because each number will have a record to provide an answer if it is dialled correctly, and a record which for if it is over-dialled (i.e. additional digits dialled).

PSTN records will consist of the number, prefixed with a PSTN Destination Group. This is an eight digit long prefix of the form 7xxxxyyy, where xxxx identifies the CP and yyy the location within the CP network. The Destination Group will provide sufficient information to allow other CPs to route calls to the appropriate point in the terminating network. IMS records will be of the form 01234567890@a123.dg.jobloggs.uktel.org.uk, where the information to the left of the “@” is the number dialled, information to the right of the “@” the IMS Destination Group.

In order to provide a transition path with existing approaches, special Destination Groups will be provided to allow embedding of Intermediate Routeing Numbers (IRNs) as used in the current MNP solution.

Where a query to the database yields no result, the default PSTN Destination Group (72000000) will be applied by the querying network, which will provide an indication to downstream networks that the database has been checked and there is no record for the number in question.

4.2 Routeing Concept

The requirements of the routeing concept and its use of the database are focussed on NGNs, but can also be applied to existing mobile and fixed networks. The approach is that UK originating networks (or transit networks acting on their behalf) will query the common numbering database.

- Where the network acting on this information is IMS-capable, the IMS-related record should be used, which will provide the <IMS_destination_group> that will be used to route the call
- Where the network acting on this information is non-IMS-capable, the PSTN related record will be used to route the call.

With the proviso that a CP can choose which record to act upon, the information received from the database is "objective"; the answer is the same regardless of the querying CP. The response to a query does not return the routeing to be used for the number, but sufficient information for the CP to properly identify the terminating CP hence allow the querying CP to select an appropriate route.

The next stage in routeing is "subjective"; the CP's policy is applied to determine how best to route the call. From the destination CP's NIPP (or equivalent) data, the originating NGN will be able to apply its own routeing policy to determine the optimal callserver and media handover to route the call to. This could be the terminating network, but could equally be a transit network. Similarly, even if the call is sent to the terminating network, this may or may not be directly to the terminating callserver, according to the CPs' connectivity and specific interconnect agreement applicable.

Note : The CP could carry out the transformation from Destination Group to next hop in real-time, or could have this information pre-provisioned within their network.

As an analogy, in TDM networks for geographic numbers the relationship that a given number range is terminated on a particular switch is objective, and the Destination Group would relate to that switch. Whether an originating network, knowing this homing, sends the call directly to that switch, to a suitable tandem switch, or to a tandem network, is subjective, and determined by the second stage.

The information from the database record, as received above, will be passed to subsequent networks to ensure that (except in the case of the following paragraph) the database isn't consulted multiple times. If the database didn't contain a record for the number in question, then a default PSTN DG will be applied to indicate that the database has been checked. This is necessary because it would be undesirable for copies of the database to be checked multiple times, lest they be out of step.

The database would be consulted a second time only in one specific circumstance. Where an originating CP network is IMS-capable and hence acted upon the IMS-related record, if the call subsequently needs to transit a non-IMS-capable network, then the database would be re-queried for the PSTN record. This is because the non-IMS-capable network wouldn't be able to correctly process the information derived from the IMS-record, meaning the call would either fail or worse, circular route. For the avoidance of doubt, the only situation where a second query should be undertaken is where the destination is in the IMS form above, but it is known that the only appropriate connectivity to the terminating IMS network is via a non-IMS network.

4.3 Security of Data

It is intended that the database will not be publicly available. An authentication scheme will exist so that CPs can request changes to information in the database securely. Each record will have an access key, known only to the CP of record to that number. Arrangements will be in place to allow CPs to release records for usage by other CPs, e.g. for portability purposes. Also, CPs can have multiple access keys each of which have access to specific groups of numbers, hence facilitating sub-allocation.

4.4 Porting Numbers

Although the definition of processes falls outside the scope of NICC, process “hooks” have had to be designed into the architecture. It is not intended that usage of the database should necessarily result in wholesale changes to porting processes, but its introduction may facilitate a review of existing approaches. The database neither precludes nor mandates the usage of a clearing house or bilateral processes. Conceptually, in a bilateral scenario the process would be:

- The customer approaches the Recipient network to have number ported.
- The Recipient network approaches the donor network with details of number(s) to be ported, together with authentication information, as agreed by process groups.
- Having verified that there are no obstacles to the port taking place, the donor network unlocks the record associated with the number to the recipient and carries out any activities associated with existing onward routeing processes (for example preparing the Rangeholder network to subsequently activate the port: this is necessary as the two schemes must co-exist until querying of the database is mandated).
- The Recipient network builds data on its network in readiness for port.
- The number port is activated by Recipient network using the access key to take ownership of the number and change the contents of record to point calls to their network.
- Within 8 minutes the database provider broadcasts a notification to CPs that have downloaded the data, and changes the database contents so that real-time queries receive the recipient network identity rather than donor.
- Within 12 minutes CPs that have local copies of the data respond to the notification, request the updated records from the database and modify their routeing tables.

4.5 Numbering Management

Nothing within the technical solution mandates that numbers be allocated in different size blocks to that currently used. However, if NICC’s proposals that all numbers be contained within the database are adopted, calls will be routed on an individual number basis. This would facilitate Ofcom issuing numbers in smaller block sizes should this be justified by a future cost-benefit analysis.

Note : It is acknowledged that this may not be feasible in the mobile space because of relationships with overseas carriers.

Similarly, nothing in this solution mandates a change to the current process whereby when a ported number is subsequently ceased, it is returned to the range-holder. However, if the database were to contain all numbers, there would then no longer be any technical need to do this, and the number could equally be returned to the pool of the CP that served the number when the customer ceased service. The management transactions specified to the database facilitate either the model of returning numbers to range-holder, or the last provider retaining rights.

Until such time as the use of the database is universal for all numbers, the absence of any number from the database **does not** indicate that it is not in use and it cannot be assumed to be unassigned.

The solution described is able to support continued use of block routeing techniques for numbers which have not been ported, but the ability to individually route non-ported numbers would not be realised.

5 Requirements and Architecture Overview

5.1 Overall Requirement

Green Release Requirement L, as modified, stated;

> For the purposes of this requirement Destination Group represents a serving communication provider's node or set of nodes for which another provider would make the same routeing decisions for all numbers within it.
>
> Separation of numbering administration and routeing **must** be supported, meaning that on an individual number basis it **must** be possible to determine the Destination Group for the number. If the customer has changed provider (number portability), this determination **must** be possible without the real-time involvement of provider to whom the number was originally allocated. Further, the architecture **must not** rely on aggregation of numbers (i.e. number ranges) to determine a particular Destination Group. Note: this requires a common shared database of number to Destination Group associations.
>
> It **must** be possible for originating call control functions, having determined the Destination Group, to route calls to the call control function serving the Destination Group. Where transit network functionality is necessary following determination of the Destination Group, sufficient information about the Destination Group **must** be provided to that transit network so that it is not necessary to query the common shared database a second time. Note: this requires publication of information by individual communications providers providing the mapping of Destination Group(s) to relevant serving nodes.

The subsequent requirements of General Condition 18 are met by Requirement L.

5.2 Functional Entities and Interfaces

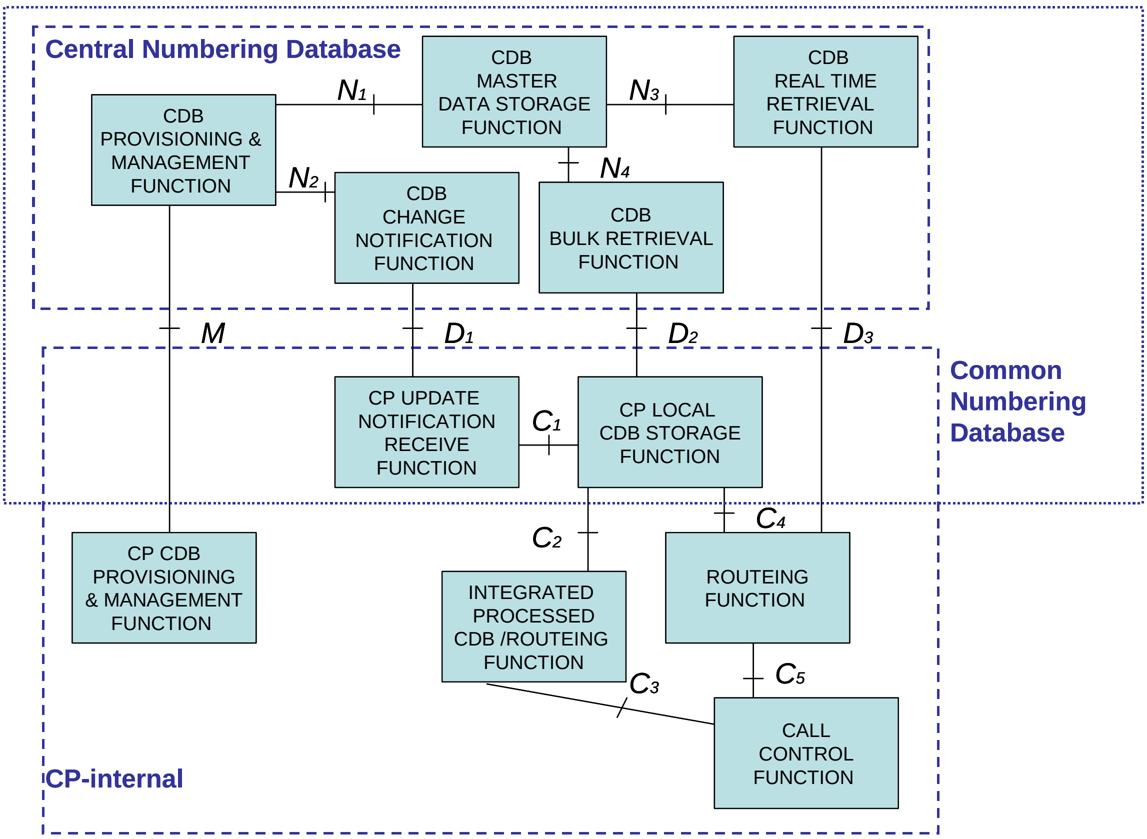

Figure 1 : Architecture

The functional architecture is depicted in Figure 1. The Central Numbering Database supports both the approach of real-time access and the approach of download of the database locally to CP networks ("reference access"). The database **will** be structured to allow CPs to choose to utilise both models according to number range.

- CPs provision and manage data in the central numbering database via Reference point M
- CPs query the data in the central numbering database via reference point D_1 and D_2 or D3
- Reference points $N_1 - N_4$ are internal to the central numbering database provider
- Reference points $C_1 - C_5$ are internal to the CP
- A CP may query the Central Numbering Database in Real Time via implementation of a Routeing Function applying local policy in real time to the result of queries to the Central Numbering Database (hence reference points C_5 and D_3),
- A CP may maintain a local copy of the Central Numbering Database and via implementation of a Routeing Function applying local policy in real time to the result of queries of the local copy of the CDB (hence reference points C_1, C_4, C_5, D_1 and D_2),
- A CP may maintain a local copy of the Central Numbering Database via implementation of a processed version of the CDB incorporating routeing information (hence reference points C_1, C_2, C_3, D_1 and D_2) or
- A CP may utilise a combination of these e.g. for different number ranges.

It could be decided that the Central Numbering Database should provide only bulk access; for example the provision of real-time access could be subject to competition between multiple providers. In this scenario, the Reference Point N_3 would not be implemented. Bulk access to the data in the Central Numbering Database in order to subsequently provide a real-time interface would be via Reference Points D_1 and D_2. This is depicted in Figure 2.

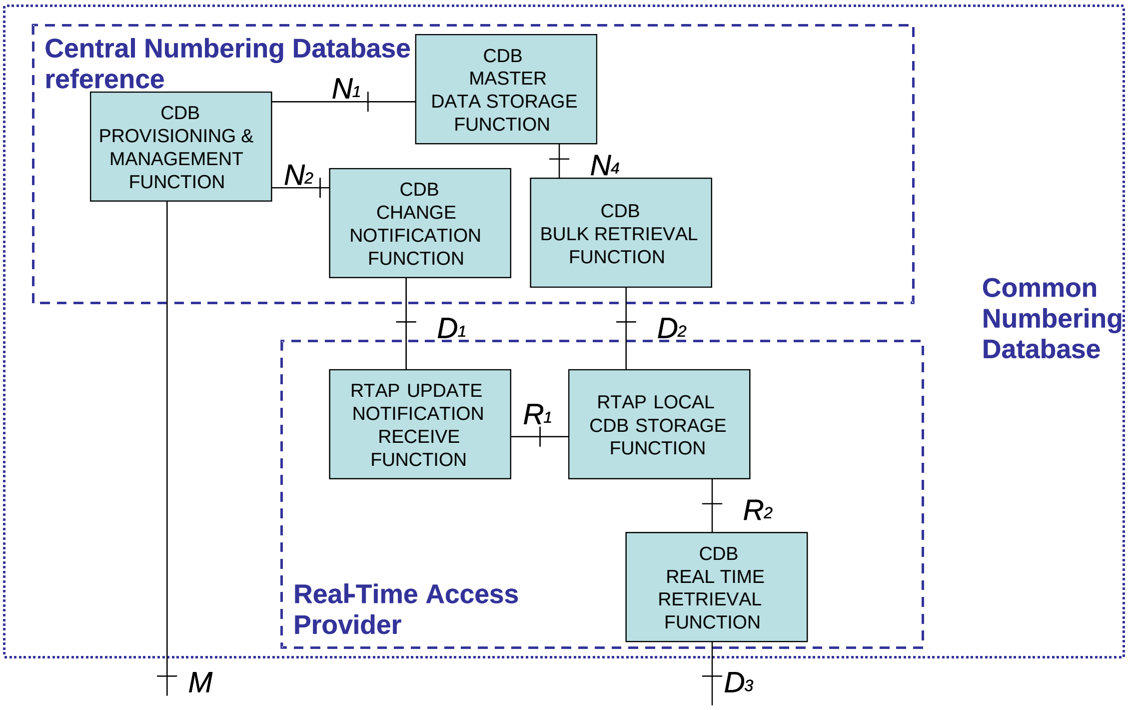

Figure 2:Splitting of read functionality of Central Numbering Database

The functions of the Real Time Access Provider Local CDB Storage Function and the Real Time Access Provider Update Notification Receive Function will be the same as those of the CP Local CDB Storage Function and the CP Update Notification Receive Function respectively.

5.3.1 Functions

5.3.1.1 CP Functions

5.3.1.1.1 Call Control Function

The Call Control Function queries the Routeing Function with the destination number, and if it has it the Destination Group, to determine the route towards the destination. It indicates to the Routeing Function whether the Common Numbering Database has previously been consulted via the presence of the Destination Group information (if present it has been consulted, if not present it hasn't).

Using the response provided by the Routeing Function, the Call Control Function passes the call onto the next hop in the call path, including the destination number and Destination Group information. As an alternative, the Routeing Function can be integrated with a processed version of the Common Numbering Database.

5.3.1.1.2 Routeing Function

The Routeing Function accepts requests from the Call Control Function, and responds with the preferred route(s) towards the Serving Control Function and sufficient information to populate signalling to subsequent Call Control Functions. Rules have been developed for the behaviour of this routeing function; see Section 8.

As an alternative, the Routeing Function can be integrated with a processed version of the Common Numbering Database so it can provide a direct mapping from the destination number to the CP's specific routeing for that number.

5.3.1.2 Common Numbering Database

The Common Numbering Database is a generic term used to describe the Central Numbering Database and all of the copies of it that are held locally by CPs.

Provision is made in the standards for the Common Numbering Database such that if it is queried with insufficient digits, a response is provided of how many additional digits are required before a valid record can be returned. Conversely, if a query is made with two many digits, the superfluous ones are discarded.

5.3.1.2.1 Central Numbering Database

The Central Numbering Database is the central reference holding mapping of telephone numbers to Destination Group. It supports both the approach of real-time access and the approach of download of the database to local CP networks. The database is structured to allow some or all of the data to be downloaded.

Bulk download : The Central Numbering Database provides a copy of its contents /changes to its contents since a given time to CPs on request.

Real-Time Access : The Central Numbering Database accepts requests from the Routeing Function and within a specified time respond with the associated Destination Group for the number. Where the Routeing Function has provided insufficient information (e.g. number length too short), the Central Numbering Database indicates the minimum number of digits required to fulfil the query. Where the Routeing Function has provided too many digits, extra digits are discarded.

Within the Central Numbering Database, separate functions deal with the management of the data (including making changes) and the download of the data via either model.

5.3.1.2.2 CP Copy of Central Numbering Database

The CP may hold a copy of central numbering database. Two approaches are possible;

- It accepts requests from the Routeing Function and responds with the associated Destination Group information for the number in real-time, or
- The CP copy of the Central Numbering Database is processed to include the CP's routeing policy and is integrated with the Routeing Function.

The Central Numbering Database must be monitored (whether by polling or receiving notifications from the CDB) and the CP copy amended within 12 minutes in case of a notified update (e.g. number activation, number port).

Use Cases have been developed for the architecture, which map the functions onto typical physical implementations. Annex B illustrates these Use Cases.

6 What the Architecture means in practice

When the database is queried, the information received is “objective”; the answer is the same regardless of the querying CP. This is achieved by the use of Destination Groups which indicate *where* a call should be terminated.

The next stage in routeing is “subjective”; the CP’s policy is applied to determine *how* best to route the call. From the destination CP’s NIPP (or equivalent) data, the originating CP will be able to apply its own routeing policy to determine the optimal callserver and media handover to route the call to. Note that this could be the destination network, but could equally be a transit network. Similarly, even if the call is sent to the terminating network, this may or may not be directly to the terminating callserver, according to the CPs’ connectivity. The CP could carry out the transformation from Destination Group to next hop in real-time, or could have this information pre-provisioned within their network.

6.1 Destination Groups

The architecture is centred around the concept of Destination Groups. A Destination Group represents the set of numbers that would ordinarily be routed in the same way. CPs may choose to route multiple Destination Groups similarly, for example a mobile CP might choose to route all Destination Groups associated with geographic calls via a common transit network, but the Destination Group represents the smallest set of destinations that can reasonably combined together in a single routeing plan.

A typical process for routeing based upon a Destination Group could be as follows;

1) A customer of CP-A dials a number (say 01234567890).

2) CP-A looks up the number in the common database and gets back a Destination Group.

For example, for a PSTN destination the response might be:

Tel : 7234567801234567890

…this indicates that 72345678 is the Destination Group.

An IMS destination may yield the response

Tel : 723456781234567890

(record to be used where the querying node / next hop is not IMS capable)

Sip : 01234567890@dg0086.dg.cp.uktel.org.uk

(record to be used where querying node / next hop is IMS capable)

…this indicates that the Destination Group to be used by PSTN networks is 72345678, that to be used by IMS networks is dg0086 within the network owned by provider “cp”.

3) CP-A then uses the Destination Group to determine the call routeing.

Assuming that the media gateway is determined by the call server IP/port pair, it will have a routeing database something like:

Destination Group	**Next hop**
dg0086.dg.cp1	cs43.cp1.uktel.org.uk : 4111
0087.dg.cp2	ssw9.cp2.uktel.org.uk : 51222
73456789	cs43.cp1.uktel.org.uk : 4133
75456789	ssw9.cp2.uktel.org.uk : 51888
76466789	cs43.cp1.uktel.org.uk : 4111

Note : Although the above table is IP-technology centric, the same principles would equally apply for the selection of TDM routes.

Note : Although a single entry for the next hop is shown, there could be a series of routes which are selected on an Automatic Alternate Routeing basis.

This table would be populated on the basis of interconnect agreements, their routeing policy and knowledge of where the Destination Groups are hosted.

The knowledge will be built from a combining a series of administrative databases. For example the first step would be to determine which CP owned each Destination Group using Ofcom (for PSTN) and NICC (for IMS) information. For example;

Destination Group range	**CP**
IMS cp1	Communications Provider 1
IMS cp2	Communications Provider 2
PSTN 73456	Communications Provider 1
PSTN 75456	Communications Provider 2
PSTN 76466	Communications Provider 3

NIPP files within a given CP could then highlight the optimal interconnect point for their Destination Groups, for example;

CP1 NIPP		
DG	**Callserver**	**Media Handover**
dg0086	cs43	cs43.cp1.uktel.org.uk : 4111
734567890	cs43	cs43.cp1.uktel.org.uk : 51888

CP2 NIPP		
DG	**Callserver**	**Media Handover**
dg0087	ssw9	ssw9.cp2.uktel.org.uk : 51222
754567890	ssw9	ssw9.cp2.uktel.org.uk : 51888

As the querying CP doesn't have a direct relationship with CP3, they may have determined via contractual arrangements that CP1 can provide a path to CP3, and the handover to use.

4) From their routeing database table, CP-A will make a SIP connection to the server and port and pass the original response onwards, for example:

   [connect to ssw9.CP-B.uktel.org.uk port 51222]

   INVITE

   To: Tel : 7234567801234567890

- If CP-B is the terminating operator, they will use the number to make internal routeing decisions.

- If CP-B is a transit operator, they will have their own routeing database and will use the Destination Group to trigger that to determine the next call server to contact.

It should be noted that with some very specific exceptions, CP-B must **not** re-query the database, because they should route based upon the Destination Group rather than the originally dialled number. This is important, because in the critical phase where a number is being ported, it is possible that individual CP copies of the database will be out of step; if this constraint was not put in place circular routeing could occur. If CP-B has chosen to integrate their routeing function with a copy of the common numbering database (processed to incorporate their routeing policy), then it should be stressed that because the routeing function will be acting upon a Destination Group (versus the destination number that CP-A used), the common numbering database wouldn't be consulted twice.

7 Database considerations

7.1 What numbers should be in the Database

NICC considered which numbers should be populated in the database, considering the following aspects;

- Only ported numbers (or all)
- UK only (or beyond?)
- E.164 only (or not)
- If E.164 only which number ranges ?

The conclusion was reached that NICC recommends that the scope of the common database should include both ported and non ported UK E.164 numbers of the following types -

- geographic numbers
- mobile numbers
- personal numbers
- 08/09 , including those (eg 0800) that are not strictly "true" E.164
- other non-geographic numbers

This means that for each number the CP currently responsible for that number must be specified. In addition for ported numbers the destination group must be specified. For non-ported numbers either of the following options may be used:

- The destination group is specified whether the number is active or not (equivalent to block routeing)
- The destination group is specified only for active numbers (equivalent to Individual number allocation).

In reaching this conclusion, the various aspects were considered, as outlined in Annex C. It should be noted that the NICC standards do not preclude populating the database with a subset of numbers (e.g. just ported ones) for routeing, but this would lead to network inefficiencies. Were a sub-set of numbers to be populated, it must be highlighted that at least administrative details of all the remaining ones would be required, in order that it could be established that a subsequent attempt to populate a number for routeing (e.g. when a number became ported) was legitimate.

7.2 Contents – What should be in the records

Depending upon the nature of the destination, the database will contain either one or two records for each number. The record format will be different according to whether the number is served by a PSTN call model (whether traditional TDM network or PES on an NGN), or within an IMS network;

- For numbers served by PSTN, a single record will be provided.
- For numbers served by an IMS network, up to two records will be provided, one format for use where the database is being queried in the knowledge that the call must be routed by a PSTN network, and another for use where the database is being queried by an IMS network hence potentially allowing routeing of the call end-to-end across IMS networks.

Note : A technical detail is that there will actually be 2 or 4 records, because each number will have a record to provide an answer if it is dialled correctly, and a record for if it is over-dialled (i.e. additional digits dialled).

Also, in order to match with existing approaches, the format of records for mobile numbers will vary slightly from those for other number ranges.

The standards assume a DNS implementation of the database, but this does not preclude database providers proposing alternatives that meet the functional requirements.

The formal specification of the records is set out in Annex A of ND1631 [1]. Using DNS, the format of the records will be as follows;

<u>Calls to PSTN/PES customers:</u>

(Fixed lines)

tel: uri, containing <PSTN_destination_group><number>
e.g. tel : 72345678 01234 567890

(Mobile numbers)

In order to provide legacy support for Mobile Number Portability according to ND1208, a range of Destination Groups starting 72007 have been designated. These align with the IRN prefixing scheme such that the Destination Group72007pqr may be directly mapped to the IRN digits pqr.

tel: uri, containing <PSTN_destination_group><number>
e.g. tel : 72007678 07957 123456

Where legacy IRN support is not required, Destination Groups other than those starting 72007 should be assigned. These will be identical in format to those used for fixed lines and the fixed/mobile differentiation ceases to exist.

<u>Calls to IMS customers</u>:

(Fixed lines)

TWO RECORDS… tel: uri, containing <PSTN_destination_group><number>

e.g. tel: 72007678 1234 567890

(note the PSTN_destination_group - 72007 is reserved for embedding mobile IRNs)

AND sip: uri, containing <number>@<IMS_destination_group>.dg.<CP>.uktel.org.uk

e.g. sip:1234 567890@a123.dg.jobloggs.uktel.org.uk

(Mobile numbers, ported)

TWO RECORDS… tel: uri, containing <PSTN_destination_group><number>

e.g. tel : 72007678 07957 123456

AND sip : uri, but format to be determined

(Mobile numbers, not ported)

TWO RECORDS… tel: uri, containing <PSTN_ destination_group><number>

e.g. tel : 72000000 07957 123456

AND sip : uri, but format to be determined

where

<PSTN_destination_group> is as set out in the definition

<number> is the queried telephone number

<IMS_destination_group> is as set out in the definition

and <CP> identifies the communications provider.

Where a query to the database yields no result, a specified PSTN_destination_group (72000000) will be applied by the querying network, which will provide an indication to downstream networks that the database has been checked and there is no record for the number in question.

It should also be noted that for mobile a special range of Destination Groups, 72007xxx, is used. These embed the existing IRNs and hence allow a translation to this form should network equipment not support the longer number lengths necessitated by Destination Groups..

7.3 Should the database be accessed in real-time or only administrative

NICC considered whether the database should be solely a real-time one, solely an administrative one, or a hybrid between the two.

Larger CPs agreed that real time queries to any routeing database should take place within their own networks. This is because it retains responsibility for network performance (e.g. availability, response time) in the hands of the CP itself. This implies that the common database should support an administrative model of downloading the whole dataset.

However, it was acknowledged that for small CPs, the task of downloading the full database could be onerous, in which case it would be preferable to either query the database in real-time, or potentially query a copy of the database held by a third party. This implies that the common database should support a real-time model.

Further, some CPs could wish to query part of the database in real-time, but have a local copy of other parts. For example a mobile CP could choose to locally download the 07 part of the database, but not the other parts.

It was therefore agreed that a hybrid model be adopted, allowing individual CPs to download the database or to query it in real time.

7.4 Potential physical implementations of the database

Independent of the considerations of which numbers to populate and whether the database should be real-time or administrative, NICC also examined various physical architectures for the database. The architectures considered and the associated advantages and disadvantages are summarised in Annex D.

NICC concluded that the preferred model would consist of having a central UK level common numbering database, authoritative for all numbers. This means that all records for the numbers would be centrally located and managed. Such an approach did not preclude storing local copies of the database for reference access, but it was important to note that these acted as slaves of the central one.

The database will initially be populated with records for each number, providing Destination Groups relating to the rangeholder. A number could then subsequently be ported by modifying the Destination Group to relate to the recipient network.

7.5 How the database is structured, managed and queried in practice

7.5.1 Access

CPs may access the database by:

- Querying a copy of the data hosted locally on their systems. This model is sometimes called the ‘admin’ or ‘push’ database;
- Querying the central database as and when required. This model is sometimes called the ‘real-time’ or ‘pull’ database;

Or a combination of the two (for example by different number types).

7.5.2 Sections

The database is partitioned into sections to allow the data to be transferred in parts and to allow the operator to only keep a copy of that data they believe is most relevant to them. These sections are divisions of the number space. NICC has concluded that the sections will be at the C-D boundary (e.g. 0121 5) level; this was a decision driven by the need to balance the size of the individual partitions against the number of partitions. The logic NICC used in making this decision is described in Annex E.

7.5.3 Querying a copy of the data hosted locally on CP systems

CPs may wish to take copies of data in one of two ways, an incremental transfer of changes or a full load.

Incremental

The CP specifies a section and a checkpoint reference and asks the central database for all the changes since then.

- Each set of changes contains a new checkpoint reference.
- When a new set of changes is available, the central database notifies the CP so that the CP can then request those changes if they wish. However the operator may choose to ignore the notification if it so wishes.
- The address to which the central database sends the notification need not be the same address as the system that requests the incremental transfer.
- Incremental changes will only be available for download for a limited period (T14a) after which a full load must be carried out.

Full load

The CP requests a full copy of a section of the database and receives it in reply. Each full copy uses the same set of checkpoint references as incremental loads so that the operator can then switch to using incremental loads to keep the data up to date. The central database only needs to provide the current copy of the database at the time of request, not any historical information.

7.5.3.1 Format and compression of bulk download data

All initial work assumed a DNS/ENUM implementation and that the bulk download would be achieved using standard DNS functionality, as a DNS/ENUM type query is the most likely format a CP will use.

Some CPs identified that an XML format bulk download was more suited to data manipulation prior to population of their local ENUM server and it was agreed that both DNS and XML options should be supported. ND1631 [1] describes a use case for each, reproduced in Annex B of this document.

By splitting the number sections of the CDB on C-D Digit boundary, each section can have up to 1,000,000 numbers stored within it. Given the potentially verbose nature of XML, NICC considered the best practices that should be adopted in order to keep the size of the bulk XML download of a fully populated number section to a manageable size, that is from the perspective of the bandwidth required by each CP to download the XML and the subsequent parsing of the downloaded XML by the CP. For incremental downloads, NICC considered that the file sizes would not be of a significant enough size to warrant specific analysis but would apply the same best practices to them to ensure consistency with the bulk download.

NICC considered two distinct aspects of XML file compression, namely transport compression (the application of a compression mechanism agnostic to the file contents themselves) and format compression (the design of the file content such that repetitive information is kept to a minimum).

In addition to using no format compression at all, NICC considered the following two types of format compression

Range Compression

Ranges are denoted by the first and last number of a contiguous block of numbers that have the same destination group. With this compression technique the destination group can be repeated multiple times within the XML file.

Batch Compression

Numbers are batched together by grouping numbers that have the same destination group, and then denoting ranges within the batch by the first and last number of a contiguous block of numbers. With this compression technique the destination group appears only once within the XML file.

The conclusions reached by NICC were that transport compression (using a compression utility like gzip [www.gzip.org]) provided significant benefit when downloading the bulk XML file and that format compression should also be used to limit the size of the XML file that the CP must parse. Providing the ranges were in ascending order, NICC decided on the use of range compression for the following reasons.

- It is easier to detect missing entries with range compression.
- It is easier to detect duplicate numbers with range compression.
- Because of the above, range compression is easier for humans to check (always important for debugging).
- It is easy to extract a number sub-range using range compression.
- The range compression presentation is more suitable for the purpose of number--> destination group translation, and easier to store in most forms of real-time

To further limit the size of the bulk XML file before (and hence after) transport compression, NICC also recommend that the use of white space and long tag names are kept to an absolute minimum.

To assist NICC in reaching these conclusions, a few porting scenarios were simulated in order to gauge the impact of both transport and format compression on the bulk XML download of a fully populated number section. The results of these simulations can be found in Annex F.

7.5.4 Querying the central database as and when required.

CPs may wish to query the central database for information as and when required.

- The CP issues a query to central database for single number.
- Each response contains with it a time to live that tells the CP how long the data is valid for. This is necessary so that the CP may cache data, but still route calls to a new CP within appropriate time should a routeing change be made to the database.
- Every query must be replied to within 20 milliseconds (excluding transport delays). It is necessary to constrain the time taken for the database to respond in order that Post Dial Delay is not impacted.

7.6 Provisioning the database

Regardless of any decision that the CDB will support all or a sub-set of numbers, it is necessary that the CDB is provisioned with all numbers from an administrative view. This is necessary so that a CP may e.g. relinquish control of a number such that a recipient CP can import it.

It is a separate decision as to whether Destination Group information is provisioned for all numbers or some sub-set.

7.6.1 The management Interface

Processes which require changes to the database can be accommodated by using transactions defined at the Management ("M") reference point. These transactions are specified in Section 9 of ND1631 [1].

A secure B2B communications interface is required between the CP and the Central Numbering Database.

The B2B interface might be achieved using the framework (Automated Business to Business (B2B) Transactions: Architecture and Principles) and NICC ebXML profile or by an alternative protocol capable of providing the necessary functionality, security, resilience and non-repudiation. However only one protocol is to be used for the B2B interface.

The B2B interface will be specified in a separate document that is specific to that protocol in which the Information flows across Reference Point M are elaborated. It is expected that design of the data structures needed is most likely to be done using XML. Consideration of the information flows in NICC documents will allow a set of XML schemas to be designed that implement the information elements and their associations in such a manner as to allow efficient implementation.

The signatures process used for digital signing of messages should be based on a known robust end to end digital signing process

7.6.2 Number Groups

A CP may define a number to be a member of a specific number group. Any number not assigned to a specific number group will be contained in the CP's Default Number Group.

Number Groups are private to each CP and provide a mechanism for a CP to manage access to its numbers by different administrative groups within the CP. A CP with a series of numbering departments could choose to associate a number group with each, each having its own access privileges. Similarly, a CP that sub-allocates numbers e.g. to resellers could choose to have a number group associated with each of the resellers

7.6.3 Access Keys

Access Keys will be required for the digital signatures applied to each transaction. Each CP will be able to create a series of access keys.

These will include Root Keys, Intermediate Key Signing Keys - typically allocated to each numbering group within a CP - which permit the signing of Application Keys. It is the Application Keys which permit the operation of the database provisioning functions.

Key management functions are included in the Central Numbering Database.

7.6.4 Typical Processes

NICC is not responsible for defining process and the CDB has been defined to be, so far as possible, process agnostic. By considering present processes and possible future requirements, appropriate "process hooks" are supported.

Processes which require changes to the database can be accommodated by using transactions on the Management ("M") interface. These transactions are specified in Section 8 of ND1631 [1]. The following functions are provided by the CDB (not exhaustive)

Transaction	Summary of Usage
Allocate	Used to reserve numbers as being under the control of a CP
Permit change of ownership	Used to allow a recipient CP to be given access rights to a number by a donor CP
Take Change of ownership	Used by a recipient CP to take access rights to a number from a donor CP
Upload	Used by a CP to change the contents of records on a non-urgent, urgent or timed basis
Cancel transactions	Used to cancel all timed transactions against a telephone number.
Delete number	Used to delete a number from the active database on either a temporary or permanent basis
Access rights query	Used to determine the CP(s) with access rights to a given number
Status and History query	Used to determine the current status and transaction history of a given number
Set Notified Servers	Used to register the servers which will receive copies of the CDB
Subscribe	Used by a CP to subscribe to receive notifications of changes to a section of the CDB
Unsubscribe	Used by a CP to unsubscribe from receiving notifications of changes to a section of the CDB.
List subscriptions	Used by the CP to retrieve a list of the sections of the CDB they are currently subscribed to receive changes for.
Register key	Used by a CP to register security keys.
Revoke key	Used by a CP to revoke a security key.
Retrieve Key	Used by a CP to retrieve other CPs' public key details
List sections	Used by a CP to retrieve a list of all sections of the CDB
Audit	Used to audit contents of database – note that this function is not intended to meet or replace CP obligations for providing Ofcom Number Audit data.

These transactions may be incorporated into processes as required and the following shows how they might be used in order to achieve the desired functionality.

7.6.4.1 Numbers are allocated by Ofcom

Numbers assigned by Ofcom will be accompanied by a an allocation certificate. The CP will then use this to set the access rights for some or all of these numbers in the Central Numbering Database to one of their number groups using the ALLOCATE transaction. This number group will be under the control of a set of access keys, as described in Section 7.6.3. If the CP chooses to associate only a subset of the numbers allocated by Ofcom to the number group in question, then the Central Numbering Database will provide a residual certificate to the CP for it to subsequently use to associate the remaining numbers with number groups.

7.6.4.2 CP provisions numbers

CPs can choose whether to add Destination Group information to numbers in the database as part of the cycle of provisioning customers, or pre-provision to route traffic to a node in their network in advance of assigning the numbers to customers. Either way, the CP adds the Destination Group information using the UPLOAD transaction.

7.6.4.3 CP reconfigures their network

Where a CP chooses to reconfigure their network so that a number is hosted on different nodes hence different Destination Groups, the Destination Group information in the database is changed using the UPLOAD transaction.

7.6.4.4. Number is ceased

Where a CP ceases a number (e.g. because the customer no longer wishes to have service on that number), then the DATA UPLOAD transaction may be used to remove the Destination Group information. This removes the Destination Group information from the database, but leaves the access rights untouched.

If it is required to completely remove a number from the database, the DELETE transaction is used.

7.6.4.5 Number is ported between networks

It is not intended that usage of the database should necessarily result in wholesale changes to porting processes, but its introduction may facilitate a review of existing approaches. Conceptually, the process would be:

- The customer approaches the recipient network to have number ported.
- The Recipient network approaches the donor network with details of number to be ported, together with authentication information.
- Having verified that there are no obstacles to the port taking place, the donor network unlocks the record associated with the number to the recipient. To accomplish this, the donor network will issue a PERMIT CHANGE OF OWNERSHIP transaction, which will unlock the records associated with the number to the recipient network.
- The Recipient network builds data on its network in readiness for port.
- The number port is activated by recipient network using the access key to change the contents of record to point calls to their network. The recipient network accomplishes this via a TAKE CHANGE OF OWNERSHIP transaction, which will change the Destination Group to which the number is routed. This TAKE CHANGE OF OWNERSHIP could be implemented immediately, or at a point in time dictated by the recipient.
- Within 8 minutes, the database provider broadcasts a notification to CPs which have subscribed to the database section, and changes the database contents so that real-time queries receive the recipient network identity rather than donor.
- Within 12 minutes, CPs that have local copies of the data respond to the notification, request the updated records from the database and modify their routeing tables.

Figure 3 shows the time line for the porting of a number and how the process functions (shown in red) align with the database changes.

Time T_2 (2 minutes) defines the period during which database processes the Permit Change of Ownership instruction from the donor CP; by the expiry of T_2, the database is ready to receive a Take Change of Ownership instruction from the recipient CP. The point in time at which the recipient CP actually issues the Take Change of Ownership is a process consideration and may, if appropriate to the process, be many days after the Permit Change of Ownership.

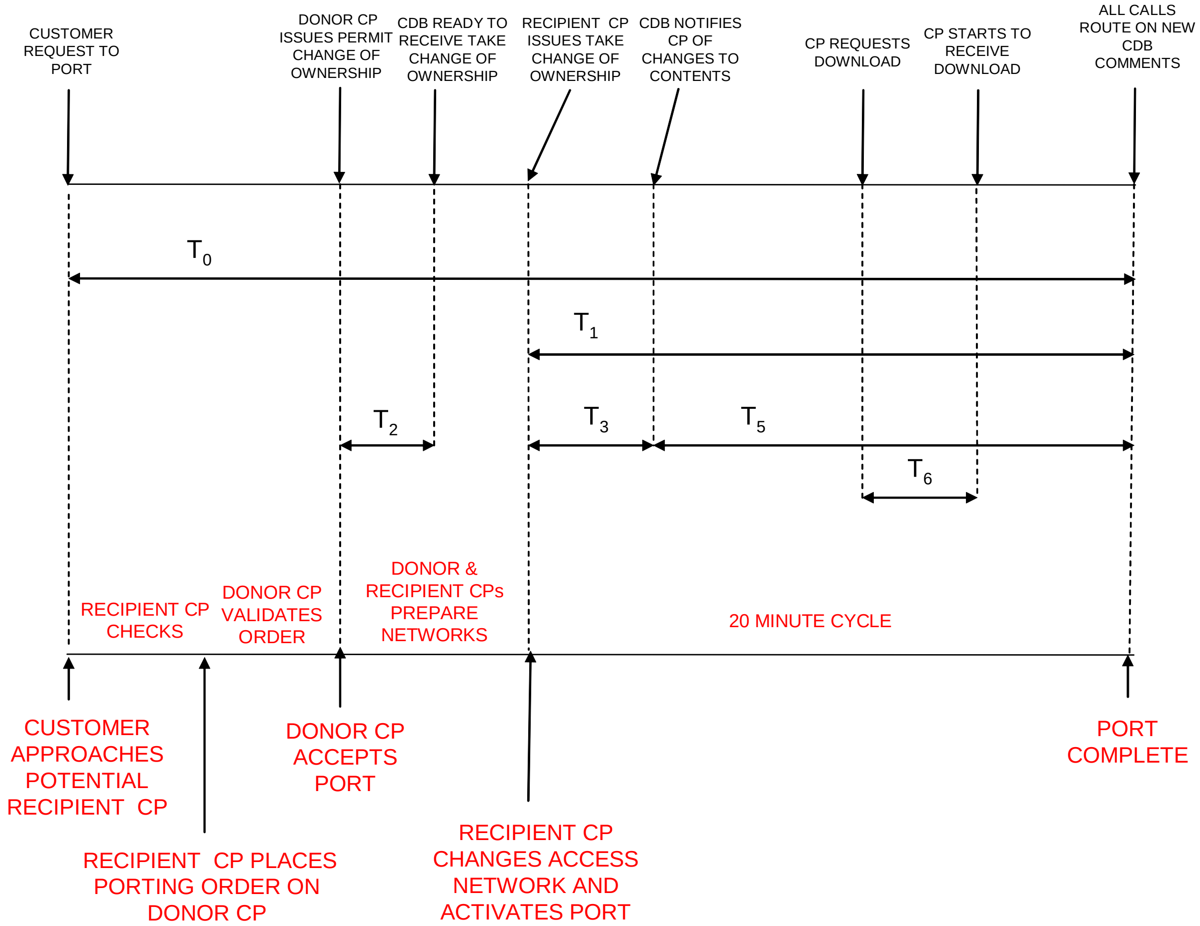

Figure 3: Porting time line.

7.6.4.6 Ported number is ceased

Where a number that has been ported is ceased. The architecture supports both a model that the number is returned to the CP originally assigned it by Ofcom, or that the number is retained by the last CP that provided service on it.

For the model where the number is returned to the CP originally assigned it, then PERMIT CHANGE OF OWNERSHIP and TAKE CHANGE OF OWNERSHIP transactions are used to return the number(s) to the range-holder.

For the model where the number is retained by the current CP, then the UPLOAD transaction may used. This removes the Destination Group information from the database, but leaves the access rights untouched.

7.6.4.7 Numbers are reclaimed by Ofcom

Where numbers are reclaimed by Ofcom, this is accomplished by a DELETE transaction being issued either by the CP or Ofcom. This removes all data associated with the numbers from the database (except audit information).

As the ultimate number administration authority, Ofcom is provided the capability to forcibly recover numbers should a CP refuse or be unable to cooperate.

8 Routeing Rules

Whilst NICC would not seek to restrict commercial flexibility, it is critical that CPs adhere to an agreed set of routeing practises in order that circular routeing is avoided. Therefore, a set of routeing rules has been devised, setting out when the database should and should not be queried, and the actions to be taken for various responses from the database.

These rules defined in Annex B of ND1631 [1].

Annex A : Studies that determined that NICC Standards should incorporate individual number routeing, i.e. All Call Query

A1 Introduction

This Annex sets out the architectures for routeing which could be implemented to support number portability between NGNs. It should be noted that this Annex is a historic record of NICC discussions, hence some of the number formats depicted are not as finally agreed.

At this stage, E.164 naming schemes are assumed, with SIP-I / SIP addressing between networks; NICC will examine the use of other addressing schemes in the future. The advantages and disadvantages of each scheme are examined, with reference to routeing efficiency, equipment capability requirements, ability to manage network resources, accounting requirements and regulatory requirements.

A2 Background

A multitude of routeing architecture models can be envisaged. In legacy networks, a choice had to be made between All Call Query, Call Drop-back, Query on Release and Onward Routeing techniques: UK network operators agreed to use Onward Routeing for portability of geographic and non-geographic numbers, and a variation on this, Signalling Relay, for portability of mobile numbers. With the introduction of NGNs, there is the opportunity to revisit these decisions based upon new capabilities via SIP signalling and the usage of DNS.

A3 Number Portability Techniques

A series of models for the treatment of number portability can be envisaged, which will be examined in this section. Fundamentally, the choice of number portability technique is one which must be agreed at a national level.

A3.1 Block routeing techniques

A3.1.1 Onward Routeing : continued usage of prefixes

A3.1.1.1 Solution overview

This model foresees a continuation of existing routeing based on the significant digits of the dialled E.164 numbers. Support of number portability under this model would continue using existing arrangements; as such, calls would route to the donor network and then be onward routed to the recipient network. Work within the NICC has already standardised the message format which would be used, in particular for the carriage of the 5xxxx prefix in the second leg of the call. Figure A..1 illustrates the call routeing.

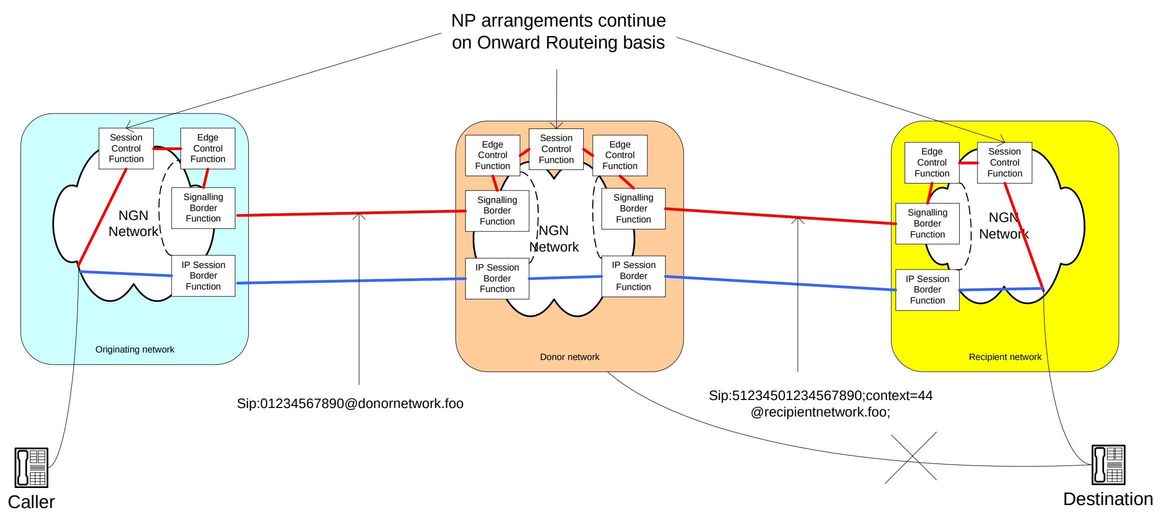

Figure A.1: Number Portability Model A : Onward Routeing

For MNP, 5xxxxx prefixes are not used, rather a prefix, for example of the form 0799x (where x defines the recipient network operator) is used, with the S-digit 7 omitted from the destination number. The call flow would therefore be as depicted in Figure A..2.

Note : other Mobile Routeing Codes are used as well as 0799.

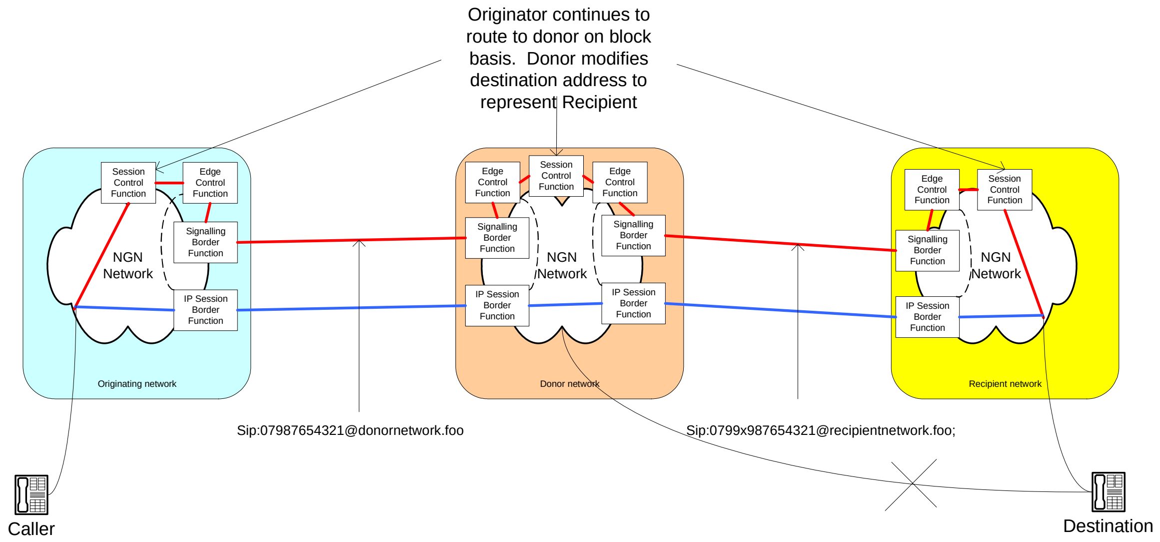

Figure A.2 : Number Portability Model A for mobile numbers

A3.1.1.2 Advantages & Disadvantages

Advantages	**Disadvantages & Issues**
• Allows "lift and shift" replacement of existing networks with minimum change to databuild and commercial arrangements • Bilateral interchange of porting data as per existing regime – so no change to processes • Since it is based around existing TDM mechanisms, would work well in mixed TDM/NGN environment • Does not require establishment of third party to administer central database	• Sub optimum routeing • Does not address the "telco failure" argument • Does not align with IMS architecture • Presents potential QoS difficulties in mixed TDM/NGN architectures where e.g. originating and terminating networks are NGN, but donor is TDM. This is particularly the case where a call path is "complex", e.g. concatenating GNP and NGNP, or CPS and number portability • The MNP solution has been devised to cope with portability between mobile networks, and hence can only be used for numbers from the 07 range.

A3.1.1.3 Conclusions

This approach emulates the current TDM solution, hence represents no change to the existing number portability solution both in the context of processes and legacy network support. Given the short timescales in which NGN interconnects must be introduced, it represented the most realistic solution for the "purple" release for NGN interconnect launch.

A3.1.2 Onward Routeing : Enhanced Onward Routeing via Modification of Destination SIP header by Donor Network

A3.1.2.1 Solution overview

This model enhances the simple block routeing architecture by utilising SIP techniques to enable number portability. Under this model, rather than using a 5xxxxx prefix to route the call to the recipient network, the donor network would simply target the correct recipient callserver. As such, the association between the E.164 number to the left and callserver name to the right of the "@" in the SIP address would yield that the number is ported: no prefixes would be required. Figure A.3 and Figure A.4 illustrate for fixed and mobile numbers respectively.

NB it must be understood that in this model, calls (both media and signalling) would still flow through the donor network. In theory, it may be possible for the final media stream to go directly from the originating to recipient network; however, it is not clear if this will be possible in practise and further study is required.

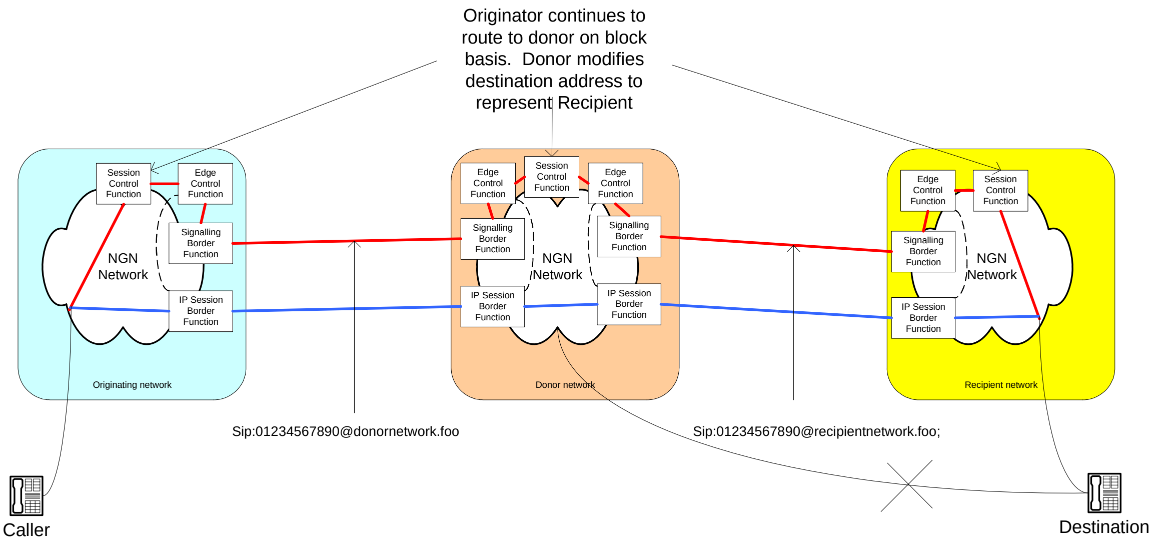

Figure A.3 : Number Portability Model B : Enhanced Onward Routeing

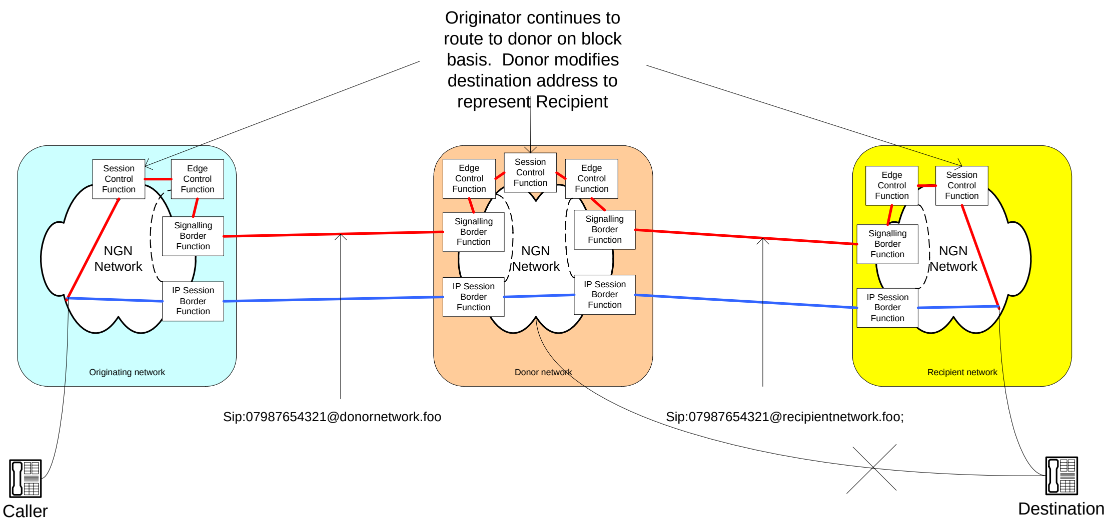

Figure A.4 : Number Portability Model B applied to mobile networks

A3.1.2.2 Advantages and disadvantages

Advantages	**Disadvantages & Issues**
• Allows minimal changes to commercial arrangements • Removes the need for potentially unnecessary porting prefixes • Bilateral interchange of porting data as per existing regime – so no change to processes and does not require establishment of third party to administer central database	• Sub optimum routeing • Risk of loops if routeing data in donor and recipient are not aligned. However, there is potential to overcome this if some form of indicator is introduced to show that the server identity has been modified for NP purposes. • Lack of NP prefix in signalling could have commercial implications (it is assumed that if this model were to be implemented, the commercial model would remain "as is", i.e. onus is on donor to reroute calls but with some scope for cost recovery) • Does not address the "telco failure" argument. It may be possible to overcome this by having an escrowed database of numbers associated with failed telcos, but this would require administrative and commercial arrangements to implement. • Solution only applicable where there is known to be an NGN path between donor & recipient. • Does not align with IMS architecture • As portability information is contained solely in the SIP header (and not within the encapsulated ISUP), it is unclear which should prevail where there is a discrepancy.

A3.1.2.3 Conclusions

For an existing porting relationship, this approach offers no material advantage over the existing onward routeing solution. However, for new entrants, there is the potential advantage that Service Establishment processes would be simplified as there would (in principle) be no need for setting up of NP prefixes.

There are significant issues that would arise with the introduction of this technique, in particular that there is scope for call looping in a mixed TDM/NGN environment, unless a "ported" indicator was introduced.

Given that the solution does not result in any more efficient routeing than the existing onward routeing approach to offset these issues, it is not recommended for any further study.

A3.1.3 Block Routeing enhanced with SIP-redirect techniques

A3.1.3.1 Solution overview

In a similar manner to Model B, this model builds upon conventional block routeing techniques to make use of the capabilities of SIP signalling. In this case, rather than modifying the signalling through to the recipient, the donor network sends a SIP redirect message to the originating network, which results in the call being sent to the recipient network. As such, it can be considered to be a next generation derivative of call-dropback. Compared to Model B, this model does not result in the final call routeing via the donor network. Figure A..5 illustrates.

Two variations to this model exist, the first being that the answer received from the donor network would have the <user> component set as the destination number (as depicted in Figure A..5), the other being that the <user> component has both the number portability prefix and destination number concatenated, and that this information is additionally mapped by the originating network into the ISUP payload.

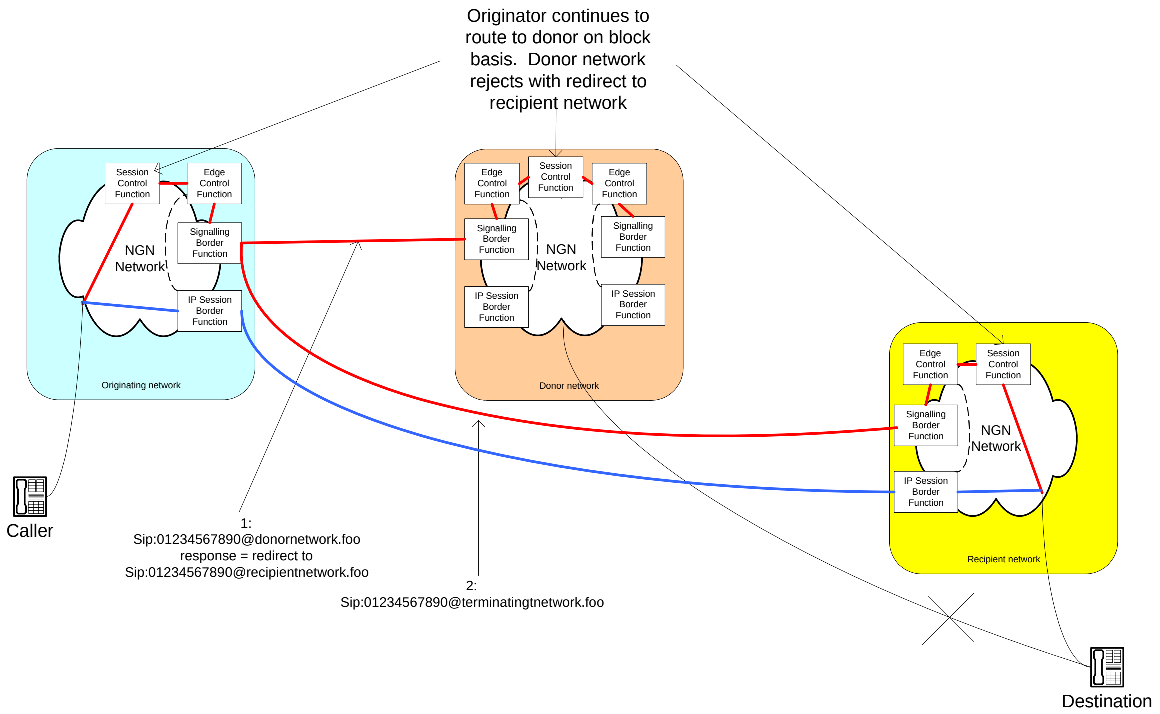

Figure A.5 : Number Portability Model C, SIP redirects

A3.1.3.2 Advantages & Disadvantages

Advantages	**Disadvantages & Issues**
• More efficient routeing • Bilateral interchange of (per customer) porting data as per existing regime – so no change to processes and does not require establishment of third party to administer central database • If data was escrowed, potential for some resolution of "telco failure" since all originating networks would be able to route to all terminating/recipient networks.	• Risk of loops if routeing data in donor and recipient are not aligned. However, there is potential to overcome this if some form of indicator is introduced to show that the server identity has been modified for NP purposes. • Does not completely address the "telco failure" argument • Changes required to commercial arrangements to incentivise/mandate originating operator to reroute traffic • Model where <user> field contains only the destination number will not work in mixed TDM/NGN environment – donor network would need to know that originating, recipient, and any intervening networks were NGN capable. Model where <user> field additionally contains the portability prefix and modifies the ISUP payload *could* work in mixed TDM/NGN environment – but donor network would need to know that originating network was NGN capable, and all networks would require knowledge of portability prefixes. • Does not align with IMS architecture • As portability information is contained solely in the SIP header (and not within the encapsulated ISUP), breaks the paradigm that where the two differ the ISUP information should prevail

A3.1.3.3 Conclusions

This approach presents some advantages in the area of telco failures, and results in marginally more efficient routeing (the final call uses the most efficient routeing, but a call attempt is made to the donor network first).

However, there are significant disadvantages; once again there is scope for calls to loop, and networks will need to have knowledge of the capabilities of upstream networks (i.e. whether they are able to process a SIP redirect message).

Therefore, once again this solution is not recommended for further study.

A3.2 Shared Database Techniques

This section examines routeing architectures which make use of shared databases.

Two approaches utilising shared databases are possible, namely that there is a single common central database, or alternatively that there are multiple shared databases. This section examines the issues associated with each approach.

A3.2.1 Single Common Database

A3.2.1.1 Solution overview

This model requires the presence of an industry-standardised common database. This database would provide authoritative data for the mapping of E.164 telephone numbers to the callserver that hosts the associated number. Ideally, the central database could be standardised on a global basis. However, if this is not possible, then standardisation at a European level would be appropriate, or as a default, agreement between UK NGN providers. A routeing model using a common database would be based on the originating network querying this database to get an objective response of where the number is hosted, then using its internal routeing logic to make the subjective decision of the best way to route to that location. As such, routeing would be based upon the terminating callserver identity, not the E.164 number itself. Figure A..6 illustrates this routeing model, where the central database was implemented via Carrier ENUM.

It should be noted that because NAT is assumed at network boundaries, although the response back from the central database that a given number is hosted on a given callserver would be universal, each NGN network would need to do their own DNS lookup to map that callserver SIP address to an IP address, because the particular edge point to target in their network would be network specific.

Two variations to this model exist, the first being that the answer received from common database having the <user> component set as the destination number (as depicted in Figure A..6), the other being that the <user> component has both the number portability prefix and destination number concatenated, and that this information is additionally mapped into the ISUP payload.

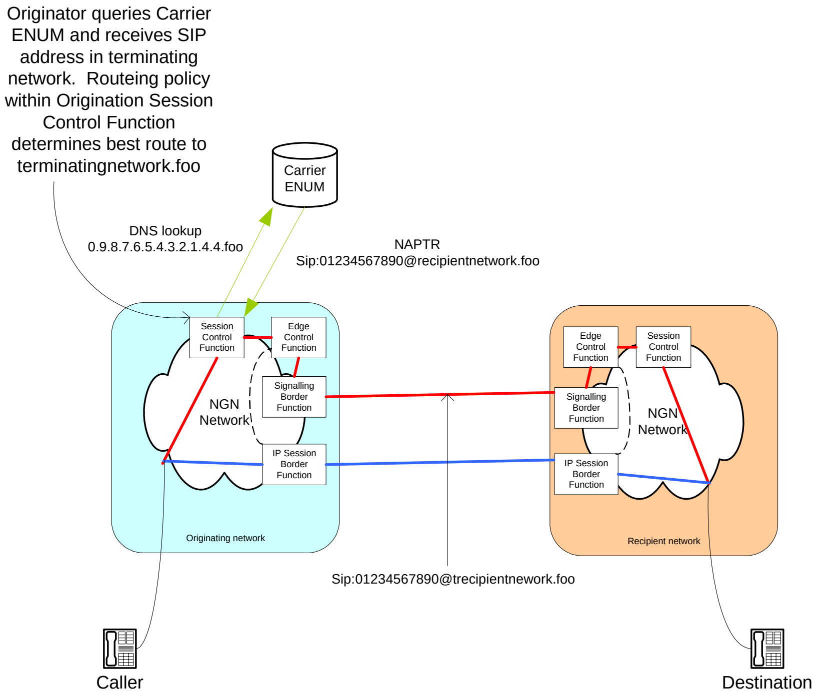

Figure A.6: Number Portability Model D, Common Database (illustrated using Carrier ENUM)

A3.2.1.2 Advantages & Disadvantages

Advantages	**Disadvantages & Issues**
• More efficient routeing • Potentially allows global scheme to develop • Avoids "telco failure" problem • Aligns with IMS architecture	• Requires 3rd party infrastructure to be put in place to common standards; this must be implemented in advance of even a small-scale rollout. • Requires complementary legal / commercial framework to be put in place • Performance of common database outside control of individual operators • Agreements would be required with respect to the ownership of data in the common database • Whole set of new NP operational processes required

A3.2.1.3 Conclusions

On an initial assessment, this approach appeared to offer the most optimal routeing. However, NICC concluded that this would require confirmation by an economic review, the particular issues being transition costs and the incremental processing required on originating callservers versus savings made via more efficient routeing.

A3.2.2 Multiple Shared Databases

A3.2.2.1 Solution Overview

This model does not assume a single authoritative central database. Rather, multiple shared databases would arise, representing federations of NGN operators that agree to co-operate. As such, the response from a shared database in this model would not give an authoritative mapping to the relevant terminating callserver: instead, it would give a mapping to a callserver which would provide a routeing towards the terminating line. Different shared databases would result in different answers for a given E.164 number. Indeed, an originating NGN operator could choose to query multiple shared databases in parallel and use the result which best fitted their commercial/routeing policy.

Figure A.7 illustrates the concept, where the shared databases are implemented using ENUM-like technologies;

- Caller 1 is connected to Originating Network 1. The shared database that it queries has had the entry for 01234567890 populated by the donor network, in the role of transit network. Based upon the response to shared database, Originating Network 1 routes to the donor, which then uses standard number portability techniques to route the call to the recipient (this could either be Model A, or Model B).
- Caller 2 is connected to Originating Network 2. The shared database that it queries has had the entry for 01234567890 populated by Transit Network 2, indicating that it can offer a route to the number. In reality, Transit Network 2 is unaware that the number is ported, hence routes the call to the donor, which is responsible for routeing the call to the recipient. Although the call paths for Caller 1 and Caller 2 are similar, the accounting arrangements between networks may well differ, because in the first case the donor has actively marketed that it can provide a routeing, whereas in the second it has been selected by default.
- Caller 3 is connected to Originating Network 3. The shared database that it queries has had the entry for 01234567890 populated by the Recipient Network. As such, by acting on the resultant record the call is routed directly to the recipient network.

As with the central database approach, two variations to this model exist, the first being that the answer received from the shared database having the <user> component set as the destination number the other being that the <user> component has both the number portability prefix and destination number concatenated, and that this information is additionally mapped into the ISUP payload.

It should be noted that although this is presented as a single number portability model, this architecture brings in elements of how transit networks will be dealt with in NGNs.

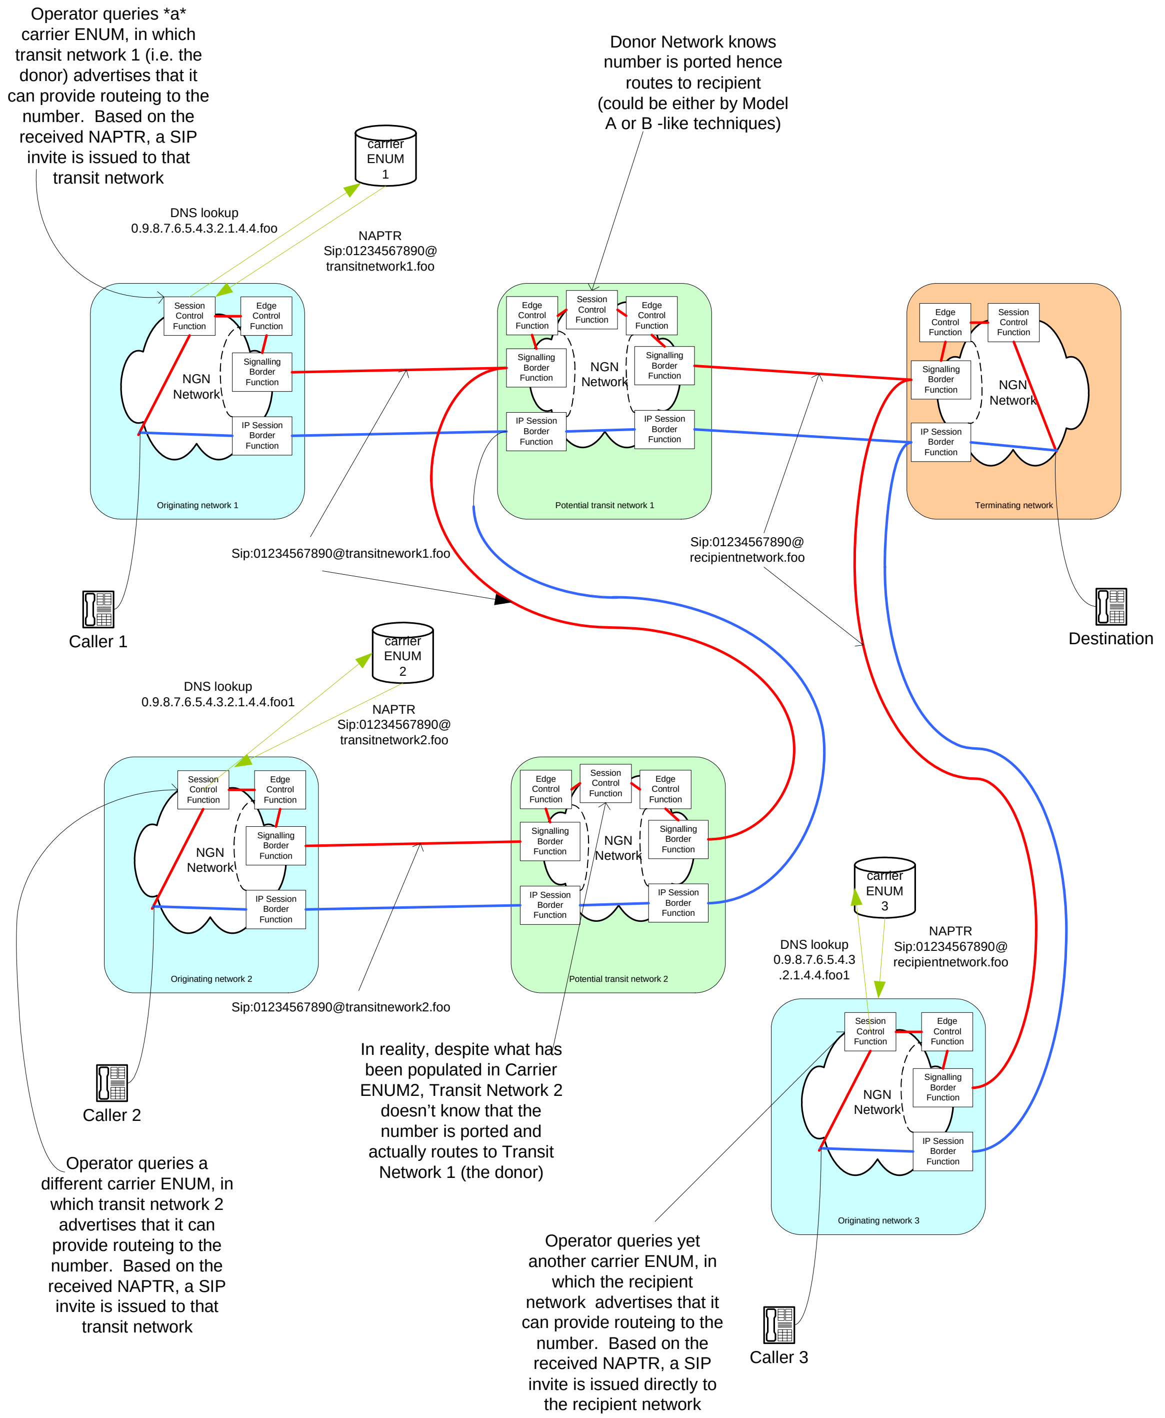

Figure A.7: Number Portability Model E, competing shared databases (illustrated using ENUM databases)

A3.2.2.2 Advantages & Disadvantages

Advantages	**Disadvantages & Issues**
• Avoids "telco failure" problem • Reflects anticipated "organic growth" of NGNs • Allows some flexibility as standards evolve • Arguably a step on the way to full central database	• Requires 3rd party infrastructure to be put in place • Similarly requires legal / commercial framework to be put in place • Performance outside control of individual operators • Whole set of new NP operational processes required • A network querying a shared database has no knowledge whether the result received represents a potential route to the terminating callserver, or the identity of the callserver itself. The database will represent a combination of objective (the ultimate destination of the call) and subjective information (potential means to get there). As such, there is no way to prioritise the results received. • As the decision to populate a given shared database will be made by the terminating network (or transit network in the case of information advertising a route to the destination), an originating network may have to query multiple shared databases to receive an answer for a given number. • As the multiple shared databases would not be standardised, potentially different technologies would be used. • If the multiple database providers choose to co-ordinate their data, how is porting data shared between different shared databases if they operate on different models? • The approach only helps with the telco failure situation if the failing telco crossloads its porting data to shared databases – and this is unlikely to be high on their list of priorities in a business failure scenario.

A3.2.2.3 Conclusions

This approach introduces a series of issues, in particular relating to which answer should prevail where querying of multiple databases results in different answers. As such, it is not recommended as a long term solution. However, operators with networks that are capable of querying e.g. carrier ENUM style technologies, may wish to co-operate to utilise such an approach as a stepping stone to the common database solution. However, this would be via bilateral agreement rather than standardisation, and participants would be responsible for transitional issues that may arise were a common database solution to be subsequently deployed.

A4 Conclusions

After considering the various options, NICC considered that the optimal solution is one where a single Common Numbering Database is adopted. The need for economic analyses was addressed by the cost-benefit study carried out for Ofcom.

Annex B : Use Cases – from ND1631

B1 Reference Access using DNS to NGN

Figure B.1.shows the case where the CP downloads a local copy of the reference database using DNS techniques. For the scenario depicted in Figure B.1, the CP then processes this into a DNS-based server which incorporates data of both which terminating node a given number is hosted, and the routeing that this particular CP uses to direct calls towards that node.

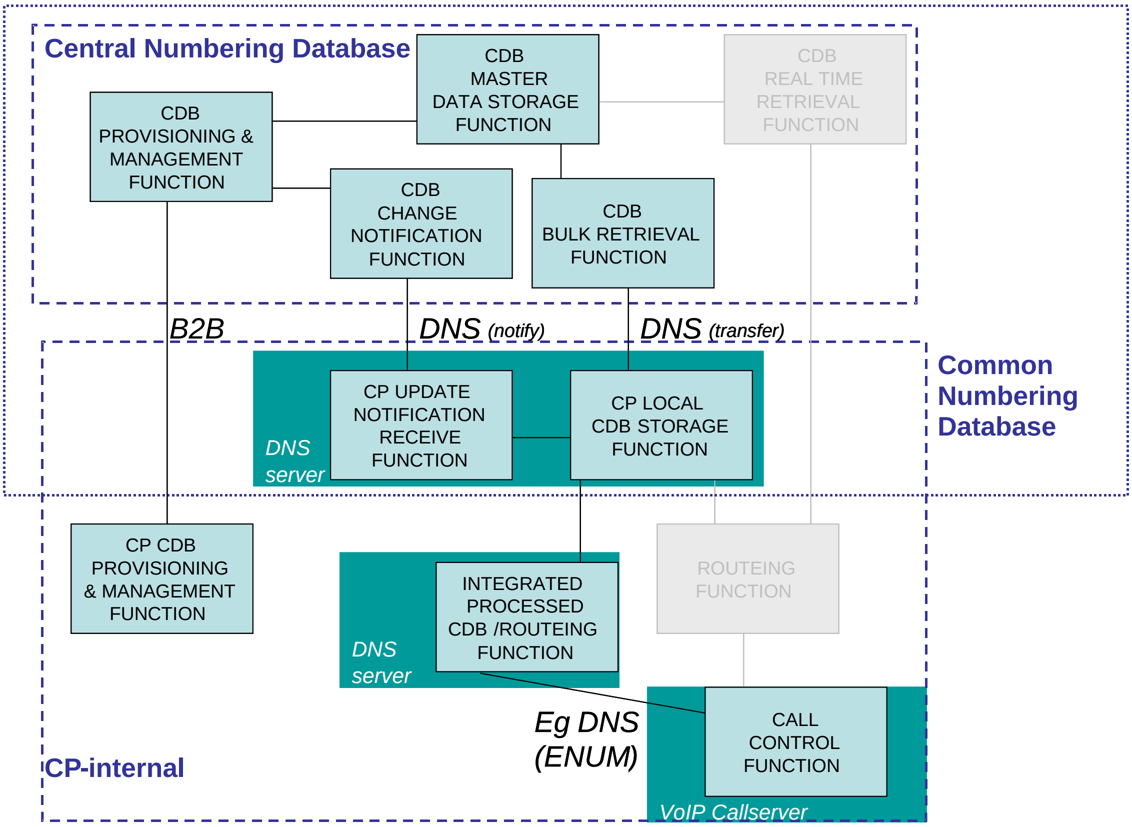

Figure B.1 : Use Case for Reference Access to database using DNS for NGN applications

B2 Reference Access using XML to NGN

Figure B.2 shows the case where the CP downloads a local copy of the reference database using XML. As in the previous case, for the scenario depicted in Figure B.2, the CP then processes this into a DNS-based server which incorporates data of both which terminating node a given number is hosted, and the routeing that this particular CP uses to direct calls towards that node.

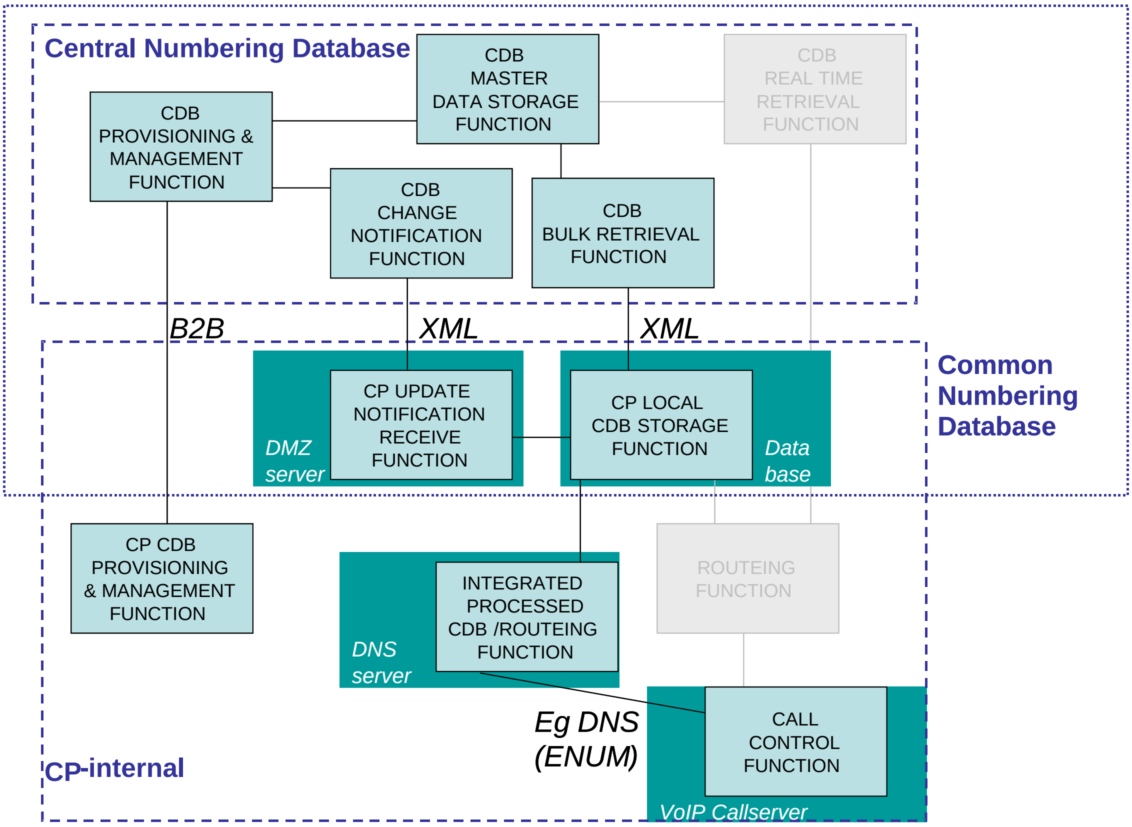

Figure B.2 : Use Case for Reference Access to database using XML for NGN applications

B3 Reference Access using XML to traditional mobile network

Figure B.3 shows the case where the mobile CP downloads a local copy of the reference database using XML. Given they are specifically interested in Mobile Number Portability data and wish this to be in a form suitable for processing by traditional technology nodes, the relevant contents of the database are manipulated into the Intermediate Routeing Number (IRN) form set out in ND1208. This data is then used to populate the Signalling Relay Function (SRF) which is queried by the mobile switching nodes.

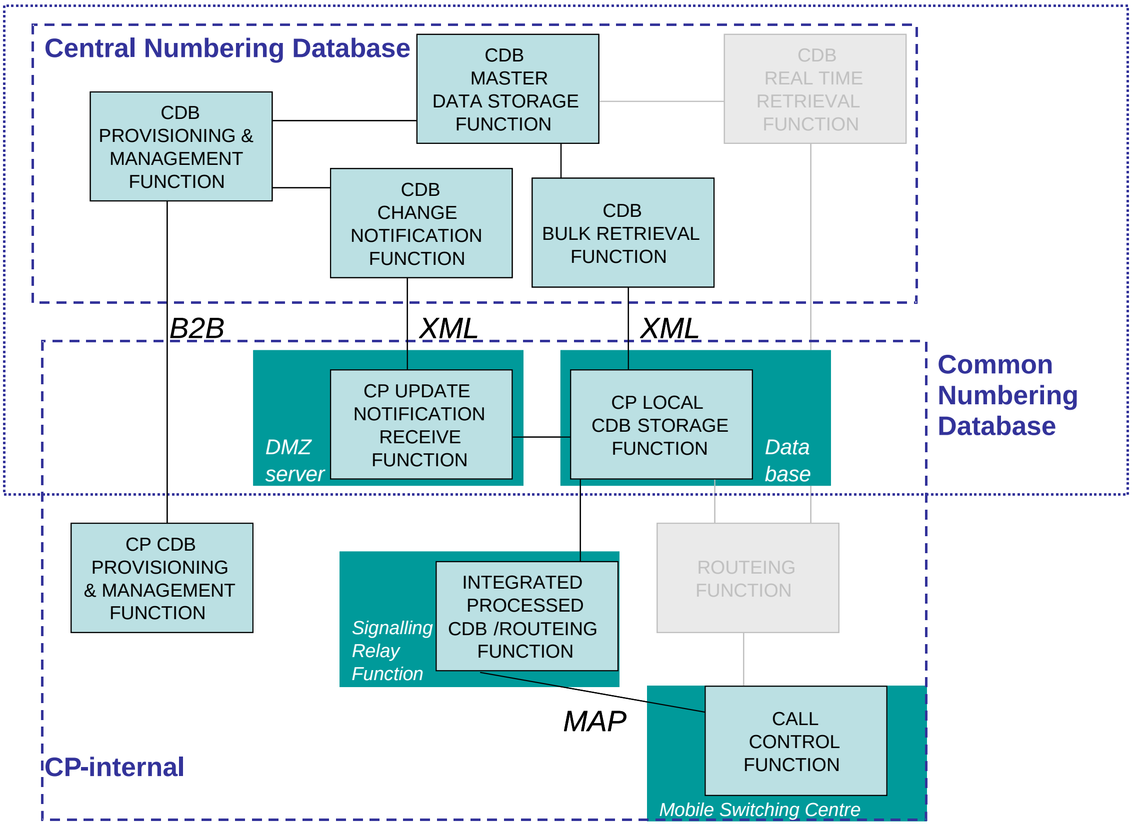

Figure B.3 : Use Case for Reference Access to database using XML for mobile applications

B4 Real Time Access using DNS to NGN

Figure B.4 shows the case where the CP queries the database in real-time using DNS.

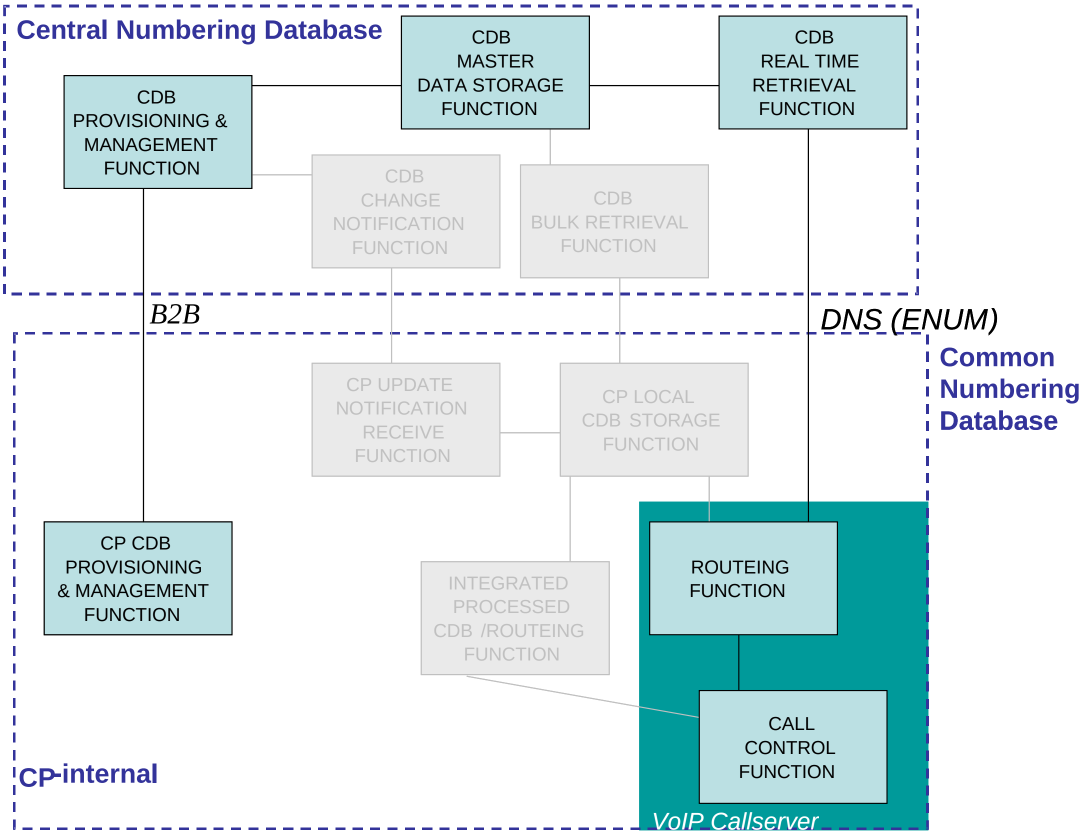

Figure B.4 : Use Case for Real-Time Access to database using DNS

Annex C : Issues considered in determining which numbers to include in the database

In reaching this conclusion, various aspects were considered by NICC.

C1 Defining the Question

Before investigating the options surrounding the database holding “all numbers” or “ported only” it is necessary to first clarify the question being addressed, as considerable misleading views and information exists.

The debate seems to have arisen because databases in some other countries contain only ported numbers. In practice, they contain (or access) ownership information as well for security reasons, but not necessarily in the same form as it would be held in a DNS-type database. This has created the misleading impression that a much smaller database could be used. If the ownership information is not in the database then it has to be held elsewhere and checked for each change to the database or else the system will not be sufficiently secure.

C2 Administrative Information.

Before a CP can be permitted to manipulate information in the database relating to any number, it is essential that the database can establish the authority of that CP to request the change to that data. This is necessary to prevent unauthorised porting of a number. There are a number of requirements for this information:

1. The source must be reliable and authoritative
2. It must support sub-allocation to other Service Providers or resellers
3. It must be secure
4. Must readily integrate with database data operations

NICC has considered the “ownership” information requirements extensively and concluded that this authority must be established with the database based directly upon the authoritative allocation of numbers to the CP by Ofcom and prior to the CP manipulating data associated with those numbers e.g. to release a number for porting. The use of external sources of “ownership” data (e.g. the Ofcom number allocation lists) has been discounted because no sufficiently reliable source could be identified and those considered cannot provide ownership information for sub-allocation of numbers to Service Providers or resellers.

For the purposes of administration of the data, NICC concluded that the database must hold CP number allocation/authorisation information for all numbers for which the database might be required to act.

Conceptually, this might extend to all potential ranges with Ofcom being the “owner” of numbers not allocated to a CP. In such a model, Ofcom would allocate numbers to a CP by “porting” the numbers to them. This would provide a ready mechanism to manage Individual Number Allocation, or permit Ofcom to easily reclaim spare numbers.

C3 Destination Information

The “all numbers” or “ported only” question really relates to the “destination” information which the database provides in response to a query or download request for a number. It is here the options exist:

1. The database responds with destination information only if the number is ported.

2. The database responds with destination information for all numbers (for which the controlling CP has provided data).

We need then to look closer at the meaning of “All Numbers”. This could mean:

A. Every possible number

B. Every number assigned by Ofcom

C. Every number potentially active with a CP

D. Every number with a live customer

Considering each of these:

“A” exceeds the scope of the administrative ownership by a CP so is inappropriate and not considered further.

“B” aligns directly with ownership data, but a CP cannot provide destination information for numbers which are assigned but not yet in service. Hence, this is inappropriate and not considered further.

“C” aligns with today’s practice of block activation where a CP will advise all other CPs that an assigned number block has been activated, and which switch/network is supporting the range. This option may require more database storage of destination information than “D” below. However, this option avoids potential disclosure of any sensitive information about any individual number.

“D” is essentially Individual Number Allocation at a customer level. If this is used where numbers are actually allocated to the CP in blocks, then by querying the database, a third party may determine which numbers of a block the CP has “sold”. This could be viewed as commercially sensitive information.

Should any given number range move to Individual Number Allocation, the database content for this option and “ported only” are equivalent.

In the following sections of this document, “all numbers” will be taken to mean either “C” or “D” above.

NICC examined the advantages and disadvantages of including destination information for “all numbers” or just ported numbers, and drew the following conclusions;

Advantages of “all numbers”

- A CDB containing ported and non-ported permits number allocation at finer granularity than may be supported at present
- If, as predicted, the amount of porting increases driven by increasing volumes of LLU (and potentially VLA/WBC-C), the amount of data in a "ported only" database is likely to increase to the point where a significant volume of fixed line numbers are included by default
- If ported and non-ported information is held, the need for each CP to separately query or merge a database of objective rangeholder information to handle non-ported calls is obviated.
- A database containing all numbers provides a single definitive source of information rather than requiring resolution of porting and non-ported rangeholder data from multiple sources and removes one potential source of errors
- Where a database is ubiquitous the requirement for CPs to circulate number range and NR activation information and databuild should be able to be simplified, potentially reducing activation times

“Advantages of ported only

- A database containing ported numbers only should (initially at least) be simpler to implement and smaller than the alternative

Disadvantages of “all numbers”

- Larger amount of work on CP OSS systems (development of pan-industry systems and rework of intra-CP ones) will be required

Disadvantages of ported only

- Database will still require administrative knowledge of all numbers even if “only ported” are to be routed.
- If only ported information is held by the CDB, then each CP must separately hold a database of rangeholder information which is either queried or merged with the CDB information prior to determining subjective routeing
- Does not aid allocation of numbers at finer granularity than today
- Does not provide opportunity to speed the activation time for new numbers

C4 Conclusions

Having considered the above advantages and disadvantages, NICC concluded that – subject to confirmation that the incremental costs of the larger CDB would not be significant – the advantages of including destination information for all numbers far outweighed the disadvantages.

Annex D : Examination of possible physical database architectures

This annex summarises the database architectural models that were examined.

Model 1 – Central authoritative database

This model consists of having a central UK level common numbering database, authoritative for all numbers. This means that all records for the numbers would be centrally located and managed.

The main disadvantage of this model is that the size of the central database could be very large. If implemented in DNS, the widely used ISC BIND DNS server application would more than likely not be able to cope with such a data-fill for this solution. However, there are high capacity ENUM/DNS solutions commercially available and it should be possible to devise a more flexible internal structure so that for example number entries are grouped together, with exception records for e.g. ported numbers. As the records would be stored centrally, it would mean that when a CP wished to re-groom their network to e.g. move a number from one call control function to another, it would be necessary to involve changing data external to the CP network.

The model has the advantage though that all subscriber numbers are stored centrally and so can be centrally controlled and administered, possibly by one O&M facility. For a real-time DNS-based approach, it also has the advantage in that it reduces the number of DNS requests that an ENUM/DNS resolver has to perform compared to the other models. For a bulk download approach, this model is optimal as all data is held in a single location thus facilitating such an exercise.

Model 2 – Data distributed, held by range-holder

This model would have a central numbering database held only at the number range level. Pointers from this database would refer to individual CP databases which would hold the actual record of associated SIP address for the number. Where a number is ported, the record in the range-holder CP database would be modified to reflect the recipient CP.

This model has the disadvantage that all UK CPs with an NGN would have to provide a database, accessible by other CPs. Another disadvantage is that the recipient network is reliant upon the range-holder network to not only make the changes at the time of porting, but to also support later additions and modifications to the records; possibly relating to services that may not be offered by the range-holder network. Set against this, Ofcom's requirement of decoupling portability from range-holder network involvement could be achieved as the CP-held databases should be of a standard form hence more easily escrowed than is the case for the existing NP solutions.

If a bulk download approach is adopted for the common numbering database, then this would be complicated with this model as to gain access to all records, it would be necessary to download a series of databases from multiple CPs.

The principle advantage of this solution is that it would reduce the size of the central database.

Model 3 – Redirection between CP databases

This model is similar to Model 2 from the perspective of the central numbering database. However, at the individual CP database level, where a number is ported the range-holder network's database would contain an entry for each ported-out number, redirecting queries to the recipient CP database.

In a DNS implementation, this redirection could be realised using a single NS record which redirects the ENUM/DNS Resolver to the recipient CP's ENUM/DNS server by returning a new DNS address to query. Such functionality could also be realised using non-terminal NAPTR records. However, support for non-terminal NAPTR records in real world implementations of ENUM resolvers is not always present; some support only one NAPTR record in a single resolution procedure, some don't even support them at all! Therefore, the use of NS records and not non-terminal NAPTR records would probably be specified.

As with Model 2, a disadvantage of this option is that all UK NGNs would need to have launched numbering database.

Another disadvantage is that the recipient CP is still reliant upon the range-holder network to make updates in their numbering database. However, unlike Model 2, the update has to be done only once (or at least, only when the customer ports their number) and encompasses ***all*** services. Whether they are supported by the range-holder network or not.

An explicit disadvantage over Model 2 though is that a total of three databases need to be consulted where the number is ported. In a DNS implementation, this would require many more resolutions..

If a bulk download approach is adopted for the common numbering database, then this would be complicated with this model as to gain access to all records, it would be necessary to download a series of databases from multiple CPs.

Set against this, compared to Model 2 once ported, the recipient CP is in control of all of their data. Compared to Model 1, the size of the central database is smaller.

Model 4 – Central re-direction database

This model consists of the central numbering database containing entries for all numbers, but rather than containing actual communications records like Model 1, the database would contain pointers to individual CP databases containing this data. This means that all records for customers are located and managed by the actual serving CP of each number, rather than the central database or the rangeholder CP.

As with Model 1, the disadvantage of this approach is that the central numbering database could be very large. However, it would not be as large as in Model 1, because although there would be records for each number, these would merely point to the relevant CP database rather than containing the actual reference data.

Although favoured in many implementations of user-ENUM, it must be noted that in those applications it is likely that a plethora of records would exist for each number. This contrasts with the NGN requirements, which foresee only one or two records.

A disadvantage when compared to Model 1 though is that two databases need to be consulted before an answer is returned.

If a bulk download approach is adopted for the common numbering database, then this would be complicated with this model as to gain access to all records, it would be necessary to download a series of databases from multiple CPs.

The advantage of this model is that it puts the serving CP in full control of the records returned for a particular number. An explicit advantage over Models 2 and 3 is that the a recipient CP is ***not*** reliant upon the rangeholder CP to make any updates in their databases, only the provider of the central numbering database.

Summary of Advantages & Disadvantages

Model	Advantages	Disadvantages
1 : Central Authoritative Database	• All data held centrally – consistent O&M and performance • Eases implementation where bulk downloads are envisaged • Only one database need be consulted.	• Size of central database • Central database involved in day to day operation of network (provides, regrooms etc)
2 : Data distributed, held by rangeholder	• Reduces size of central database • CPs have control of their own data	• Relies on concept of rangeholder • Two database lookups for each number (central then CP) • Involves rangeholder in recipient CP service (albeit with greater scope for data escrow) • Ill-suited to approaches where bulk download is envisaged
3 : Redirection between CP databases	• Reduces size of central database • CPs have control of their own data • Compared to Model 2 decreases role of rangeholder	• Relies on concept of rangeholder • Involves rangeholder in porting process (albeit with greater scope for data escrow) • Two database lookups involved for each number (central then CP). Requires three database lookups for ported numbers (central then rangeholder then recipient) • Ill-suited to approaches where bulk download is envisaged
4 : Central re-direction database	• CPs have control of their own data • No rangeholder involvement • Volume of data in central database reduced versus Model 1	• Size of central database • Two database lookups for each number (central then CP) • Ill-suited to approaches where bulk download is envisaged.

It was concluded that there seemed little merit in pursuing Models 2 and 3.

If the approach adopted involves bulk download of the database to individual CP networks, the Model 1 appeared to be the best approach. If the approach was purely that the database is to be queried in real-time, then further study would be required to select between Models 1 and 4. Since, as outlined in Section 7.3, it was recognised that the database would need to support both administrative and real-time approaches, then Model 1 was selected.

Annex E : Partitioning of the database

Summary

NICC concluded it was best to split the database at the C-D boundary of the number (e.g. 0121 2). This gives 1,397 zones, all of which can expect an update every working day. At 50 bytes of data per number, that makes a full download of a database Section up to 50Mb with an average size of 28Mb.

Methodology

NICC used the Ofcom CSV files to be found on their website. The data was as of the morning of 2007-04-26. These files include all numbers except 0500 and BT six-digit 0800 numbers, for which Ofcom don't provide any allocation data.

NICC examined only those blocks with a status of "Allocated". The block prefixes vary from 5 to 7 digits and the number lengths from 7 to 10. Where a block length isn't given, 10 digits was assumed. This happens with:

all of 03

a few blocks in 07 and 09

80 blocks in 0844 and 0871

a number of blocks in 01481 and 01534 (Jersey and Guernsey)

For 01 and 02, where the files say "F-digit 3 4 6" (for example) this was then split that into three separate blocks.

This gave:

S-Digit	Total assigned blocks	Blocks of number length			
		7 digits	8 digits	9 digits	10 digits
01	28503		2	323	28178
02	3824				All
03	665				All
05	309				All
07	4029				All
08	7371	2		2757	4612
09	8284				All

NICC then split all the blocks up into sub-blocks along the F digit. That is, a block such as 122345 was split into 1223450, 1223451, ..., 1223459.

This gives a total of 774,254 F blocks:

S-Digit	1k Blocks of number length				Total possible numbers
	7 digits	8 digits	9 digits	10 digits	
01		20	3230	281780	200,230,200
02				341360	34,136,000
03				6650	6,650,000
05				3090	3,090,000
07				402000	402,000,000
08	11		2757	39613	39,888,711
09				82840	82,840,000

For a grand total of 768,834,911 potential numbers.

Each of seven possible boundaries for database Sections, from S-A to F-G were then considered;

Split	Number of Sections	Maximum	Minimum	Mean	Big
S-A	7	402,000,000	3,090,000	109,833,558	1
A-B	34	98,400,000	370,000	22,612,791	1
B-C	189	10,000,000	20,000	4,067,909	25
C-D	1,397	1,000,000	10,000	550,347	436
D-E	9,763	100,000	1,000	78,749	5,965
E-F	83,740	10,000	10	9,181	75,954
F-G	77,254	1,000	1	993	768,236

The second column is the number of Sections. The next three are the maximum, minimum, and mean Section size. Finally the last column shows how many of the Sections are the maximum size.

To look at how many updates a day to expect, a mathematical model was used where each number has an update with equal probability. For a given split, all the Sections were assumed to be of the same size (to do otherwise is too complex, and it turns out to be a reasonable approximation). Four scenarios were examined:

(A) each number gets an update every 250 working days

(B) each number gets an update every 833 working days

(C) each number gets an update every 2,500 working days

(D) each number gets an update every 8,333 working days

These scenarios correspond to the following rates of update and number block utilisation:

Average updates	Average utilisation			
	100%	30%	10%	3%
Once per year	(A)	(B)	(C)	(D)
Once per 3.3 years	(B)	(C)	(D)	
Once per 10 years	(C)	(D)		

With splits from S-A to D-E all zones see at least one update per day. With the smaller splits, we see the following probability of an update on a given day:

Scenario	E-F split	F-G split
(A)	100%	98%
(B)	100%	70%
(C)	97%	33%
(D)	67%	11%

Conclusions

Incremental updates are precisely that, so the amount of data to be transferred daily is going to be the same except for the framing of the messages themselves. Thus the split is a trade-off between full downloads - requiring whole Sections to be transferred - and the number of update messages. It was considered that the latter was far more significant and it's better to get (say) 1,397 update messages a day rather than 56,105 (67% of 83,740). This means few large zones are preferable. On the other hand, full downloads shouldn't be too big; they should be megabytes rather than gigabytes. Assuming 50 bytes per telephone number, the size of a zone file would be:

Split	Maximum zone size	Mean zone size
S-A	30.1GB	4.5GB
A-B	4.9GB	1.1GB
B-C	500MB	203MB
C-D	50MB	28MB
D-E	5MB	4MB
E-F	500KB	460KB
F-G	50KB	50KB

Thus it was concluded that the optimal balance lay splitting at the C-D boundary. This gives 1,397 zones, all of which can on average expect an update every working day. At 50 bytes of data per number that makes a full download be up to 50Mb with an average size of 28Mb.

Annex F : Impacts of compression on the bulk XML download of a fully populated number section

In order to assist the evaluation of the benefit of various compression techniques on the bulk XML download of a fully populated number section, various scenarios were simulated by varying the overall porting rate and the number of other CPs to which a number could port. It should be noted however, that in the simulation each other CP was denoted by a single destination group, so in effect the other number of CPs becomes equivalent to the total number of unique destination groups within a number section for the purposes of the simulation.

A fully populated number section was created to which was applied a random number of ports to a random number of CPs. In order to limit the benefit of the format compression techniques a pessimistic approach was taken to simulating ports whereby numbers to port were selected individually (akin to Consumer porting). Selecting contiguous blocks of numbers to port (akin to Business porting) would have biased the results in favour of the format compressions under investigation. As such number sections that contain many contiguous blocks of numbers could expect to obtain much better compression rates than those found during the simulation.

<u>No format compression.</u>

The following is an excerpt from a file that has no format compression.

```
<entry id="000000">
    <pstn>72000799</pstn>
    <ims>dg.a799.uktel.org.uk</ims>
</entry>
<entry id="000001">
    <pstn>72000799</pstn>
    <ims>dg.a799.uktel.org.uk</ims>
</entry>
<entry id="000002">
    <pstn>72000798</pstn>
    <ims>dg.a798.uktel.org.uk</ims>
</entry>
<entry id="000003">
    <pstn>72000799</pstn>
    <ims>dg.a799.uktel.org.uk</ims>
</entry>
```

Range format compression

Under range format compression the same excerpt would take the following form.

```
<range first="000000" last="000001">
    <pstn>72000799</pstn>
        <ims>dg.a799.uktel.org.uk</ims>
</range>
<entry id="000002">
    <pstn>72000798</pstn>
        <ims>dg.a798.uktel.org.uk</ims>
</entry>
<entry id="000003">
    <pstn>72000799</pstn>
        <ims>dg.a799.uktel.org.uk</ims>
</entry>
```

Batch format compression

Lastly, under batch format compression the same excerpt would take the following form.

```
<batch>
    <pstn>72000799</pstn>
    <ims>dg.a799.uktel.org.uk</ims>
    <range first="000000" last="000001" />
    <entry id="000003" />
…
</batch>
<batch>
    <pstn>72000798</pstn>
    <ims>dg.a798.uktel.org.uk</ims>
    <entry id="000002" />
…
</batch>
```

The first simulation was carried out by setting the overall port rate to 15% with 30 potential CPs to which a number could port. The following file sizes were observed.

format compression	file size	gzip compressed file size
no format compression	84.88MB (-)	3.35MB (0.21MB)
range compression	24.76MB (0.21MB)	1.64MB (0.14MB)
batch compression	8.12MB (0.01MB)	1.22MB (0.04MB)

The number in brackets is the amount by which the file size reduced when the number of potential CPs to which a number could port was reduced to 5.

The next simulation increased the overall port rate to 40% with 30 potential CPs to which a number could port. The following file sizes were observed.

Format compression	file size	gzip compressed file size
no format compression	84.88MB (-)	3.92MB (0.42MB)
range compression	55.16MB (2.08MB)	3.51MB (0.41MB)
batch compression	16.97MB (0.26MB)	2.47MB (0.09MB)

Again, the number in brackets is the amount by which the file size reduced when the number of potential CPs to which a number could port was reduced to 5.

The results suggested that the number of CPs (equivalent to unique destination groups within a single number section) has little impact on the file sizes so no further variations in number of other CPs were carried out.

To further reduce the uncompressed file sizes the white space within the XML file was reduced to a minimum and information relating to a single entry or range was limited to a single line.

<u>No format compression</u>

With no format compression the previous excerpt would then look as follows.

<entry id="000000"><pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims></entry>

<entry id="000001"><pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims></entry>

<entry id="000002"><pstn>72000798</pstn><ims>dg.a798.uktel.org.uk</ims></entry>

<entry id="000003"><pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims></entry>

Range format compression

Under range format compression the previous excerpt would then look as follows.

```
<range first="000000" last="000001"><pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims></range>
<entry id="000002"><pstn>72000798</pstn><ims>dg.a798.uktel.org.uk</ims></entry>
<entry id="000003"><pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims></entry>
```

Batch format compression

Under batch format compression the previous excerpt would then look as follows.

```
<batch>
<pstn>72000799</pstn><ims>dg.a799.uktel.org.uk</ims>
<range first="000000" last="000001" />
<entry id="000003" />
…
</batch>
<batch>
<pstn>72000798</pstn><ims>dg.a798.uktel.org.uk</ims>
<entry id="000002" />
…
</batch>
```

For 15% overall porting levels, the following files sizes were observed

format compression	**file size**	**gzip compressed file size**
no format compression	77.25MB	3.24MB
range compression	23.17MB	1.59MB
batch compression	7.85MB	1.22MB

For 40% overall porting levels, the following files sizes were observed

format compression	file size	gzip compressed file size
no format compression	77.25MB	3.78MB
range compression	51.5MB	3.38MB
batch compression	16.36MB	2.45MB

Lastly, to reduce the size of the uncompressed files further still the most repetitive tags within the XML were replaced with single characters.

No format compression

With no format compression the previous excerpt would then look as follows.

```
<e id="000000"><p>72000799</p><i>dg.a799.uktel.org.uk</i></e>
<e id="000001"><p>72000799</p><i>dg.a799.uktel.org.uk</i></e>
<e id="000002"><p>72000798</p><i>dg.a798.uktel.org.uk</i></e>
<e id="000003"><p>72000799</p><i>dg.a799.uktel.org.uk</i></e>
```

Range format compression

Under range format compression the previous excerpt would then look as follows.

```
<r f="000000" l="000001"><p>72000799</p><i>dg.a799.uktel.org.uk</i></r>
<e id="000002"><p>72000798</p><i>dg.a798.uktel.org.uk</i></e>
<e id="000003"><p>72000799</p><i>dg.a799.uktel.org.uk</i></e>
```

Batch format compression

Under batch format compression the previous excerpt would then look as follows.

```
<batch>
<p>72000799</p><i>dg.a799.uktel.org.uk</i>
<r f="000000" l="000001" />
<e id="000003" />
…
</batch>
<batch>
<p>72000798</p><i>dg.a798.uktel.org.uk</i>
<e id="000002" />
…
</batch>
```

For 15% overall porting levels, the following files sizes were observed

format compression	**file size**	**gzip compressed file size**
no format compression	60.08MB	3.06MB
range compression	17.72MB	1.48MB
batch compression	6.07MB	1.17MB

For 40% overall porting levels, the following files sizes were observed

format compression	**file size**	**gzip compressed file size**
no format compression	60.08MB	3.47MB
range compression	39.6MB	3.08MB
batch compression	12.94MB	2.36MB

Further savings in range compressed file size can be achieved by better handling of the common part of the ims destination group ('.uktel.org.uk') which is repeated many times when using that format compression technique.

History

Document history		
<Version>	<Date>	<Milestone>
1.1.1	April 2008	Initial issue
1.1.2	August 2008	Addition of hyperlink to ND1610 release sheet with document version numbers

End of Document