Document filename: ITK 2 0 Addressing and Routing Requirements v1.0.docx			
Directorate / Programme :	HSCIC - Architecture	**Project**	Interoperability
Document Reference :		HSCIC-ITK-ARCH-107	
Project Manager :	Rob Shaw	**Status :**	Final
Owner :	George Hope	**Document Version :**	1.0
Author :	George Hope	**Version issue date :**	23/06/2014

ITK Addressing and Routing Requirements

Document Management

Revision History

Version	Date	Summary of Changes
1.0	31/05/2014	First version issued by HSCIC

Reviewers

This document was reviewed by the following people:

Reviewer name	Title / Responsibility	Date	Version
George Hope	ITK Architecture Lead	30/04/2014	1.0
Sanjay Paul	ITK Architect	30/04/2014	1.0
Richard Dobson	ITK Accreditation Manager	30/04/2014	1.0
David Barnet	ITK Communication and Messaging	30/04/2014	1.0
Nigel Saville	ITK Accreditation	30/04/2014	1.0

Approved by

This document was approved by the following people:

Name	Signature	Title	Date	Version
Shaun Fletcher		Head of Architecture	31/05/2014	1.0
Rob Shaw		Director Operational Services	31/05/2014	1.0

Reference Documents

Ref no	Doc Reference Number	Title	Version
1.			
2.			
3.			
4.			

Document Control:

The controlled copy of this document is maintained in the HSCIC corporate network. Any copies of this document held outside of that area, in whatever format (e.g. paper, email attachment), are considered to have passed out of control and should be checked for currency and validity.

Contents

1 Introduction

This document forms part of the overall document set for ITK Architecture.

1.1 Purpose of Document

This document defines a set of requirements for ITK Addressing and Routing.

1.2 ITK Architecture Documentation Set

The position of this document in relation to the document set is shown below.

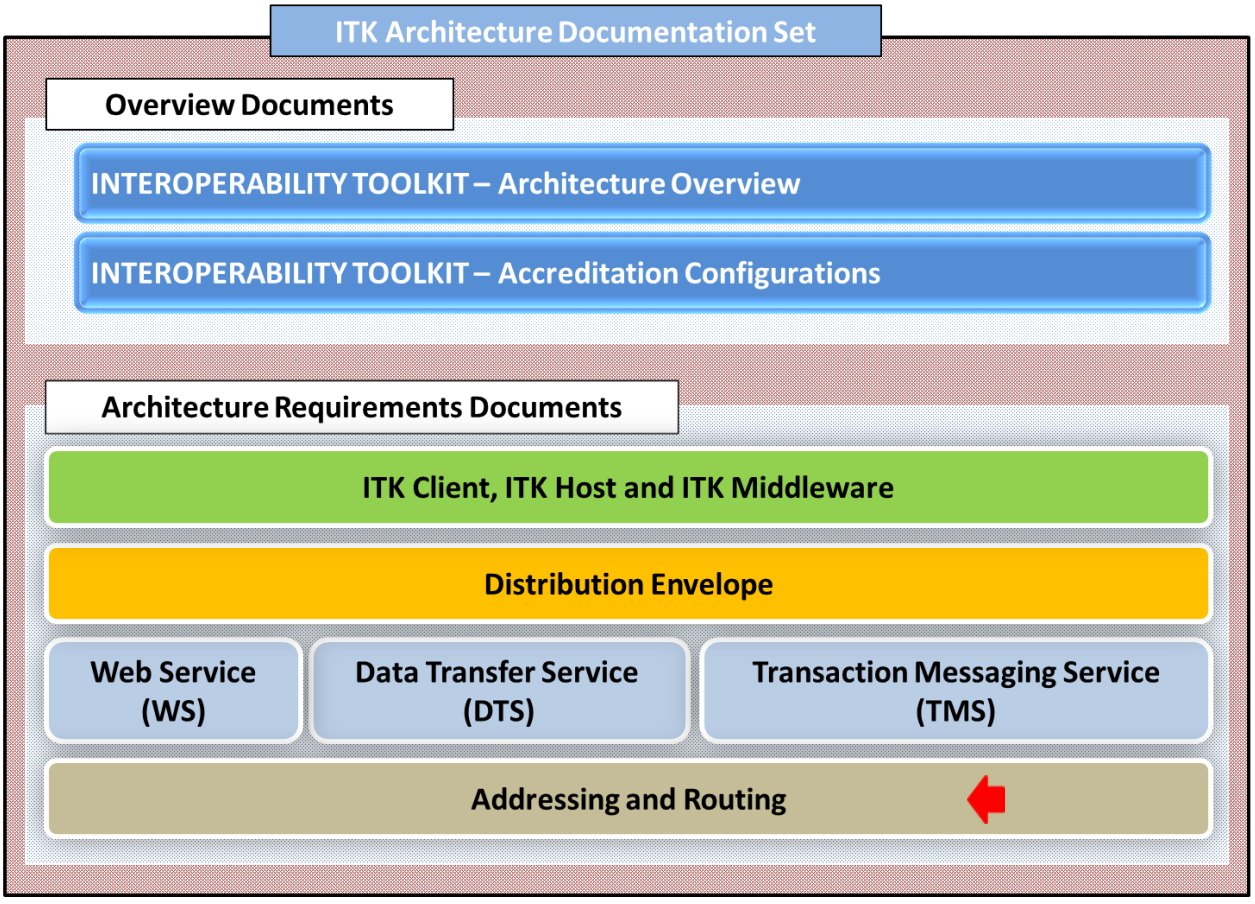

Figure 1 – ITK Architecture Documentation Set

1.3 Audience

The primary audience are supplier technical and product development staff who are interested in developing a Toolkit Implementation.

1.4 Document Scope

The document covers the ITK Addressing and Routing Interoperability requirements required for accreditation.

1.5 Document Overview

The rest of this document covers a number of areas of functionality. Within each area the functionality is described, and a number of formal requirements are listed in bold type, with additional detail provided in smaller type below this.

1.6 Requirements Presentation

The requirements are presented in the format given below:

Ref (1)	Description (2)	Client (3)	Host (4)	MW (5)	SMSP (6)
COR-REL-03	**Toolkit Implementations MUST retain responsibility for processing until a request completes**	**Y**	**N**	**Y**	**N**
NB (7)	Specifically, any response returned from the initial part of the asynchronous invocation does NOT indicate a transfer of responsibility. It is only a transport acknowledgement, and it does NOT imply that the message has necessarily been persisted, nor does it indicate a transfer of responsibility, nor promise that subsequent application processing will be completed.				

Clarification Notes

(1) The requirement reference

(2) The Description of the requirement

(3), (4), (5) and **(6)** Shows the requirements applicability for accreditation

(7) Provides further details relating to the requirement and supplementary notes

Colour Coding Notes

- The fill colour of the Reference relates to a particular document from the document map.
- Where requirements are universally applied the fill colour will always be blue. Where requirements are conditional and may impact accreditation the fill colour will be Orange.
- See the Accreditation Configuration spread sheet for related details.

1.7 Reference Implementation

An ITK reference implementation pack is available as a training and development aid and it contains example code snippets for typical Healthcare Interoperability scenarios.

http://developer.nhs.uk/library/interoperability/nhs-interoperability-framework/

2 Addressing and Routing Concepts

The addressing and routing architecture assumes that ITK distribution envelopes are used, and can be queried by routers to obtain addressing information but are otherwise transferred intact.

An **address** is a label for a communications end point that is meaningful in a business context. **Addressing** is a collective name for business and user processes that use addresses. These include:

- The discovery of the address of a business entity.
- The use of that address in sending a message (and in the user declaring their own address so that the message can be acknowledged or failures notified).
- The management of one or more addresses for an organisation, allocating new addresses, and ensuring that each is unique.

Addresses and addressing are considered entirely “logical” things independent of the physical transports that are used to ship messages:

Discovery

In many cases it is likely that a recipient address will already be known to a sender because, for example, the message is a discharge summary to a local GP. This is an example of a “local address book”. ITK middleware suppliers SHOULD demonstrate how client-side discovery of non-local addresses can be supported although ITK does not mandate any particular technology for doing so.

How such discovery services are populated with national data is out of scope for this specification.

2.1 Addresses

An address is just a label, a string that represents a communications end point. ITK addresses are intended for use in two ways:

- As a user-friendly, or at least readable, way to describe something that there is a business requirement to communicate with. “Something” is most likely a role, team or organisation – for example an organisation’s Caldicott Guardian, a neonatal unit in a particular Trust or a GP practice. But it may be a service such as a central archive or PACS, or the Spine PDS.
- As something that routers can work on to do their job.

As such an address must be:

- **Unique**. When an address is used to resolve a path, there can be no ambiguity as to the identity of the intended ultimate recipient.
- **Readable**. Human users provide a great deal of sanity checking: but they can only do so if information is presented in a readable form. An address which takes the form of a sequence of numbers or other meaningless string has to be taken on trust. One that is readable invites far less user error.

- **Manageable**. Addresses must be easy for administrators to handle. They must have a clear correlation with organisational or other structures in the real world.
- **Processable**. Routers must be able to use the address.

ITK addresses are structured strings which support each of these four requirements. The address representation is a printable ASCII string, conformant with the XML Schema anyURI data type (http://www.w3.org/TR/2004/REC-xmlschema-2-20041028/#anyURI).

ITK routers MUST be able to recognise and process addresses in the format described here. Suppliers of ITK routers MUST be able to demonstrate administration of ITK addresses, and SHOULD be able to demonstrate support for local directory services or address books.

2.2 Address Structure

To distinguish the subset of anyURI values that correspond to the NHS ITK address space, the following prefix is used:

urn:nhs-uk:addressing:

The address representation's structure is the prefix, followed by an arbitrarily-long sequence of colon-delimited tokens:

urn:nhs-uk:addressing:a:b:…:N

The tokens represent a hierarchy with the highest level on the left ("a" in this example) and the lowest level on the right ("N" here). There are no explicit semantics attached to any of the tokens at any level. A token is simply a label.

In the ITK addressing scheme, the highest-level token implicitly represents a naming authority. The next token implicitly represents some form of organisation that is self-contained or autonomous at least to some degree. However there is no restriction on what the token itself actually is.

Examples of logical addresses might be:

A GP practice	urn:nhs-uk:addressing:ods:B83019
A particular service within that GP surgery	urn:nhs-uk:addressing:ods:B83019:practicenurse
An acute trust	urn:nhs-uk:addressing:ods:RHM
A very specific in-tray within a particular team within that acute trust	urn:nhs-uk:addressing:ods:RHM:oncology:labreports:histology:in-tray
A social care team in a local authority	urn:nhs-uk:addressing:ods:V999:childservices:referrals
A team in another government department	urn:nhs-uk:addressing:govuk:ABCD:freeschoolmeals

**All these addresses are only examples by pattern.*

2.3 Management

While the Logical Address format is very flexible, a constraint is that every address MUST be a globally unique string. This will become particularly important as the ITK evolves and ITK systems potentially become increasingly interconnected beyond purely local boundaries.

However it is clearly both impractical and undesirable to micro-manage addressing down to the level of teams within organisations. A delegated approach similar to that used for internet domains is used. The logical address is structured, as shown in the diagram below:

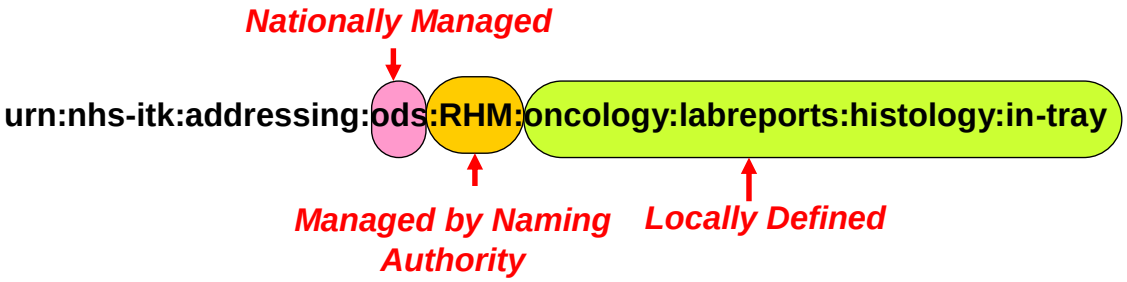

Figure 2 – ITK Address Management

The first element of the address identifies a Naming Authority. It is expected that this will be a very short and stable list. In the first instance it is envisaged as consisting of one entry for the NHS Organisational Data Service

The second element of the address is managed by the naming authority. Organisations must obtain a code to represent their organisation, unique within the naming authority.

The remainder of the address is locally defined. Organisations themselves are responsible for organising their own internal address details. (Large organisations may of course decide to further delegate responsibility internally, based on further sub-elements).

The first "organisational" address element will be allocated and managed using, as far as possible, existing organisational identifiers.

- ***ODS Codes**. Most NHS organisations already have a unique ODS code, and the allocation of these is already managed by the Organisational Data Services (ODS) team. ODS Codes will be reused as the first element of a Logical Address.*
- ***GP Practice Codes**. In a similar way, GP practice codes are already allocated and managed. GP practice codes will be reused as the first element of a Logical Address.*
- ***Other**. The above options involve no new centralised processes, and should cover the vast majority of circumstances. However there may be unusual circumstances or special cases where an organisation wishes to allocate Logical Addresses but does not qualify for an ODS code. In this circumstance the organisation may contact the ITK enquiries team (toolkit.enquiries@hscic.gov.uk)- who can verify the validity of the request and arrange an alternative addressing hierarchy.*

3 Requirements Specification

ITK addressing and routed infrastructure provides a context in which various system actors collaborate to perform the function of delivering content.

The specifications are based on the performance of addressing and routing functions only: they do not cover any requirements specific to the implementation of any particular physical route types. Any such conformance requirements are stated against the specification of that type of transport.

3.1 Routing Requestor conformance

Ref	Description	Client	Host	MW	SMSP
MOD-RTS-01	**Routing requestors MUST provide an externally-routable sender address for an end-to-end acknowledged message**	**N**	**N**	**N**	**N**
NB	Messages routed outside an organisation MUST include a sender address that is externally routable if they are to use any “end-to-end acknowledged” routing service. This requirement does not apply to unacknowledged “Request” configuration.				

MOD-RTS-02	**Routing requestors SHOULD provide an externally-resolvable audit identity**	**N**	**N**	**N**	**N**
1	Messages to be routed outside an organisation SHOULD include identifiers that either: - Are ITK address-style identities showing the scoping organisation, or - Are sets of identities which MAY include local identifiers but which MUST include a recognisable originating organisation scope, such as SDS identifier or ODS code.				
2	Routers will reject messages which fail this specification. Routing requestors SHOULD be configured to avoid a router so rejecting. However as a sender does not necessarily know the path that a message will take the sender cannot know if that path will involve an external route. So it is not a conformance violation on the part of the sender to include an “illegal” identity, although deployment configurations should attempt to avoid it.				

MOD-RTS-03	**Routing requestors MUST support the “urn:nhs-itk:services:201005:SendInfrastructureAck-v1-0” service**	**N**	**N**	**N**	**N**
1	An exception MAY be granted where the routing requestor can demonstrate that routing request has requested no acknowledgement.				

MOD-RTS-04	**Routing requestors MUST be able to handle multiple end-to-end infrastructure returns**	**N**	**N**	**N**	**N**

1	ITK message routing supports addressing a single message to multiple recipients. As the infrastructural acknowledgment is end-to-end each receiver will return its own acknowledgment. The routing requestor MUST be able to handle multiple acknowledgments for a single message transmission.
2	Although ITK routing is consolidated, delivery to each recipient is independent of delivery to any others. As such different paths can fail independently of one another, and so a requestor may receive a mixture of successful infrastructure acknowledgments from receivers, and error reports. A routing requestor MUST be able to support this scenario.
3	An exception MAY be granted where the routing requestor can demonstrate that routing request has requested no acknowledgement.

3.2 Routing Receiver conformance

Ref	Description	Client	Host	MW	SMSP
MOD-RTR-01	**Routing receivers MUST acknowledge using the urn:nhs-itk:services:201005:SendInfrastructureAck-v1-0" service**	**N**	**N**	**N**	**N**
1	ITK message routing includes an end-to-end infrastructural acknowledgment message that is returned to the original requestor, as identified by the "senderAddress" element of the distribution envelope. A receiver MUST return a message consisting of a distribution envelope wrapping a Generic Infrastructure Response to confirm technical receipt, or to provide information about a failure. This acknowledgement or NACK is sent using a request to the "urn:nhs-itk:ns:201005:SendInfrastructureAck" service on the receiver's local router.				
2	An exception MAY be granted where the routing requestor can demonstrate that routing request has requested no acknowledgement.				

3.3 Router conformance

Ref	Description	Client	Host	MW	SMSP
MOD-RTI-01	**Router Implementations MUST support configurable service-based routing**	**N**	**N**	**N**	**N**
NB	The Toolkit Implementation MUST be able to route calls to an endpoint based on the service being called.				

MOD-RTI-02	**Routing Implementations MUST support configurable content-based routing**	**N**	**N**	**N**	**N**
1	The Toolkit Implementation should be able to route calls to an endpoint based on the service being called AND the content of the message.				
2	The fields to be routed on MUST be configurable based on service type.				

3	The values to be routed on MUST also be configurable.

MOD-RTI-03	**Routing Implementations SHOULD support dynamic binding**	**N**	**N**	**N**	**N**
NB	Dynamic binding means that endpoint addresses are resolved dynamically from the registry at runtime. This allows configuration changes to take immediate effect.				

MOD-RTI-04	**Routing Implementations MUST be able to convert a Toolkit Address URI + service name into a Physical Route for the next messaging hop**	**N**	**N**	**N**	**N**
1	The goal of routing is to resolve the Toolkit Address uri for the ultimate end destination into physical details of the next messaging hop. (It is thus possible that this physical address is not the final destination, and further routing of subsequent hops may be needed) This mapping of Toolkit Address URI to physical route MUST be configurable.				

MOD-RTI-05	**Routing Implementations MUST support redirection and blocking routing table entries**	**N**	**N**	**N**	**N**
NB	The redirection and blocking entries are important elements in change management. Toolkit routing implementations MUST support these in addition to any other routing facilities they may be provided.				

MOD-RTI-06	**Routing Implementations MUST optimise the number of message instances transmitted**	**N**	**N**	**N**	**N**
NB	For messages with multiple recipients, the Toolkit Implementation MUST optimise the messages transmitted, so that it only sends a single message to each distinct physical route. For example, if 3 recipients all resolve to the same system endpoint then just one copy of the message (not 3 copies) should be sent there.				

MOD-RTI-07	**Routing Implementations MUST support configuration of one or more Authoritative Domains**	**N**	**N**	**N**	**N**
1	Toolkit Implementations MUST allow an Authoritative domain for its routing to be defined via configuration.				
NB	An Authoritative Domain for a Toolkit Implementation is the set of Toolkit Addresses which lie within its control for routing purposes. *(This is as opposed to "external" addresses which the router may forward to but does not directly control).* Those addresses within a router's Authoritative Domain are considered to be "inbound" for routing purposes, whilst all other addresses are considered to be "outbound".				

MOD-RTI-08	**Routing Implementations MUST support inbound vs outbound routing**	**N**	**N**	**N**	**N**

NB	In order to constrain loops and duplication, a message must know, at any moment in time, whether it is “outbound” or “inbound”. This is deduced based on configuration of the endpoint through which the message enters the router. There are three possible modes to choose from when configuring an endpoint: **IN – the message is inbound.** Therefore only addresses which fall into the router’s Authoritative Domain must be routed to. **OUT – the message is outbound.** Therefore only addresses which fall outside of the router’s Authoritative Domain must be routed to. **RELAY – used when a message is first posted into the network**. It therefore needs to be routed in both directions initially – both inbound and outbound. Note that this approach has been selected as it avoids the need for routing state to be stored in the message. What a router elects to do with a resolved address is dependent on its operating mode, but is otherwise stateless. No state is held in the message itself.

MOD-RTI-09	**Routing Implementations MUST be able to route based on processing of a Toolkit Distribution Envelope address list**	**N**	**N**	**N**	**N**

MOD-RTI-10	**Routing Implementations MUST reject requests for external routing that are only locally identified**	**N**	**N**	**N**	**N**
NB	Messages routed outside an organisation MUST include identifiers that either: - Are ITK address-style identities that show the scoping organisation, or - Are sets of identifiers which MAY include local identifiers but which MUST include a recognisable originating organisation scope, such as an SDS identifier or ODS code. Routing service requestors SHOULD be configured to avoid routers having to reject messages based on this requirement. However it is in the nature of the routing system – especially in light of redirect entries - that a requestor does not necessarily know the path a message will take. So, ultimately, it is the responsibility of the router to reject.				

MOD-RTI-11	**Routing Implementations SHOULD be able to use auditIdentity content to make access control decisions**	**N**	**N**	**N**	**N**

MOD-RTI-12	**Routing Implementations MUST consider it an error if it receives a message containing a Toolkit Address which it cannot handle**	**N**	**N**	**N**	**N**
NB	Routers MUST reject messages containing a toolkit address which is un-routable.				

MOD-RTI-13	**Routing Implementations MUST reject requests for end-to-end acknowledged external routing that have no routable sender address**	**N**	**N**	**N**	**N**

NB	Messages routed outside an organisation MUST include externally routable sender addresses for services that use end-to-end acknowledgment, as the acknowledgment process depends on it. The requirement does not apply to the unacknowledged “Request” message configuration.

MOD-RTI-14	**Routing implementations MUST support the “urn:nhsitk: ns:201005:SendInfrastructureAck” service**	**N**	**N**	**N**	**N**
1	The “urn:nhs-itk:ns:201005:SendInfrastructureAck” service is unacknowledged and routing errors MUST be logged but a router MUST NOT generate a NACK. The router is expected not to try to interrogate the content of an acknowledgment nor to correlate with any previous message it may have handled, the acknowledgments are routed based on their address alone				
NB	For “Response” message configuration, ITK message routing includes an end-to-end infrastructural acknowledgment message that is returned to the original requestor either by the receiver (to confirm technical receipt), or by an intermediary router or the receiver to provide information about a failure. This acknowledgement is sent using the “urn:nhs-itk:ns:201005:SendInfrastructureAck” service, and routers requestors MUST be able to route it.				

MOD-RTI-15	**Routing implementations MUST signal routing failures using the "urn:nhs-itk:services:201005:SendInfrastructureAck-v1-0” service if Infrastructure Ack is requested**	**N**	**N**	**N**	**N**
1	For “Response” message configuration, ITK message routing includes an end-to-end infrastructural acknowledgment message that is returned to the original requestor, as identified by the “senderAddress” element of the distribution envelope. A router MUST NOT return an acknowledgment when a message is successfully relayed.				
NB	A router MUST return a message consisting of a distribution envelope wrapping a Generic Infrastructure Response to signal and to provide information about a failure. This NACK is routed via a request to the “urn:nhs-itk:services:201005:SendInfrastructureAck-v1-0” service using the router’s own routing tables.				

4 ITK Addressing and Routing in Practice

The following example is intended to demonstrate all the various features of ITK addressing and routing, described to date. It shows a multi-organisation routing problem, the various routing tables, and physical routes.

4.1 ITK Addressing and Routing Scenario

A sender “urn:nhs-uk:addressing:ods:RHM:A”, writes a CDA discharge summary addressed to:

urn:nhs-uk:addressing:ods:RHM:B
urn:nhs-uk:addressing:ods:B83019:C
urn:nhs-uk:addressing:ods:ABC:D
urn:nhs-uk:addressing:ods:RHM:F

The actors are:

Recipient	Description
urn:nhs-uk:addressing:ods:RHM:B	Internal recipient, “colleague” team or department within RHM. Can process CDA.
urn:nhs-uk:addressing:ods:B83019:C	External recipient – GP or care home. Unable to process CDA natively but is a subscriber to an “agency” document management service that can render CDA into a form that “C” can use. The rendering service is proven via the IHE Connectathon process and natively implements a listener for the IHE XDR protocol.
Organisation ABC	Some external organisation involved in the follow-on care for the patient. Implements a native ITK listener for the SendCDADocument-v2-0 service.
urn:nhs-uk:addressing:ods:ABC:D	External recipient in organisation ABC, can process CDA.
urn:nhs-uk:addressing:ods:ABC:E	A second external recipient in organisation ABC, can process CDA. This is the old “team F” from RHM.
urn:nhs-uk:addressing:ods:RHM:F	Department in RHM which has been transferred to organisation ABC, and is known as “team E” there, with an address “urn:nhs-uk:addressing:ods:ABC:E”
Rendering Service	A notional service which, for example, runs a document management system that converts CDA clinical content to PDF and manages access for clients.

Table 1 : ITK Addressing Examples

For A to send the message is straightforward. All routing requests from A are sent via an ITK request to the URL https://rhm-internal-relay.rhm.nhs.uk/relay. This configuration is similar to the way Spine services are set up, with a fixed URL providing the service, or to the way email clients are set up with an “SMTP host” to which outbound messages are sent.

Because “RHM-internal-relay” is configured in the routing “relay” mode, it will route all messages from those requestors authorised to use it. Its routing table looks like:

Address	Service	Type	Route Data
urn:nhs-uk:addressing:ods:RHM:A	*	ITK-Queue	Local:A
urn:nhs-uk:addressing:ods:RHM:B	*	ITK-Queue	Local:B
*	*	ITK-Webservice	https://itk.rhm.nhs.uk/outbound

Table 2 : ITK Routing Table Examples

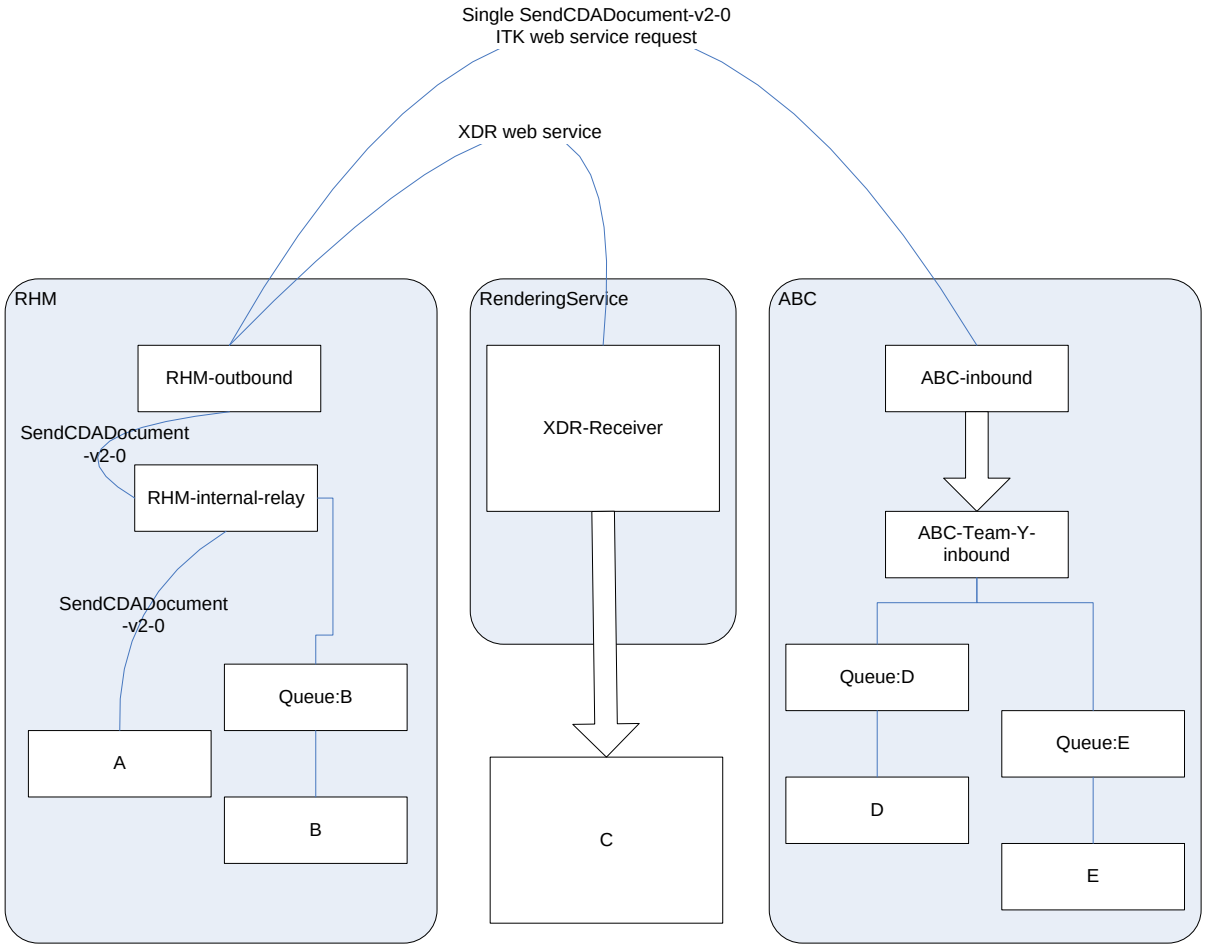

Figure 3 – Fully Worked Example

The relay will resolve “Queue:B” for the recipient “urn:nhs-uk:addressing:RHM:B” and deliver to that queue. Recipient B will collect the message from that queue, the next time it polls.

All the other addresses match the “wildcard” only, and so the RHM-internal-relay re-lays the SendCDADocument-v2-0 to the RHM-outbound router. On receipt, the outbound router will not attempt to deliver to recipient “urn:nhs-uk:addressing:RHM:B” again, because it is authoritative for that address and outbound routers only route to non-authoritative addresses.

The RHM-outbound routing table is:

Address		Service	Type	Route Data
urn:nhs-uk:addressing:ods:RHM:*	**A**			
urn:nhs-uk:addressing:ods:RHM:F	N	*	Redirect	urn:nhs-uk:addressing:ods:ABC:E
urn:nhs-uk:addressing:ods:ABC:*	N	*	ITK-Webservice	https://itk.abc.nhs.uk/inbound
urn:nhs-uk:addressing:ods:B83019:*	N	SendCDADocument-v2-0	XDR	https://renderingservice.nhs.uk/xdr
urn:nhs-uk:addressing:ods:B83019:*	N	SendDocument-v1-0	ITK-Webservice	https://itk.b83019.nhs.uk/inbound
*	N	*	Block	Unresolvable address

Table 3 : ITK Routing Examples

It is best to look at how the router handles addresses. The routing table entries can be classified by target organisation:

1. For addresses that no other route matches, mark the address as blocked and return an error message saying that the address is unresolvable[1].

2. Apart from the special case of “urn:nhs-uk:addressing:ods:RHM:F”, all RHM addresses are “inside” the RHM router and it is authoritative for them. So, being an “outbound-configured” router, these are ignored so “B” does not get a second copy of the message.

3. The address “urn:nhs-uk:addressing:ods:RHM:F” is replaced by “urn:nhs-uk:addressing:ods:ABC:E” by the re-direction special route type, and this new address is resolved to an ITK webservice at the URL https://itk.abc.nhs.uk/inbound - and the address “urn:nhs-uk:addressing:ods:ABC:D” resolves to the same URL. So a single ITK web service call is made to the SendCDADocument-v2-0 on organisation ABC’s inbound router[2].

[1] This is an example where the “universe” of addressable organisations is very small, so it is reasonable to block anything we don’t match. In practice, it is to handle cases such as this “everything else” type of entry, that the physical route over Spine transport, attached to a “common content” message, is envisaged. “Route data” in this case would be resolved by a function call that performed a lookup on an SDS cache.

[2] In this “small universe” example, RHM’s outbound router would be configured with knowledge of ABC’s inbound router by mutual agreement. More scalably, each organisation that hosts an entry point or inbound router for ITK routed messages would publish an accessible route to it. Such an accessible route would either be via ITK web services over N3, or possibly some other physical transport. There are many advantages to using the existing SDS as a vehicle for publishing such routes.

4. The original service request is for SendCDADocument-v2-0, and the address “urn:nhs-uk:addressing:ods:B83019:C” matches two routing table entries. So the RHM outbound router checks the service and finds that to send SendCDADocument-v2-0, it must convert the ITK distribution envelope into an IHE XDR format, and make an XDR protocol call to the URL https://renderingservice.nhs.uk/xdr. The rendering service is considered responsible for getting the content to recipient “C” in a suitable form[3].

On arrival at organisation ABC’s inbound router, the routing table is:

Address		Service	Type	Route Data
urn:nhs-uk:addressing:ods:ABC:*	A	*	URL	https://ABC-team-y-inbound.internal.abc.nhs.uk/inbound
urn:nhs-uk:addressing:ods:RHM:F	A	*	Redirect	urn:nhs-uk:addressing:ods:ABC:E
urn:nhs-uk:addressing:ods:RHM:*	N	*	ITK-Webservice	https://itk.rhm.nhs.uk/inbound
urn:nhs-uk:addressing:ods:B83019:*	N	SendCDADocument-v2-0	XDR	https://renderingservice.nhs.uk/xdr
urn:nhs-uk:addressing:ods:B83019:*	N	SendDocument-v1-0	ITK-Webservice	https://itk.b83019.nhs.uk/inbound
*	N	*	Block	Unresolvable address

Table 4: ITK Routing Examples

As a router configured “inbound”, all “non-authoritative” routes are ignored. So only the matches for “urn:nhs-uk:addressing:ods:ABC:*” and “urn:nhs-uk:addressing:ods:RHM:F” are processed. The “RHM:F” re-direction works as before, and a single ITK web service call is made to the organisation ABC inbound router at the URL https://ABC-team-y-inbound.internal.abc.nhs.uk/inbound.

[3] Note that had the content been some format other than CDA, and the SendDocument-v1-0 ITK service had been called, the RHM outbound router would have been able to use ITK web service forwarding.

This final router has the table:

Address		Service	Type	Route Data
urn:nhs-uk:addressing:ods:ABC:D	A	*	ITK-Queue	Queue:D
urn:nhs-uk:addressing:ods:ABC:E	A	*	ITK-Queue	Queue:E
urn:nhs-uk:addressing:ods:RHM:F	A	*	Redirect	urn:nhs-uk:addressing:ods:ABC:E
*	N	*	ITK Webservice	https://itk.abc.nhs.uk/outbound

Table 5: ITK Routing Examples

Again this is an inbound-configured router, so any matches that are not authoritative are ignored. There are three authoritative routes. The message address list is still addressed to “urn:nhs-uk:addressing:ods:RHM:F” but this is re-directed to “urn:nhs-uk:addressing:ods:E”, for which there is an explicit authoritative route. Similarly, an explicit authoritative route exists for “urn:nhs-uk:addressing:ods:ABC:D”. Both these routes are handled by ITK queues, so the messages are collected by final recipients D and E the next time they poll.

* * * End of Document * * *