

Document filename: ITK 2 0 Architecture Overview v1.0.docx			
Directorate / Programme :	HSCIC - Architecture	Project	Interoperability
Document Reference :		HSCIC-ITK-ARCH-100	
Project Manager :	Rob Shaw	Status :	Final
Owner :	George Hope	Document Version :	1.0
Author :	George Hope	Version issue date :	23/06/2014

Interoperability Toolkit - Architecture Overview

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	5
1.1	Purpose of Document	5
1.2	ITK Architecture Documentation Set	5
1.3	Audience	5
1.4	Document Scope	6
1.5	Document Overview	6
1.6	Reference Implementation	6
2	What is NHS Interoperability Toolkit	7
2.1	Architecture Specifications and Guidance Documents for Health and Social Care Interoperability	7
2.2	Domain Message Specifications for Healthcare Interoperability	7
2.3	Trust Operating Model	7
2.4	Resources to aid Development and Deployment of ITK Accredited Software	8
3	Interoperability Contexts	9
3.1	Local Intra-Trust Context	9
3.2	Local Health and Social Care Community Context	9
3.3	National Context	10
4	ITK Architecture Principles and Architecture Design Decisions	11
4.1	ITK Architecture Principles	11
4.2	ITK Architecture - Key Design Decisions	11
5	ITK Logical Architecture	13
5.1	ITK Architecture – Logical Application Components	13
6	ITK Message Structures	15
6.1	Distribution Envelope - Messaging Configuration	16
6.2	Use of Vocabularies	16
6.3	Multiple Addressees	16
7	ITK Messaging Configurations	17
7.1	ITK Messaging Configurations - Request	17
7.2	ITK Messaging Configurations – Request / Response	18
7.3	Queue Collection	19

8	ITK Acknowledgement Framework	22
8.1	Introduction	22
8.2	Infrastructure Acknowledgements	23
8.3	Business Acknowledgements	23
8.4	Enabling the Acknowledgement Framework	23
9	ITK Capabilities	24
9.1	Discovery	24
9.2	Message Versioning	24
9.3	Monitoring and Management	24
9.4	Reliability and Error Handling	24
9.5	ITK Architecture – Addressing and Routing	25
9.6	Security	26
9.7	Throttling	26
9.8	Transformation and Mediation	27
9.9	Validation	27
10	ITK - Components of a Service	28
10.1	ITK Service Layers	28

1 Introduction

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

1.1 Purpose of Document

The purpose of this document is to provide a scene setting overview of NHS's Interoperability Toolkit (ITK).

1.2 ITK Architecture Documentation Set

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

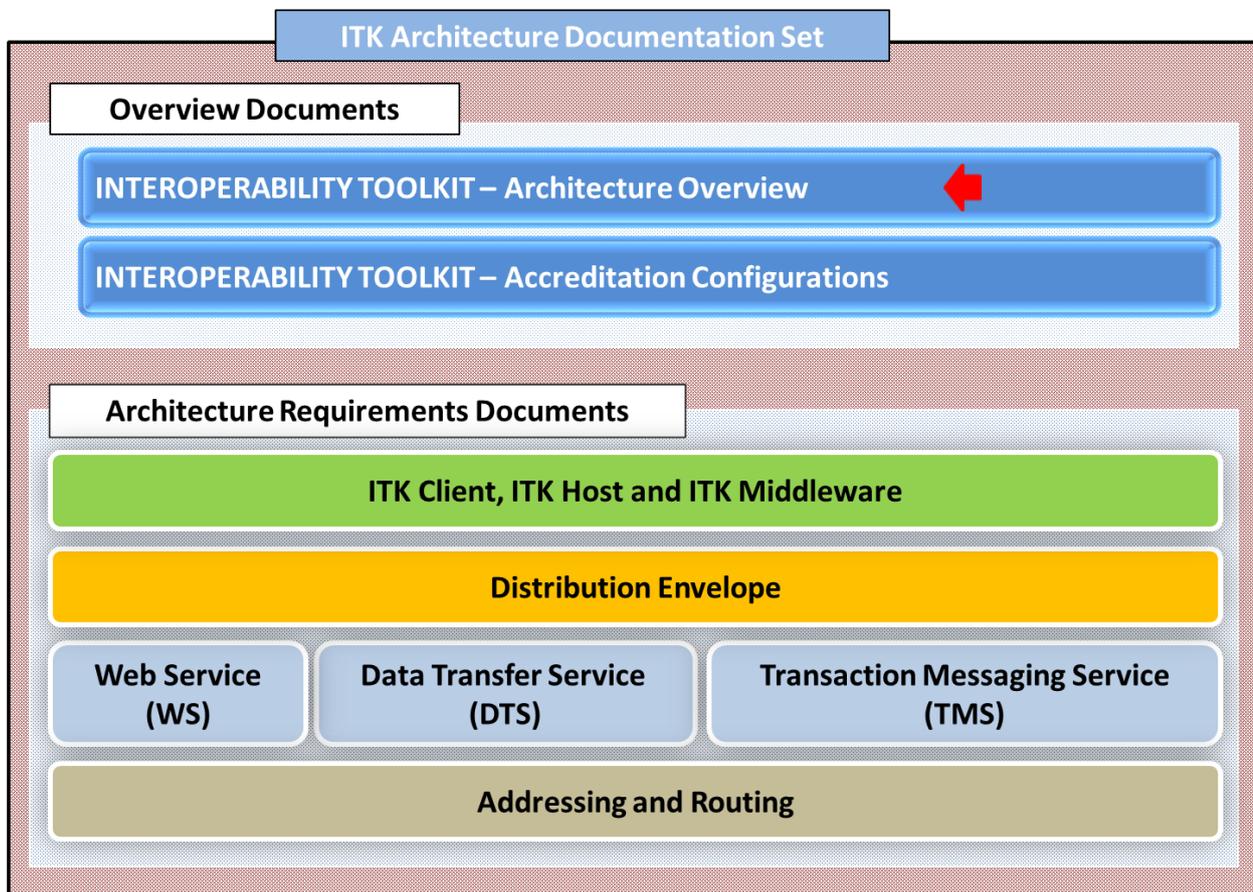


Figure 1 – ITK Architecture Documentation Set

1.3 Audience

The document is intended for a mixed technical audience of individuals within an organisation of all levels intending to develop ITK Accredited Systems, as well as those in the Health and Social Care Architecture communities.

1.4 Document Scope

The document covers the ITK Architecture in its entirety.

1.5 Document Overview

This document covers the following topics:

- Toolkit Logical Architecture
- Toolkit Messaging Configurations and Structures
- Toolkit Acknowledgement Framework
- Toolkit Capabilities and Service Definitions

1.6 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 What is NHS Interoperability Toolkit

The interoperability toolkit comprises of centrally provided four sets of documents and associated resources. The details of the said sets are as follows.

2.1 Architecture Specifications and Guidance Documents for Health and Social Care Interoperability

The ITK Architecture Specifications (documents) define the mandatory and optional architecture requirements of an ITK compliant implementation. These documents are message payload agnostic and have no clinical meaning. Their collective aim is to assured and auditable delivery of Health and Social Care Information between other Health and Social Care IT Systems securely and safely.

Typically they include requirements regarding how systems connect and interoperate, requirements for data security, requirements for payload distribution and routing.

2.2 Domain Message Specifications for Healthcare Interoperability

Domain Message Specifications (DMS) specify message payloads and interaction details for healthcare domains. A DMS contains all the relevant artefacts to build message payloads, which are based on international standards such as HL7 CDA. The payload provides a means to share healthcare information in a clinically safe manner. DMS message payloads are transport agnostic.

DMSs also contain interaction details, which describe behaviours for sending and receiving messages. For example requesting and receiving acknowledgements and accept/reject message configurations.

2.3 Trust Operating Model

The ITK Trust Operating Model provides best-practice guidance which Trusts need to consider when assuring their own Architectures. It provides assistance with the implementation of integrated systems within a local environment. In particular it lays out local responsibilities when connecting to Spine compliant systems. The ITK Trust Operating Model contains a guidance document and best practice checklist which covers the critical operational aspects such as: resilience, performance, installation, configuration etc.

2.4 Resources to aid Development and Deployment of ITK Accredited Software

In order to support the accreditation and deployment of clinically safe interoperable IT systems, there is a need to supply resources that cater both Vendors and Trusts. Typical resources being:

- Accreditation Support
- Automated Test Workbenches
- Sample Code
- Technical Guidelines.

For Accreditation and Deployment support materials, please see here:

<http://developer.nhs.uk/testcentre/>

For reference implementation, please see here:

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

3 Interoperability Contexts

For the purposes of this document, Interoperability is defined as a Health and Social Care application's ability to interact with other Health and Social Care systems and applications using common standards and specifications.

The following, are the common contexts for using the ITK Architecture Specifications:

3.1 Local Intra-Trust Context

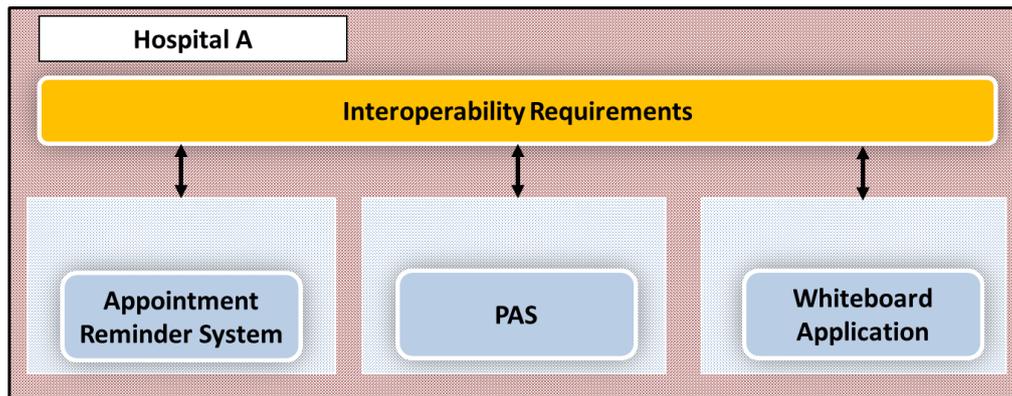


Figure 2 - Local Intra-Trust Context

Here the interoperability requirements may not require any sophisticated levels of service access security, as the services are located, for example, within a Trust's fire wall secured infrastructure and primarily externalise non-sensitive data within trusted systems.

3.2 Local Health and Social Care Community Context

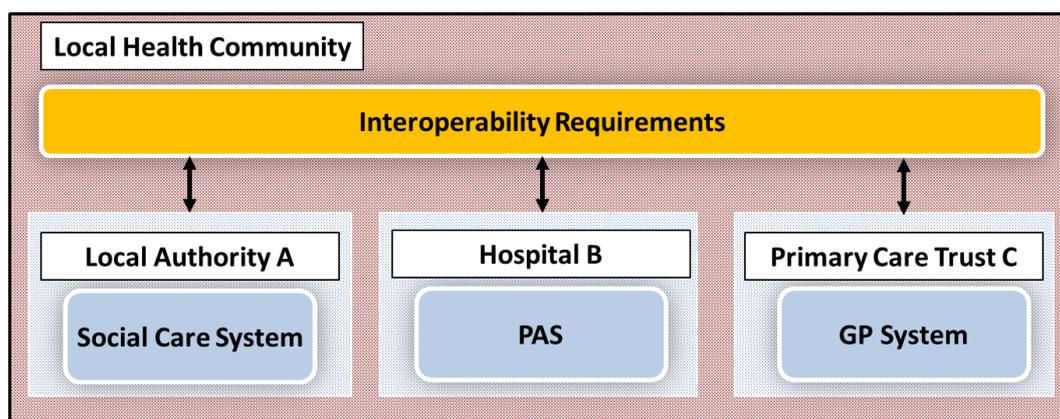


Figure 3 - Local Health and Social Care Community Context

This context is where functionality and data are shared across a local health economy involving a wider range of care settings, for example, acute and primary care providers. This requires, among other things, greater levels of security and more sophisticated routing mechanisms.

3.3 National Context

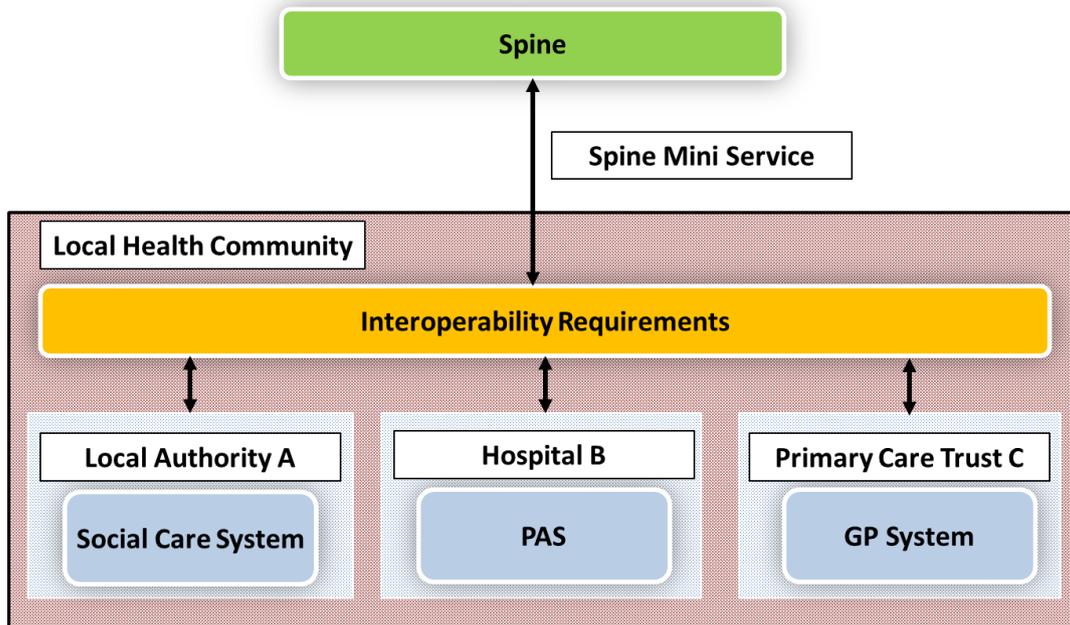


Figure 4 - National Context

Finally, in cases requiring integration with national systems, Spine Mini Service, or the establishment of national services, the ITK requirements are (and will be) at their most comprehensive.

4 ITK Architecture Principles and Architecture Design Decisions

The following section lists the principles and key design decisions used when defining the Interoperability Toolkit Specifications.

4.1 ITK Architecture Principles

- **Change and Existing Systems** - reduce the need for system change by loosely coupling new services and existing systems.
- **Infrastructure Services** - Non-service specific elements of a service (e.g. security, error reporting, routing, and validation) should be developed as generic and reusable functionality.
- **Minimum Standards** -The ITK will provide a minimum standards approach to interoperability.
- **Modular and Extensible** -Toolkit implementations will be modular and offer clear interfaces for extensibility of standard functionality.
- **Service Gateway** -Based on National Interoperability Standards ITK implementations will provide a common gateway for all services offered by a Trust.
- **Service Implementation** -An ITK service interface ensures that the client application does not require knowledge of the host service location or implementation.
- **Standards Based Specifications** -The Toolkit will use Open Standards where ever possible, unless there are good reasons why these standards should not be used.

4.2 ITK Architecture - Key Design Decisions

The following section provides at a more technical level a list of architecture Key Design Decisions (KDD).

- **Application Based Security** - Initial security scenarios are based on Application security. This means that services are invoked by trusted applications – as opposed to end-users themselves invoking services in their own name.
- **Common Messaging Headers** - The Toolkit will define a set of message headers based on SOAP and WS-* standards.
- **Configuration** - Local Configurability, based on context: One Trust might choose to deploy within the context of a secure data centre, to connect two locally hosted applications. By way of contrast, a Local Health Community might connect a range of clinical applications across different organisations and care settings.
- **Context Sensitive Security** -It is important to determine (based on risk assessment) what security mechanisms are appropriate for a given implementation.
- **Flexible Payload Content** - The Toolkit will be flexible in terms of the standards used for payload content.
- **Flexible and Extensible Security** -The Toolkit will define a flexible security framework based on a combination of both transport-level security (e.g. TLS) and message-level security (e.g. WS-Security).

- **Integration Services** - An ITK Client will hand off how service calls are routed. This approach decouples senders and receivers, allowing complex scenarios to be supported, and isolates the Client and Host applications from change.
- **Separation of Messaging Layers** - The layers of a messaging solution are clearly identified and segregated, with explicit allowance for layers covering transport e.g. (HTTP), Service Request (SOAP), Distribution (Distribution Envelope) which includes Acknowledgement, Workflow (via Interaction Id) and payload. Separating concerns through layering enables a modular and extensible design.
- **Use of Open Standards for Interoperability** - The ITK will use of widely-adopted International Healthcare, Transport and Messaging standards e.g. HL7 and Web Services.

5 ITK Logical Architecture

The following diagram offers a logical view of the ITK Architecture. It should be noted that these logical components need not be physically separate. The following section contains a brief description of each component.

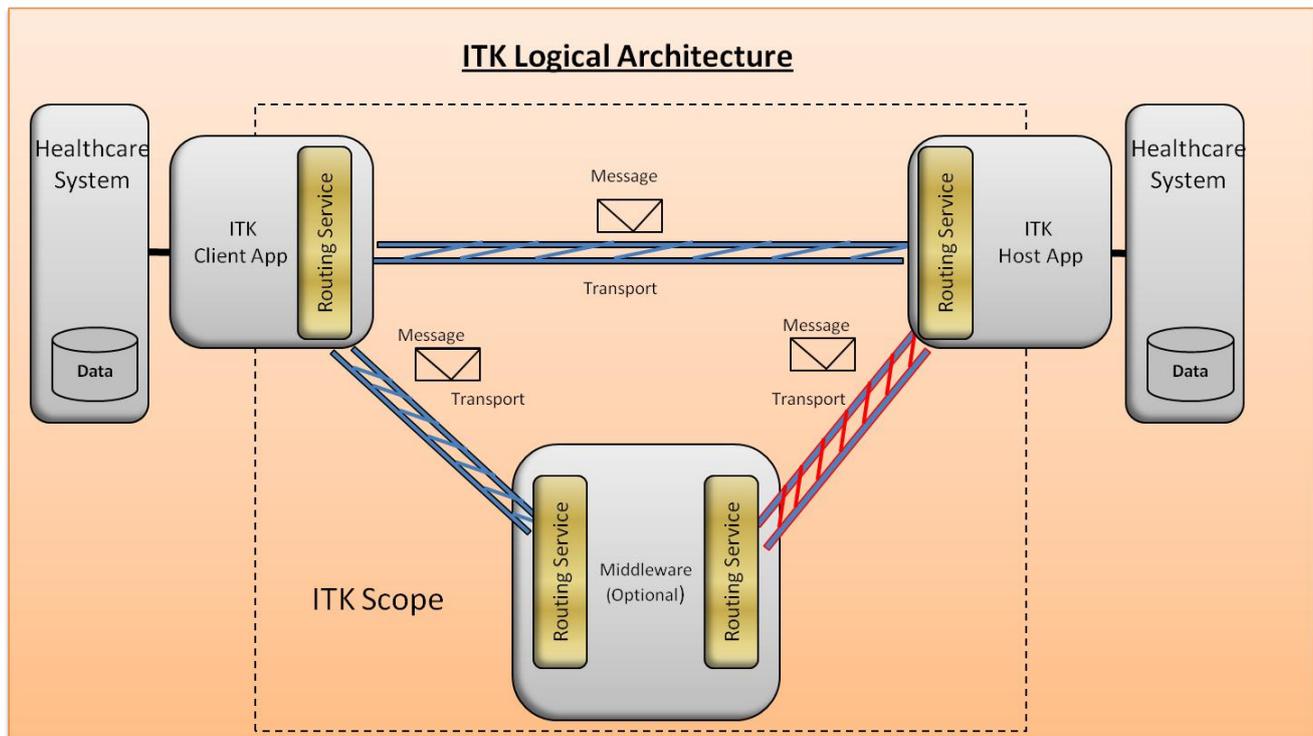


Figure 5 - The ITK Logical Architecture

5.1 ITK Architecture – Logical Application Components

5.1.1 ITK Client

The key characteristics of the ITK Client are: this component must be accredited to be called an ITK Client. The ITK Client is independent of the location of the service it requests. The Client is responsible for creating the ITK Compliant message.

5.1.2 Routing Service

The key characteristics of the Routing Service are: to receive the message produced by the ITK Client Application and to route the message to the ITK Host Application and vice versa. This component is not accredited separately but as part of either the ITK Client, ITK Host or ITK Middleware accreditations.

5.1.3 ITK Host

The key characteristics of the ITK Host are: this component must be accredited to be called an ITK Host. The ITK Host is independent of the location of the Client. The Host is responsible for consuming the ITK Compliant message.

5.1.4 ITK Middleware

The key characteristics of the ITK Middleware are: this component must be accredited to be called ITK Middleware, the ITK Middleware is independent of the location of the Client, or Host, is responsible for message routing, transformation, mediation, orchestration etc.

5.1.5 ITK Message

The key characteristics of the ITK Message are: it includes the Distribution Envelope which contains the payload and all the information required regarding: the address of a service, the payload and the payload processing requirements, allowing payload delivery to be transport independent.

5.1.6 ITK Payload

The ITK Payload is the business information/data which is passed between Health and Social Care organisations. These are fully defined within the individual Domain Message Specifications.

5.1.7 ITK Transport

The underlying technology used for transporting messages. The currently defined ITK transports are Data Transfer Service (DTS), Transaction Messaging Service (TMS) and Web Service (WS).

6 ITK Message Structures

The diagrams below shows the outline structure of an ITK Message, note that since the Distribution Envelope contains the ITK addressing and routing information required, it enables messages to flow across [multiple](#) transports protocols.

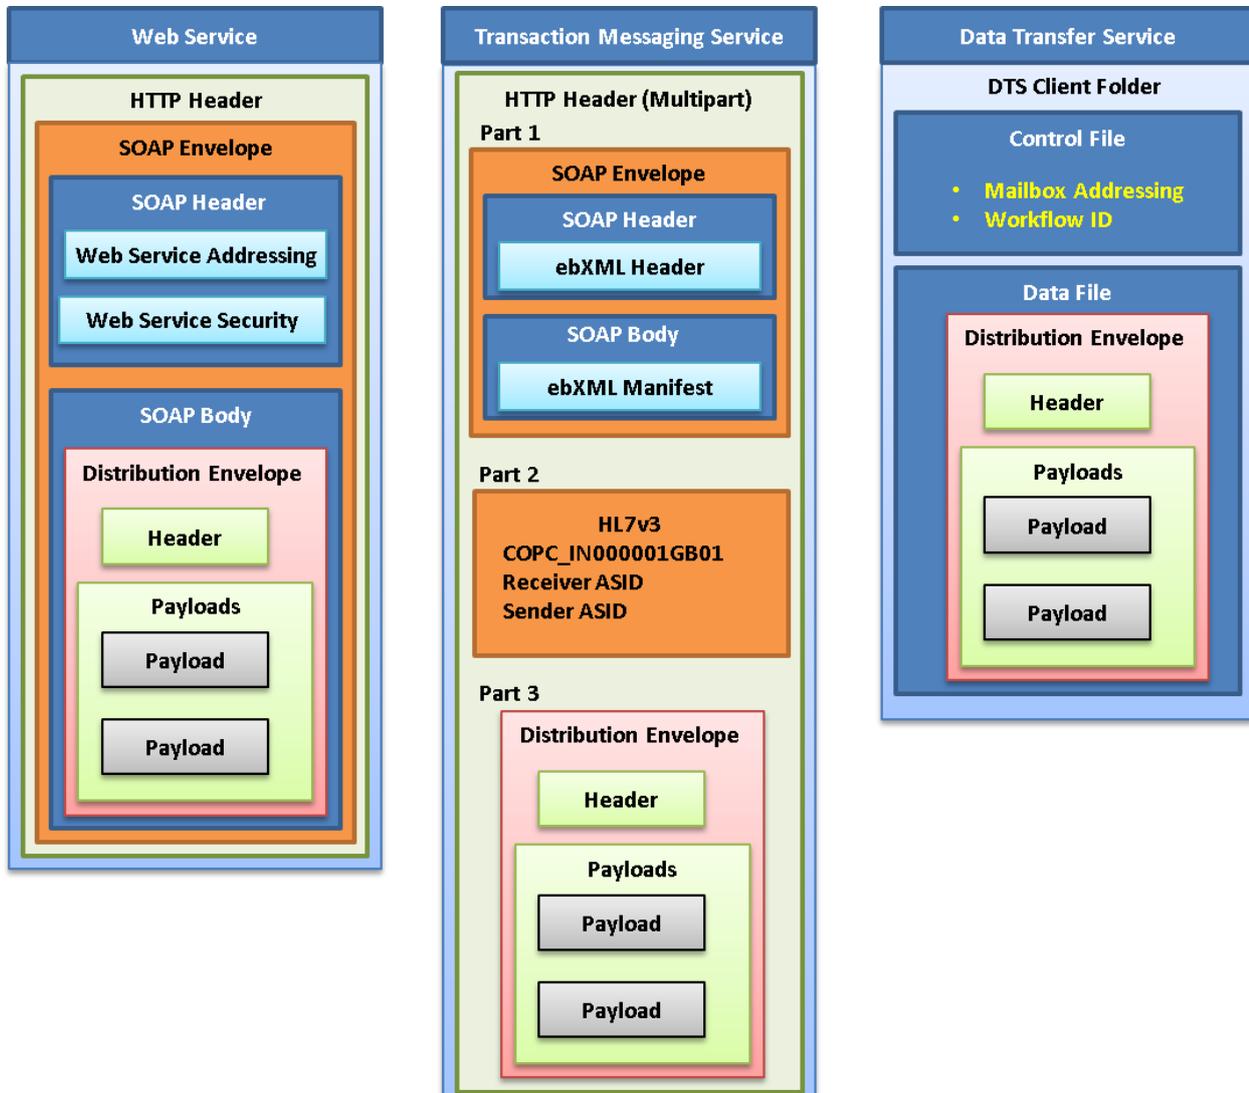


Figure 6 - The ITK Message Structures

Future ITK Transports (i.e. XDS) will where appropriate endeavour to use the Distribution Envelope as the common messaging component.

It should be noted that the intention will always be that the DE is consistent component across all transports.

6.1 Distribution Envelope - Messaging Configuration

The distribution envelope is used to define the behaviours for messaging. There are three flags in the handling specification sections.

The following table shows the generalised messaging configurations in relation to various messaging formats.

Messaging Format	Distribution Envelope Handling Specification Configuration		
	Infrastructure Ack	Business Ack	Business Response
HL7v2	true or false**	true or false**	false
HL7v3	true or false	true or false	true or false*
CDA	true or false	true or false	true or false*
Other Formats	true or false	true or false	true or false

* Defined in the Domain Message Specification's Interactions

** HL7v2 acknowledgement (acts as both infrastructure and business acknowledgement)

Table 1: Distribution Envelope Handling Specifications for Messaging Formats

6.2 Use of Vocabularies

Data items, which are restricted in terms of sets of values they may contain, require vocabularies to define and manage the allowable values. The required vocabularies will be managed by HSCIC as part of service definition. Examples of vocabularies include

- Standard Technical Errors
- Service specific error codes
- Service specific business vocabularies

6.3 Multiple Addressees

One ITK Message contains one and only one Distribution Envelope (DE). The ITK distribution envelope supports multiple addressees. The routing service must send only one copy of the message to each address.

7 ITK Messaging Configurations

There are a number of configurations that can be used in an ITK implementation, and the invocation styles can be Synchronous or Asynchronous.

ITK Messaging Configurations support both stateless and stateful business processes.

- Stateless implementations support high throughput and low latency requirement scenarios
- Stateful implementation supports more complex integration scenarios, such as message sequencing, subscription, integration process orchestration, or long running integration processes.

The Service Listings provides the messaging configurations that align with the DMS and can be found in the accreditation pack.

Requirements for Queue Collection are defined where an asynchronous delayed response is required, useful for example for batch processing.

7.1 ITK Messaging Configurations - Request

The Requestor provides a message for processing where no business response is required. An example might be a “notification” message. The Toolkit supports the Request configuration by the appropriate setting of flag within the Distribution Envelope. This is used for “one-way” calls where a response is not expected.



Figure 7 – Request Configuration

Handling Specification	Value
urn:nhs-itk:ns:201005:infackrequested	false
urn:nhs-itk:ns:201005:ackrequested	false
urn:nhs-itk:ns:201005:busresponserequested	false

Table 2: Request Configuration flag settings

Note: a “dummy” response may be returned by the provider to satisfy the completeness requirements of the underlying transport protocol.

For example, the Web Service implementation uses a dummy “SimpleITKResponse” to enable all ITK interactions to be implemented using widely supported request/response SOAP calls.

Where a request fails – due for example to a transport fault such as inability to resolve a remote provider – the requestor must throw an exception which is defined in the appropriate transport specification.

7.2 ITK Messaging Configurations – Request / Response

The Requestor provides a message for processing and the Provider provides information back as a result. An example might be a query which returns the requested information. Or it might be confirmation that an update will be processed.

Here the Toolkit supports the Request/Response configuration by setting the the appropriate flag within the Distribution Envelope.

A requestor calls a provider and gets either a response, or an exception.

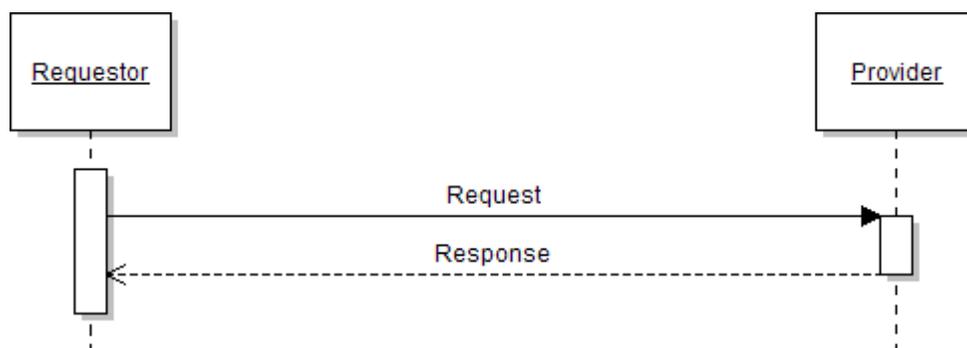


Figure 8 - Request / Response Configuration

Handling Specification	Value
urn:nhs-itk:ns:201005:infackrequested	true or false
urn:nhs-itk:ns:201005:ackrequested	true or false
urn:nhs-itk:ns:201005:busresponserequested	true

Table 3: Request / Response Configuration flag settings

Exceptions may be raised for technical reasons – malformed requests, security or access control rejections and the like.

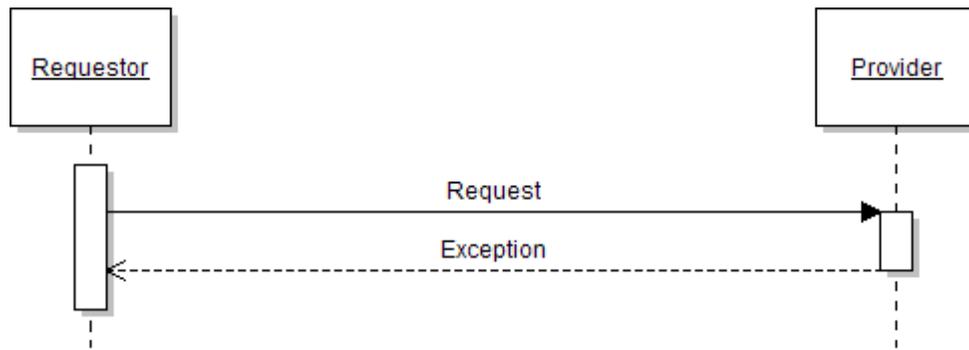


Figure 9 - Exception returned by provider

It is an error for neither a response nor an exception to be returned, however it is the requestor’s responsibility to detect this – for example due to request timeout, inability to resolve or contact the provider and so on.

7.3 Queue Collection

The ITK specifies a “queue collection” capability where a provider maintains a store of asynchronous messages for a requestor, and from which the requestor can retrieve the messages.

The queue may be populated either as a means of delivering asynchronous responses, or “unsolicited” content such as the delivery of clinical documents.

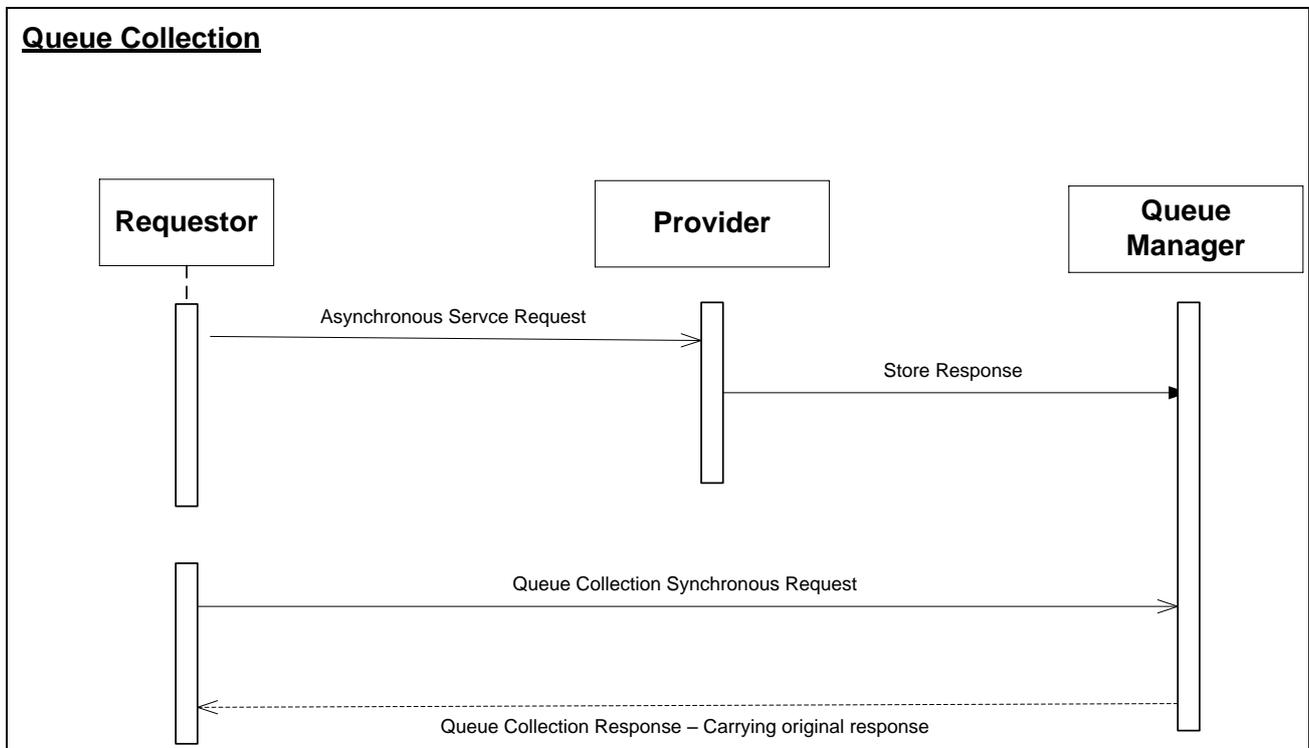


Figure 10 - Queue collection

When populated by an asynchronous response, either the requestor can specify the queue as a response address, or the provider can intercept the response based on knowledge of the requestor preferences. The ways in which these are done is transport-specific¹. Delivery to a queue completes service request handling, as the queue acts as a proxy for the requestor.

7.3.1 Queue Names

Queues are generally expected to be local and no explicit ITK queue naming rules are enforced. It is not expected that ITK implementations outside of an organisation should be able to reference a queue directly – that task is handled by the optional ITK routing specification.

It is RECOMMENDED that queue names be valid URIs as this meets the requirements of WS-Addressing.

7.3.2 Queue Properties

Queue implementations SHOULD support message sequencing, and MUST support reliable retrieval where a message is removed from the queue only on receiving of confirmation that the requestor has accepted it.

¹ For example: In SOAP 1.1 the asynchronous reply address is given in the WS-Addressing “ReplyTo” header of the request. This reply address may be a URL to call on the requestor, or a queue name. If a queue name, the provider would write the response to the queue. Alternatively the provider may implement a rule saying “anything for URL ‘X’ should go to queue ‘Y’”.

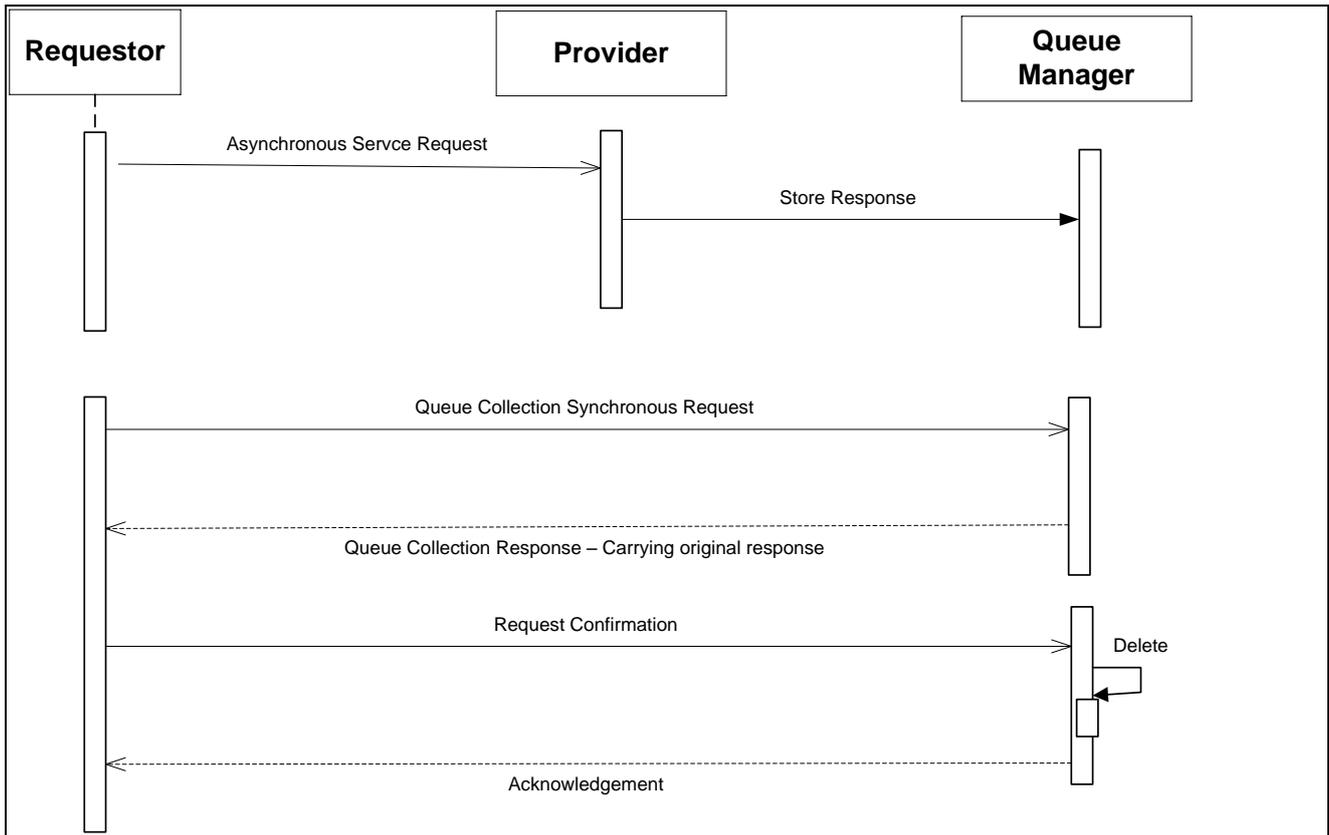


Figure 11 - Queue collection for general delivery

8 ITK Acknowledgement Framework

8.1 Introduction

The ITK Acknowledgement Framework is simply implemented by ITK Components returning messages to the message sender, to inform the sender of the status of the message they have sent.

The ITK Acknowledgement Framework is primarily designed for a multi-hop routing, and provides a sender with a reliable view of the state of a payload’s transmission.

Acknowledgement Framework

- Note: 1. Only the Endpoint can send a Business Acknowledgement
- 2. Infrastructure Acks can be generated by End Points or Intermediaries

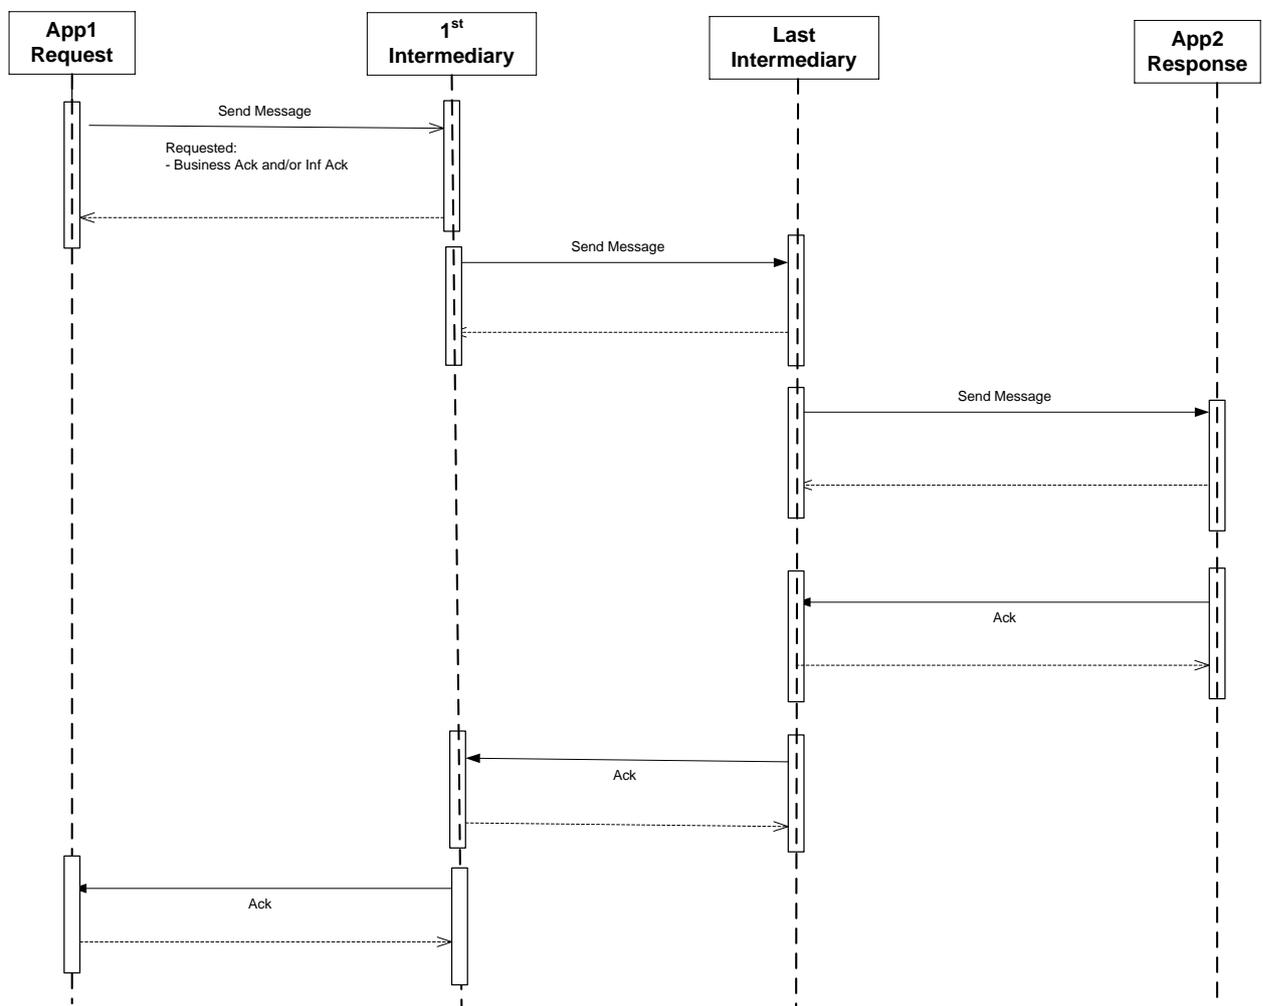


Figure 12 - ITK Acknowledgement Framework

To facilitate this capability the ITK defines the XML messages to be used for the Infrastructure and Business Acknowledgements.

End-to-end acknowledgments (ACKs) or negative acknowledgments (NACKs) are delivered in a Distribution Envelope(s) as separate messages.

8.2 Infrastructure Acknowledgements

The purpose of the Infrastructure Acknowledgement is to inform the sender that the payload has been delivered to its final destination.

8.3 Business Acknowledgements

The purpose of the Business Acknowledgement is DMS/payload/context sensitive and can for example inform the sender that the payload has been processed and/or delivered to a clinician for review. A Business Acknowledgement may implement the following:

- Received document at the correct practice
- Patient matched
- Correspondence seen/not seen by GP/Clinician
- Correspondence filled

8.4 Enabling the Acknowledgement Framework

To enable the Acknowledgement Framework the ITK Distribution Envelope:

- Requires a routable sender address, routers and receiver(s) MUST use this address for the ACKs/ NACKs.
- Requires the `urn:nhs-itk:ns:201005:infackrequested` handling specification entry, be set to true/false for an Infrastructure Acknowledgement.
- Requires the `urn:nhs-itk:ns:201005:ackrequested` handling specification entry, be set to true/false for a Business Acknowledgement.
- Requires the `urn:nhs-itk:ns:201005:busresponserequested` handling specification entry, be set to true/false for a Business Response.
- Requires the `urn:nhs-itk:ns:201005:interaction` Interaction Id entry, to be set.

In all cases should the infrastructure fail to deliver or return messages the message sender needs to be able to handle a time-out situation.

9 ITK Capabilities

The ITK defines a set of capabilities that the accredited systems need to implement. These capabilities ensure that from an architecture perspective ITK Accredited Systems are: secure, extensible, scalable, resilient and ready for use and deployment in any Health and Social Care setting.

9.1 Discovery

Service discovery is an optional part of a Toolkit implementation and may be achieved using a Universal Description, Discovery and Integration (UDDI) registry.

9.2 Message Versioning

Message versioning is a mandatory requirement of a Toolkit implementation. ITK message headers provide an attribute to hold the message versions.

9.3 Monitoring and Management

Monitoring and management of the interoperability infrastructure and eco system is a responsibility of Toolkit implementations.

As a minimum, the applications participating in the interoperability scenarios must support some form of message/error logging to support end-to-end operational management capabilities.

9.4 Reliability and Error Handling

9.4.1 Reliability

The reliability of information/message exchange between an ITK Client and an ITK Host is based on implementation of acknowledgments at Infrastructure and Business levels. The receipt of the acknowledgement(s) allows the receiver to assume responsibility for processing and possible onward delivery.

The requirements for handling duplicate messages, for example resulting from an application's incorrect processing are defined in the Toolkit. For example, it is the responsibility of the ITK Host's application logic to manage duplicates.

9.4.2 Error Handling

Errors that occur during end-to-end integration scenarios fall into two categories.

- Application errors, components in the end-to-end ITK implementation must be capable of receiving and handling exceptions.
- System errors, typically the system administration layer will provide error capture and handling capabilities.

9.5 ITK Architecture – Addressing and Routing

The ITK specifies the means by which addressing and routing are implemented.

This covers “addresses”, “addressing” and “routing”. It is critical to an understanding of the ITK that these related terms are defined carefully.

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 resolution 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.

Actually to transfer a message from sender to recipient involves finding a path, which may be a sequence of physical routes, between the sender’s and the recipients’ systems. Routing is the overall process by which this happens, and the systems which perform the routing functions are routers.

9.5.1 Routing

The Toolkit defines responsibilities regarding routing. ITK Client routing requests should be agnostic to the service being called.

The Toolkit defines how to route service requests. This decouples the ITK Client from the ITK Host, insulating service implementation changes, and allowing complex service scenarios to be designed.

A generic routing component (e.g. in middleware) may then use the address (in a standard URI format) thereby resolving a logical address across multiple physical “hops”.

9.5.2 Sequencing

From an interoperability perspective it is helpful to consider sequencing in two distinct parts:

- **Sorting messages into the correct order (Business Sequencing)**

This requires an understanding of the sequencing context – for example that events for the same patient are related, but those for different patients are independent. It may also require knowledge of business rules. The necessary handling for out-of-sequence events will be use-case specific.

- **Preserving the order once sorted (Technical Processing)**

By contrast this involves preserving an orderly flow of messages, such that, once sorted, they do not become “jumbled” in transit. This is independent of any business context or rules, and consists purely of preserving the technical message flow. This feature must be configurable - as it generally involves serialising processing and therefore implies a trade-off in terms of throughput.

- **Reliable Messaging**

Currently the ITK uses the simpler approach of awaiting acknowledgement of one call before beginning the next.

Sequencing is an end-to-end, system consideration, with potential responsibilities for all components involved in the message flow.

9.6 Security

Security is clearly an essential aspect of the Toolkit, and standards-based security requirements are provided as an enabler to interoperability. The ITK approach to security is based at the Transport Level and Application Level.

Initially the approach taken by the ITK to security is based on providing capabilities for two common scenarios:

- Scenario 1 (Trusted Network)– minimal security based on low-sensitivity data exchanges within a secure environment. This is intended to be quick and simple to implement, and to be suitable for a limited, but significant, range of scenarios within a Trust.
- Scenario 2 (Un-trusted Network) – more extensive security, based on exchanging potentially sensitive data and transiting un-trusted infrastructure (e.g. N3). Using digital certificates to secure the exchange, with options for controls to be applied at both the infrastructure level (TLS mutual authentication) and the messaging level (XML Encryption).

The ITK approach is extensible and aims to accommodate other more detailed scenarios as required.

Toolkit implementations must include functionality to secure service invocation against unauthorised access.

- The ITK Client is wholly responsible for performing authentication, authorisation, and audit of the requester.
- The ITK Host is responsible for performing authentication, authorisation, and audit of the invoking application.
- Toolkit implementations must also include functionality to secure messages in transit:
- The ITK Client and ITK Host are jointly responsible for agreeing how a message will be secured in transit to ensure any confidentiality and integrity requirements are met.

9.7 Throttling

In an ITK context throttling describes a system's ability to manage the rate at which it sends or receives messages. This is important for situations where, due to its design or deployment, a system is unable to process more than a certain number of instructions during any given time period.

9.8 Transformation and Mediation

Mediation describes the function to act as the go-between or broker between applications. This includes the ability to do any combinations of translation, transformation, conversion and enrichment of information between network protocols, messaging formats and message structures. The ITK defines a set of mediation requirements for successful interoperability.

9.9 Validation

The appropriate level of validation will always depend on the specifics of the interoperability scenario and the extent to which the participating applications are under the control of the system integrator. The general principle is that as much validation as possible is carried out close to the initial service invocation. This reduces redundant processing resulting from message handling for invalid invocation details

10 ITK - Components of a Service

The Toolkit's view of integration is defined by a series of services that are invoked by ITK Clients and implemented by ITK Hosts.

To date integration in the NHS, especially integration across system and organisational boundaries, has been dominated by message based integration concepts and mechanisms. As a result, most integration requirements are expressed as messages being exchanged between systems and applications.

10.1 ITK Service Layers

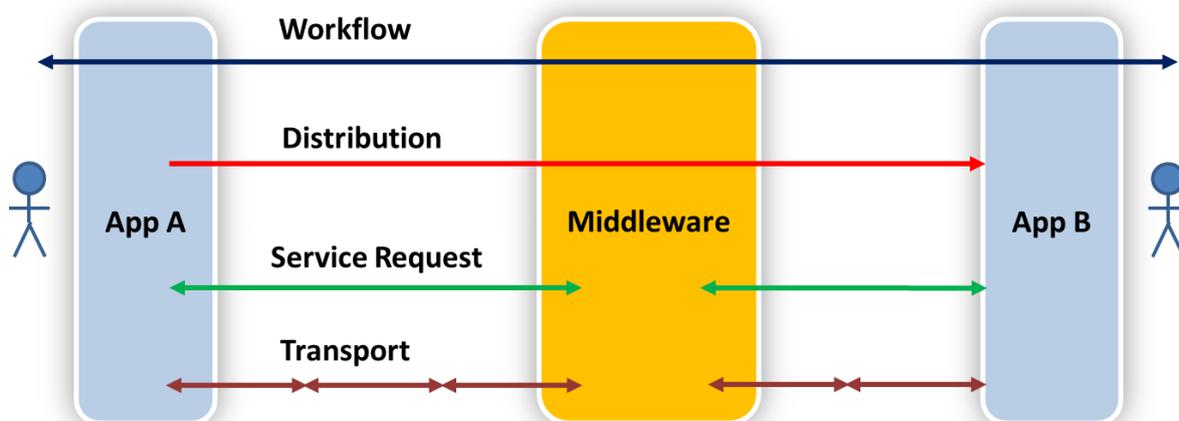


Figure 13 - ITK Service Layers

The diagram above shows the layers in the ITK Service Stack. These are:

- Transport**
 This is the low-level technical mechanism for moving the message from the source to the destination(s). There might potentially be many hops between low-level transport devices.
- Service Request**
 This layer is covered by widely accepted services standards, e.g. SOAP Headers. The SOAP Headers contain addressing, identity and security information relating to a Service Request “hop”. To go from A to B via the Middleware is considered as two “hops” – i.e. from A to the Middleware is a Service Request in its own right.
- Distribution**
 This is the end-to-end distribution of a message, from a sender to receiver(s).
- Workflow**
 This layer represents the business workflow. For example, there may be a long-running orchestration where the original message triggers a dialogue of further related messages (in both directions).

* * * End of Document * * *