
Document filename: ITK 2 0 Trust Operating Model Architecture Guidance v1.0.docx			
Directorate / Programme :	HSCIC - Architecture	Project	Interoperability
Document Reference :		HSCIC-ITK-ARCH-201	
Project Manager :	Rob Shaw	Status :	Final
Owner :	George Hope	Document Version :	1.0
Author :	George Hope	Version issue date :	23/06/2014

ITK Trust Operating Model Architecture Guidance

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 Barnett	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	4
1.1	Purpose of Document	4
1.2	TOM Documentation Set	4
1.3	Audience	5
1.4	Document Scope	5
1.5	Document Overview	5
2	Integration Architecture Context, Configurations and Services	6
2.1	Business Context	6
2.2	Message Configurations	6
2.3	Logical Architecture	6
3	Introduction to Architecture Principles	7
4	Security Principles	8
4.1	Integrity	8
4.2	Chains of systems	9
4.3	Malware	10
5	Integration Processing Principles	11
5.1	Technical Orchestration	11
5.2	Validation	12
5.3	Transformation	12
5.4	Exception Handling	13
6	Configuration and Management Principles	15
6.1	Quality of Service	15
6.2	Service Management	15
6.3	Version Management	16
7	Information Architecture Principles	18
7.1	Use of Data Standards	18
7.2	Mapping from alternative standards	18
7.3	Migration from alternative standards	19
8	Architecture Guidance Summary	20
9	Appendix – Architecture Principles Summary	22

1 Introduction

This document forms part of the overall document set for the Interoperability Toolkit (ITK).

1.1 Purpose of Document

This document is part of the Trust Operating Model component of the Interoperability Toolkit. See the document “Trust Operating Model – Overview” for a more complete description of the document set.

This specific document provides guidance on the architecture-related aspects of Local Application Integration. It does this by providing a set of Architecture Principles to guide key design decisions in this space.



This document is intended to provide information to support the “Architecture” tab of the Self-Evaluation Checklist. Having read this document, the Architecture tab should be completed.

1.2 TOM Documentation Set

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

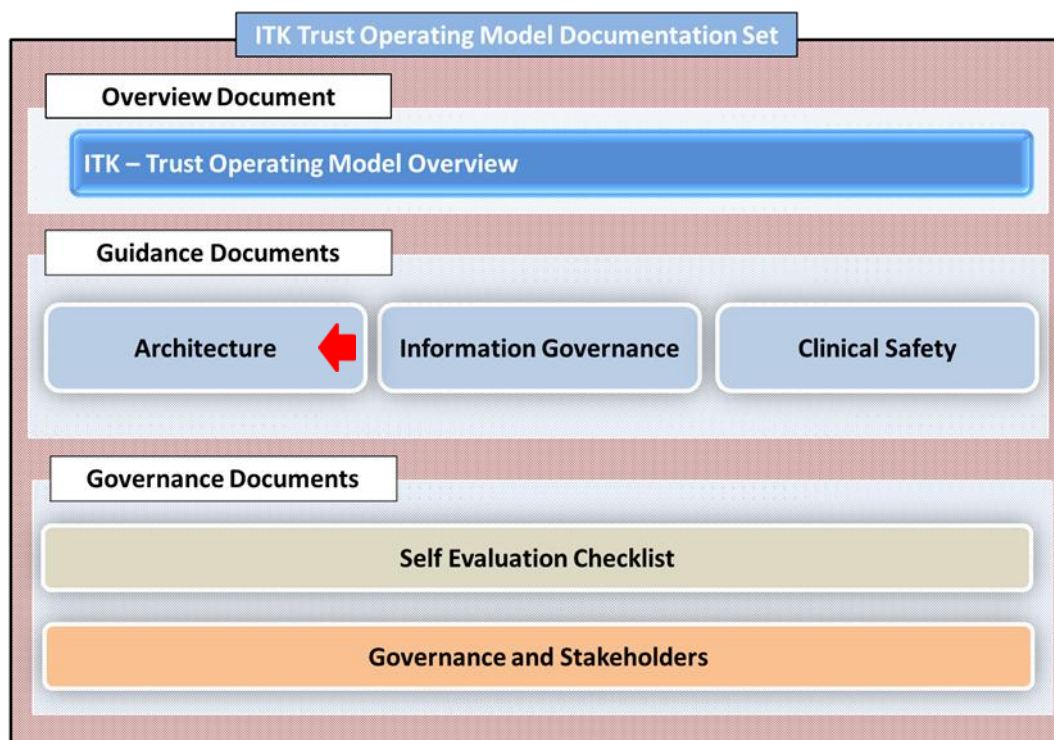


Figure 1 - The ITK Trust Operating Model Document Set

1.3 Audience

The primary audience for the Trust Operating Model is project teams within a Trust who are responsible for implementing a local integration project.

This document will be of particular relevance to architects and technical members of the project team.

Secondary audiences may include 3rd parties such as supplier and HSCIC architects

1.4 Document Scope

The Trust Operating Model focuses on integration between Local Trust Systems and Spine Compliant systems, and also on integration between Local Trust Systems and / or Non-NHS Systems within a Local Health Community environment. Please see the Overview document for further explanation of these concepts.

It does not cover integration at a National level through the Spine – existing Compliance documentation is already available on this topic.

Also note that the focus is on the integration-specific aspects of a project. General topics necessary for any successful project (e.g. training, communications, service management etc) are not covered.

1.5 Document Overview

The rest of this document covers the following topics:

- **Architecture Context, Configurations and Services**
A brief overview of key architectural decisions to be considered at the start of any integration project
- **Architecture Principles**
The rest of this document provides principles relating to various aspects of Local Application Integration. For each principle there is a brief discussion of the topic, which is then summarised into a statement of a principle.

The final chapter is a summary of the key points, and the Appendix provides a consolidated list of principles, for ease of reference.

Toolkit Specifics

Note that this document provides generic architecture guidance that is relevant to any local integration solution.

However for projects making use of the Interoperability Toolkit, the Toolkit standards provide further explicit guidance and solutions. In addition the middleware provides pre-built services to assist. Boxes labelled “Toolkit Specifics” throughout the text highlight relevant points.

2 Integration Architecture Context, Configurations and Services

Before embarking on detailed design work, it is recommended to clarify key architectural decisions in a number of areas:

2.1 Business Context

- **Impacted Organisations**
What organisation(s) will be impacted by the integration work? For example, who will drive the requirement and be ultimately become responsible for the new interface? Also what other organisations may be affected, and thus may need to signoff any risks involved?
- **Business Scenarios**
What are the Business Scenarios that underlie the technical interfacing work? Is business change also required to support the technical implementation?
- **Messaging Scenarios**
What are the proposed technical messaging scenarios? How do they support the Business Scenarios? Is everything covered?

2.2 Message Configurations

- **Integration Style**
Would messaging or file transfer be more appropriate?
- **Messaging Configurations and Combinations**
What messaging configurations are most appropriate for each interface? What are the message configuration combinations needed across the overall end-to-end deployment?

2.3 Logical Architecture

- **Identification and allocation of Integration Services**
What integration services will be required? How will they be allocated across the various components of the logical architecture?

It is expected that the answers to these questions would be captured as part of the architecture and design documentation of a local integration project - thus providing the high-level background and context needed to assure the appropriateness of the solution.

3 Introduction to Architecture Principles

An Architecture Principle is a statement of belief, approach or intent which directs the formulation of the architecture. Architecture Principles are guidelines for the development of the architecture, as they underpin analysis and investigation of architecture options, and provide a structured set of justifications for the decisions that were made about the components in the architecture

In terms of interoperability between Non-Spine Compliant Systems, there is a varied existing architectural and contractual model with increasingly complex flows to be supported. However the problems to be addressed are common across service providers, and therefore guidance can be defined based upon the architectural principles to be adhered to.

These architectural principles are derived from the different architectural areas defined within the NHS HSCIC Enterprise Architecture Integration Conceptual Services model. This model defines the different services that are needed to be provided by an integration capability. The architectural principles relate to the usage of these services. In addition, there are further principles relating to the information architecture.

The diagram below shows the Level 1 and Level 2 conceptual services only

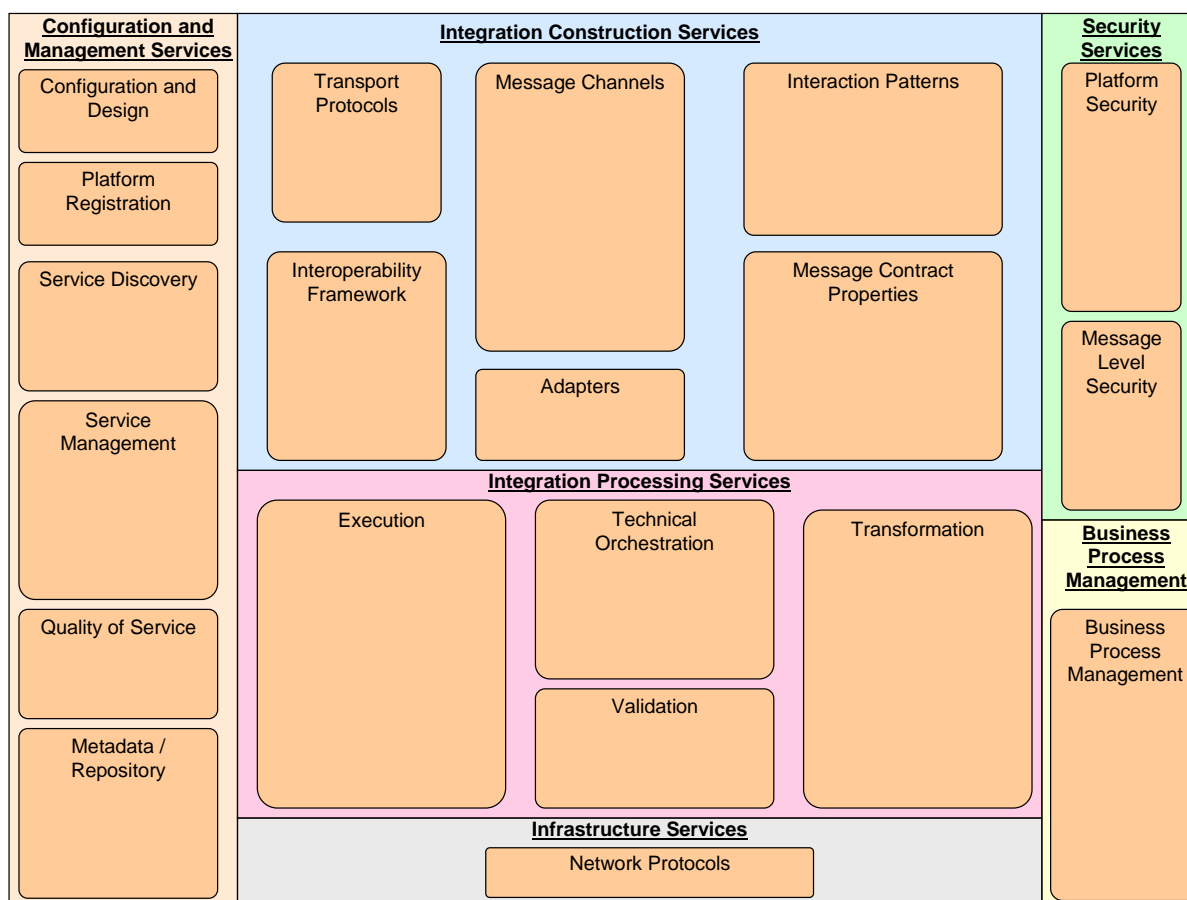


Figure 2 - Integration Conceptual Services - Level 1 and level 2 services

The rest of this document describes the Principles, based upon this model

4 Security Principles

The security services cover both platform level security and message level services including services such as integrity, authentication and authorisation.

4.1 Integrity

Within Local Application Integration one of the key concepts is around the integrity of the data passed within the messages and files transferred as this flows between Spine Compliant and Non-Spine Compliant systems.

Since the Spine Compliant system is not only connected to Non-Spine Compliant systems but also the spine, both

- the flow of information into the Spine Compliant system from a Non-Spine Compliant system needs to be managed (e.g. inbound demographic updates from a maternity system) as at a subsequent point the Spine Compliant system will update the national services such as PDS
- the flow of information from the Spine Compliant system down to a Non-Spine Compliant system needs to be managed to ensure no data leakage from the initial Non-Spine Compliant system to other Non-Spine Compliant systems.

This first architectural principle clarifies that manner of update to the spine. Through Local Application Integration, a Spine Compliant system may receive an update from a Non-Spine Compliant system. The updates could either come from direct messaging between systems or through scheduled batch updates. However, in all cases, any subsequent interaction with the spine must come from an accredited¹ user (or appropriate data flow owner²) irrespective if that information was generated by a Spine Compliant system or came from a Non-Spine Compliant system.

Architectural Principle 1 – Data Flows to the Spine must always be attributable to a known user

A spine update (immediate or scheduled) from a Spine Compliant system must always be actioned via an accredited user. If information into that Spine Compliant system is from a Non-Spine Compliant system, this must be authorised from an accredited user prior to update to spine. The user can either be the end user or the appropriate data flow owner.

¹ “Accredited” implies that the user is registered in the Spine Directory Service, and has been through a suitable authentication process. (For details see the IG Framework document).

² The “Data Flow Owner” is an accredited user who agrees to be accountable for a data flow with no direct user interaction – for example an automated feed.

There are specific requirements regarding information governance that need to be met both in managing data inbound from Non-Spine Compliant systems as well as outbound from Spine Compliant systems to ensure the integrity, and confidentiality of data in line with information governance existing codes of practise such as the NHS CRS This is summarised in Principle 2. The IG Framework document outlines which requirement categories need to be addressed for different business flows and their implications. These will need to be addressed for each deployment activity.

Architectural Principle 2 – Maintenance of Integrity and Confidentiality of data between Spine Compliant and Non-Spine Compliant systems

Integrity and confidentiality of demographic and clinical data between Spine Compliant and Non-Spine Compliant systems must be maintained by use of appropriate Information Governance controls. These are defined further in the IG Framework.

4.2 Chains of systems

It is essential to ensure there are the commensurate security controls in place locally that cover the receiving Non-Spine Compliant system but also subsequent Non-Spine Compliant systems that may be connected to this one. This is to ensure that data flows between the systems do not in any way impact existing code of practice relating to information governance such the NHS Care Record Guarantee. ***Consequently, the same level of rigour of checks needs to be applied across the chain of Non-Spine Compliant systems.***

Architectural Principle 3 – The same level of rigour of checks defined for Locally Assured systems needs to be applied across the chain of Non-Spine Compliant systems.

Appropriate security controls must be applied across all Non-Spine Compliant Systems in a “chain” of systems. Specifically, the security controls around the Spine Compliant system must not be impacted when interfacing with Non-Spine Compliant systems and subsequently linked Non-Spine Compliant systems.

4.3 Malware

The core ITK specifications do not mandate requirements for protection against malware (Viruses, Trojans, Spyware, Spam). It is therefore essential that the ITK connected applications have built in measures to deal with such threats.

Architectural Principle 3.1 – The Trust shall ensure that ITK connected applications have appropriate malware protection in place

Appropriate malware protection measures should be in place such that any ITK connected application is both protected from adverse reactions due to inbound malicious messages and is prevented from sending malicious messages.

5 Integration Processing Principles

The integration construction services focus upon defining which components are required within a messaging configuration. The integration processing services focus on the capabilities required when executing these configurations. Consequently, this includes services such as technical orchestration, validation, exception handling and transformation. Note that the validation, orchestration and exception handling services also apply to file transfer mechanisms.

5.1 Technical Orchestration

Within Local Application Integration the business processes e.g. admission, discharge and transfer of patients or the provision of interim and final results may require the concept of sequencing to be supported.

Sequencing can occur at two levels – business and technical.

- Sequencing at a business level relates to managing the steps in the overall business process irrespective of the technical method of providing information in sequence (e.g. message ordering). Consequently, business level sequencing can only be at the application layer and not at the messaging layer.

Furthermore, the sequencing at a business level should be done within the context of the patient (as opposed to per message type level) to ensure that all activities in relation to that patient are within sequence as opposed to all admission messages being in sequence. The principle related to business level sequencing is defined below.

- Technical level sequencing relates to the managing of the sequence at a messaging layer – that is keeping all messages in a FIFO order, regardless of business content.

Note that even if the messages are provided within sequence order from the Spine Compliant system data centre, this does not necessarily mean that these messages will arrive at the receiving organisation within sequence order. To ensure this, mechanisms must be put in place on every transmission hop and every processing step. For example, gaps in message sequence leading to sequence failure can be exhibited due to:

- Exceptions thrown in message handling at any of the previous handlers (or 'hops') in the integration flow leading to the message not being transported to that handler.
- Transport delays between any of the previous handlers e.g. network blips leading to messages being overtaken across the messaging handlers.

Preserving a technical sequence is thus possible but difficult to achieve in the general case, and requires a single-threaded processing approach which may not be suitable for higher volume flows.

Architectural Principle 5 – Applications must provide business level sequencing, where sequencing is required

Technical level sequencing cannot be assumed in all scenarios. Applications must therefore, where required, provide business sequencing capabilities to handle messages received out of sequence.

5.2 Validation

Whilst validation itself could occur at each hop within the messaging processing flow, the sending system knows the exact version being sent and can therefore perform tight and version-specific structural / syntactic validation. Structural validation at subsequent hops may need to allow flexibility for handling different versions.

Therefore, it is the responsibility of the sending system to ensure that the messages generated are syntactically valid as this is the point of generation of the message. Run-time validation is preferable although this has performance implications.

Additional validation can be put in place at the receiving system but this should be in addition to the validation at the sending system. This may include both structural validation and validation of additional business rules.

Architectural Principle 6 – Syntactical message validation is the responsibility of the sending system

It is the responsibility of the sending system to ensure that the messages generated are syntactically valid and in accordance with data structures and profiles regarding what is valid vocabulary and data items to be used

5.3 Transformation

Integration work needs to consider at least three different types of transformation:

- 1) Protocol Translation** – for example converting from Web Services to ebXML
- 2) Structural Translation** – for example converting from proprietary internal formats to HL7
- 3) Content Translation** – for example converting between different code lists and vocabularies, e.gg Pathology Result Codes

Types (1) and (2), while still potentially non-trivial to implement, are generally well understood and covered by established industry standards (e.g. ebXML, WS-*, HL7).

Type (3) – Content Translation – is typically more challenging – and may in some cases be impossible. Despite good progress on coding standards such as SNOMED, there are still a wide range of vocabularies and coding mechanisms in common use throughout NHS Local Trust systems. This topic is addressed further in section “7 Information Architecture” below.

Errors in content are potentially more difficult to detect during testing – with the system appearing to work during early technical testing, and the invalid data only later being detected by clinicians. Therefore it is extremely important that compatibility and / or translation between vocabularies and coding schemes are considered during the early architectural design phases.

Architectural Principle 7 – Compatibility of data content across systems must be ensured

Structural compatibility of the message definitions is a necessary but not sufficient condition for safe and successful integration. The message content (e.g. coding and vocabulary lists) must also be aligned.

5.4 Exception Handling

The triggering of application behaviour using application acknowledgements provides the mechanism for end to end reliability in that the sending system will be aware of a failure within message processing due to the receipt of negative acknowledgement or due to the lack of receipt of a positive acknowledgement within a defined time period.

However, this is reliant upon the sending system providing the following:

- An explicit mechanism for acting upon the receipt of a negative acknowledgement
- An explicit mechanism for acting upon the lack of a positive acknowledgement within a pre-defined time period. This time period will be dependent upon the business process enacted.
- Execution of defined retry behaviour as part of acting upon indication of failure in messaging processing. This could include system retry of the same message or generation of a new message or manual process as appropriate.

Architectural Principle 8 – End to End reliability is reliant upon the implementation of business level application acknowledgments

Both the source and recipient systems are required to ensure usage and handling of business level application acknowledgements in order to support end to end reliability.

This includes implementation of the relevant retry behaviour as appropriate for messages not processed successfully.

6 Configuration and Management Principles

These principles relate to the managing of the overall Service Capability provided as well as managing changes to the overall Service Capability. This includes principles covering Service Management, Quality of Service and Version Management.

6.1 Quality of Service

Within messaging between the Spine Compliant systems and the National services, there are defined service level agreements in place to ensure that the performance is appropriate to the business processes being supported.

Within Local Application Integration between the Spine Compliant and Non-Spine Compliant systems, there are no such service level agreements in place across the national programme. However, the end to end performance between the Spine Compliant and Non-Spine Compliant systems must meet the needs of the business process. Consequently, the principle regarding end to end performance has been included.

Note: there is no barrier to the definition of service level agreements at a local level with the relevant regional programmes.

Architectural Principle 9 – End to End performance must be provided to meet needs of the business process

As part of the design process, the requirements for end to end performance between the Spine Compliant and Non-Spine Compliant systems must be considered, and the solution engineered such that it supports the needs of the business process. (This does not preclude the local definition of SLAs with LSPs)

6.2 Service Management

In addition to defining the configurations, protocols and execution capabilities, the end to end service management implications of Local Application Integration need to be considered.

6.2.1 Audit

In order to provide a basis for end to end service management and uphold our legal and NCRG obligations audit logging across all systems within the end to end integration flow needs to be robust and sufficient to meet the needs of Audit. This includes auditing of the

Spine Compliant system and associated components but also auditing at the Non-Spine Compliant system and its associated components to ensure a complete and joined up audit. This audit approach needs to be consistent with the overall audit responsibilities defined for suppliers within the NHS HSCIC IG guidance documentation (see the IG Framework for details).

These documents generically describe the responsibilities of each party that creates, reads or updates or interacts with any of the systems provided within the NPfIT. These are applicable to all parties in a local application integration scenario and parties wishing to connect to Spine Compliant systems or exchanging information with them. For these scenarios, they must demonstrate how they meet the guidelines and requirements.

Architectural Principle 10 – Audit compliance with IG guidance is required across each component to enable provision of a complete end to end audit

Audit is required across the Spine Compliant system and Non-Spine Compliant system(s) landscape to enable a complete end to end audit of messaging across application and infrastructure. Each system is responsible for a complete audit of its transactions.

In addition to providing the necessary audit capabilities, all parties need to consider and agree how they will support the service. This will include co-ordination of:

- Help Desks at the Spine Compliant system and Non-Spine Compliant system estates
- Error Reporting at the Spine Compliant system and Non-Spine Compliant system estates
- Alerting at the Spine Compliant system and Non-Spine Compliant system estates

6.3 Version Management

Across the Local Application Integration estate, there will be a variety of different interface versions being supported. This will include several different versions of HL7v2 messaging e.g. HL7v2.3, HL7v2.4. Suppliers are expected to ensure backwards and forwards compatibility in the handling of these different messaging versions. This is summarised in the principle below.

Architectural Principle 11 – Provision of backwards and forwards compatibility in handling of different messaging versions

Suppliers should focus on ensuring the backwards and forwards compatibility in the handling of different messaging versions to ensure that differences between application versions supported across local organisations can be catered for.

7 Information Architecture Principles

While the preceding sections are concerned primarily with the processing of messages, for true interoperability the information content must also be compatible. This section therefore defines a number of Information Architecture Principles.

7.1 Use of Data Standards

HSCIC has defined Data Standards to be used to enable consistent information exchange throughout the NHS. These include:

- **SNOMED** – for reference terminology
- **NHS Data Dictionary** – for administrative codes
- **ICD-10 and OPCS** – for reporting

By using these standards, Trusts and suppliers can be confident that their systems will be aligned with others in the NHS, and will be able to exchange information in a meaningful and clinically safe manner. They will also be able to benefit from the extensive research which underpins these standards, as well as from pre-defined mappings between them where appropriate (e.g. from SNOMED clinical codes to ICD-10 for reporting).

It is also important to note that these standards are very much seen as a living entity and are undergoing constant active development based on the needs of the NHS. Therefore if an apparent “gap” is discovered, Trusts are strongly encouraged to contact the HSCIC Data Standards Team for advice – as opposed to developing a bespoke solution. (Email to DataStandards@nhs.net). The Data Standards team will be happy to advise on best use of the existing standards, and/or to provide guidance on initiating a request for new standards development.

Architectural Principle 12 – Use of NHS Data Standards (SNOMED, NHS Data Dictionary, ICD-10 / OPCS)

***Systems should use the established NHS wide data standards.
Use of these standards enables meaningful and clinically safe
exchange of information throughout the NHS.***

7.2 Mapping from alternative standards

It is recognised that in reality there are many existing systems in the NHS that may not necessarily use the above standards. Therefore mappings may be necessary. Indeed, in some cases the solution may not be clinically safe unless mapping is done to a consistent set of standards. (For example, consistency in use of measures such as “Full Blood Count” and “Whole Blood Count”).

However it is important to realise that mappings between data standards is not necessarily a simple technical translation exercise (for example as might be found between product codes in a logistics system). Rather it is likely to involve an intricate semantic conversion between different healthcare concepts. Specific examples include:

- Conversion from fine-grained to course-grained. This will generally be possible, but will result in information loss (including degradation to text). Depending on the circumstances this may or may not be acceptable and clinically safe.
- Conversion from course-grained to fine-grained. This may be impossible to automate. Alternatively it may involve a judgement call to pick the “most likely” of the fine-grained options. There may be clinical safety implications of this.
- Conversion between incompatible classification schemes. In some cases, different classification schemes categorise the world in a different way. It is thus not possible or meaningful to convert between them.

This is not to say that translation cannot be done, and mapping tables between well known standards do exist. However in using such mappings, Trusts must remain aware of their limitations, and take responsibility for considering the appropriateness and clinical safety implications for each specific scenario.

Architectural Principle 13 – Local Responsibility for mappings

Mapping between clinical standards is not necessarily straightforward. Trusts must take responsibility for maintaining any mappings they may wish to use, and for ensuring they are appropriate and clinically safe for the specific scenario where they are applied.

7.3 Migration from alternative standards

As explained above, it may sometimes be impossible to map between different coding standards, and in this case at least one of the two end systems must be changed – so that a common standard can be used by both. System change such as this should always be used as an opportunity to migrate towards the National NHS Data Standards.

Architectural Principle 14 – System change always migrates to NHS Data Standards

If a system’s information coding approach needs to be changed, then this should always be used as an opportunity to migrate towards the NHS Data Standards.

8 Architecture Guidance Summary

This chapter provides a summary of the areas described in the document that need to be assured with regard to Trust and supplier implementations for Local Application Integration. These are listed below, with further description and rationale of the principles in the appendix.

Business Context
○ Identify impacted organisation(s), Identify business Scenarios to be implemented
○ Define relevant messages scenarios

Message Configurations
○ Identify appropriate Local Application Integration messaging configuration per messaging interface for the required business scenarios
○ Identify appropriate file transfer mechanisms for the required business scenarios
○ Define the different configuration combinations required across the overall deployment

Logical Architecture
○ Appropriate Non-Spine Compliant system logical services provided
○ Appropriate Non-Spine Compliant system integration layer logical services provided
○ Appropriate Spine Compliant system integration layer logical services provided
○ Appropriate Spine Compliant system logical services provided

Architectural Principles Adherence	
○	Data Flows to the Spine must always be attributable to a known user
○	Maintenance of Integrity and Confidentiality of data between Spine Compliant and Non-Spine Compliant systems
○	The same level of rigour of checks defined for Locally Assured systems needs to be applied across the chain of Non-Spine Compliant systems
○	Applications must provide business level sequencing, where sequencing is required
○	Syntactical message validation is the responsibility of the sending system
○	Compatibility of data content across systems must be ensured
○	End to End reliability is reliant upon the implementation of business level application acknowledgments
○	End to End performance is provided to meet needs of the business process
○	Audit compliance with IG guidance is required across each component to enable provision of a complete end to end audit
○	Provision of backwards and forwards compatibility in handling of different messaging versions
○	Use of NHS Data Standards (SNOMED, NHS Data Dictionary, ICD-10 / OPCS)
○	Local Responsibility for mappings
○	System change always migrates to NHS Data Standards

9 Appendix – Architecture Principles Summary

ID	Principle	Rationale	Implications
LTI 01	Data Flows to the Spine must always be attributable to a known user	<p>Since the Spine Compliant system is not only connected to Non-Spine Compliant systems but also the spine, both</p> <ul style="list-style-type: none"> the flow of information into the Spine Compliant system from a Non-Spine Compliant system needs to be managed the flow of information from the Spine Compliant system down to a Non-Spine Compliant system needs to be managed 	A spine update (immediate or scheduled) from a Spine Compliant system must always be actioned via an accredited user. If information into that Spine Compliant system is from a Non-Spine Compliant system, this must be authorised from an accredited user prior to update to spine. The user can either be the end user or the appropriate data flow owner.
LTI 02	Maintenance of Integrity and Confidentiality of data between Spine Compliant and Non-Spine Compliant systems	Regarding integrity and confidentiality of data, there are specific requirements regarding information governance that need to be upheld both in managing data inbound from Non-Spine Compliant systems as well as outbound from Spine Compliant systems to ensure that the Care Record Guarantee is upheld	Integrity and confidentiality of demographic and clinical data between Spine Compliant and Non-Spine Compliant systems must be maintained by use of appropriate Information Governance controls. These are defined further in the IG Framework
LTI03	The same level of rigour of checks defined for Locally Assured systems needs to be applied across the chain of Non-Spine Compliant systems.	It is the responsibility of the Trusts to ensure the appropriate security controls are in place locally at the receiving Non-Spine Compliant system but also subsequent Non-Spine Compliant systems that may be connected to this one so that data flows between the Non-Spine Compliant systems do not in any way impact the Care Record Guarantee. Consequently, the same level of rigour of checks needs to be applied across the chain of Non-Spine Compliant systems.	<p>Responsibility for appropriate security controls across all Non-Spine Compliant systems locally lies with Trusts. This should not impact the NCRS agreement.</p> <p>The security controls around the Spine Compliant system must remain intact when interfacing with Non-Spine Compliant systems.</p>

ID	Principle	Rationale	Implications
LT104	Business Level Sequencing Handling is at the Application Layer	Sequencing at a business relates to managing the steps in the overall business process irrespective of the technical method of providing information in sequence (e.g. message ordering). Consequently, business level sequencing can only be at the application layer and not at the messaging layer.	Business level sequencing is the responsibility of the source and target application and not any intermediary hops. Sequencing should be provided within the context of the patient.
LT105	Applications must provide business level sequencing, where sequencing is required	Technical level sequencing requires FIFO mechanisms to be put in place on all transmission hops and processing steps in the chain. This cannot be assumed in all scenarios.	Technical level sequencing cannot be assumed in all scenarios. Applications must therefore, where required, provide business sequencing capabilities to handle messages received out of sequence.
LT106	Syntactical message validation is the responsibility of the sending system	Whilst validation itself could occur at each hop within the messaging processing flow, the sending system knows the exact version being sent and can therefore perform tight and version-specific syntactical validation. Validation at subsequent hops may need to allow flexibility for handling different versions.	It is the responsibility of the sending system to ensure that the messages generated are syntactically valid and in accordance with data structures and profiles regarding what is valid vocabulary and data items to be used
LT107	Compatibility of data content across systems must be ensured	Structural compatibility of the message definitions is a necessary but not sufficient condition for safe and successful integration. The message content (e.g. coding and vocabulary lists) must also be aligned.	Investigation of vocabularies and code lists during the early architectural design phases. Provision of translation services to map between different vocabularies. (Note that this may not always be possible and / or clinically safe)

ID	Principle	Rationale	Implications
LTi08	End to End reliability is reliant upon the implementation of business level application acknowledgements	The use of application acknowledgements provides the mechanism for end to end reliability in that the sending system will be aware of a failure within message processing due to the receipt of negative acknowledgement or due to the lack of receipt of a positive acknowledgement within a defined time period.	Both the source and recipient systems are required to ensure usage and handling of business level application acknowledgements in order to support end to end reliability. This includes implementation of the relevant retry behaviour as appropriate for messages not processed successfully.
LTi09	End to End performance must be provided to meet needs of the business process	Within Local Application Integration between the Spine Compliant and Non-Spine Compliant systems, there are no such service level agreements in place across the national programme. However, the end to end performance between the Spine Compliant and Non-Spine Compliant systems needs to meet the needs of the business process.	As part of the design process, the requirements for end to end performance between the Spine Compliant and Non-Spine Compliant systems must be considered, and the solution engineered such that it supports the needs of the business process. (This does not preclude the local definition of SLAs with LSPs)
LIT10	Audit compliance with IG guidance is required across each component to enable provision of a complete end to end audit	In order to provide a basis for end to end service management, sufficient audit and logging across all the components within the end to end integration flow needs to be ensured. This includes auditing of the Spine Compliant system and associated components but also auditing at the Non-Spine Compliant system and its associated components to enable a complete and joined up audit.	Audit and tracing is required across the Spine Compliant system and Non-Spine Compliant system(s) landscape to enable a complete end to end audit of messaging across application and infrastructure. Each system is responsible for a complete audit of its transactions.
LTi11	Provision of backwards and forwards compatibility in handling of different messaging versions	Across the Local Application Integration estate, there will be a variety of different interface versions being supported. This will include several different versions of HL7v2 messaging e.g. HL7v2.3, HL7v2.4.	Suppliers should focus on ensuring the backwards and forwards compatibility in the handling of different messaging versions to ensure that differences between application versions supported across Trusts can be catered for.

ID	Principle	Rationale	Implications
LTI12	Use of NHS Data Standards (SNOMED, NHS Data Dictionary, ICD-10 / OPCS)	Systems should use the established NHS wide data standards. Use of these standards enables meaningful and clinically safe exchange of information throughout the NHS.	Use of the NHS Data Standards
LTI13	Local Responsibility for mappings	Mapping between clinical standards is not necessarily straightforward and may, depending on the scenario, introduce clinical risk	Trusts must take responsibility for maintaining any mappings they may wish to use, and for ensuring they are appropriate and clinically safe for the specific scenario where they are applied.
LTI14	System change always migrates to NHS Data Standards	There are advantages to the NHS Data Standards, as described above. Therefore if systems are being changed anyway, then the opportunity should be taken to migrate to them	If a system's information coding approach needs to be changed, then this should always be used to migrate towards the NHS Data Standards

*** End of Document ***