

# nehta

---

## **Connectivity**

### **Introductory Guide**

Version 1.1 — 30 June 2010

---

**National E-Health Transition Authority Ltd**

Level 25

56 Pitt Street

Sydney, NSW, 2000

Australia.

[www.nehta.gov.au](http://www.nehta.gov.au)**Disclaimer**

NEHTA makes the information and other material (“Information”) in this document available in good faith but without any representation or warranty as to its accuracy or completeness. NEHTA cannot accept any responsibility for the consequences of any use of the Information. As the Information is of a general nature only, it is up to any person using or relying on the Information to ensure that it is accurate, complete and suitable for the circumstances of its use.

**Document Control**

This document is maintained in electronic form. The current revision of this document is located on the NEHTA Web site and is uncontrolled in printed form. It is the responsibility of the user to verify that this copy is of the latest revision.

**Copyright © 2010, NEHTA.**

This document contains information which is protected by copyright. All Rights Reserved. No part of this work may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without the permission of NEHTA. All copies of this document must include the copyright and other information contained on this page.

# Table of contents

<b>Table of contents</b> .....	<b>iii</b>
<b>Document information</b> .....	<b>iv</b>
Change history .....	iv
<b>1 Executive overview</b> .....	<b>1</b>
<b>2 Preface</b> .....	<b>2</b>
2.1 Document purpose .....	2
2.2 Intended audience .....	2
2.3 Definitions, acronyms, abbreviations .....	2
<b>3 Introduction</b> .....	<b>3</b>
3.1 Connectivity .....	3
3.1.1 Structured communications .....	3
3.1.2 Different parties .....	3
3.1.3 Business processes .....	3
3.2 Connectivity architecture .....	3
3.3 Examples .....	4
<b>4 Services</b> .....	<b>5</b>
4.1 Use of services .....	5
4.2 Design starts with business services .....	5
4.3 Decentralised services .....	5
4.4 Advantages of services .....	6
<b>5 National Infrastructure Services</b> .....	<b>7</b>
5.1 Healthcare Identifiers .....	7
5.1.1 Individual Healthcare Identifier .....	7
5.1.2 Healthcare Provider Identifier .....	8
5.1.3 Other types of identifiers .....	8
5.2 National Authentication Service for Health .....	8
5.2.1 Benefits .....	9
5.3 Endpoint Location Service .....	9
<b>6 Interactions</b> .....	<b>10</b>
6.1 Bootstrapping interaction .....	10
6.2 Identification interaction .....	10
6.3 Core connectivity interaction .....	11
6.3.1 Service invoker .....	11
6.3.2 Service provider .....	13
6.4 Further information .....	14
<b>Appendix A: References</b> .....	<b>15</b>
<b>Appendix B: Change log</b> .....	<b>16</b>

# Document information

## Change history

Version	Date	Comments
1.0	2008-12-01	Release
1.1	2010-06-30	Release

# 1 Executive overview

Connectivity is the establishing and conducting of communications in a reliable and secure manner.

E-health related communications are to be conducted using methods identified as being appropriate for these interactions. The primary method identified in the architecture for this purpose is Web services.

Security is provided by appropriately applying the use of digital signing and encryption. This is supported by the keys and certificates issued by the National Authentication Service for Health, which is one of the national infrastructure services.

Reliability is provided by using appropriate mechanisms in the design of the services to ensure reliability capabilities. They will be designed to support acknowledgements at the business, operation and network levels.

This document provides an introduction to the key elements of the connectivity architecture. Those elements are:

- **Services** expose function in a reusable fashion that is useful for conducting electronic business interactions. Services will be distributed without a central service bus or messaging hub. Therefore, each interaction involving a service is responsible for handling its own communications.
- **National infrastructure services** will be available to provide identification and authentication services. These are shared services that everyone can use. They provide national coordination as well as assisting each service interaction to handle its own communications.
- **Interactions** have been defined for how these national infrastructure services are used to support the process of communicating healthcare information.

## 2 Preface

### 2.1 Document purpose

This document is an introductory guide to the connectivity architecture.

### 2.2 Intended audience

This is a non-technical document that is intended for a general audience.

### 2.3 Definitions, acronyms, abbreviations

HI	Healthcare Identifier
HPI	Healthcare Provider Identifier
IHI	Individual Healthcare Identifier
NASH	National Authentication Service for Health
NEHTA	National E-Health Transition Authority
PKI	Public Key Infrastructure
ELS	Endpoint Location Service

# 3 Introduction

## 3.1 Connectivity

Connectivity in this context is the establishing and conducting of e-health related communications in a reliable and secure manner.

To understand how connectivity addresses improving healthcare, it is useful to examine the three key points in the definition:

- Structured communications;
- Different parties; and
- Business process.

### 3.1.1 Structured communications

Structured communications involves exchanging data that can be meaningfully processed by the other party. Structured data is a requirement for application-to-application communications between computer systems. The aim is to achieve a shared understanding between the two systems. This involves the use of common data structures and common terminologies.

### 3.1.2 Different parties

This architecture focuses on communications with external organisations. Organisations which have different ownership, work under different policies, and have implemented their systems using different technologies.

The techniques used by this architecture can be used inside an organisation for internal communications, but that is not the main focus of the architecture.

The connectivity architecture uses a service oriented approach, where functionality is defined by well defined units of work with standardised interfaces. These services can be implemented and deployed by different organisations, and one organisation can use the services provided by another organisation.

### 3.1.3 Business processes

The focus is on supporting the healthcare business processes. This involves identifying the requirements of the healthcare process and mapping them into services and their associated interactions.

## 3.2 Connectivity architecture

The connectivity architecture describes how connectivity is achieved in the national e-health environment.

At a high level, the connectivity architecture contains three main points:

- The use of services;
- National infrastructure services to provide common functions that many different parties will use; and
- Interactions that describe how those national infrastructure services are used together to achieve connectivity.

These three points are reflected in the structure of this document. Chapter 4 describes services, chapter 5 describes the national infrastructure services, and chapter 6 describes the interactions.

The architecture also describes how communications will be conducted using appropriate protocols. One of the protocols used by the architecture is Web services.

Security is provided by appropriately applying the use of digital signing and encryption. This is supported by the keys and certificates issued by the National Authentication Service for Health, which is one of the national infrastructure services.

Reliability is provided by using appropriate mechanisms in the design of the services to ensure reliability capabilities. They will be designed to support acknowledgements at the business, operation and network levels.

### **3.3 Examples**

The example of pathology result reporting will be used to provide illustrations of how connectivity is used. Although details can be found in the NEHTA pathology result reporting package, it is sufficient to know that this example is about delivering a pathology report from a pathology laboratory to a general practice. The pathology report is data. The pathology laboratory is the sender. The general practice is the receiver.

The connectivity architecture can also be applied to many other tasks, some of which are much more complex than pathology results reporting. The example of pathology was chosen because it is simple to understand and serves as a good illustration of how the architecture is used.

## 4 Services

### 4.1 Use of services

The connectivity architecture has a service view of the business environment. Where a service is a unit of well-defined functionality offered in a standard form.

When this document talks about “business,” it is referring to the function of healthcare. For example, the business is performing a pathology test, or referring a patient to another provider. It focuses on the healthcare or clinical outcome. In this document, it does not refer to the finance or administration side of healthcare.

### 4.2 Design starts with business services

This service-oriented approach organises the business of healthcare in terms of services and interactions between those services. It attempts to bring order to the understanding of the environment by breaking up a complex system into a set of simple services which are aligned to the functions of business. Each service provides a reliable well tested capability that can be leveraged by additional participants and possibly within other business processes.

The design process should first identify business services instead of technical services. A business service defines what function needs to be performed, without defining how it is done. A technical service is a mechanism that implements a business service; it is a realisation of that service.

For example, delivering a pathology result is a business service that could be realised in a number of different technical means (e.g. physical delivery, fax, or Web services). Obviously, some technical services are outside the scope of connectivity as defined here (e.g. provision of a shop front service) because the focus is on electronic communications.

The distinction between business services and technical services help reinforce the practice that the design needs to be driven by the business needs and not by technology. The focus should be on delivering healthcare, and technology is a means to achieving that end.

### 4.3 Decentralised services

Services can be operated by different parties and distributed across the nation. Each organisation can choose the services they provide and where to host them.

Organisations can, but do not have to, host the services themselves. Some organisations will have the need and resources to host services. While other organisations might choose not to host services or contract a third party to host the services on their behalf. For example, a pathology laboratory might host their own services, but a general practice might contract a third party to run services for them off-site.

Services implementations are responsible for managing their own communications. They need to handle every aspect of operating a service, such as identification, addressing, authentication, encryption and ensuring reliability. When services are implemented inside an organisation, there is often an Enterprise Service Bus (ESB) that handles many of these functions. But in the national e-health environment there is no ESB, so each service needs to perform these functions themselves.

Services can use the national infrastructure services to assist them to perform many of these functions. This makes implementing services easier, since the

service can reuse the functions provided by the national infrastructure services instead of implementing it all themselves. The national infrastructure services help with identification, addressing, authentication and encryption. They are described in section 5.

## 4.4 Advantages of services

The services based approach has many advantages that are important in an interoperable national e-health environment.

Some of these advantages are:

- Services can be operated by different parties. This reflects the decentralised nature of the healthcare system in Australia. Without a services oriented approach, one party will have to operate all components of the e-health environment: although this is possible within a single organisation, it is impossible to do nationwide across multiple independent organisations.
- Different types of services can be defined. The e-health environment needs to change over time to meet new healthcare needs. The services approach allows some services to change without affecting the other services in the system. If a service oriented approach was not used the entire system will have to change whenever a change occurs, which is impractical in a large nationwide system.
- Services can be standardised. These standards can enable different vendors to implement and deploy products which can interoperate with products from other vendors.
- Services can be reused. A service deployed by one party can be used by others.

# 5 National Infrastructure Services

The connectivity architecture has identified some services that need to be available on a national basis. They provide infrastructure that can be reused by every party in the e-health environment.

The national infrastructure services are:

- Healthcare Identifiers;
- National Authentication Service for Health; and
- Endpoint Location Service(s).

This chapter describes what these services are. The next chapter, chapter 6, will then describe how these services are used together to achieve connectivity.

## 5.1 Healthcare Identifiers

The Healthcare Identifier (HI) services are a set of services providing identifiers for different types of party in the national e-health environment.

There are two needs: having identifiers, and having the same type of identifiers.

- Identifiers are important in communications. Without identifiers, it is difficult to determine who a message is coming from, who it is going to, and who it is about. For example, a pathology report could be sent to the wrong GP if the GP could not be identified, or the pathology report could be associated with the wrong patient if the individual could not be identified.
- In addition to having identifiers, the same type of identifiers needs to be used by the communicating parties. If the pathology identifies individuals using their own private identifier, then the GP will not know who they are referring to. In a national environment, a nationally common identifier is required—which is what HI provides. For a nationwide system, the same type of identifiers needs to be used by parties across the nation.

The HI services satisfy these two needs. They provide an identification mechanism and support those identifiers for parties across the whole nation.

The HI is needed because local identifiers cannot be used across the nation. For example, patient identifiers allocated by a hospital cannot be used outside that hospital.

The HI services provide several different types of identifiers, including:

- Individual Healthcare Identifier (IHI) for identifying individuals (also known as patients or consumers).
- Healthcare Provider Identifier (HPI) for identifying organisations and individuals which provide healthcare (e.g. pathology laboratories and doctors).

### 5.1.1 Individual Healthcare Identifier

An Individual Healthcare Identifier (IHI) is used to identify individuals. Individuals are people who receive healthcare. They are sometimes referred to as patients or consumers.

NEHTA will put in place the national IHI system. This system includes services and processes. Services will be available to lookup IHI numbers. Processes will be available for issuing and managing IHI numbers.

### 5.1.2 Healthcare Provider Identifier

A Healthcare Provider Identifier (HPI) is used to identify healthcare providers.

There are two types of healthcare providers:

- Healthcare Provider Identifier for Individuals (HPI-I) are for people who provide healthcare. For example, doctors and pharmacists.
- Healthcare Provider Identifier for Organisations (HPI-O) are for organisations who employ healthcare providers. For example, general practice clinics, hospitals, and pharmacies.

NEHTA will put in place the national HPI system. This system includes services and processes. Services will be available to lookup HPI numbers and information about HPI holders. Processes will be available for issuing and managing HPI numbers.

### 5.1.3 Other types of identifiers

National identifiers can work along side with local identifiers. Although national identifiers should be used instead of local identifiers, there is no requirement to eliminate local identifiers. For example, a hospital can still use their local patient identifiers inside their systems as well as use IHI numbers for communications to outside parties.

## 5.2 National Authentication Service for Health

The National Authentication Service for Health (NASH) is designed to be a set of services and processes for issuing and managing security credentials.

These security credentials can be used to secure communications:

- They can be used to sign data. This is used to authenticate parties, to prove that they are the holder of the security credential.
- They can be used to encrypt data. This is used to preserve confidentiality so that other parties cannot see or use the data.

NASH supports the Public Key Infrastructure (PKI) standard for security credentials. The NASH issues, manages, and revokes these artefacts:

- Key pairs for use with public key cryptography. These are private and public keys that are used with cryptographic algorithms to sign and encrypt data. The owner of the private key must not reveal it to any other party. The corresponding public key, on the other hand, can be shared with all other parties for them to use.
- X.509v3 public key certificates. This is an industry standard format of representing public keys by binding them to a subject name. The format allows parties to check who issued the certificate and to validate that it has not been forged.
- Hardware tokens such as smartcards. These devices are used to store the private keys in a secure manner. They are secure because they protect the private key from copying, and therefore reducing the risk of it being revealed to other parties.<sup>1</sup>

---

<sup>1</sup> Not all private keys will be stored on smartcards. NASH will also support "soft keys," which are private keys stored as data files. These files can be copied and therefore the applications that use them must take precautions to protect them. Whether hardware tokens or soft tokens are used will depend on the application they will be used in and the issuing policy for that type of credentials.

The NASH will provide services that support the use of these PKI keys. As users of PKI, the services that connectivity will use are:

- A directory of public X.509v3 certificates. For obtaining a copy of the public certificate when the identity of the owner or the certificate is known.
- A Certificate Revocation List (CRL) service. This service publishes a list of certificates that have been revoked. Certificates can be revoked for a number of reasons, such as the issuer cancelling the owner's right to have the certificate or the private key has been accidentally revealed to other parties.
- A certificate checking service. This is a service for finding out the status of a particular certificate, and whether it has been revoked or not. This is used to achieve the same outcome as using a CRL, but can provide more up to date information since the status is obtained when a certificate is used instead of when the CRL was last updated. The NASH will support the industry standard Online Certificate Status Protocol (OCSP) for checking certificates.

### 5.2.1 Benefits

NASH provides a PKI deployment that can be used by parties in the national e-health environment.

This reduces the cost of implementing PKI, since organisations do not have to provide their own deployments. They simply use the service provided by NASH, instead of having to issue their own certificates.

The NASH is a national service, so certificates can be used across the nation. This can be more beneficial than locally issued certificates, which can only be used within an organisation or jurisdiction.

Having PKI allows parties to secure their communications using digital signing and encryption to provide authentication, integrity and confidentiality.

## 5.3 Endpoint Location Service

The Endpoint Location Service (ELS) is a service for looking up information needed to connect to other services. It is a directory that contains technical information for using services.

For example, if a pathology laboratory wants to send a pathology report to a general practice, they would look up the ELS for the pathology services supported by that particular general practice. The ELS will return the information needed to connect to the GP, including: the network address of where the service is, who is operating the service, and how to secure requests to it.

The ELS is an essential component of the national e-health environment. Without it, there is no way for clients to know where the services are or how to connect to them. It is not practical to hardcode this information into the clients, because there are too many services and they are not controlled by any central authority. The ELS provides a directory where this information can be looked up when it is needed.

There are a number of possible deployment models for the ELS. In the decentralised model, each provider hosting services will also host their own ELS. In the centralised model, there is one national ELS which contains entries for every service provider. There can also be hybrid models, where a ELS can serve a group of providers (e.g. a ELS could be hosted by a division for all its members).

## 6 Interactions

The Connectivity Architecture has identified three interactions that are used to achieve connectivity. These interactions show how the national infrastructure services are used together to achieve connectivity.

The three interactions are:

- Bootstrapping interaction
- Identification interaction; and
- Core connectivity interaction.

In this chapter these interactions are described by using an example of a pathology laboratory system delivering a pathology result report to a general practice.

### 6.1 Bootstrapping interaction

The first interaction used is the bootstrapping interaction, which is used to obtain the information required to connect to the national infrastructure services.

1. The pathology laboratory system contacts the Base-ELS and queries it for the IHI service.

It can do this because the location of the Base-ELS and its public certificate has been configured into the pathology laboratory system.

2. The Base-ELS returns the address (a URL) and a reference to the public certificate for the IHI service.

The pathology laboratory repeats these steps for the HPI-I service, the HPI-O service and the NASH service.

Since these national infrastructure services rarely change, this information can be cached by the pathology laboratory for future use. The bootstrapping interaction is only needed when an application first starts up. It should not be used for every transaction that it conducts.

### 6.2 Identification interaction

Using the national identification services (obtained from the bootstrapping interaction) the pathology laboratory can then identify the parties involved in the delivery of the pathology report.

1. The pathology laboratory uses the IHI service to lookup the patient's IHI number.
  - a. Send a request to the IHI service with a set of identifying traits about the patient (e.g. their name and date of birth).
  - b. Receive back a response with the patient's IHI number if the traits successfully match.

The IHI number is used in the report so that the receiver (the GP) knows exactly which patient the report is about.

Often the IHI number is already known to the pathology laboratory. For example, it was provided with the pathology request or the patient tells the pathology laboratory their IHI number. However, if it is not available (e.g. the request was an old paper request or the patient did not know their IHI number) this process can be used by the pathology laboratory to look up the IHI number. This process can also be used by other provider organisations, such as when the patient first visits a general practice.

2. The pathology laboratory uses the HPI-I service to lookup the HPI-I number for the general practitioner (the person).
  - c. Send a request to the HPI-I service with a set of identify traits about the general practitioner.
  - d. Receive back a response with the GP's HPI-I number if the traits successfully match.

The HPI-I number is used to identify which GP the report is intended for.

Often the HPI-I number is already known to the pathology laboratory. It could have been included in the pathology request.

3. The pathology laboratory uses the HPI-O service to lookup HPI-O number for the general practice (the organisation).
  - a. Send a request to the HPI-O service with a set of identify traits about the general practice.
  - b. Receive back a response with the practice's HPI-O number if the traits successfully match.

The HPI-O number is used to identify which organisation to send the report to. It is also used to lookup the ELS in the core connectivity process (see section 6.3).

Often the HPI-O number is already known to the pathology laboratory. It could have been included in the pathology request, or the pathology laboratory has communicated with them before.

These identifiers can be cached for future use. The identification interactions only need to be used if these identifiers are not known to the system.

## 6.3 Core connectivity interaction

The main interaction that is used is the core connectivity interaction.

Using the previous interactions, the pathology laboratory system knows where the national infrastructure services are and the identifiers for the parties involved. It then uses the core connectivity interaction to obtain the information needed to connect to the general practice's pathology report service and use it.

The core connectivity interaction uses the national infrastructure services to obtain the address of the desired service instance and the certificates needed to invoke it. It then invokes the operation it wants to perform.

There are two sides to the interaction: one from the service invoker (the pathology laboratory), and one from the service provider (general practice).

### 6.3.1 Service invoker

In brief, the interaction for the service invoker consists of the following steps:

1. Use the HPI-O service to determine the Endpoint Location Service (ELS) to use;
2. Use the Endpoint Location Service (ELS) to determine the service interface and obtain the address of where the service instance is operating; and
3. Use the National Authentication Service for Health (NASH) to obtain the digital certificates used to sign and encrypt the data.
4. Invoke the service instance.

In detail, the interaction consists of the following steps:

1. Query the HPI-O service to obtain the location of the general practice's Endpoint Location Service (ELS).

All provider organisations that provide services will also operate a Endpoint Location Service. In this case the general practice will have a ELS. This ELS could be operated by the general practice or by a third party on behalf of the general practice. Provider organisations could have their own ELS or could share a ELS with other provider organisations. Since there could be many ELS instances, the first step is to locate the ELS instance that contains entries for the general practice the report needs to be sent to.

The pathology laboratory looks up the HPI-O service and retrieves the location of the ELS for that general practice. It queries the HPI-O service using the HPI-O number of the general practice (which was obtained in the identification interaction, see section 6.2). The HPI-O service returns back a record which contains the address of that HPI-O's ELS.

2. Query the general practice's ELS to obtain the interaction patterns.

Using the information obtained from the previous step, the pathology laboratory queries that ELS to find out what technical interaction patterns it supports for pathology report delivery.

- a. Send a request containing the service category that it wants to perform. In this case the service category is for delivering pathology result reports.
- b. Receive back a response containing a list of interaction patterns that the general practice supports.

3. Choose which interaction pattern to use.

If the ELS returns only one interaction pattern, then the client does not have a choice to make. The general practice only supports one method for receiving pathology result reports.

If the ELS returns more than one interaction pattern, then the pathology laboratory will need to choose which one it wants to use. The pathology laboratory will pick an interaction pattern out of the ones it supports and in the order of what is most efficient for it to use.

If the ELS returns zero interaction patterns, then the general practice does not support that interaction category, so the pathology laboratory will have to find a way to deliver the pathology report.

4. Obtain the necessary public certificates from NASH.

The ELS entry for the chosen interaction pattern will indicate which certificates are needed to invoke the service instance. It will contain references to X.509v3 certificates.

This step takes those references and queries the NASH service to obtain the actual X.509v3 certificate. For example, if the ELS entry indicated that the data needs to be encrypted with a certificate that has a Subject Key Identifier (SKI) of "375347234057057347", then this step obtains the actual certificate that matches that SKI number.

5. Sending a service request to the service provider.

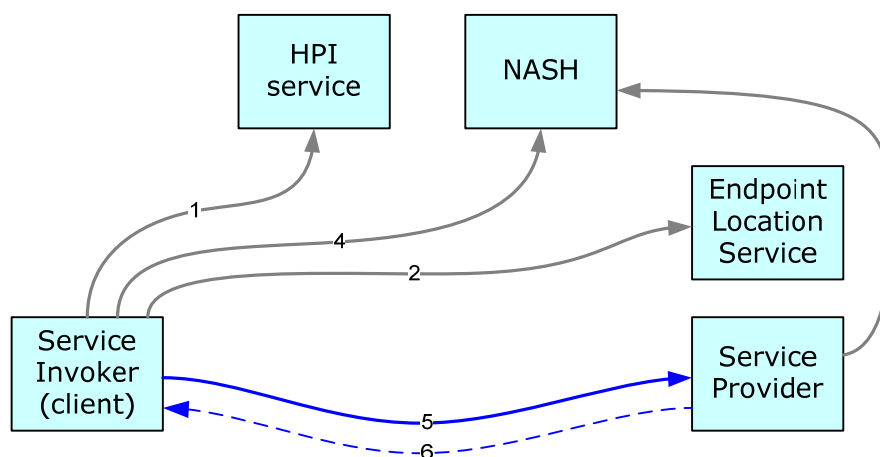
The pathology laboratory now has all the information it needs to invoke the service.

It uses that information to invoke the service instance.

- The payload is created and populated with the request details. This will contain the pathology report and the IHI of the subject and the HPI-I number of the general practitioner.
- The request is digitally signed using the pathology laboratory's private key.

- The request is encrypted using the general practice's public certificate (which was indicated in the ELS entry and obtained from the NASH in the previous step).
  - The request is sent to the address that was obtained from the ELS entry.
6. Receiving a service response from the service provider.  
A response comes back from the service instance. The Web service operations will use a request-response message exchange pattern. So for every request sent, there will be a response (or fault) returned by the service.
  7. Validating the service response.  
The response is checked for validity. The digital signature on the response is validated to ensure that it has not been forged or tampered with.
  8. Verifying the security credentials used by the service provider.  
The response needs to have come from the expected general practice. This step checks that this by checking that the certificate used for signing the response belongs to the intended general practice.

The steps in the core connectivity interaction are illustrated in Figure 1. The arrows correspond to the service invocations.



**Figure 1: Core connectivity main process**

### 6.3.2 Service provider

In brief, the service provider uses the NASH service to check the service request before it processes it.

In detail, the interaction consists of the following steps:

1. Receiving the service request from the client.  
The general practice runs a computer program that continuously waits for a request to arrive. The pathology laboratory's system connects to it and sends it a request.
2. Validating the service request.  
The request is checked for correctness. This involves checking the digital signature on the request is correct. This ensures that the request has not been forged or tampered with.
3. Verifying the security credentials used by the service invoker.  
The security credentials presented by the pathology laboratory are checked to ensure that they are acceptable.

This step involves contacting the NASH service to check that the certificate key pair has not been revoked. Alternatively, the Certificate Revocation List (CRL) published by the NASH can be checked.

4. Performing the operation.

Having passed the tests, the general practice processes the request.

It can validate the pathology report for technical correctness and then store it for the intended GP doctor (as identified by their HPI-I number) to view.

5. Sending a service response back to the client.

A response is generated and sent back to the client to indicate the result of the operation. In some cases, when the operation has failed, a fault is sent instead.

Figure 1 shows step 3 as the unnumbered line between the service provider and NASH.

## 6.4 Further information

This chapter has provided examples of the interactions used to achieve connectivity. For a more formal description of the interactions, see the *Connectivity: Architecture* [CA2010].

In a real operating environment, the interactions usually involve fewer operations than has been described here. This is because a real implementation will cache values from previous transactions and reuse them. Caching and other implementation issues are discussed in the *Connectivity: Implementation Guide* [CIG2008].

# Appendix A: References

- [CA2010] NEHTA, *Connectivity: Architecture v1.1*, 30 June 2010.
- [CIG2010] NEHTA, *Connectivity: Implementation Guide v1.1*, 30 June 2010.

# Appendix B: Change log

Version 1.1

- Service Instance Locator (SIL) renamed to Endpoint Location Service (ELS).