Document Title: Audit Data Collection & Storage High Level Design

Document Type: High Level Design

Release: (Release specific/Release Independent/Not Applicable)

Abstract: This document describes the Audit data collection & storage facilities within HNG-X

Document Status: APPROVED

Author & Dept: Alan Holmes

Internal Distribution:

External Distribution:

Approval Authorities:

Name	Role	Signature	Date
Steve Evans	Solution Design		
Graham Allen	Development		
Alan Holmes	Architecture		

Note: See Post Office Account HNG-X Reviewers/Approvers Role Matrix (PGM/DCM/ION/0001) for guidance.

0 Document Control

0.1 Table of Contents

0.2 Figures and Tables

0.3 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1	25/01/2007	Initial draft. Content taken from the equivalent HNG-X document SD/HLD/001	
0.2	11/05/2007	Second draft. Includes changes resulting from comments received on version 0.1 and PCI changes CR957	
0.3	08/01/2009	Changes resulting from review of version 0.2 Changes resulting from E2E group review of the parent architecture document ARC/SVS/ARC/0001 Changes resulting from the following CPs • HNG-X CP0157 – CP4623 • HNG-X CP0299 – CP4810 • HNG-X CP0258 – CP4751 • HNG-X CP0284 – CP4791 Extended migration coverage	
1.0	05/10/2009	Minor updates resulting from review. Issued for Approval	

0.4 Review Details

Review Comments by :	N/A
Review Comments to :	Alan Holmes
Mandatory Review	
Role	Name
Solution Design	Andy Williams
Development	Steven Meek
Architecture	Alan Holmes
Solution Design	Steve Evans
SSC	Mik Peach
Business Continuity	Tony Wicks
System Test	John Rogers
Optional Review	
Role	Name
Security & Risk Team	GRO
Programme Manager	Phil Day
Applications Architecture	David Johns
System Qualities Architecture	Dave Chapman
Architect	Jason Clark
Security Architect	Jim Sweeting

Test Design	Peter Robinson
Test Design	George Zolkiewka
Head of Service Management	Steve Denham
Head of Service Change & Transition	Graham Welsh
Service Support	Peter Thompson
Service Network	Alex Kemp
Data Centre Migration	Caroline Montgomery
Data Centre Migration	Peter Okely
Testing	Peter Dreweatt
SV&I Manager	Sheila Bamber
Tester	Hamish Munro
RV Manager	James Brett (POL)
VI & TE Manager	Mark Ascott
HNG-X Acceptance & Risk	Wayne Roberts (POL)
Integrity Testing	Alan Child
Integrity Testing	Michael Welch
Core Services	Ed Ashford
Core Services	Andrew Gibson
Business Architect	Gareth Jenkins
Development	Gerald Barnes
Development	Graham Allen
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	Name

(*) = Reviewers that returned comments

0.5 Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Document Template	Dimensions
ARC/SVS/ARC/0001			HNG-X Architecture – Support Services	Dimensions
DES/APP/HLD/0029			Audit Data Retrieval High Level Design	Dimensions
DEV/INF/ION/0001			Audit Server Configuration	Dimensions
IA/PRO/004			Audit Data Extraction Process	PVCS
ARC/GEN/REP/0001			HNG-X Glossary	Dimensions
SVM/SDM/SD/0017			Security Management Service - Service Description	Dimensions
CR/FSP/006			Audit Trail Functional Specification	PVCS

Reference	Version	Date	Title	Source
SD/IFS/014			Audit to Tivoli Cluster Information Interface Specification	PVCS
SD/HLD/001			Audit Data Collection & Storage High Level Design	PVCS
DES/PPS/PPD/0037			HNG-X Audit Workstation (AUW) - Physical Platform Design	Dimensions
DES/PPS/PPD/0035			HNG-X Audit Server (ARC) - Physical Platform Design	Dimensions
DEV/INF/ION/0008			Audit Server Migration Configurations	Dimensions
DEV/APP/SPG/0016			Audit Extraction Client Support Guide	Dimensions
DEV/INF/ION/0009			Audit Server Schedule Design	Dimensions
ARC/APP/RTM/0005			Requirements Traceability Matrix for Support Services	Dimensions
DEV/GEN/MAN/0015			Audit Extraction Client User Manual	Dimensions
DEV/APP/SPG/0016			Audit Extraction Client Support Guide	Dimensions
DES/APP/HLD/0020			Branch Database High Level Design	Dimensions
DES/SEC/HLD/0011			HNG-X Anti Virus High Level Design	Dimensions
DES/SEC/HLD/0002			HNG-X Crypto Services High Level Design	Dimensions
DEV/APP/HLD/0071			Audit Data Retrieval Low Level Design	Dimensions
DEV/APP/SPG/0020			Audit Server Support Guide	Dimensions

Unless a specific version is referred to above, reference should be made to the current approved versions of the documents.

0.6 Abbreviations

Abbreviation	Definition
ARQ	Audit Record Query
AS	Audit Server
ATD	Audit Track Deleter
ATR	Audit Track Retriever
ATS	Audit Track Sealer
Baseline HNG-X	the existing solution being re-architected
Branch	Post Office outlet identified by a unique Branch Code. Within the HNG model, a Branch is a logical entity that can be composed of several physical locations at which business is transacted.
Sensitive Cardholder Data	The PAN, Cardholder name, Service code & Expiration date information relating to a

Abbreviation	Definition
	payment card
COTS	Commercial off the Shelf
CSV	Comma Separated Variables
DRDB	Branch Database
EMC	Company that supplies resilient disk technology
FS	Fujitsu Services
FTMS	File Transfer Managed Service
HNG-X	HNG-X Next Generation – Plan X
KEL	Known Errors Log
NBS	Network Banking Service
NBSC	National Business Support Centre
NPS	Network Banking Persistence Service
NSP	Atalla Network Security Processor
OBC	Operational Business Change
Operational Services	Those services that are needed to run the HNG-X system that are not directly supporting the Post Office business. Examples include software distribution, audit, security management etc.
PCI	Payment Card Industry. A set of security controls defined by the Payment Card Industry organisation.
PCI DSS	Payment Card Industries Data Security Standard. A set of security controls defined by the Payment Card Industry organisation.
PO	Post Office
Post Office	Post Office Limited
RAD	Real time Active Dashboard
RDDS	Reference Data Distribution System
RDMC	Reference Data Management Centre
RDS	Post Office Reference Data System
RDT	Reference Data Team - the Post Office and Fujitsu Customer Services teams use the RDT environment to validate and verify the reference data associated with business changes.
Sensitive Authentication Data	The full contents of any track from the magnetic stripe (on the back of a card, in a chip, etc.), Any data element extracted from the magnetic stripe other than Cardholder Name, PAN, expiry date and service code, and Encrypted PIN blocks.
SLT	Service Level Target
SYSMAN	The systems management environment.
TES	Transaction Enquiry Service
XML	Extensible Markup Language

0.7 Glossary

See *HNG-X Glossary (ARC/GEN/REP/0001)*

Term	Definition

0.8 Changes Expected

Changes

This document will be changed as a result of comments received.

0.9 Accuracy

Fujitsu Services endeavours to ensure that the information contained in this document is correct but, whilst every effort is made to ensure the accuracy of such information, it accepts no liability for any loss (however caused) sustained as a result of any error or omission in the same.

1 Introduction

Within the HNG-X system, Fujitsu Services are required to provide facilities to produce, store and present to (customer) auditors for analysis Audit Track data in support of the security policy and audit requirements laid down for the system.

Within HNG-X, audit data is collected from a number of subsystems. The basic types of audit data that is collected are:

- Counter application messages as received by the Branch Access Layer and stored in the Branch Database message journal table. This will include counter transactions and events
- Data transferred across HNG-X system boundaries. E.g. Bulk file transfers to and from Post Office and their clients
- Host database systems audit and archive data. In this context database audit data refers to the saving of logs of updates applied to the databases, and database archive data refers to the saving of old data that has been purged from the primary databases.
- HNG-X system events – including security events
- Logging of activities undertaken by administrative users during maintenance of the system
- System scheduler logs

Audit data may be requested by a number of different end users, for a number of different reasons. These include:

- Post Office Auditors in connection with Fraud investigations – in which case the data may be presented as evidence in court
- Fujitsu Services Post Office Account Security staff monitoring compliance with security requirements
- Post Office users handling enquiries regarding banking transactions
- Fujitsu Services System Support Centre for diagnostic information

It is essential that the Audit System can both maintain the integrity of data under its management and subsequently be able to prove that integrity if and when the data is retrieved for analysis. This applies to both data gathered from HNG-X sources and existing Horizon audit data.

The architecture for the audit sub-system within the HNG-X system is described in *HNG-X Architecture – Support Services* (ARC/SVS/ARC/0001). This Audit Data Gathering & Storage High Level Design Specification is consistent with that architecture.

This High Level Design (HLD) specifies the components required to be integrated to provide the Audit Data Gathering & Storage facilities together with their interfaces and functionality.

The level of detail in this HLD is intended to be adequate to enable detailed design, implementation, integration and test work packages to be specified.

The Audit Data Gathering & Storage facilities have been designed to be generic and extensible; in particular any new applications introduced into the HNG-X system should interface to the Audit Server as specified in Section 6.2.

The Audit Data Extraction & Filtering facilities are provided for the Fujitsu Services Internal Auditors to provide extracts of the audit data from the Audit Archive in response to information requests from external auditors. These are described in *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

This document also includes details of the changes to the Audit system required for CP4305. The only impact of this change to the Audit system is that PANs, which are classified within PCI DSS as

Cardholder Data, are no longer stored in clear text form within Audit Tracks. Thus the Audit system requires a mechanism to handle hashed and/or encrypted PANs where they form a part of retrieved Audit data.

This document is based upon the Horizon equivalent document SD/HLD/001.

2 Scope

This High Level Design Specification covers:

- The interfaces between the applications creating the Audit Tracks and the Audit Data Storage facilities
- The mechanisms for the storage of the Audit Archive
- The structure of the data stored in the Audit Archive

The standard interface defined for data gathering gives the ability to gather & store audit files from new HNG-X subsystems & additional files from existing HNG-X subsystems with minimum impact on the Audit Design. All new Audit Points and files will need to be defined in the *Audit Server Configuration* (DEV/INF/ION/0001) as it provides details of the actual interfaces configured for each instance of the Audit Server.

The data that will be stored in the Audit Archive and hence the data that needs to be retrieved is defined in *Audit Server Configuration* (DEV/INF/ION/0001).

The scope of this HLD does not cover:

- Audit Track Generation
- Access to the current data on production systems
- Auditing of any information beyond the HNG-X external gateways
- Specification of Information Requests, this is defined in *Audit Data Extraction Process* (IA/PRO/004)
- Online access to live data to support Internal Audit
- The analysis/interpretation of Audit Tracks in order to identify specific Audit Trails

The following functionality is covered in detail in the document *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

- Online extraction of Audit data from the Audit archive
- Seal Checking to ensure extracted files have seals intact
- Maintenance & Monitoring of Audit Record Queries (ARQs)
- Presentation to Auditor, tools to present data in required format
- The interfaces between the applications restoring the Audit Tracks and the Audit Data Extraction facilities

3 Design Principles

The main principle of the Audit design has been to provide the required audit data storage and retrieval facilities while minimising the impact on development activities and timescales.

The Audit Architecture as defined in *HNG-X Architecture – Support Services* (ARC/SVS/ARC/0001) identifies the need to be able to cope with change as the usage of the HNG-X system develops, especially as new applications and services are introduced. Thus a significant design principle is for the Audit system to be able to support the introduction of such new facilities with minimum impact.

The extraction mechanisms must provide only that data that is appropriate to the External Auditor Role requesting the audit data.

Other principles are to provide facilities that can be built upon and developed to provide enhanced audit facilities for future releases.

4 Requirements

Incoming requirements will be captured in the Support Services Topic Architecture SRS. A skeleton copy of this document, containing section numbers only, will be lodged in the requirements system. For each requirement in the SRS, the skeleton document will identify in each section where the requirement is explicitly addressed or allocated to a topic or component architecture.

4.1 Design Assumptions

This section records the explicit assumptions relating to requirements that have been made during this design.

1. A failure of an Audit Server on one data centre has not been rectified until it has caught up with the gathering of the Audit Tracks from all sources of Audit Tracks in that data centre. A failure of the second Audit Server on the other data centre during this period is deemed to be a second point of failure.
2. The facilities provided for auditor access will be adequate for any other access needs. e.g. in support of the investigation of processing failures.
3. The maximum time to fix an Audit Server following an outage for whatever reason (while the rest of the data centre operates correctly) will be five days. All points within the HNG-X system that generate Audit Tracks have enough disk space to hold at least 5 days of Audit Tracks.
4. It is acceptable to duplicate Audit Tracks in the Audit Archive and duplicate records within Audit Tracks.
5. It is assumed that all sources of the Audit Track data will use synchronised clocks.

5 System Overview

HNG-X Architecture – Support Services (ARC/SVS/ARC/0001) describes the overall audit system design for HNG-X. This section provides an overview of the design of the Audit System.

This overview provides more detail, in its focused area, than the Audit Architecture while acting as an introduction to the detail in the rest of the high level design.

5.1 Major Components

Figure 1 shows the major components of Audit Solution at a single data centre. The configuration is duplicated on both centres.

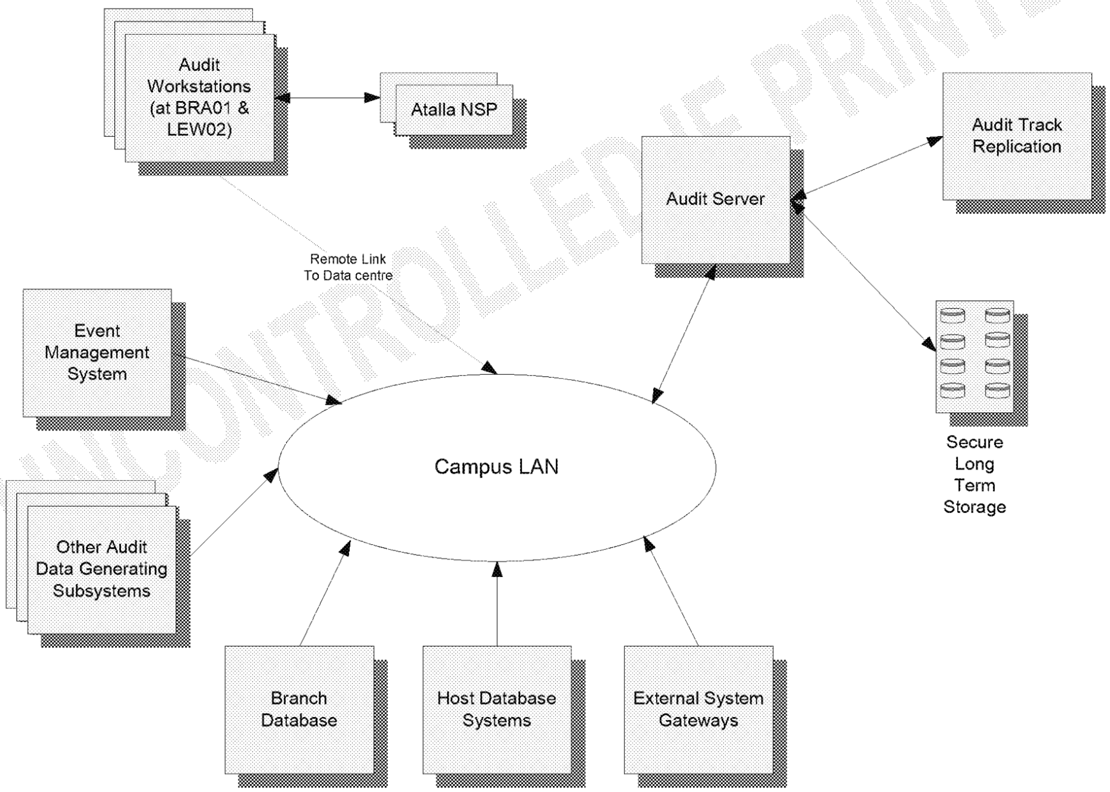

Figure 1 - Main Components of Audit Data Storage and Retrieval

The Audit server is responsible for gathering Audit Tracks generated from a wide range of components of the HNG-X systems including: -

- Post Office Counters
- Systems Management Facilities
- Database Hosts (including the Reference Data System)
- FTMS Gateways

- System Scheduler logs

Audit Tracks (including the External Gateway Audit Tracks) are not automatically duplicated at each data centre. Files for these tracks are produced by each Audit Server and are exchanged between the two sites utilising the Inter-data centre link.

An Audit server exists at each data centre, i.e. IRE11 & 1RE19, they operate in an Active/Active configuration, each audit server gathering audit tracks from their local data centre. In normal post Hydra operation, audit tracks will only be generated and gathered on the Active campus. If there is a complete failover of the HNG-X service to the DR data centre, the audit server at the DR data centre will take over responsibility for gathering audit data from all failed over machines. During such a failover it will not be possible for the IRE19 audit server to replicate to the IRE11 audit server.

During Hydra, the IRE11 Audit server will be responsible for gathering Audit data from retiring platforms located at the Bootle data centre & the IRE19 audit server for platforms located at Wigan.

Each Audit server is dependent on the other to support the replication of audit tracks.

As well as gathering and storing audit data on EMC Centera all of the Audit Tracks, the Audit Server provides facilities to retrieve data from the Audit Archive.

Tools to extract and prepare data for analysis are provided together with basic facilities to support internal Fujitsu Services data retrieval activities (see *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029) Access, by Fujitsu Services staff, to the retrieval and extraction facilities is via the user interface provided on the Audit Workstation.

The function of the Audit Retrieval & Extraction system is to: -

- Effect extraction of branch transaction records via any Audit Workstation
- Extract the appropriate records using specified extraction criteria to give a manageable number of files/records. Extracting a subset of records from an audit archive file is not provided for all file formats.
- Provide facilities to filter and subsequently browse the extracted data to meet the criteria of the 'Information Request'
- Seal check retrieved files
- Maintain & Monitor Retrievals by ARQ

5.2 Audit Server Overview

Figure 2 shows the major logical components of the Audit Server.

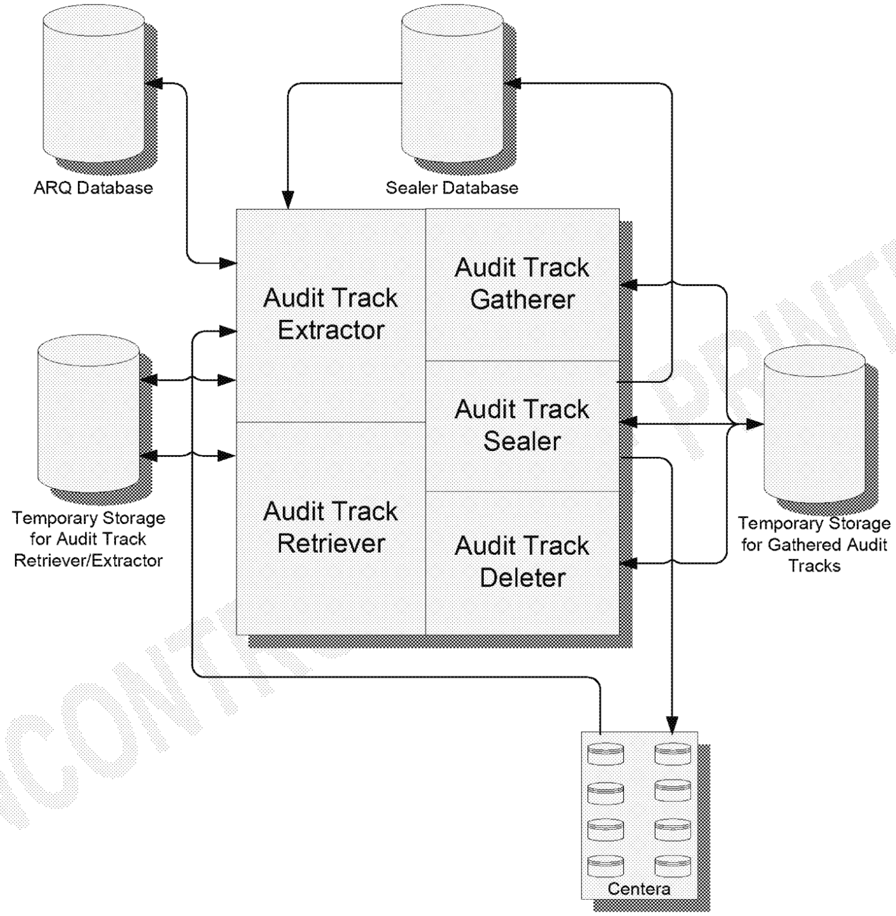

Figure 2 - Audit Server Logical Structure

The functions of the major logical components of the Audit Server are described below:

5.2.1 Audit Track Gatherer

The Audit Track gatherer is responsible for Collection of Audit Tracks that have been generated within the HNG-X system. The majority of these tracks are created on different platforms and are gathered onto temporary disk storage on the Audit Server.

The Audit Track Gatherer is responsible for any renaming of the gathered files.

Gathering is implemented using Windows shares. This requires that Samba is installed on UNIX systems.

The Audit Tracks need to be gathered at regular intervals. The Scheduling of the transfers varies with the type of Audit Point and the locations from which the tracks are gathered and is controlled via the scheduling facilities of the HNG-X system and frequency / time period facilities within the Audit solution.

Multiple instances of the Audit Track Gatherer can be configured on a single Audit server.

5.2.2 Audit Track Sealer

The Audit Track sealer is responsible for calculating a seal value for the file prior to writing the audit track to the archive.

The Audit Track is then compressed (if so configured) and written to Centera and the Audit Track details stored in the Audit Server database.

When an Audit Track is retrieved its seal is recalculated and checked against the value in the database. This confirms the integrity of the retrieved file. See *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029).

5.2.3 Audit Track Deleter

The Audit Track Deleter is responsible for the deletion of Audit Tracks from the machines on which they were generated after they have been gathered. The point in the processing of an Audit Track (by the Audit Server) at which the original copy of each gathered file is deleted is configurable. Audit Track Deletion takes place between the completion of Audit Track Gathering and some (configurable) time after the completion of Audit Track Sealer for any particular Audit Track file.

The Audit Track Deleter is also responsible for regularly producing a list of files processed by the Audit Server. The details of processed files are archived as part of the Audit Servers own Audit Trail.

5.2.4 Audit Track Retriever

See *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

5.2.5 Audit Track Extractor

See *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

5.3 Major Data Flows

The distributed nature of the HNG-X system combined with the relatively long timescales for Audit Server operations and the need for it to operate with minimal operator intervention mean that the Audit Server must be designed to be robust in operation and to automatically cope with and recover from unsuccessful operations and a wide range of failure modes.

The flow of files and other data through the AS is a key feature of its design. Full details of the flows through the components are given in the interface definitions of each component in section 6.1 below.

Figure 3 below provides a view of the key flows.

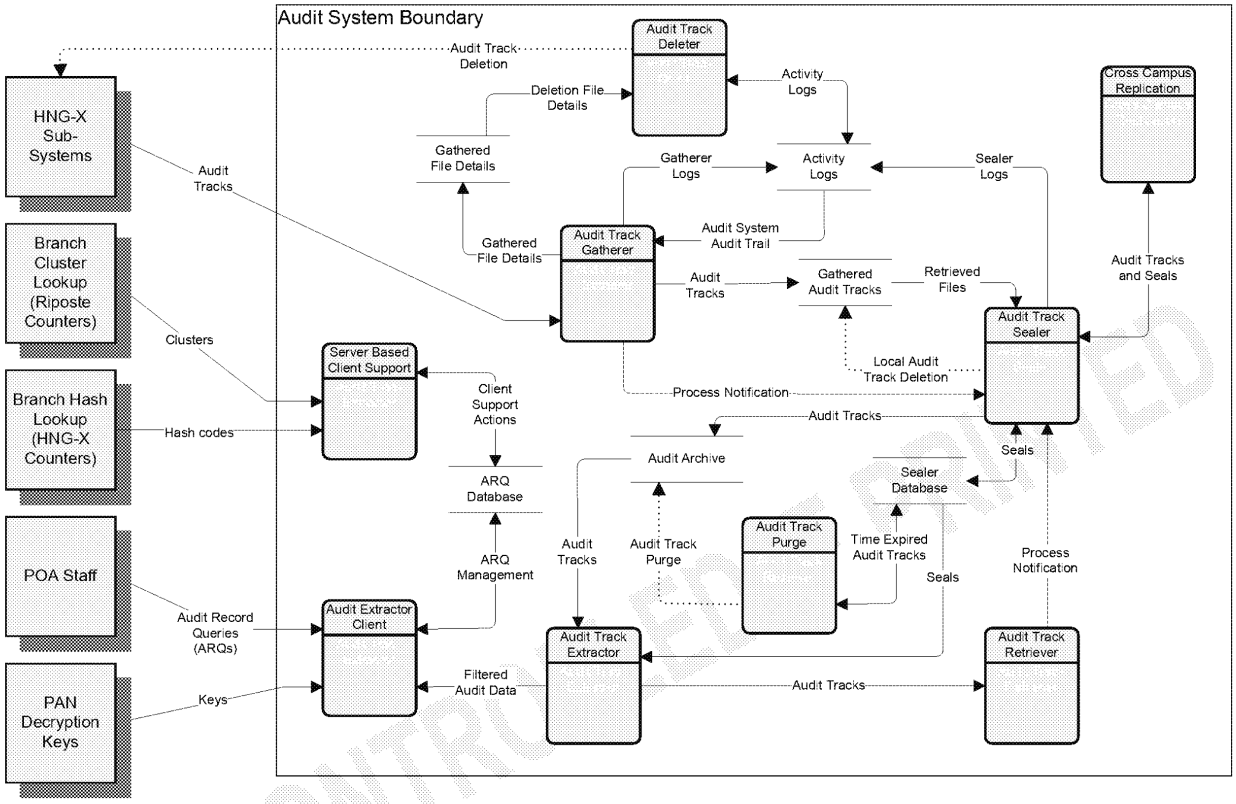

Figure 3 - Major Data Flows between the Audit Server Components

Where the inter component communication relates to control information that information is held in text form. This is designed to allow administration staff to manually control the flow of data in emergency circumstances.

For performance reasons, whenever possible file movement is achieved by “moving” directory entries rather than copying all of the files.

A number of the activities carried out by the Audit Server can take a relatively long time and can utilise a very high percentage of system resources in addition to requiring large amounts of data to be passed between many of the components. In order to support individual components being started and stopped independently, persistent storage is used to hold inter component communications. This also has the advantage that the loss of a single component for whatever reason will allow the other components to continue processing any input data in their persistent input buffers.

The Audit Server employs WinZip to compress files prior to them being written to Centera. This compression is optional[1] and is controlled by the Audit configuration file. See *Audit Server Configuration*

[1] HNG-X CP0299 introduced Pre-Compressed audit data sources. These audit tracks are not compressed again by the Audit server.

(DEV/INF/ION/0001). Typically an audit Sub-Point will not be configured for compression if the audit tracks within it are already compressed by the generating sub-system.

The design of the Audit Server does not include any synchronisation between the Audit Track Sealer and the Audit Track Extractor

Parts of the Audit Server, which are controlled via configuration files, e.g. the ATG, should all (at a configurable time) be scheduled to read the configuration information and implement any changes to their configuration.

5.4 Audit Workstation

See *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

5.5 Audit Archive

A copy of the Audit Archive is kept at each of the data centres, each one is held on a separate EMC Centera. Audit Tracks generated on each data centre are added to the local copy of the archive..

A copy of all the stored Audit Tracks is written to an Export directory and subsequently transferred to an Import directory at the other data centre using Robocopy. The files are subsequently added to the copy of the Audit Archive held there from this directory.

In addition to the transfer of Audit Tracks (via copying over the inter-data centre link) between the two sites it is necessary to transfer copies of the seals (relating to the transferred files) in a secure manner. This enables the remote machine to verify the integrity of the transferred files. The seals are exchanged using the same robust file copy mechanism as used for transferring Audit Tracks.

When Audit Tracks are introduced on an Audit Server by the copy over the inter-data centre link the files containing the names of the tracks are subsequently introduced into the Sealer database. Thus each Audit Server maintains a full index of all of the files held in its part of the Audit Archive.

5.6 Audit Track Replication

Audit tracks that are gathered at one data centre are replicated to the Audit server at the remote data centre. This replication process is managed by the Audit Track Sealer. As Audit tracks are secured to the Audit archive, they are moved to an export area awaiting transfer to the remote campus. A second file, containing the calculated seal value for the audit track is also stored in the export area.

Audit tracks & seals are copied, using robocopy, to the equivalent import area on the remote audit server as part of Audit server overnight schedule. On arrival, the sealer on the remote audit server recalculates the seal value of the imported audit track and compares it with the original value in the imported seal file. Assuming they match, the file is then written to the remote Audit archive. If the seals do not match, the Audit track & seal file are moved to a holding area & an event is raised. Manual investigation is necessary to investigate the cause of the discrepancy.

Figure 4 illustrates this process.

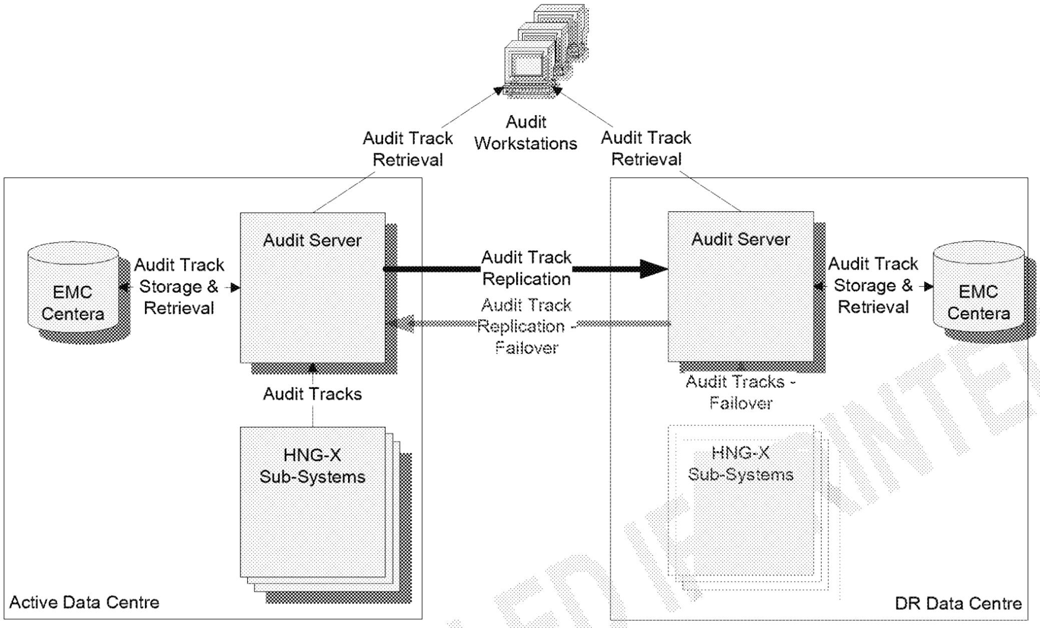

Figure 4 - Audit Track Replication

If, during a site failover, the IRE11 audit server is not available & the IRE19 server has taken over responsibility for gathering audit data this replication will not take place

5.7 Audit Points

An Audit Point is a logical concept introduced in this design to minimise the linkage (as seen by users of the Audit Workstation) with the physical design of the HNG-X system. This is intended to help reduce the knowledge that the end users need of the details of the way HNG-X has been implemented.

The term Audit Point is used, in a number of places within this design, to refer to the logical location at which a particular Audit Track is generated, e.g. where the TMS Audit Track is generated. Due to the distributed and resilient nature of the HNG-X system an Audit Point is actually realised at a number of different physical locations.

The specific locations at which the Audit Track of a particular Audit Point is generated are identified as Audit Sub-Points. An Audit Sub-Point maps onto a single sub-directory on a single component in the HNG-X system. It is however possible for an Audit Sub-Point to map onto a (finite) set of such sub-directories. Where there are a number of sub-directories they will be nested beneath a single top level sub-directory.

The files stored in the Audit Archive are named in terms of an Audit Point and an Audit Sub-Point. These logical concepts are designed to be stable across the life of the HNG-X system and to assist in locating the files containing the relevant data to support any particular audit activity. For example, all TMS journal (i.e. Correspondence Server Audit Track) files will be identified by the same Audit Point and all files generated on a given Correspondence Server cluster will be identified by the same Audit Sub Point.

Audit Track data collected from an Audit Sub-Point can be held in multiple files (of multiple data types).

The files containing Audit Track data collected from a single Audit Sub-Point shall all have the same retention period.

Migration of systems that generate Audit data and the consequential, incremental changes to the audit configuration will need to be carefully coordinated during the migrations of systems to the new data centres and to the HNG-X environment. As retiring Horizon systems go off-line, their corresponding audit configuration entries will need to be disabled, and as new systems appear they will need to be configured in line with the HNG-X Audit and Security Policies. This is further described in section 10.

6 System Components

The following conventions are used in this section:

Interface Naming - all interface names have the form “I-“<TLA>-<n><*>, where <TLA> is a three (or four) letter acronym identifying the logical component, <n> is an Id number, <*> is optional but if present indicated an interface which can have multiple instantiations. Examples are I-ATG-4*. If necessary instances of an instantiated Interface are identified by the addition of “-“<n> on the end, e.g. I-ATG-4*-3. Interfaces to common infrastructure items e.g. operating system are not named.

Note that the interfaces are per component and links between components of the Audit Server have two, identified by interface names at both ends of the link, e.g. I-ATG-3/I-ATS-4.

There are two instances of the Audit Server one at each data centre.. Unless otherwise specified all discussions apply to each instance.

6.1 Application Components

6.1.1 Audit Track Gatherer

Figure 5 identifies the main interfaces to the Audit Track Gatherer (ATG)

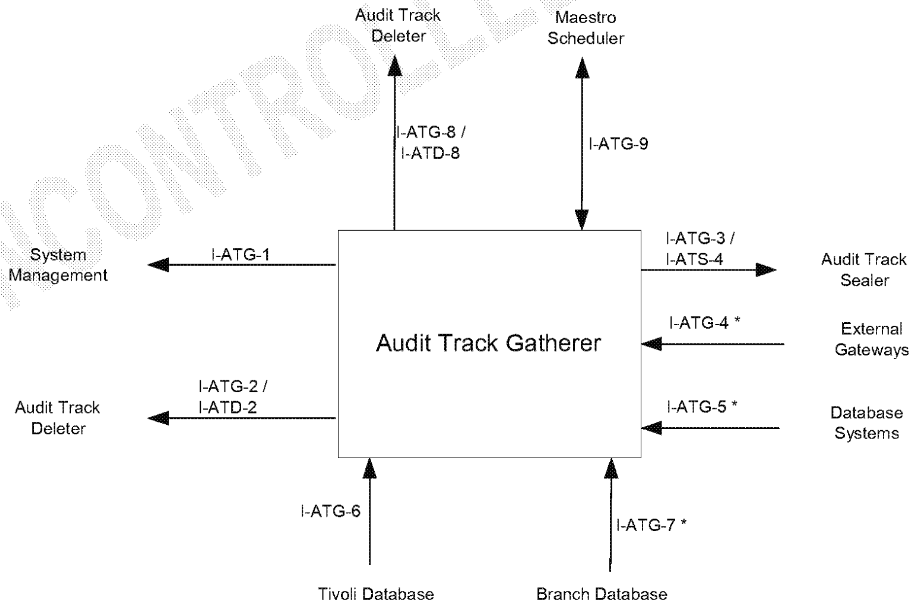

Figure 5 - Interfaces to the Audit Track Gatherer

6.1.1.1 ATG Functionality

The Audit Track Gatherer can be configured to run multiple instances on a single Audit Server, each instance collecting Audit Tracks from different Audit Sub-Points. Each instance of the ATG will be configured with a unique id.

The ATG performs regular checking for new Audit Tracks on remote machines and their gathering by the methods specified in the interface definitions. The scheduling of individual instances of the ATG is under the control of the HNG-X wide system scheduler. The frequency of checking for Audit Tracks needs to be configurable based on Audit Sub-Points. In order to enable load balancing to be matched with slack periods in the system the gathering times shall be configurable and definable as times, at which instance of the ATG shall run.

Within the overall schedule further control of the gatherers is required, as frequencies within those time periods at which the ATG instance shall attempt to gather. The level of control required is:

- Frequency of gatherer activity
- Duration of those period of activity

All locations from which Audit Track files are to be collected shall be configurable.

All Audit Tracks are subject to renaming before being stored in the Audit Archive. The renaming is carried out by the Gatherer based on the rules of the renaming policy associated with the Audit Sub-Point. See *Audit Server Configuration* (TD/ION/011) for details of the renaming policies.

Once the gathered files have been renamed the ATS is notified via I-ATG-3

A record of activities is to be passed to the Audit Track Deleter via I-ATG-2.

Since the interval between the gathering of files from a given location can be substantially shorter than the interval between a particular file being gathered and that file being deleted (in normal operation) it is necessary for the ATG to be able to identify which files have already been gathered and which have not. The method of marking will be decided as part of the detailed design but must have the following characteristics

- Work over an ATG crash
- Allow files marked as being gathered to be unmarked (without loss of information) by administrators in an emergency situation
- Cause, at most, minor increases in network traffic. See section 9.2 for an analysis of the failure modes. Once a file has been successfully copied to the Audit Server the ATG will mark it as gathered and then signal to the ATD. The ATG shall not attempt to gather files that have been marked as gathered.

6.1.1.2 ATG Interfaces

6.1.1.2.1 I-ATG-1

Interfaces to the system administration staff for monitoring and control purposes.

Monitoring of the system is through the event log.

The ATS will be under the control of the HNG-X scheduler. In addition command line interfaces are to be provided to start and stop instances of the ATG. Facilities to stop an instance of the ATG shall provide an option to stop an instance immediately (and leave the current file ungathered) or to stop the instance when it has finished gathering the current file.

6.1.1.2.2 I-ATG-2

For every Audit Track file successfully (or unsuccessfully) gathered by the ATG it will inform the Audit Track Deleter (ATD) of:

- The time and date of the gathering
- Id of the ATG instance
- A meaningful success/failure code
- The path name (including remote share details) of the remote (gathered) file
- (If successful) The name the gathered file was renamed to by the ATG
- (If successful) the date and time it was passed to the Sealer
- The size of the file
- The elapsed time taken to gather the file

A file has been successfully gathered when it has been renamed and made available to the Audit Track Sealer.

The ATG will pass the details of the files gathered to the ATD via an agreed directory. Details of the gathered files will be put in a Record File in the shared directory for collection by the ATD. Each Record File must contain the details of at least one gathered file, but may contain many, or an indication that no files were available for gathering from the remote directory. Details of gathered files batched together in one Record File must always be held on persistent storage media (e.g. disk) and must never be over written. Record Files must regularly (e.g. every hour) be passed over to the ATD. The Record File shall be a text file.

6.1.1.2.3 I-ATG-3

This interface makes gathered Audit Tracks available to the Audit Track Sealer. The files are written to a common directory and are picked up by the Audit Track Sealer. Files in the directory are immediately available for sealing. The Audit Track Sealer is responsible for removing the files from the directory once they have been sealed.

6.1.1.2.4 I-ATG-4

This interface represents the gathering of Audit Tracks from the FTMS implementation on the local external gateways.

There are multiple external gateways. An external gateway requires two systems one on a HNG-X data centre (the local gateway system), one on the appropriate external site (the remote gateway system).

This design only supports the gathering of transferred files from Local Gateway Systems (both primary and secondary).

For each Logical FTMS External Connection, control files and copies of transferred files are gathered from different directories. The term audit directory is used in this HLD to refer to a directory containing copies of the transmitted files. The term control file directory is used to refer to a directory containing the control files.

Files to be gathered are accessed using remotely mounted windows shares

Copies of transferred files are available for gathering as soon as they have been placed in the agreed audit directory for each Logical External Connection. All files placed in an audit directory are given a unique name.

It is potentially possible if difficulties are encountered in the transmission for the same file to be collected by the local gateway machine more than once. If this event does occur the copies of the files written to an audit directory are still given unique names.

When a local gateway machine receives a transmission from a remote gateway machine a copy of the transmitted file is written to the same audit directory.

Control files, written to a control file directory, are built up each day with details of the transfer being appended to the same file. The file name identifies the date of creation hence each file name is unique. The control files are only available for gathering after the day on which they have been created.

The Audit Track Gatherer must gather Audit Tracks from the gateway machines within 5 days of their generation. See section 9 for details of the resilience.

A gatherer instance only gathers files from local gateways at its local data centre. This means that related Audit Tracks may be gathered on different sites and will only be brought together when the Audit Tracks are exchanged by secure copy via the inter-data centre link.

6.1.1.2.5 I-ATG-5

A range of applications run on a number of database servers in the HNG-X system from which audit data is gathered.

6.1.1.2.6 I-ATG-6

This interface represents the gathering of Audit Tracks from the Tivoli Database on the local data centre. The Tivoli Database generates Audit Tracks of security relevant events that have been collected by the Systems Management Facilities.

6.1.1.2.7 I-ATG-7

The ATG on an Audit Server gathers Audit Tracks from the Branch Database at the local data centre. A large volume of data is generated on this interface.

6.1.1.2.8 I-ATG-8

Audit Tracks generated by components of the Audit Server are gathered on this interface. The files are written to a common directory and are picked up by the Audit Track Gatherer. Files in the directory are immediately available for gathering. The Audit Track Deleter is responsible for removing the files from the directory once they have been gathered in the same way remote files are deleted after gathering.

6.1.1.2.9 I-ATG-9

Instances of the ATG will be scheduled via the scheduler. The ATG shall implement the standard interfaces. Individual instances of the ATG will be scheduled to gather files from predefined (configurable) Audit Sub Points.

6.1.2 Audit Track Sealer

Figure 6 identifies the main interfaces to the Audit Track Sealer

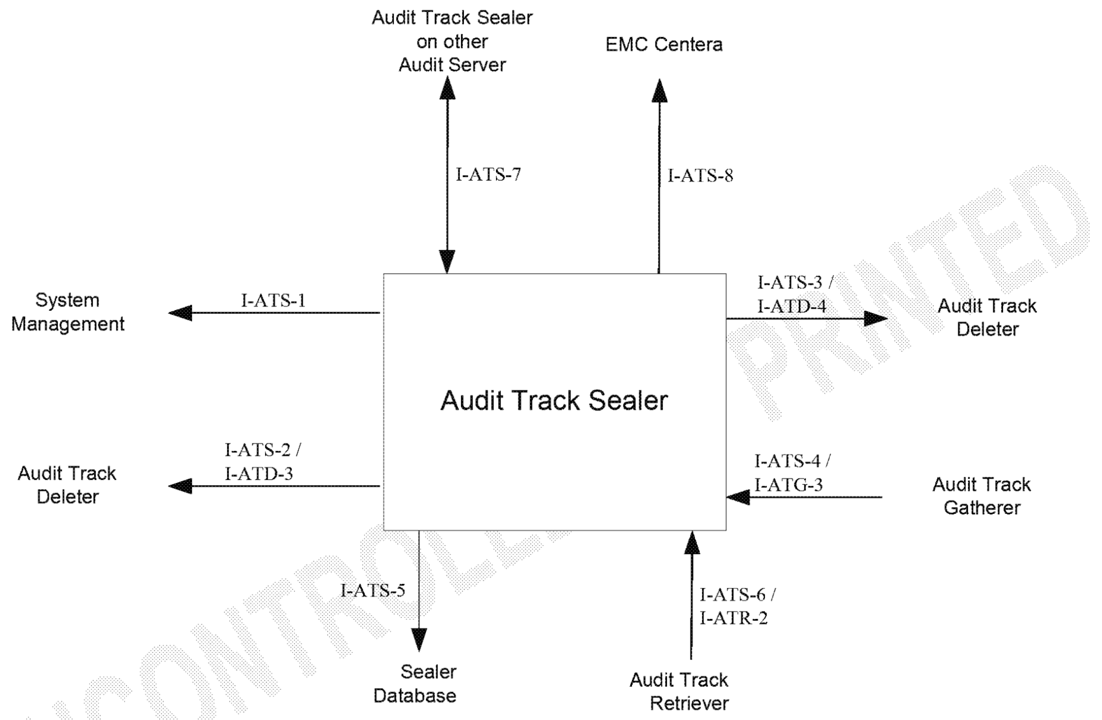

Figure 6 - Interfaces to the Audit Track Sealer

6.1.2.1 ATS Functionality

There will be a single instance of the ATS that concurrently accepts files for sealing/seal checking from ATG and ATR and notifies sealed files to the ATD and into the Sealer Database for subsequent use by the Audit Track Extractor.

The ATS shall collect files for sealing via I-ATS-4 and shall write a log of its activities to the ATD via I-ATS-2. In sealing a file the seal shall be generated using a secure hash algorithm, the MD5 algorithm has been selected.

Once a file has had a seal calculated the file will be written to Centera & details will be stored in the Audit Track Seal Database via I-ATS-5.

Only once a sealed file has been written to Centera can any local copies be deleted.

The ATS shall give priority to checking the seals as requested by the ATR over generating seals for newly gathered data.

The seals and audit tracks which are to be copied to the remote Audit Server shall be configurable per Audit Sub-Point.

The Seal Database will be backed up by the Audit Server at the alternate site as part of the regular daily backup, which will be retained for 2 weeks.

The size of the Sealer DB should not grow indefinitely hence the Audit Track Purge process will occasionally remove any seals with a retention period that have expired and the corresponding Audit Tracks from Centera – see section 6.1.5.

The ATSDB will need to support not only the storage and eventual deletion of the seals but also the querying of seal data from the Audit Track extraction facilities.

6.1.2.2 ATS Interfaces

6.1.2.2.1 I-ATS-1

Interfaces to the system administration staff for monitoring and control purposes.

Monitoring of the system is through the event log.

The ATS will be under control of the HNG-X system scheduler scheduling.

In addition command line interfaces will be provided to allow control of the ATS outside of the schedule to start and stop instances of the ATS. Facilities to stop an instance of the ATS shall provide an option to stop an instance immediately (and leave the current file unsealed) or to stop the instance when it has finished sealing the current file.

6.1.2.2.2 I-ATS-2

For each file that is successfully or unsuccessfully sealed or a seal checked a record is written to the Audit Track Deleter. The record shall contain:

- Date and time of sealing
- File Name (as received from the Audit Track Gatherer/Audit Track Retriever)
- An indication of whether the operation was sealing or checking
- A meaningful success/failure code
- The id of the instance of the ATS
- The size of the file
- The elapsed time to seal the file
- An indication of the checksum algorithm used

A file has been successfully sealed when a seal has been calculated, details of the seal stored and the file made available to the Audit Track Deleter (ATD).

The ATS will pass the details of the files gathered to the ATD via an agreed directory. Details of the sealed files will be put in a Record File in the shared directory for collection by the ATD. Each Record File must contain the details of at least one sealed/checked file, but may contain many. Details of sealed/checked files batched together in one Record File must always be held on persistent storage media (e.g. disk) and must never be over written. Record Files must regularly (e.g. every hour) be passed over to the ATD. The Record File shall be a text file.

6.1.2.2.3 I-ATS-3

When a file has been successfully sealed and details of the seal have been written to the Audit Track Seal Database the file shall be made available to the Audit Track Deleter. Such sealed files are placed in a shared directory to be picked up by the ATD.

6.1.2.2.4 I-ATS-4

See I-ATG-3. The ATS picks up files to be sealed from this interface.

6.1.2.2.5 I-ATS-5

For each file that is sealed and written to Centera or has its seal checked an entry is written to the Sealer database. Each entry shall contain:

- Date and time of seal calculation
- File Name (as received from the Audit Track Gatherer/Audit Track Retriever)
- An indication of the checksum algorithm used
- An indication that it is the generation of an initial seal or the checking of a seal
- The value of the seal(s), i.e. the original seal in the case of a generation or the recalculated seal in the case of a check.
- An indication of whether the seal was generated on the local Audit Server or whether it was imported from the remote Audit Server.
- (If an initial seal calculation) The retention period for the Audit Track.
- Centera ID

6.1.2.2.6 I-ATS-6

See I-ATR-2. Files whose seals are to be checked are notified to the ATS by the Audit Track Retriever (ATR) via this interface. Each request is in the form of a marker file that uniquely identifies the filename of the file to be seal checked. Files to be seal checked are placed in an agreed directory. The ATS regularly checks the directory and removes the first entry and checks the seal. The ATS is responsible for removing files from the common directory.

6.1.2.2.7 I-ATS-7

This is a bidirectional interface that is used to support the replication of audit tracks to the remote audit server.

As files are sealed and written to the local Centera a pair of files, one containing the original audit track and one containing the calculated seal are written to an export directory. As part of the audit server overnight schedule, these export files are copied to an import directory on the remote audit server using robocopy. The schedule must ensure that both sealer processes are inactive while this copy takes place.

Upon reactivation, the remote audit server will scan the import directory picking up the pairs of files. It will recalculate the seal on the imported audit track and compare it with the seal within the imported seal file. This confirms the integrity of the data as received on the interface. Assuming the seals match, the remote sealer will then write the file to its Centera and record the details in its ATSDB.

If the seal cannot be verified, the pair of files are written to a holding area and an event raised. Manual intervention will be necessary to resolve this.

6.1.2.2.8 I-ATS-8

Once the seals have been calculated the audit tracks are compressed if the associated Sub-Point is configured for compression and written to Centera.

6.1.3 Audit Track Deleter

Figure 7 identifies the main interfaces to the Audit Track Deleter (ATD).

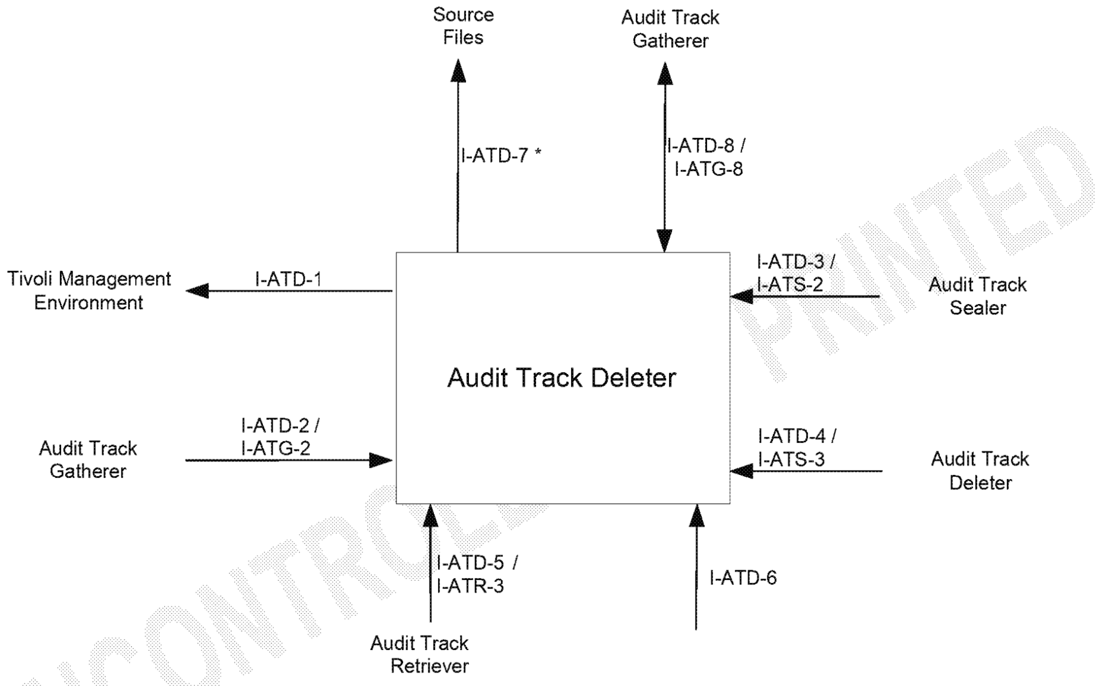

Figure 7 - Interfaces to the Audit Track Deleter

6.1.3.1 ATD Functionality

The Audit Track Deleter has two main functions:

- Deletion of the remote copies of Audit Track files which have been gathered from remote systems
- Writing of a single audit log of the main actions of the Audit Server.

There is only one instance of the ATD per Audit Server.

The point at which the remote copies of audit files are deleted needs to be:

1. Once the ATG has signalled that it has completed the gathering
2. Once the ATS has signalled that it has sealed the file and written it to Centera
3. A configured time (typically hours) after the above events.

6.1.3.2 Source Audit Track Deletion

The queue of audit tracks to be deleted, maintained by the ATD, must be stored on persistent storage so that the queues can be reused after ATD outages.

The delay before the source file is deleted should be configurable per Audit Sub-Point.

All attempts to delete files (whether successful or not) should be auditable. Where a file can not be successfully deleted the ATD shall retry a number of times. The number of times and the interval between the files shall be configurable. If after all of the retries the file has not been successfully deleted an exception report shall be generated.

The file names supplied by components other than the ATG will potentially have been transformed by the ATG; part of the deletion function will be to match the transformed file name to the original file (and path) name.

6.1.3.3 Auditing

The ATD shall generate two files for audit (and other administration purposes):

- An Activity File
- A Deletion Exception File

The ATD shall write out all messages it receives across interfaces I-ATD-1 to I-ATD-6 to the Activity File. The file shall be a text file formatted to ease reading/scanning. The results of file deletion attempts via I-ATD-7* shall also be recorded to the Activity File.

The ATD shall also calculate and record the number of files and the total amount of data archived each day. This information shall also be written to the Activity File.

Details of any file which is not successfully deleted after the maximum number of retries shall be recorded in the Deletion Exception File.

The Activity File and the Deletion Exception File will be written daily. A copy shall be passed to I-ATD-8 and a copy stored for diagnostic usage. These will be retained online for a period of 14 days.

The files shall both be text files, formatted for ease of reading.

6.1.3.4 ATD Interfaces

6.1.3.4.1 I-ATD-1

Interfaces to the system administration staff for monitoring and control purposes.

Monitoring of the system is through the event log.

The ATD will be under control of the HNG-X system scheduler. In addition command line interfaces should be provided to start and stop the ATD outside of scheduler control.

6.1.3.4.2 I-ATD-2

See I-ATG-2

6.1.3.4.3 I-ATD-3

See I-ATS-2

6.1.3.4.4 I-ATD-4

See I-ATS-3

6.1.3.4.5 I-ATD-5

See I-ATR-3 & *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

6.1.3.4.6 I-ATD-6

For each file that is subject to a status change within the ARQ Database, a log is written to the ATD. Each record will contain

- ARQ Reference
- Filename
- Status
- Date/Time of Status Change

6.1.3.4.7 I-ATD-7*

This interface is used by the ATD to delete files. These are the files collected by the ATG on I-ATG-4*, I-ATG-5*, I-ATG-6, I-ATG-7*, and 8.

6.1.3.4.8 I-ATD-8

See I-ATG-8

6.1.4 Audit Track Hoarder

This component is now redundant.

6.1.5 Audit Track Purge

The Audit Track Purge process runs on a daily basis and deletes Audit tracks that have reached their expiry date – either 18 months or 7 years after they were gathered. The process will both delete the audit track from Centera and delete the corresponding entry in the Sealer database.

The process maintains an audited log of its actions.

6.2 Interfaces

All current interfaces are defined in section 6.1. Any newly identified applications where audit files need to be gathered should comply with the standard interface between the Audit Track Generating Applications and the Audit Track Gatherer as detailed below.

- The Audit Track files generated at each Audit Sub-Point shall be placed in a dedicated (single) sub directory for the Audit Track Gatherer to collect. The files placed in each such sub directory shall either (i) be moved into the directory such that they are available for collection immediately they have been placed there (ii) be renamed to indicate their availability for gathering (e.g. the file is renamed from xxx.tmp to xxx.dat. Exclusion from gathering based upon file suffix is defined in the Exclude parameter in the audit server configuration file). The Audit Track directory must be shared and permissions set such that the Audit system can gather and deleted the files within it.
- The sub directories should not be used to hold any other files.
- The name of each Audit Track file shall conform to one of the naming standards specified in *Audit Server Configuration* (DEV/INF/ION/0001)
- Access to the Audit Track files for gathering shall be via Samba (for Unix systems) or NTFS (for Windows systems). Access to the sub directory shall be limited to the application generating the Audit Track and the Audit Track Gatherer. Audit track files should be written in write-append mode.
- The Audit Track Deleter shall delete files from the sub directories a configurable period of time after they have been successfully secured in Centera. Once the Audit Track files are available for gathering the generating application should not make any assumptions about their ongoing availability in the sub directory.
- In order to minimise the need to re-gather files following a network or other failure the size of the Audit Track files should be controlled. Typically they should not exceed 100MB. Within that limit the number of files should be minimised.
- The share containing the Audit Tracks should have sufficient capacity to withstand a 5-day outage of the audit system.
- Permission must be set on the share such that the audit system can read and delete files contained within the share

6.3 Distributed Application Services

Figure 8 show the distributed application interfaces between the components of the Audit system

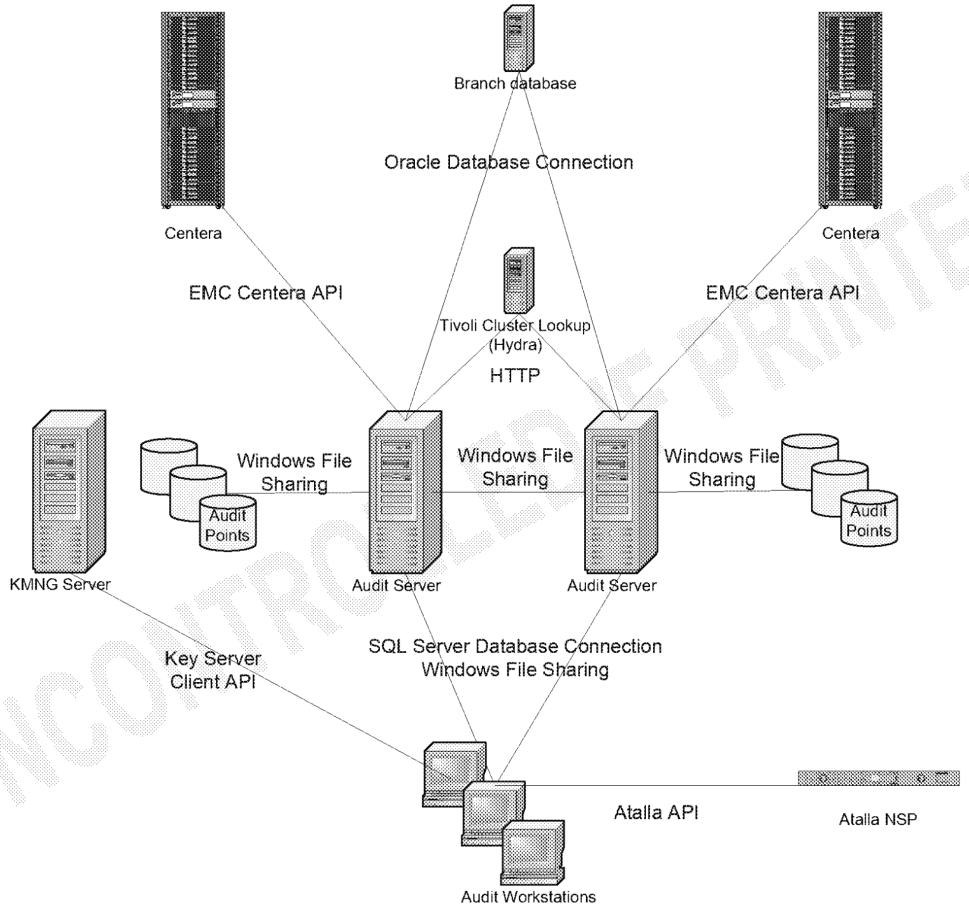

Figure 8– Audit Distributed Application Interfaces

The ARQ SQL Server 2000 database is used to support client-Server communication between the Audit workstations & servers. Service requests are initiated by writing to an interface table in the database. Triggers on the interface table are used to initiate server based processes. The client monitors the progress of server actions via a status table in the database.

The Audit servers access the EMC Centera units using the EMC Centera API.

The Tivoli cluster lookup service is used to determine which Correspondence Server cluster a particular branch belonged to at a particular point in time. This is only relevant to Horizon Riposte transactions, but access to this data must be maintained for the 7 years that Riposte data exists in the Audit archive. The cluster lookup service will be retained. As part of SYSMAN2, while Windows NT4 based platforms still exist. When the service is eventually retired, the cluster look up data will be imported into the Audit system and accessed directly. See *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

Usage of other interfaces is described in *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029).

6.4 Networking Services

The Audit Server will use the standard HNG-X network services. Figure 9 shows the networked interfaces between components of the Audit system. All application services utilise TCP and/or UDP protocols on a predictable set of ports.

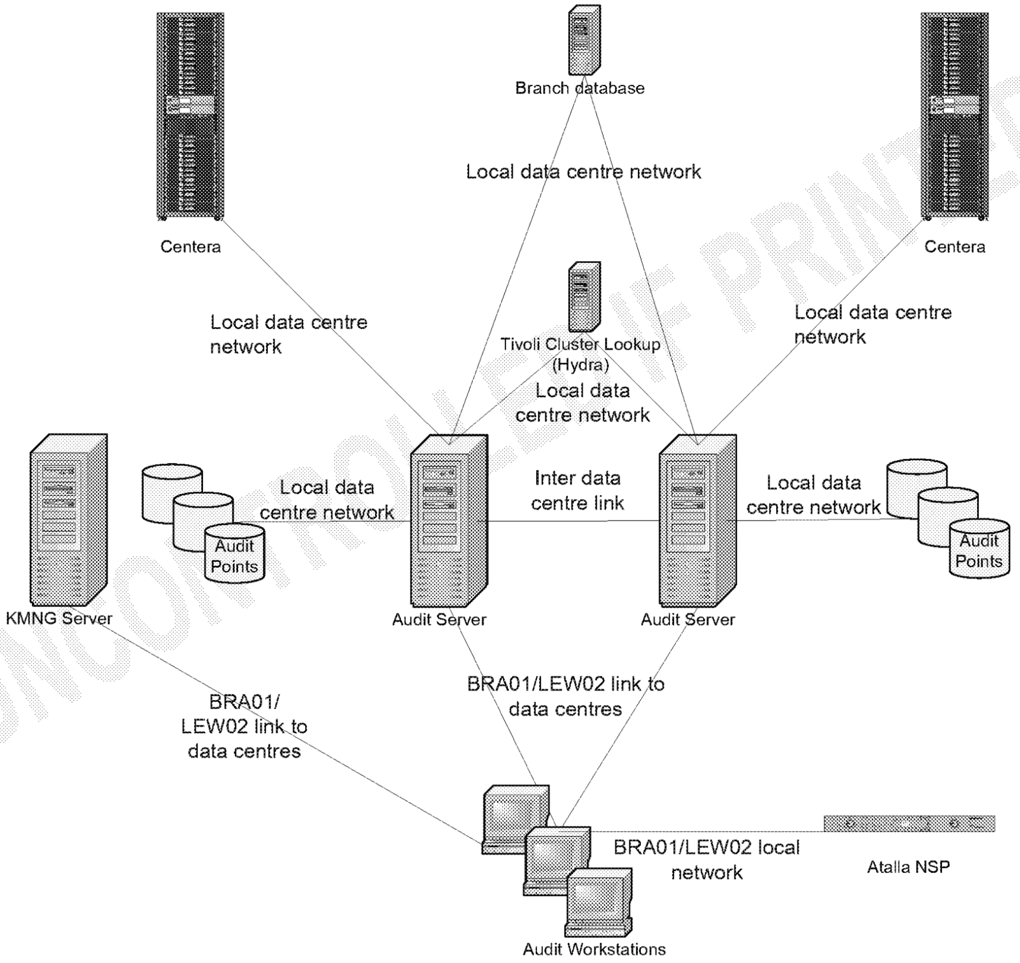
Centera
Branch database
Centera
Local data centre network
Local data centre network
Tivoli Cluster Lookup (Hydra)
Local data centre network
Local data centre network
Local data centre network
Inter data centre link
Local data centre network
Audit Points
Audit Points
KMNG Server
Audit Server
Audit Server
BRA01/ LEW02 link to data centres
BRA01/LEW02 link to data centres
BRA01/LEW02 local network
Atalla NSP
Audit Workstations

Figure 9 – Network Services

The Audit system can move large volumes of data around the network. On a peak day, it is estimated that the audit system will gather 10 GB of data at the active data centre. This data will flow as follows

1. Local data centre network - HNG-X systems ➔ Audit server 1
2. Local data centre network – Audit server 1 ➔ Centera 1
3. Inter data centre link – Audit server 1 ➔ Audit server 2
4. Local data centre link – Audit server 2 ➔ Centera 2

In addition, during Hydra, the following additional flows will exist

1. Bootle/IRE11 link - Horizon systems → Audit server 1
2. Wigan/IRE19 link - Horizon systems → Audit server 2
3. Inter data centre link – Audit server 2 → Audit server 1

During Hydra, a peak load of 35GB per day will flow from Bootle → IRE11 and Wigan → IRE19. The volume of this traffic will decrease as the HNG-X counter rollout progresses.

The bulk of Audit data (by volume) is generated by the Branch database. This data will be made available to the audit system by an overnight extract taken shortly after midnight. Thus the audit system will utilise the local network quite heavily during this period.

Replication of audit data must take place when the audit system is quiescent. This is scheduled to take place after the Audit backups have completed. During this period the audit system will be heavily utilising the inter data centre link. This may require some traffic management on the network to avoid the Audit replication traffic flooding the link.

If the Audit server is in recovery mode (i.e. it has been out of action for a period of time) it will put a considerably higher load on the network than stated above.

6.5 Platforms

The Audit servers run in an Active/Active configuration. They are housed in the BladeFrame.

The Audit servers will run Windows Server 2003 Standard Edition.

The servers will be of the following minimum specification:

- 2 Dual Core CPUs
- 4GB RAM
- 1.1 TB of non replicated storage

The Physical Platform Design for the Audit server is *HNG-X Audit Server (ARC) - Physical Platform Design (DES/PPS/PPD/0035)*

7 Systems Management

7.1 Monitoring

The Audit applications use standard Windows event logging to record non-fatal application events. Events are classified using the standard mechanism and are marked as Information, Warning or Error as appropriate.

The SYSMAN3 Event Management system is responsible for analysing this event stream and taking suitable action.

The Audit servers will be monitored by the Real time Active Dashboard (RAD) the following spreadsheet categorizes the events that are raised by the audit processes.

Audit_EventsforRAD
.xls

The Audit gathering system automatically manages the availability of free local disk space to hold temporary copies of Audit Tracks. If the amount of free disk space falls below a given threshold, the gathering system will suspend until sufficient free space becomes available. This is normally an automatic process as other audit processes secure Audit Tracks and delete the local copy.

7.2 Schedule Requirements

The audit system is dependent on the Tivoli Workload Scheduler to manage its operation. The scheduler is used to:

- Detect fatal application failures & raise alerts
- Schedule (run and shutdown) individual Audit application processes
- Manage the cross campus replication of Audit Tracks
- Manage the Audit server backups
- During Hydra, manage the failover state of audit gathering processes

The Audit system schedule requirements are fully documented in *Audit Server Configuration (DEV/INF/ION/0001)*

8 Application Development

The Audit data collection system is developed using the following development tools and technologies

- Microsoft Visual Studio 6.0
  - Visual Basic 6.0
  - Visual C++ 6.0
- Microsoft SQL Server 2000

9 System Qualities

9.1 Availability

The Audit service is normally available 24x7 with the exception of a scheduled overnight slot where the following activities take place:

- Database & filestore backup
- Cross campus Audit Track replication

During this slot all normal Audit processes are stopped.

The need to ensure that the Audit Server does not fall significantly behind in the collection of Audit Tracks on a single data centre places further constraints on the design. Following the failure of an Audit Server, on restoration of the server the audit applications must be able to catch up with a five days outage within one day. The Audit Server and links to other components must be able to cope with such a load.

9.2 Resilience and Recovery

9.2.1 Failure Scenarios

9.2.1.1 Campus Level Disaster

Audit Tracks gathered at one data centre are replicated to the remote data centre by a scheduled daily job. Thus up to a days audit tracks could be lost in the event of a complete campus level disaster.

In normal, post Hydra, operation, the IRE11 audit server is responsible for gathering all audit tracks and replicating them to the IRE19 Audit server. If the IRE11 data centre is failed over to IRE19 it will be necessary for the IRE19 Audit server to assume this responsibility. This will be achieved stopping all Audit Schedules and manually switching in an alternative DR configuration file to the IRE19 Audit server. This process is described in *Audit Server Support Guide* (DEV/APP/SPG/0020).

During Hydra, the IRE11 Audit server will gather all HNG-X audit tracks but Horizon audit tracks will be gathered by both Audit servers in the same manner as in Horizon. The IRE11 Audit server will gather Bootle audit tracks and the IRE19 audit server will gather from Wigan.

9.2.1.2 Audit Data Source Platform failure

It is the responsibility of the appropriate generating software and host system to ensure that the Audit Track is maintained in a robust form to survive errors on that system.

The time at which the source file(s) of an Audit Track are deleted is configurable and can be delayed until a (configurable) time after a copy of the Audit Track has been copied onto the Audit Archive. These times are held in configuration files on the audit server and will only be changed on direction from the appropriate design authority.

9.2.1.3 Audit Server Failure

The Audit Data held on an Audit Server is held on resilient SAN storage so that a single disk failure will not result in loss of the data. Audit server filestore is local to each audit server and is not replicated between the IRE11 & IRE19 storage systems.

In the event of a complete failure of an Audit server, the server will need to be rebuilt using the Audit Server platform build & the previous days backup (see section 9.2.2) restored. Reference to any Audit tracks gathered since the backup was taken will be lost. The Audit tracks will still be present on the source platform. When the Audit server is rebuilt it will re-gather the lost Audit tracks. This places a requirement on the Audit configuration that all deletion policy values are set to at least 1 day. One consequence of this approach is that reference to Audit tracks that have been stored in the archive since the last backup will be lost.

If an Audit server is out of action for a period in excess of 5 days, source platforms may run out of disc space to store audit tracks. Manual intervention will then be necessary to temporarily relocate audit tracks until the audit server is recovered.

9.2.1.4 Network Failure

In the event of a local data centre network failure, the Audit Gatherer processes will not be able to gather any new audit tracks that are generated. Recovery will be automatic once the underlying problem is resolved.

Similarly, the Audit Deleter process will be unable to delete source Audit tracks. It will periodically retry the delete and after a preset number of failures will delete all marker files relating to the Audit track. Once the network is restored to operation the Audit track will be seen as new & re-gathered. Thus two copies of any affected Audit tracks will be written to the archive.

Any failure of the inter data centre network link will suspend the Audit track replication process until the problem is resolved

9.2.1.5 Centera Failure

In the event of a failure on the EMC Centera system which holds the Audit Archive, it will not be possible to secure any Audit tracks. The Audit sealer process will stop processing any new Audit tracks until the problem is resolved. The gatherers will continue to gather Audit tracks until the Audit servers working filestore free space falls below a threshold value. The gatherers will stop gathering new Audit tracks at this point. The audit server will generate error events in the Event log to highlight this problem.

If a Centera is out of action for a period in excess of 5 days, source platforms may run out of disc space to store audit tracks. Manual intervention will then be necessary to temporarily relocate audit tracks until the audit server is recovered.

Recovery will be automatic once the Centera problem is resolved.

9.2.2 Data Backup

All non transient data on the Audit Server shall be backed up on a daily basis. This backup must include:

- ARQ Database
- Sealer Database
- Audit server working filestore

A consistent snapshot of the Sealer database and Audit server working filestore must be taken. Cold backups of the SQL Server database are taken; hence it is necessary for the schedules to shut down all Audit server processes while the backup is taken.

Recovery of a complete audit server requires a rebuild and synchronised restore of both its databases and working filestore. Reference to any Audit tracks that have been gathered & written to Centera since the backup was taken will have been lost in the restored databases. Once the restored Audit system is operational again, these missing Audit tracks will be automatically regathered.

The User area on the audit server, which is used to hold working copies of retrieved data, is not backed up. If necessary, this data may be simply re-retrieved from Centera.

9.3 Usability

Human users of the Audit Server fall into two categories:

- Operations/Administrator staff
- End users of the Audit workstations.

The bulk of the operation of the Audit Server is automated and does not require user intervention. Where users are required to interact with the Audit Server such interactions are carried out, from the Audit Workstation, using standard facilities provided by COTS software including Windows Explorer and specifically developed User Interfaces.

Some of the activities require relatively long elapsed times, e.g. recovery of files from Centera. The Windows based facilities provided by the COTS allow other activities to be progressed while waiting for the longer term ones to complete.

Usability of the Audit Track Extraction facilities is addressed in *Audit Data Retrieval High Level Design* (DES/APP/HLD/0029)

9.4 Performance

9.4.1 Volumetrics

Post Hydra, the peak loading of the HNG-X system is expected to generate approximately 10 GB of audit data in one day.

During Hydra, a peak load of 35GB per day will flow from Bootle ➔ IRE11 and Wigan ➔ IRE19. The volume of this traffic will decrease as the HNG-X counter rollout progresses.

The sizing of the Audit Server has to be able to cope with a situation where the server has been out of operation for up to five days. It is a design aim that the Audit Server shall be able to archive those five days worth of data plus the current day's data on the day after the outage.

9.4.2 Impact on other parts of the HNG-X System

The need to collect the very considerable amount of audit data from around the system has impacts on many components outside of the Audit Servers.

The requirements the design places on the network are summarised in section 6.4

Impacts on other components are outside the scope of this design.

9.5 Security

9.5.1 Identification and Authentication

All users (including administrators) of the Audit Workstation and Audit Server shall log onto systems using two factor authentication in conjunction with the HNG-X Active Directory system. Each user shall be uniquely identifiable.

9.5.2 Audit

Details of the main file operations on the Audit Server shall be audited as defined in section 6.1.3.3. In addition the following operating system level events on the Audit Server will be audited via the System Management event monitoring facilities:

- Log on/Log off (including unsuccessful log on attempts)
- File Creation, Deletion and Modification (on selected files)
- Modifications to system configuration (inc software configuration and account details)
- System start up and shut down
- Recovery actions
- Exception conditions
- Change of user rights

9.5.3 Remote Directory Access

The remote directories from which the Audit Server gathers Audit Tracks will be configured so that only the Audit Server (or an administrator who has been explicitly given permission) is able to delete files in the directory.

9.5.4 Physical Access Controls

All Audit Server and Audit Workstation & Centera hardware shall be held in physically secure areas where physical access to the systems is controlled.

9.5.5 Sensitive Data

During data collection, the Audit system is totally transparent to the data being gathered. If a particular Audit track contains any sensitive data that cannot be stored in the audit archive, it is the responsibility of the originating subsystem to obfuscate that sensitive data before the Audit track is presented to the Audit system.

9.6 Roles

There shall be separate roles for:

- Audit Server (inc. Audit Workstation) Administration
- Fujitsu Services Audit Staff

The roles shall be mutually exclusive, i.e. no one individual shall be given access rights of more than one role.

9.6.1 Access Controls

The access control configuration of the Audit Server shall be compatible with the HNG-X Access Control Policy.

The Fujitsu Services Audit Staff role shall not have any write, modify or delete access to the Audit Archive. The ACL configuration on the directory structure of the Audit Server shall ensure that the Fujitsu Services Audit Staff can only recover files into a retrieval/extraction working area; it shall not be possible for the Fujitsu Services Audit Staff to recover files to other parts of the Audit Server directory structure.

The Fujitsu Services Audit Staff will have full control over the files in the retrieval/extraction working area, and in particular shall be responsible for managing the files in that area.

The Audit Server Administrator role shall have full access to manage all of the Audit Server and Audit Workstation file stores and shall be granted the necessary Windows privileges.

9.6.2 Potential for Change

HNG-X Architecture – Support Services (ARC/SVS/ARC/0001) identifies the need for the Audit Server to be able to support the introduction of new applications which generate Audit Tracks of their activities into the HNG-X system. Since the locations from which the Audit Server gathers such files are configurable and a standard interface can be defined, see section 6.2, it is possible to introduce new applications easily. However it should be noted that the implications of such new requirements on the sizing of the infrastructure will need to be assessed.

Other features of the Audit Server, e.g. addition of additional retention periods (other than 18 months or 7 years), or a need to support shorter file retrieval periods may well require changes to the HLD and the implementation.

10 Migration

10.1 Audit Server Platform

The Audit servers will be upgraded to Windows Server 2003 in line with the HNG-X platforms architecture requirements. New BladeFrame based servers will be used in the new data centres

The existing Audit server data collection components must be ported to run under the Windows Server 2003 environment. It is expected that the majority of such applications will run unchanged on Windows Server 2003. This will entail an initial development team exercise to check the compatibility of the existing client applications with Windows Server 2003 & to identify & resolve any issues encountered.

No significant functional changes are required to the data collection components of the Audit server for HNG-X.

A new Audit server platform build will need to be defined for the HNG-X environment in *HNG-X Audit Server (ARC) - Physical Platform Design (DES/PPS/PPD/0035)*.

Migration scripts will be developed to export and import the audit server's databases & files from the old to the new audit servers. Prior to migration all schedules should be stopped and the export script executed on each Audit servers to create a zip file of the databases and filestore. The Bootle zip file will be transferred to IRE11 and Wigan to IRE19. The zip files will be transferred over the network using secure copy (scp). The import script will then be executed in the Irish data centres. Database schema updates will be applied as part of the import script.

Given that platforms that generate Audit data are only required to hold up to an additional 5 days worth of Audit data, the audit server migration must complete within 5 days. If unforeseen problems occur and this time period may be exceeded a decision must be to regress back to the Bootle/Wigan Audit servers.

10.2 Centera Migration

The Audit archive Centera systems will need to be physically shipped to the new data centres. This introduces a risk of them being lost or damaged in transit. To mitigate against this risk two swing Centeras will be used. Archive data will be copied onto the swing Centeras prior to shipping of the original Centeras.

It is estimated that the copying process, which will be implemented by EMC, could take up to 3 weeks. Thus it will need to be initiated well in advance of the Audit migration phase. The outline approach planned by EMC is as follows

1. Install swing Centeras in Bootle & Wigan
2. Bulk copy the data on the original Centeras to the swing kit
3. Once the bulk copy is completed, the swing kit is kept synchronized using Centera native replication
4. Separately ship the original Bootle/Wigan Centeras to IRE11/IRE19
5. Install (at IRE11/19) & verify the integrity of the original Centeras
6. Once the Centeras in Ireland are declared to be fully operational, the data on the swing kit in Bootle/Wigan can be securely erased & the kit decommissioned

The above sequence of operations will be applied to each Centera in parallel.

Data held on the Centera systems is not encrypted and could potentially be accessed if the systems were physically intercepted whilst being moved. A suitably secure means of transport must be arranged.

10.3 Audit Configuration

Migration of systems that generate Audit data and the consequential, incremental changes to the audit configuration will need to be carefully coordinated during the migrations of systems to the new data centres and to the HNG-X environment.

As old systems go off-line, their corresponding audit configuration entries will need to be disabled, and as new systems appear they will need to be configured in line with the HNG-X Audit and Security Policies.

Audit Server Configuration (DEV/INF/ION/0001) will describe the target configuration. Changes to the configuration through the migration phases will be described in *Audit Server Migration Configurations* (DEV/INF/ION/0008)

Figure 10 illustrates the sequence of configuration changes that will be applied.

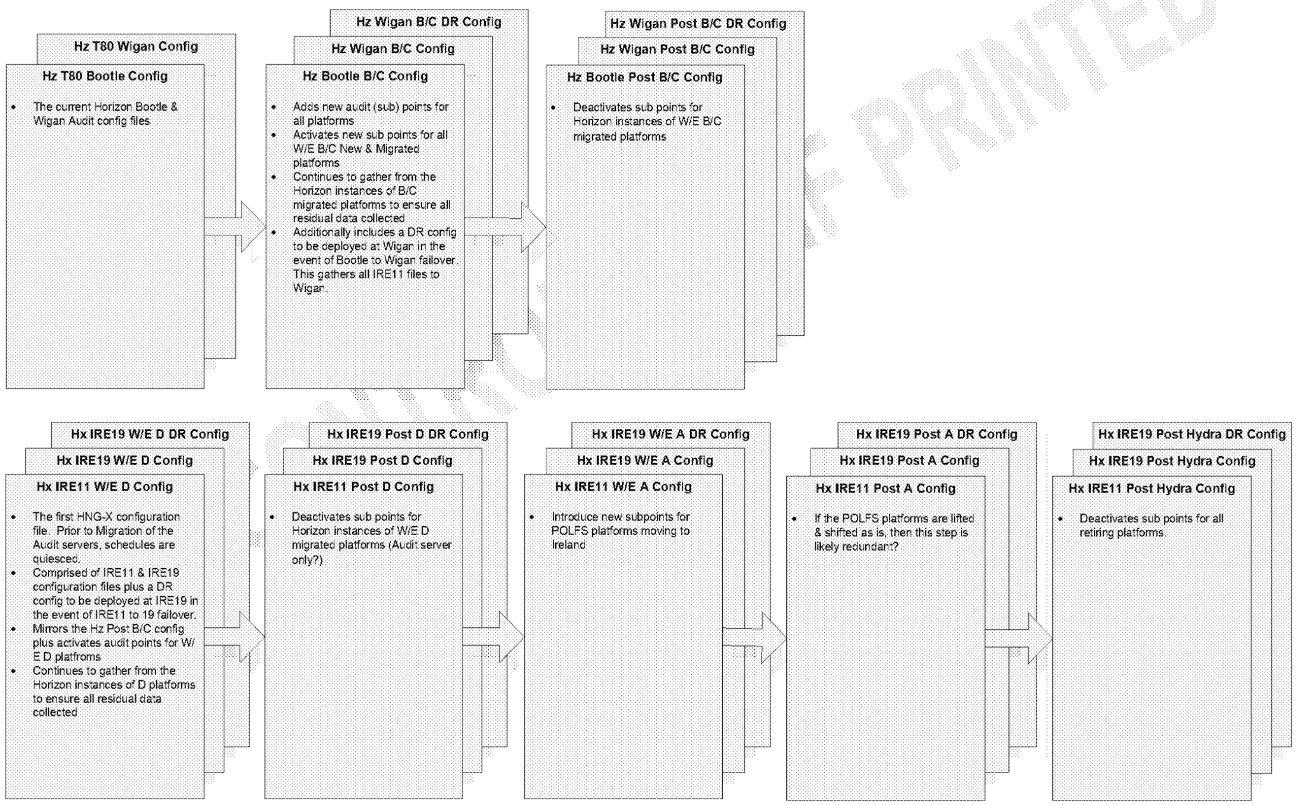

- The current Horizon Bootle & Wigan Audit config files

- Adds new audit (sub) points for all platforms
- Activates new sub points for all W/E B/C New & Migrated platforms
- Continues to gather from the Horizon instances of B/C migrated platforms to ensure all residual data collected
- Additionally includes a DR config to be deployed at Wigan in the event of Bootle to Wigan failover. This gathers all IRE11 files to Wigan.

- Deactivates sub points for Horizon instances of W/E B/C migrated platforms

- The first HNG-X configuration file. Prior to Migration of the Audit servers, schedules are quiesced.
- Comprised of IRE11 & IRE19 configuration files plus a DR config to be deployed at IRE19 in the event of IRE11 to 19 failover.
- Mirrors the Hz Post B/C config plus activates audit points for W/E D platfroms
- Continues to gather from the Horizon instances of D platforms to ensure all residual data collected

- Deactivates sub points for Horizon instances of W/E D migrated platforms (Audit server only?)

- Introduce new subpoints for POLFS platforms moving to Ireland

- If the POLFS platforms are lifted & shifted as is, then this step is likely redundant?

Hx IRE19 Post Hydra DR Config
Hx IRE19 Post Hydra Config
Hx IRE11 Post Hydra Config

- Deactivates sub points for all retiring platforms.

Figure 10 – Audit Configuration During Migration

In normal, post Hydra, operation, the IRE11 audit server is responsible for gathering all audit tracks and replicating them to the IRE19 Audit server. If the IRE11 data centre is failed over to IRE19 it will be necessary for the IRE19 Audit server to assume this responsibility. This will be achieved by manually switching in an alternative DR configuration file to the IRE19 Audit server. Prior to the Audit server migration at weekend D, this DR configuration will be applied to the Wigan Audit server.

10.4 Migration Phases

A number of alternate Audit configuration files will be delivered against various migration phases. These are described in detail in *Audit Server Migration Configurations* (DEV/INF/ION/0008). These changes are summarized in the following subsections.

10.4.1 AP Clients Migrated to EDG

The Horizon Audit server will be reconfigured to disable auditing of the APS local FTMS gateway. All APS audit data will be gathered from the EDG local gateway.

10.4.2 HNG-X Migration Enabling Upgrades for Data Centres

This phase has no impact upon the Audit system.

10.4.3 Data Centre Build

This activity installs and commissions new hardware and software in the Data Centres at IRE11 and IRE19, initially for use in the ST, SV&I, RV, V&I and LST test environments. The installation is incremental as test environments are established. The Audit components are installed using production settings and identities.

10.4.4 Move Wigan Network Management Servers

This phase has no impact upon the Audit system.

10.4.5 Data Centre Preparation

This phase involves preparing the environment for live use following Volume & Integrity testing.

The Audit server will need to be cleared down back to its initial state, clearing down all data and other changes introduced during V&I testing.

The live Centera systems are not used during V&I testing, instead a dedicated test Centera is used. Thus no clear down will be required for the live Centeras.

V&I testing will use dedicated instances of the Audit workstation platforms, so no clear down of the live Audit workstations will be required. The Live audit workstations will require access to live PAN decryption keys & the live PAN Hash salt value.

HNG-X Audit schedules must be inactive until the Migration of Audit Services phase.

10.4.6 Cutover Rehearsal

Given the time required to secure the Centera data onto the swing kit, it will not be viable to rehearse this activity during this phase.

10.4.7 Migration of POL FS

The HNG-X Audit configuration will need to be updated to reflect migration of the POLFS systems which generate Audit tracks.

10.4.8 Migration of Batch Services

The Horizon Audit system will need to be reconfigured to reflect migration of the systems which generate Audit tracks.

10.4.9 HNG-X Specific Services

This phase has no impact upon the Audit system.

10.4.10 Migration of Online Services

The Horizon Audit system will need to be reconfigured to reflect migration of the systems which generate Audit tracks.

10.4.11 Migration of Audit Services

The migrated Audit system (Audit server & workstations) becomes operational at this phase.

Prior to the commencement of this phase, the data on the Audit Centera systems must have been copied onto swing kit in the Bootle & Wigan data centres - §10.2. This activity will take an estimated 3 weeks to complete.

The Audit Centera systems are separately shipped to Ireland. Cold backups of the existing Audit server databases & working file systems are taken and shipped to Ireland to be installed onto the new Audit servers. The Audit server database schema will be upgraded to that required by the updated applications - §10.1.

The HNG-X schedules, used to control the Audit server, will be activated.

Required DNS aliases will be established.

10.4.12 Migration of Branch Services

As a result of HNG-X CP0284 Horizon Branch Services will remain in Bootle/Wigan. The Audit servers in Ireland will gather audit data from the Bootle/Wigan platform instances until they are retired - §6.4.

10.4.13 Move Bootle Network Management Servers

This phase has no impact upon the Audit system.

10.4.14 Decommission Wigan and Bootle

This phase has no impact upon the Audit system.

10.4.15 Horizon Counter Changes for PCI Compliance

This phase has no direct impact upon the Audit system. Any retrieved Audit tracks generated by Horizon PCI counters may need to be processed differently by the end user of the Audit workstation.

10.4.16 HNG-X Migration Enabling Upgrades for Counters

This phase has no impact upon the Audit system.

10.4.17 HNG-X Application Pilot and Rollout

This phase has no impact upon the Audit system.

10.4.18 Branch Router Rollout

This phase has no impact upon the Audit system.

10.4.19 Counter Event Management Changes

This phase has no impact upon the Audit system.

10.4.20 Counter XP Upgrade

This phase has no impact upon the Audit system.

10.4.21 Post-Application ADSL Changes

This phase has no impact upon the Audit system.

10.4.22 Final Decommissioning

Once all SYSMANn2 & the Tivoli cluster lookup service is retired, Horizon Branch to Correspondence Server cluster mapping data can be imported into the Audit server database.

10.4.23 Estate Management Upgrade

The HNG-X Audit configuration will need to be amended to gather any new Audit tracks generated by the Estate Management system.