FUJ00098211

Document Title: Audit Data Storage & Retrieval High Level Design Specification

Document Type: Design

Abstract: The components required to implement the audit data storage and retrieval facilities are identified and their individual roles are specified. Interfaces between the components and with other parts of the Horizon system are identified.

Status: Approved

Distribution:

Steve Doyle
Richard Long
Peter Sewell
Stephan Robson
Jan Holmes
Richard Laking
Bryan Muir
Pathway Library
Martyn Bennett
Jack Kirwan
Simon Fawkes
Gareth Jenkins
Andy Scott

Author: Frank Womack / Michele Myles

Comments to:

Comments by:

FUJ00098211

0 DOCUMENT CONTROL

0.1 DOCUMENT HISTORY

Version	Date	Reason
0.1	01/5/98	First draft -Very limited internal distribution
0.2	5/5/98	Second Draft - for Internal Review
0.3	27/5/98	Third Draft - for development information only
0.4	3/6/98	Draft for Approval
1.0	25/6/98	Version for approval
1.1	12/02/99	Updates to HLD in line with NR2 changes since approval – CP1326.
2.0	25/02/99	Version for approval, incorporating formal review comments

0.2 APPROVAL AUTHORITIES

Name	Position	Signature	Date
T. Austin	Systems Director		
R. Long	Design Manager		
C. Humphries	Development Manager		
G. Jackson	Test & Integration Manager		
A. Ward	Chief Architect		
J. Dicks	Customer Requirements Director		
S. Muchow	Customer Services Director		
M. Bennett	Quality and Risk Director		
D. Groom	Quality Manager		

FUJ00098211

0.3 ASSOCIATED DOCUMENTS

	Reference	Vers	Date	Title	Source
[1]	TD/ARC/001	4.3	17/12/98	Technical Environment Description	Pathway
[2]		Issue 2	18/3/98	Briefing Note - Audit	Pathway
[3]	CR/FSP/006	2.4	5/2/99	Audit Trail Functional Specification	Pathway
[4]	IA/REQ/001	0.7	14/12/98	Pathway Internal Audit Requirements	Pathway
[5]	IA/REQ/002	0.6	04/2/99	Audit Data Retrieval Requirements	Pathway
[6]	SD/PRP/001	1.0	21/8/97	Proposal for Auditing at Release 1C	Pathway
[7]	TD/DES/037	2.0	13/3/98	Release 1c to Release 2 Migration Strategy	Pathway
[8]	TD/DES/052	0.1	27/3/98	Release 1c to Release 2 Migration Detail of Implementation	Pathway
[9]	TD/DES/0023	4.0	14/9/98	Long Term Archiving of PAS/CMS Data	Pathway
[10]	TD/STD/001	2.0	9/11/98	Host Applications Database Design and Interface Standards	Pathway
[11]	RFC 1321		April 1992	RSA Data Security Inc. MD5 Message-Digest Algorithm	
[12]	Audit.doc	1.0	09/03/97	Systems Management Auditing for Pathway Release 2	Pathway
[13]	TD/DES/053	2.0	9/11/98	Long Term Archiving of OBCS Data	Pathway
[14]	SD/DES/064	1.0	24/08/98	Audit Data Gap Analysis	Pathway
[15]	TD/DES/027	1.2	21/09/98	File Transfer Managed Service	Pathway
[16]	TD/ION/003 – TD/ION/010			FTMS Configurations at Release 2	Pathway
[17]	SD/DES/074	1.0	03/12/98	Audit Data Filtering and Extraction HLD	Pathway
[18]	IA/SPE/008	0.1	27/11/98	Audit Data Catalogue	Pathway
[19]	RS/FSP/001	3.3	11/12/98	Security Functional Specification	Pathway
[20]	RS/SPE/005	0.2	26/5/98	Security Event Management Requirements	Pathway
[21]	RS/DES/004	0.4	16/04/98	Use of SecurID Token Authentication for Release 2	Pathway
[22]	RS/POL/003	3.0	18/12/98	Access Control Policy	Pathway

FUJ00098211

[23]	TD/DES/092	0.2	19/01/99	Audit Server Resilience & Recovery for Release 2	Pathway
[24]	TD/ARC/012	3.1	24/09/98	Technical Environment Implementation for Release 2	Pathway

0.4 ABBREVIATIONS

AS	Audit Server
AS(B)	Audit Server (Bootle)
AS(W)	Audit Server (Wigan)
ATE	Audit Track Extractor
ATD	Audit Track Deleter
ATG	Audit Track Gatherer
ATH	Audit Track Hoarder
ATR	Audit Track Retriever
ATS	Audit Track Sealer
ATSDB	Audit Track Seal Database
CD-W	CD-Writable(in the instance of the Audit Server/Workstation Write Once)
COTS	Commercial Off The Shelf
CS	Correspondence Server
CSAH	Correspondence Server Audit Harvester
DB	Database
DLT	Digital Linear Tape
DSS	Department of Social Security
FTMS	File Transfer Managed Service
GB	GigaByte
HLD	High Level Design
Mb	Megabit
MB	MegaByte
MIS	Management Information System
NFS	Network File System
OBCS	Order Book Control System

FUJ00098211

PAS/CMS	Payment Authorisation System/Card Management System
POCL	Post Office Counters Limited
RED	Reconciliation Acceptance Database
SLAM	Service Level Agreement Monitoring
TME	Tivoli Management Environment
TMS	Transaction Management Service
TOD	Tivoli Oracle Database

0.5 DEFINITIONS

The Audit Archive is the sum of audit data written to secondary storage (DLT tapes) by the Audit Servers which is available for subsequent Audit Trail Extraction and other recovery/extraction purposes. The Audit Archive is generated in two parts one on the Bootle campus and one on the Wigan campus. Both parts contain a copy of all audited information.

Audit Server is the system which is responsible for the gathering , archiving, retrieving and potential extraction for subsequent analysis of all audit information that is required to be retained by Pathway beyond normal operational use.

Audit Data Storage covers Audit Track Gathering, and Audit Track Hoarding.

Audit Data Restoration covers Audit Track Retrieval and Audit Trail Extraction.

An Audit Point identifies a logical position in the Horizon system at which an Audit Track is generated. In reality an Audit Point is distributed across a number of locations in the system each such location is identified as an Audit Sub-Point

Audit Sub-Point - see definition of Audit Point

An Audit Track is a sequential record of activities made by a particular subsystem.

Audit Track Deletion covers the deletion of Audit Tracks once they have been gathered. The Audit Server is responsible for the deletion of all Audit Tracks outside of the Audit Archive. Audit Track Deletion does not cover the removal of time expired Audit Tracks from the Audit Archive.

Audit Track Gathering covers the transfer of Audit Tracks to the Audit Server prior to Audit Trail Hoarding.

Audit Track Generation is the production of Audit Tracks by application software in formats suitable for Audit Data Storage.

Audit Track Hoarding covers the writing of Audit Tracks to offline storage in a format suitable for retrieval.

Audit Track Retrieval covers the reading of Audit Tracks from secondary storage (after having been written there by the Audit Track Hoarding).

Audit Track Sealing is the generation and independent storage of a checksum for each file before it is subject to Audit Track Hoarding and after Audit Track Retrieval.

FUJ00098211

Comparison of the before and after seals can demonstrate, to a very high degree of probability, that the Audit Tracks have not been tampered with.
An Audit Trail is one or more Audit Tracks which between them enable an auditor to follow the treatment of related data transfers, movements or accesses by named individuals.
Audit Trail Extraction includes the extraction of Audit Track information for presentation to and subsequent use by customer auditors and by Pathway internal auditors.

0.6 CHANGE HISTORY

- **VERSION 1.0 to Version 1.1**
  - Changes Forecast replaced by Change History
  - Requirements: Addition of Footnotes for Reqs 4 & 5
  - Assumptions: Removal of Assumptions 12 & 13
  - Rephrasing of some sentences in Audit Server Overview to reflect implemented functionality of Audit modules
  - Archive Server replaced by Audit Server
  - Table 5.1 and associated text removed
  - General textamendments to Section 6.1 to bring the functionality in line with implementation, including
    - I-ATG-4 Large amount of text re FTMS removed as covered in FTMS documentation
    - References to Appendix C & F replaced by ref to Audit Data Catalogue
    - I-ATG-5, file naming text removed
  - Section 6.2: Text from Appendix D, inserted here
  - Section 6.5: Text from Appendix A inserted here
  - Section 9.1.1: Now just a list of failure modes and a reference to Audit Server Backup & Recovery HLD, to avoid holding this data in 2 documents.
  - Appendices
    - A – text now in appropriate section, 6.5
    - B – Renamed to Appendix A
    - C – Removed
    - D – text now in appropriate section,6.2
    - E – Renamed to Appendix B
    - F – Removed
    - G - Renamed to Appendix C

  CP1712 – New Appendix B updated with revised Audit Server sizing details
  CP1782 –Amendments to Requirements, Assumptions & Section 9

Version 1.1 – 2.0

- Incorporation of Review Comment

FUJ00098211

0.7 TABLE OF CONTENT

1. INTRODUCTION 10

2. SCOPE 10

3. DESIGN PRINCIPLES 11

4. REQUIREMENTS 12

4.1 Source of Requirements 12

4.2 Additional Requirements 13

4.3 Assumptions 13

5. SYSTEM OVERVIEW 14

5.1 Major Components 14

5.2 Audit Server Overview 15

5.3 Audit Workstation 20

5.4 Audit Archive 21

5.5 Audit Points 22

6. SYSTEM COMPONENTS 22

6.1 Application Components 23
6.1.1 Audit Track GatherER 23
6.1.2 Audit Track Sealer 27
6.1.3 Audit Track Deleter 31
6.1.4 Audit Track Hoarder 33
6.1.5 Audit Track Retriever 36
6.1.6 Audit Trail Extractor 37

6.2 Interfaces 37

6.3 Distributed Application Services 38

6.4 Information Management 38

6.5 Networking Services 38

6.6 Platforms 38

FUJ00098211

7. SYSTEMS MANAGEMENT 39

8. APPLICATION DEVELOPMENT 39

9. SYSTEM QUALITIES 39

9.1 Availability 39
9.1.1 Failure Modes 39
9.1.2 Availability of Data ON TApe 40

9.2 Usability 40

9.3 Performance 41
9.3.1 Volumetrics 41
9.3.2 Correspondence Server Harvesting 41
9.3.3 MD5 Algorithm Performance 41
9.3.4 Impact on other parts of the Horizon System 42

9.4 Security 42
9.4.1 Identification and Authentication 42
9.4.2 Audit 42
9.4.3 Domain Structure 43
9.4.4 Remote Directory Access 43
9.4.5 Physical Access Controls 43
9.4.6 Roles 43
9.4.7 Access Controls 43
9.4.8 Tivoli Tasks 44
9.4.9 Data Backup 44

9.5 Potential for Change 44

10. MIGRATION 44

10.1 Release 1c Audit Facilities 44

10.2 Overall Migration Strategy 45

10.3 Audit Migration Strategy 45

11. SOLUTION IMPLEMENTATION STRATEGY 47

12. COSTS, RISKS AND TIMESCALES 47

APPENDIX A POSSIBLE FUTURE CHANGES 48

APPENDIX B HARDWARE SIZING AND CONFIGURATION 49

APPENDIX C INTERFACES 52

FUJ00098211

1. INTRODUCTION

Within the Horizon system, ICL Pathway is 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.

The architecture for the audit sub-system within the Horizon system is described in [1]. This Audit Data Storage & Retrieval High Level Design Specification is consistent with that architecture.

This High Level Design (HLD) specifies the concrete components required to be Integrated to provide the Audit Data Storage and Retrieval 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 mechanisms for the generation of the Audit Tracks are not within the scope of the Audit Data Storage and Retrieval facilities. The Audit Track Generation is the responsibility of the different applications within the Horizon system. The Audit Data Storage and Retrieval facilities are responsible for gathering those tracks and the subsequent storage, retrieval and extraction of relevant audit data for subsequent analysis.

The facilities for preparing selected Audit Track data for presentation to the POCL and DSS auditors together with those facilities necessary to provide the basic capabilities required by internal Pathway auditors are specified in [17].

The Audit Data Storage and Retrieval facilities have been designed to be generic and extensible, in particular any new applications introduced into the Horizon system should interface to the Audit Server as specified in Section 6.2.

This HLD is focused on New Release 2 of the Horizon system. Where possible the components have been designed to cope with full steady state workloads in order to minimise nugatory work. Where changes to the design are felt likely to be necessary beyond New Release 2 they are identified in Appendix A.

2. SCOPE

This High Level Design Specification covers:

- the audit storage and retrieval facilities for New Release 2 of the Horizon System as outlined in [2]
- the interfaces between the applications creating the Audit Tracks and the Audit Data Storage facilities
- the mechanisms for the (off-line) storage of the Audit Archive
- the structure of the data stored in the Audit Archive
- the mechanisms for retrieving data from within the Audit Archive
- the interface with the Release 1c audit trail

FUJ00098211

- the strategy for migrating from the Release 1c audit facilities to the New Release 2 facilities.

As indicated in [2] only basic retrieval capabilities will be provided at New Release 2, and hence are specified in this document. In line with [2] this HLD does not provide specific facilities to support the auditing of data from the MIS applications, e.g. SLAM and RED. However the standard interface defined in Section 6.2 is able to support the archiving of such data. The requirements [4] specify that the retention periods for some of these components is 7 years.
The components and products used to provide the facilities described herein are identified.
The scope of this HLD does not cover:

- Audit Track Generation
- Auditor access to the current data on the production system (including provision of facilities currently provided by the Extract Workstation)
- Auditing of any information beyond the Horizon external gateways (except for CAS (VME) information audited by PAS/CMS)
- The Audit Track Extraction tools for use by internal Pathway auditors

Although the design of the Audit Track Extraction tools is not included in this HLD the interface between those facilities and the Audit Data Storage and Retrieval facilities is specified as part of the Audit Track Retriever in section 0.
This HLD is not specific about the exact instances of the interfaces between the Audit Server and the rest of the Horizon system. The HLD does define the types of the interfaces and the anticipated numbers of the interfaces. It is intended that configuration of the Audit Data Storage and Retrieval facilities will cope with the changes in the number of interfaces. In this way the HLD will not have to be changed as the number of interfaces changes. Only when new types of interfaces are introduced will the changes need to be reflected in the HLD. [18] provides details of the actual interfaces configured for each instance of the Audit Server.
This design does not take into account the use of the Audit Server to provide a general backup/archive and recovery service for non audit purposes. The use of the Audit Server for back-up & recovery of the Correspondence Server and the impacts on the Auditing capacity of the Audit Server are detailed in [23].

3. DESIGN PRINCIPLES

The main principle of this design is to provide the required audit data storage and retrieval facilities while minimising the impact on the New Release 2 development activities and timescale. Consequently this HLD includes design features to interface to and support existing system features. It is expected that any new applications should interface to the Archive Server in the manner specified in Section 6.2 Interfaces.

FUJ00098211

Due to the very large amounts of audit data which have to be collected each day data (estimated maximum is estimated to be in excess of 50GB) there is a major design aim to ensure that each item of audit data which is collected only transits the campus LAN at most once. It is also intended that audit traffic should not traverse the intercampus links.
The Audit Architecture as defined in the TED, ref [1], identifies the need to be able to cope with change as the usage of the Horizon system develops, especially as new applications and services are introduced. Thus a significant design principle is for the Audit Data Storage and Retrieval system to be able to support the introduction of such new facilities with minimum impact.
The Horizon system as a whole is designed to be resilient to a single point of failure. The same philosophy applies to the Audit Data Storage and Retrieval design.
New Release 2 is currently expected to be used to support approximately 6000 outlets. In order to minimise nugatory work it is an aim of the New Release 2 Audit Data Storage and Retrieval design to be able to support the full steady state load (following New Release 2+) from 20,000 outlets except where specifically identified in Appendix A.

4. REQUIREMENTS

4.1 SOURCE OF REQUIREMENTS

The following documents include requirements on the Audit Data Storage and Retrieval design:

- Audit Trail Functional Specification, [3]
- Pathway Internal Audit Requirements, [4]
- Audit Data Retrieval Requirements, [5]
- Audit Architecture, [1]
- Security Functional Specification, [19]
- Security Event Management, [20]
- Access Control Policy, [21]

This design does not address all of the requirements raised in those documents given that they address a wider scope than audit data storage and retrieval. In particular the following areas are not covered:

- Specific support of POCL Emergency Manager/Auditor role
- Access by auditors to data on the operational system (i.e. prior to archiving)
- Auditing of components beyond the external gateways except for CAS(VME) data audited by PAS/CMS.

FUJ00098211

4.2 ADDITIONAL REQUIREMENTS

This section records additional and qualified requirements which have been identified during the production of this design and which were not recorded in the sources of requirements identified in section 4.1 above.

1. The audit data storage and retrieval design shall not use the same compression algorithm (or one with similar characteristics) as the Correspondence Server.
2. It shall be possible for retrieval of audit data collected during the operation of Release 1c to be carried out for at least 18 months after collection of the data. This implies that the audit facilities for Release 1c data must be available in parallel with the New Release 2 facilities.
3. It shall also be possible for auditors to carry out audits of activities spanning the swap from Release 1c to New Release 2. It will be acceptable for two independent systems to be used for Release 1c and New Release 2 provide the semantics of the relevant Audit Trails can be matched.
4. The Audit Archive is only to be used for audit purposes. The Audit Archive (and the Audit Server) are not designed to support recovery of data for general purposes.[1]
5. Following the failure of an Audit Server on restoration of the server it shall be able to catch up with a two day outage within one day. This means that the Audit Server and the links to the other components of the system for Audit Trail Gathering shall be able to cope with such a load. This requirement is dependent on requirement 4 above, refer associated footnote.

4.3 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 campus is deemed not to have been rectified until it has caught up with the gathering of the Audit Tracks from the Correspondence Servers and other sources of Audit Tracks on that campus. A failure of the second Audit Server on the other campus during this period is deemed to be a second point of failure.
2. The consequences of any redesign activities that are carried out for New Release 2+ which might impact the Audit Data Storage and Retrieval function will be considered in a New Release 2+ reassessment of this HLD.
3. The migration activities from Release 1c to New Release 2 do not require any more auditing than the operation of Release 1c prior to the migration.

[1] While this may have been a requirement for optimum operation of the Audit Server, at NR2 the Audit Server is used to back up the Correspondence Server, [23] refers.

FUJ00098211

4. The facilities provided for auditor access will be adequate for any other access needs. e.g. in support of the investigation of processing failures. Also see point 4 of Section 4.2
5. The maximum time to fix an Audit Server following an outage for what ever reason (while the rest of the campus operates correctly) will be two days. Also see point 5 of Section 4.2. All points within the Horizon system which generate Audit Tracks have enough disk space to hold at least 3 days of Audit Tracks.
6. It is acceptable to duplicate Audit Tracks in the Audit Archive and duplicate records within Audit Tracks.
7. The Tivoli Oracle Database server operates on only one campus at a time and its audit tracks are duplicated by tape copying at the Audit Server.
8. Loss of the previous day's non-Correspondence Server Audit Tracks should the unlikely event of a campus level disaster occur is acceptable.
9. The Riposte message numbers will remain unique across the R1C/R2 boundary.
10. The basic design of the Horizon system means that the Post Offices can operate even if their communications to the central sites are lost. Under such circumstances it is possible that the messages relating to a day's operations will not be audited on that day. Hence to recover the audit of a particular day's operations at a particular Post Office it may be necessary to search a number of days worth of Audit Tracks. Note that it is possible as a result of certain extreme failure modes, e.g. destruction by fire of a Post Office following an outage of that Post Office's ISDN link, for there to be missing data in the Audit Trail.
11. It is assumed that all sources of the Audit Track data will use synchronised clocks.

5. SYSTEM OVERVIEW

The Audit Architecture, [1] describes the overall audit system design for Pathway. This section provides an overview of the design of the parts of the Audit System supporting the Audit Data Storage and Retrieval functions. 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 5.1 shows the major components of Audit Data Storage and Retrieval on a single Campus. The configuration is duplicated on both the Wigan and Bootle sites.

FUJ00098211

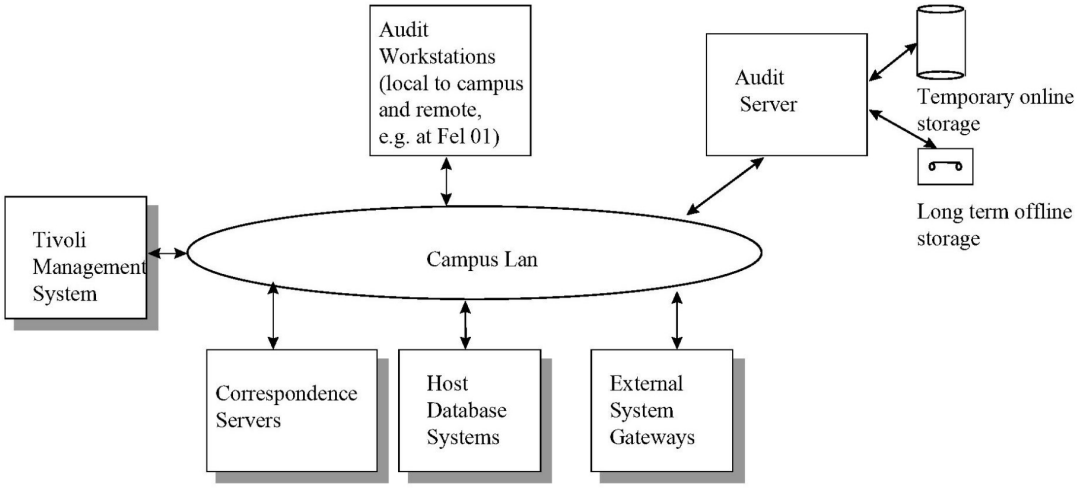

Figure 5.1
Main Components of Audit Data Storage and Retrieval

The Archive server is responsible for gathering Audit Tracks generated from a wide range of components of the Horizon system including:

- Correspondence Servers
- Tivoli Management Facilities
- Database Hosts (including the Reference Data System)
- Gateways to External Systems

The Correspondence Servers within the Horizon system have been designed to duplicate their message store data across the two sites, the way the Audit Servers on each campus operate means that a duplicate of the Correspondence Server Audit Track is automatically produced and gathered on each site. Other Audit Tracks (including the External Gateway Audit Tracks) are not automatically duplicated on each campus, duplicates of tapes containing such tracks are produced by each Audit Server and are exchanged between the two sites.

As well as gathering and storing, on offline mass storage tapes, 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 by the POCL and DSS auditors are to be provided together with basic facilities to support internal Pathway audit activities, as specified in [17]. Access by Pathway internal audit staff to the retrieval and extraction facilities is via the user interface provided on the Audit Workstation.

5.2 AUDIT SERVER OVERVIEW

Figure 5.2 shows the major logical components of the Audit Server.

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

FUJ00098211

- Audit Track Gatherer

  Collects Audit Tracks that have been generated within the Horizon 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 NTFS for Correspondence Server, Tivoli Object Database and External Gateway Audit Tracks. NFS is used to collect files from Unix systems in particular the database applications, e.g. PAS/CMS. 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 Maestro scheduling facilities of the Horizon system.

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

FUJ00098211

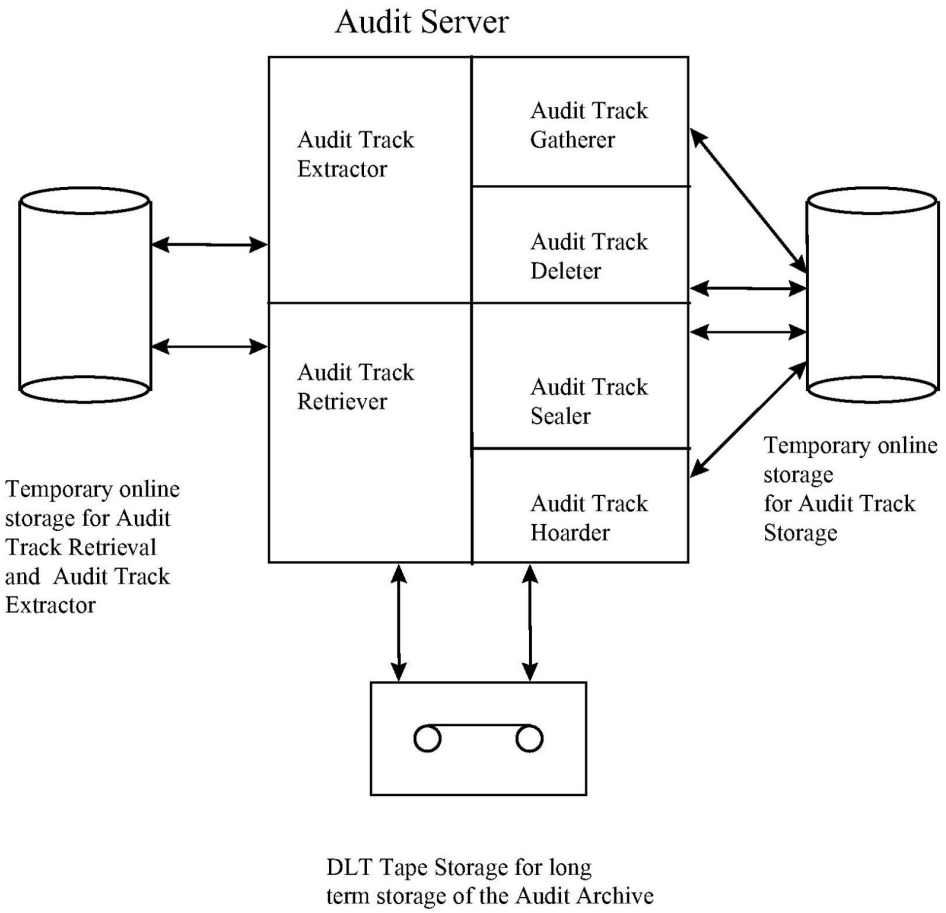

Figure 5.2
Audit Server Logical Structure

- Audit Track Sealer

Before Audit Tracks are hoarded a seal is calculated for the file. The seal is stored on the Audit Server in a database which links the seal to the file.

When an Audit Track is retrieved its seal is recalculated and checked against the value in the database.

- Audit Track Hoarder

Transfers Audit Tracks from the Disk Storage on the Audit Server onto long term storage media (DLT tapes). This component is implemented using the Legato NetWorker product.

FUJ00098211

- 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 Hoarding 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

- Audit Track Retriever

  Selectively restores Audit Tracks written to long term storage media onto temporary disk storage on the Audit Server prior to the extraction of relevant data for use by the auditors. A major portion of the functionality of this component is implemented using the Legato NetWorker product.

- Audit Trail Extractor

  Provides facilities for the Pathway Internal Auditors to provide extracts of the audit data from the Audit Archive. The design of the extraction facilities is addressed in [17].

The distributed nature of the Horizon 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 (see section 0 for details).
During the operation of New Release 2 it will be necessary for the auditors to be able to audit activities that took place during Release 1c operation and during the period of migration from Release 1c to New Release 2. Section 10 of this design specifies how that will be done.
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 5.3 below provides a view of the main flows. For the sake of clarity figure 5.3 omits the Audit Track Deleter and its associated data flows.

FUJ00098211

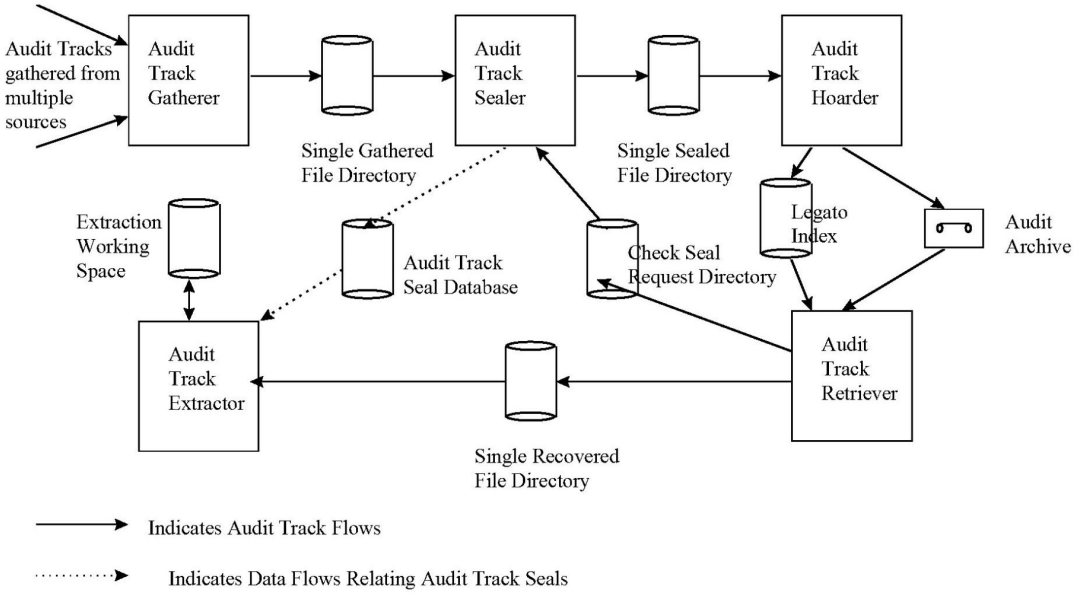

Figure 5.3
Major Data Flows between the Audit Server Components

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.
Where possible file movement should be 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 what ever reason will allow the other components to continue processing any input data in their persistent input buffers.
The Audit Server employs the data compression facilities integral to the DLT tape devices, no other compression is used in its design. However it should be noted that some of the files gathered by the Audit Server may already have been compressed prior to their gathering.

The Audit Track Retriever is the only point prior to extraction where any of the Audit Track files need to be duplicated. This approach has been adopted to allow the checking of recovered files seals to be calculated asynchronously to the selection and extraction of relevant data by the Pathway Internal Audit staff. On the grounds that there are likely to be long delays in the recovering of files from tape and the checking of seals synchronously will further significantly add to the time before the file is available for extraction of the required data. Allowing data from a file to be extracted

FUJ00098211

before its seal has been checked does not invalidate the results of the extraction as long as the validity of the seal is confirmed before the results of the extraction are used, e.g. Audit Trail information passed to POCL/DSS auditors.
The design of the Audit Server does not include any synchronisation between the Audit Track Sealer and the Audit Track Extracter
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.

5.3 AUDIT WORKSTATION

The Audit Workstation provides facilities for Pathway audit staff to access the Audit Server in order to retrieve Audit Track data from the Audit Archive and to either select and prepare Audit Track data for presentation to the POCL and DSS auditors or carry out basic online extraction of Audit Track data in support of internal audit activities.
The access to the facilities to request retrieval of the Audit Track files from the Audit Archive is provided by the Legato NetWorker User Client. This client provides facilities for the audit staff to browse the Legato Index and to select files to be recovered from the Audit Archive tapes. The client requests the Legato NetWorker Server software on the Audit Server to retrieve the files, the server in turn requests the operations staff to load the appropriate tape(s) from which the files are then retrieved.
The recovery of files from tape may take a noticeable length of time, e.g. over an hour, dependent upon the amount of data to be recovered, the number of tapes to be recovered and loaded and the other activities taking place on the Audit Server.
The files retrieved from the Audit Archive are stored on the Audit Server. Retrieved files are not normally transferred to the Audit Workstation, normally selected data is recovered from the files using the extraction facilities before being transferred to the Audit Workstation.
In addition to providing access to the retrieval facilities the Audit Workstation provides access to the extraction facilities. The design of the extraction tools is specified in [17]. The Audit Architecture, [1], indicates that any parts of the extraction tools which run on the Audit Workstation should be thin clients.
The Audit Workstation supports a Write-Once CD to which selected Audit Track data can be exported and passed to POCL and DSS auditors.
Seal Database backups to CD-W, are via the Audit Workstation.

5.4 AUDIT ARCHIVE

A copy of the Audit Archive is kept on both the Bootle and Wigan sites, each one is held on DLT tapes. Audit Tracks generated on each campus are added to the local copy of the archive. The locally generated Audit Tracks include a full copy of the TMS Audit Tracks generated on the local instances of the Correspondence Server machines and all other non-TMS Audit Tracks generated on the campus.

FUJ00098211

The Audit Server stores Audit Tracks generated on the local campus onto DLT tapes which then become part of the Audit Archive. In the case of non-TMS Audit Tracks a copy of all the stored Audit Tracks is written to another tape which is then transported to the other campus. When the other tape is received at the destination campus it is added to the copy of the Audit Archive held there. Figure 5.4 shows this structure.

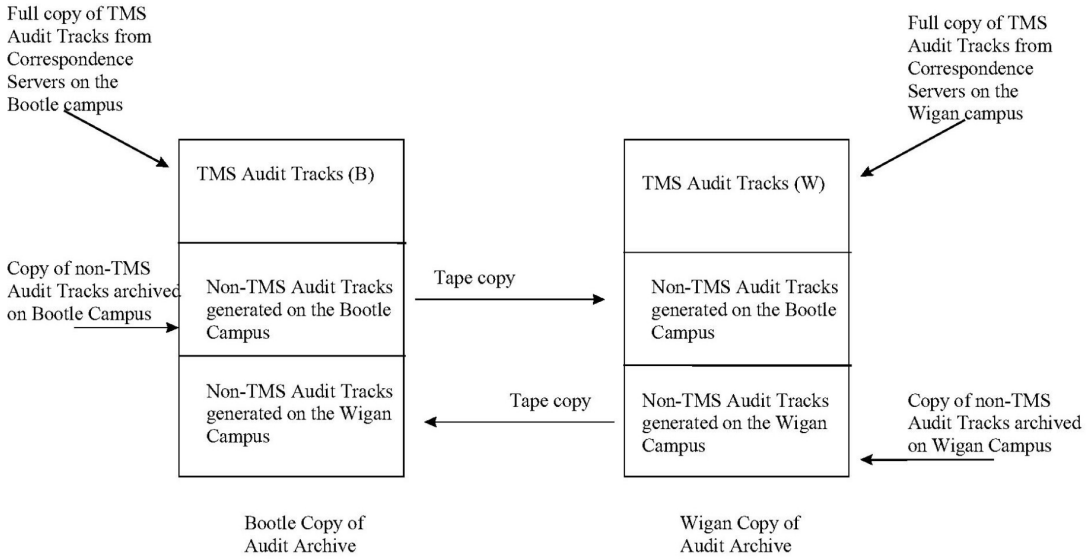

Figure 5.4
Audit Archive Structure

The correspondence server message stores are automatically duplicated on both the Wigan campus and the Bootle campus. However the duplication and auditing mechanisms do not maintain the sequence of the audited messages in the TMS Audit Tracks across both sites. Hence although the TMS Audit Tracks archived on the two sites will contain the same information the order of the information will vary.

In addition to the transfer of Audit Tracks (via tape) between the two sites it is necessary to transfer copies of the seals (relating to the transferred files) in a secure manner. The seals are exchanged using the FTMS facilities over the intercampus links.

When Audit Tracks are introduced on an Audit Server by the tape exchange mechanism the files containing the tracks are introduced into the Legato index. Thus each Audit Server maintains a full index of all of the files held in its part of the Audit Archive.

5.5 AUDIT POINTS

An Audit Point is a logical concept which is introduced in this design to minimise the linkage (as seen by users of the Audit Workstation) with the physical design of the

FUJ00098211

Horizon system. This is intended to help reduce the knowledge that the auditors need of the details of the way Horizon 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 Horizon 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 Horizon 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 Horizon system and to assist the Pathway Auditors 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 the same 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.
From a detailed design stance it may prove useful to keep data files associated with a single Audit Sub-Point separate during storage and processing so that during Hoarding a separate Legato Save Set can be used for each Audit Sub-Point. This enables a different Retention and Browse Policy to be associated with each Audit Sub-Point.

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 a number, <*> is optional but if present indicated an interface which can have multiple instantiations. Examples are I-ATG-4* and I-CSAH-1. 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 are identified by interface names at both ends of the link, e.g. I-ATG-3/I-ATS-4.
Appendix C provides an overview of all the interfaces and links between the components.

FUJ00098211

Instances of the Audit Server - There are two instances of the Audit Server one on the Wigan campus, the other on the Bootle campus. Unless otherwise specified all discussions apply to each instance. Where necessary the terms Audit Server (Wigan), AS(W), and Audit Server (Bootle), AS(B) are used.

6.1 APPLICATION COMPONENTS

6.1.1 AUDIT TRACK GATHERER

Figure 6.1 identifies the main interfaces to the Audit Track Gatherer

ATG Functionality

The functionality of the Audit Track Gatherer includes:

The ability to run multiple instances of the ATG on a single Audit Server and to configure each ATG to collect Audit Tracks from different interfaces and different Audit Sub-Points within an interface. Each instance of the ATG will be configured with a unique id.

The regular checking for 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 Horizon wide Maestro 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, (schedules by Maestro), at which instance of the ATG shall run, as frequencies within those time periods at which the ATG instance shall attempt to gather, e.g. every 4 minutes and duration.

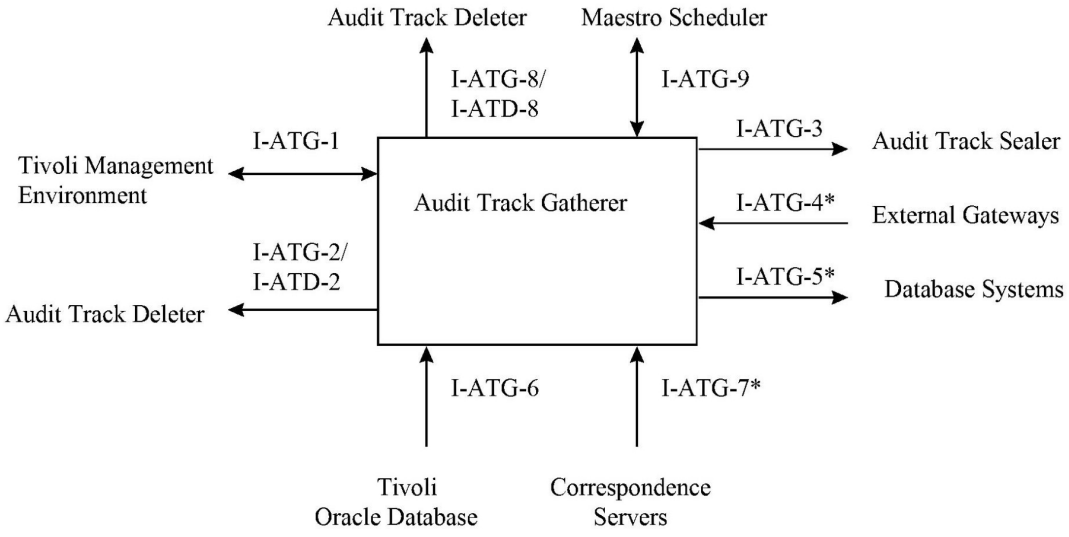

Figure 6.1

FUJ00098211

Interfaces to the Audit Track Gatherer

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

  The naming conventions used by the applications which generate the Audit Tracks are not consistent with the need to use the file name to easily identify the date and main contents of a file for recovery by the Audit Track Retriever, and more specifically the Legato product which implements much of the functionality of that module.

  All Audit Tracks are subject to a renaming policy before being stored in the Audit Archive. The appropriate renaming is carried out by the Gatherer based on the rules of the renaming policy. The renaming policies are detailed in [18], the Audit Data Catalogue.
- Once the gathered files have been renamed the ATS is notified via I-ATG-3
- Generation of a record of activities to be passed to the Audit Track Deleter via I-ATG-2.
- Marking files as gathered. Since the frequency of gathering of files 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 (i) work over an ATG crash, (ii) allow files marked as being gathered to be unmarked (without loss of information) by administrators in an emergency situation, (iii)cause, at most, minor increases in network traffic. See section 9 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 which have been marked as gathered.

ATG Interfaces

I-ATG-1 Interfaces to the system administration staff and facilities are all via the Tivoli Management Environment (TME). TME shall provide facilities via Tivoli tasks 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.

FUJ00098211

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 Audit Trail 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 frequency shall be a configurable parameter. The Record File shall be a text file.

I-ATG-3 provides complete files containing Audit Tracks of all types generated in the system 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.

I-ATG-4* gathers information from the FTMS implementation on the external gateways, as specified in [15] and [16].
There are multiple external gateways. An external gateway requires two NT systems one on a Pathway campus (the local gateway system), one on the appropriate external site (the remote gateway system).
Each external link is an Audit Point with each gateway system being an Audit Sub-Point.
This design only supports the gathering of transferred files from Local Gateway Systems (both primary and secondary).
The Audit Tracks are gathered from the various gateway systems from directories which contain tracks for only one Logical External Connection. For each Logical External Connection, control files and copies of transferred files are gathered from different directories. The term audit directory is used in this

FUJ00098211

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 NTFS 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 (on a per LEC basis) 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. Control files should only be gathered after 1200 noon on the day after they were created.
The Audit Track Gatherer is responsible for notifying the ATD of successfully gathered files.
The Audit Track Gatherer must gather Audit Tracks from the gateway machines within 3 days of their generation. See section 9 for details of the resilience.
The AS on each campus only gathers part of the set of Audit Tracks generated by all of the external gateways on the system.
It should be noted that the above mechanism means that Audit Tracks relating to an individual file transfer may be gathered on different sites and will only be brought together when the Audit Tracks are exchanged by tapes (see section 0).

I-ATG-5* A range of applications run on a number of database servers in the Horizon system from which audit data is gathered These include:

PAS/CMS
OBCS
Reference Data Management System
Roll Out Database

I-ATG-6 gathers Audit Tracks from the Tivoli Oracle Database on the same campus as the Audit Server. The Tivoli Oracle Database (TOD) generates Audit Tracks of security relevant events that have been collected by the Systems Management Facilities.
The Audit Tracks are written as files to a shared directory on the TOD system. All files written to the shared directory have unique file names. Copies of transferred files are available for collection as soon as their extension has been changed from ".tmp". The ATG collects them at regular intervals using the

FUJ00098211

remote share capabilities of NTFS. The ATG is responsible for deleting files in the shared directory.

I-ATG-7* The ATG on an Audit Server gathers Audit Tracks from the CS on the same campus as the Audit Server. Since the Correspondence Server Audit Tracks are effectively duplicated on each campus a resilient copy is automatically generated on each site.

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.

I-ATG-9 Instances of the ATG will be scheduled via Maestro. The ATG shall implement the standard Maestro 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 6.2 identifies the main interfaces to the Audit Track Sealer

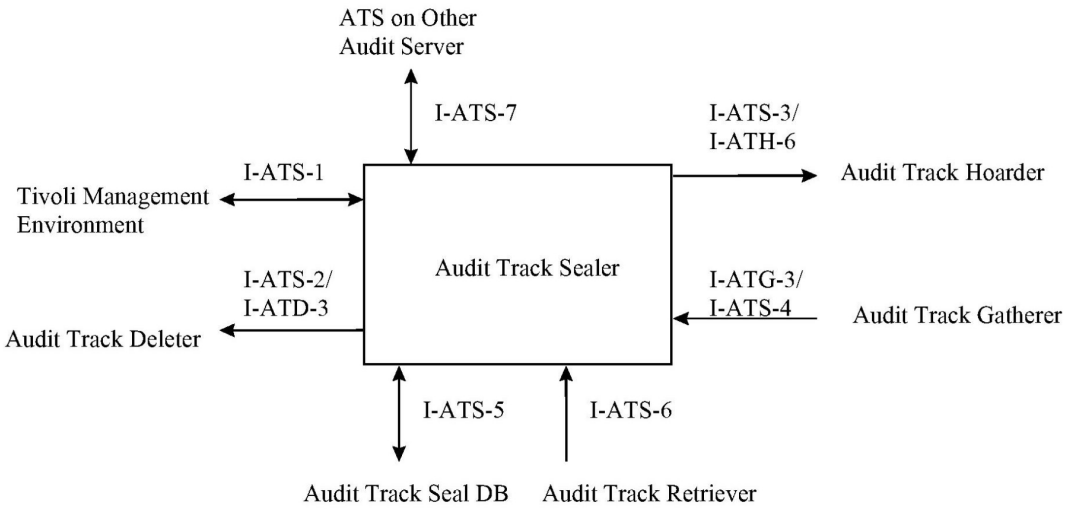

Figure 6.2
Interfaces to the Audit Track Sealer

ATS Functionality

FUJ00098211

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 ATH and into the ATSDB for 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, [11] has been selected. See section 9 for performance implications.
Once a file has had a seal calculated details (including the retention period of the seal) will be stored in the Audit Track Seal Database. Failed attempts to calculate a seal should also be recorded in the ATSDB.
Only when a sealed file has been made available to the ATH can any (unsealed) local copies be deleted. Files which have not been notified as sealed to ATH must not be deleted from I-ATS-4.
The ATS shall give priority to checking the seals as requested by the ATR (over generating seals).
It shall be possible to configure the times at which the ATS seals files from the ATG and checks files from the ATR independently.
The seals which are to be copied to the remote Audit Server shall be configurable per Audit Point.
The use of FTMS to transfer seals between the Audit Servers will generate control files and audit copies of the files transferred. They will be written to the same CD-W to which the ATSDB is regularly backed up. Each Audit Server will copy the FTMS audit and control files generated on itself. FTMS shall be configured to audit both transmitted and received files.
The ATSDB will regularly be written to a CD-W. No details of the seals will be held in the Audit Archive.
It will be necessary to keep the CDs for audit purposes. It is likely they will need to be kept for 7 years (i.e. the longest retention period for the seals on the disks).
The size of the ATSDB should not grow indefinitely hence the ATS should occasionally remove any seals with a retention period that has expired.
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 Trail Extraction facilities. Full details of the querying requirements are specified in [17

ATS Interfaces

I-ATS-1 Interfaces to the system administration staff and facilities are all via the Tivoli Management Environment (TME). TME shall monitor the existence of the expected instances of the ATS and shall report as a Tivoli event any unexpected unavailability of such instances. In addition the TME shall provide facilities via Tivoli tasks 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.

FUJ00098211

I-ATS-2 For each file that is successfully or unsuccessfully sealed or a seal checked a record is written to the Audit Trail 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 Hoarder (ATH). A seal has been successfully checked when a new seal has been generated, checked against the original seal and the results stored in the ATSDB.
The ATS will pass the details of the files gathered to the ATH 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 frequency shall be a configurable parameter. The Record File shall be a text file.

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 Hoarder. Such sealed files are placed in a shared directory to be picked up by the ATH.

I-ATS-4 See I-ATG-3. The ATS picks up files to be sealed from this interface.

I-ATS-5 For each file that is sealed or has its seal checked an entry is written to the Audit Track Seal DB (ATSDB). Where possible a entry of any attempt to calculate a seal which failed (for whatever reason) should be written to the ATSDB as well. 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. one seal in the case of a generation, the recalculated seal and the original seal value retrieved from the Audit Track Seal DB in the case of a check.

FUJ00098211

- A meaningful success/failure code which indicates
  The operations (un)successful completion in terms of the calculation
  (If a check) the success of the comparison
- 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 seal.

I-ATS-6 Files who's 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 sealed/checked. Files to be sealed/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.

I-ATS-7 The ATS shall frequently (configurable) pass details of seals calculated for the files (to be transmitted to the other AS via tape) to the other Audit Server. The records of those seals shall be exported to a flat file and FTMS shall be used to tansfer that file to the other AS.
Received files shall be imported and written to the ATSDB. The ATSDB shall clearly identify imported seals and seals which have been locally generated.

6.1.3 AUDIT TRACK DELETER

Figure 6.3 identifies the main interfaces to the Audit Track Deleter.

ATD Functionality
The Audit Track Deleter has two main functions:

- Deletion 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.

FUJ00098211

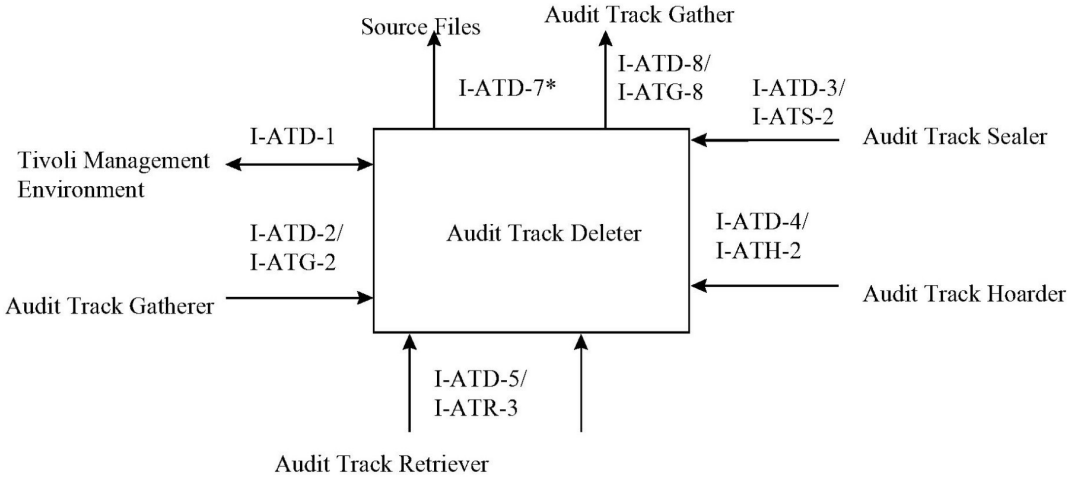

Figure 6.3
Interfaces to the Audit Track Deleter

The point at which the original audit files are deleted needs to be configurable: possible configurations should include:

- Once the ATG has signalled that it has completed the gathering
- Once the ATH has signalled that it has written the file to tape
- Once ATS has sealed file
- A configured time (typically hours) after one of the above events.

File Deletion

Any queues of files to be deleted maintained by the ATD must be stored on persistent storage so that the queues can be reused after ATD outages. This mechanism allows the gathered files to be written to appropriately resilient storage before the source file is deleted.

The configuration information should be variable 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.

Auditing

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

- An Activity File
- A Deletion Exception File

FUJ00098211

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 administrator usage. The administrator shall be responsible for checking the files regularly to ensure correct operation of the Audit Server and correction of any abnormalities identified.
The files shall both be text files, formatted for ease of reading.

ATD Interfaces

I-ATD-1 Interfaces to the system administration staff and facilities are all via the Tivoli Management Environment (TME). TME shall monitor the existence of the ATD and shall report as a Tivoli event any unexpected unavailability. In addition the TME shall provide facilities via Tivoli tasks to start and stop the ATD.

I-ATD-2 See I-ATG-2

I-ATD-3 See I-ATS-2

I-ATD-4 See I-ATH-2

I-ATD-5 See I-ATR-3

I-ATD-7* The interface via which ATD deletes files using NFS and NTFS. These are the files collected by the ATG on I-ATG-4*, I-ATG-5*, I-ATG-6, I-ATG-7*, and 8.

I-ATD-8 See I-ATG-8

6.1.4 AUDIT TRACK HOARDER

Figure 6.4 identifies the main interfaces to the Audit Track Hoarder

FUJ00098211

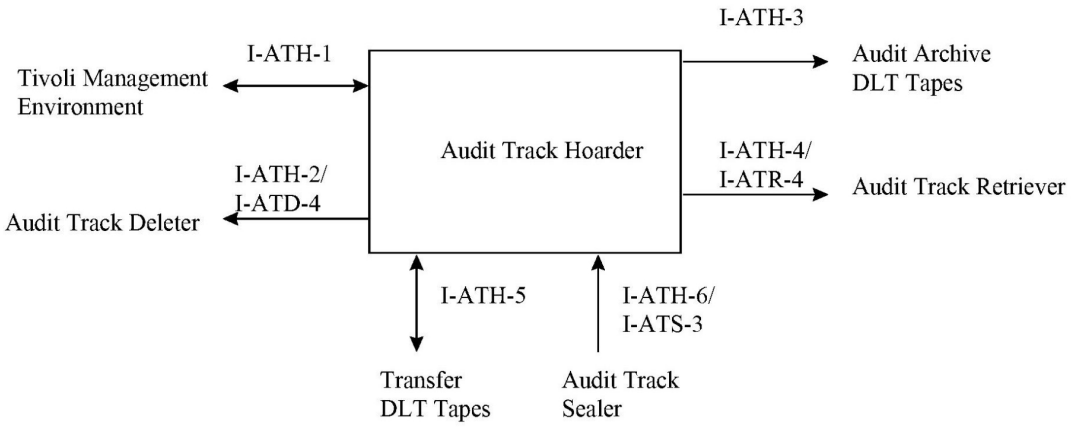

Figure 6.4
Interfaces to the Audit Track Hoarder

ATH Functionality

There shall be one instance of the Audit Track Hoarder which may be configured to support streaming to multiple parallel DLT Tapes.

The writing of the Audit Tracks to tape and the management of the tapes (and their contents) shall be carried out using Legato NetWorker.

The Backup facilities provided by Legato NetWorker shall be used, the Archive facilities shall not be used (so that the Browse and Retention Policies can be used). Legato Network does not delete files that have been backed up. The ATH will have to delete the files after they have been backed up. It will be necessary to use the FULL backup level.

Backups will take place at scheduled times. The schedule shall be configurable. See [18] for proposed schedule.

The ATD will utilise three separate tape pools:

- Tapes containing the TMS Audit Tracks exclusively, the TMS Pool
- Tapes containing all other Audit Tracks with an 18 month retention period, the Non TMS Pool
- Tapes containing Audit Tracks with a 7 year retention period, the 7 Year Pool

The Browse Policy for the TMS and Non TMS Pools shall be configured to retain Index Entries for 18 months.

The Retention Policy for the TMS and Non TMS Pools shall be configured to retain Save Sets for 18 months.

The browse and retention policies for the 7 Year pool will be 7 years.

Every Audit Point Sub Point will be associated with a single Save Set. Note this probably means that different sub-directories should be used to store Audit Tracks from different Audit Points as they are manipulated by the Audit Server. It also means that the Audit Tracks associated with each Audit Point Sub Point need to have the

FUJ00098211

same retention period. By this means Legato NetWorker will know what the retention period on the different files are.
Separate Pools will be used for the different classes of tapes identified in I-ATH-3. Each pool will be linked to a specific (configurable) set of tape devices.
The contents of the Save Sets backed up to the Non TMS Pool will also be written to tape via I-ATH-5. This copying will be done without using Legato NetWorker. The directory structure of the files on the tape should be such that files with 18 month retention periods and 7 year retention periods are kept separately.
On a daily basis all of the Non TMS Audit Tracks will be written out via I-ATH-5 and the tape will be transferred to the Audit Server on the other site.
When a tape containing a copy of the Non TMS Audit Tracks is received at the other campus it contents are read onto the Audit Server to the ATH and backed up using the Legato NetWorker facilities in such a way that the files are written to the tapes with the appropriate retention periods. By taking this approach the files on tape and the entries in the Legato index are subject to the retention and browse policies and hence are automatically managed by Legato NetWorker.
The seals for the transferred files are exchanged between the Audit Servers using an alternative path, see section 0. Note that the seals on the files transferred between sites are not checked until a file is retrieved from the Audit Archive.
The Auto Media Management facilities in Legato shall not be used. Note this means that it will be necessary to label tapes before they are used. This will be a manual procedure which will also involve physical as well as electronic labelling. If a tape is loaded on to the Audit Server which does not have a recognisable Legato label it will not be accepted by the system. This approach minimises the risk of accidentally overwriting data already on a tape.
A tape with a Legato label, e.g. a tape from the Audit Archive will not be overwritten unless all the data on it has passed its retention period.
The Legato NetWorker Power edition shall be used. Note that the Power edition is not needed because of the number of devices etc. it supports but rather to remove the need for data being streamed from Audit Server disk to local Tape to traverse the TCP/IP stack.
Events requiring administrator intervention shall be notified to System Administrators via the Tivoli Management Environment. The interface shall be configured so that any warning messages do not flood the operators while messages requiring operator intervention, e.g. tape loading, are delivered in a timely fashion.
A Network Server Bootstrap shall be produced and backed up daily.
Full Legato NetWorker facilities will be needed on the Audit Server

ATH Interfaces

I-ATH-1 Interfaces to the system administration staff and facilities are all via the Tivoli Management Environment (TME). TME shall monitor the existence of the ATH and shall report as a Tivoli event any unexpected unavailability. In addition the TME shall provide facilities via Tivoli tasks to start and stop the ATH.

FUJ00098211

I-ATH-2 For every file successfully (or unsuccessfully) written to Tape by the ATH it will inform the Audit Track Deleter (ATD) of:

- The time and date of the Writing
- A meaningful success/failure code
- The name of the file
- (If successful) The tape the file was written to

The ATH will pass the details of the files gathered to the ATD via an agreed directory. Details of the 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 file, but may contain many. 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 frequency shall be a configurable parameter. The Record File shall be a text file.

I-ATH-3 This interfaces supports three classes of tapes

- Tapes containing the TMS Audit Tracks exclusively, the TMS Pool
- Tapes containing all other Audit Tracks with an 18 month retention period, the Non TMS Pool
- Tapes containing Audit Tracks with a 7 year retention period, the 7 Year Pool

The files written out on this interface are all under the control of Legato NetWorker.

I-ATH-4 This interface is effectively the Index File maintained by Legato NetWorker.

I-ATH-5 This interface is used to transfer (via tape) non TMS Audit Tracks that need copying between the Audit Servers on the Wigan and Bootle sites. The import and export of files via this interface is not under the control of Legato Networker.

I-ATH-6 See I-ATS-3

6.1.5 AUDIT TRACK RETRIEVER

Figure 6.5 identifies the main interfaces to the Audit Track Retriever

FUJ00098211

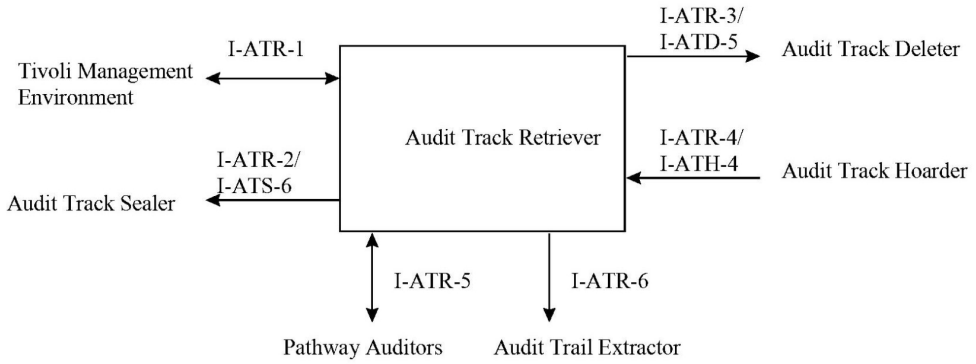

Figure 6.5
Interfaces to the Audit Track Retreiver

ATR Functionality

The Pathway Auditors use the Legato NetWorker User Program to browse the File Index and to select files for retrieval from the Audit Archive. Files retrieved from the Audit Archive are placed in a directory on the local Audit Server.

Retrieval of the files will involve the operators at the appropriate data centre having to collect the appropriate tapes from the archive and loading them on the Audit Server. This may involve what appears a considerable delay to the Auditor.

Once a file has been retrieved it needs to be automatically passed onto the ATS (for seal checking) and the ATE for extraction. This involves making an extra copy of the file so that the extraction and seal can progress asynchronously.

ATR Interfaces

I-ATR-1 Interfaces to the system administration staff and facilities are all via the Tivoli Management Environment (TME). TME shall monitor the existence of the ATR and shall report as a Tivoli event any unexpected unavailability. In addition the TME shall provide facilities via Tivoli tasks to start and stop the ATR.

I-ATR-2 See I-ATS-6.

I-ATR-3 For every file successfully (or unsuccessfully) retrieved from tape by the ATR it will inform the Audit Track Deleter (ATD) of:

- The time and date of the Retrieval
- A meaningful success/failure code
- The path name file

The ATR will pass the details of the files retrieved to the ATD via an agreed directory. Details of the 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 file, but may contain many. Details of retrieved files batched together in one Record File must always be held on persistent storage media (e.g. disk) and must never be over

FUJ00098211

written. Record Files must regularly (e.g. every hour) be passed over to the ATD. The frequency shall be a configurable parameter. The Record File shall be a text file.

I-ATR-4 This interface is effectively the Index File maintained by Legato NetWorker.

I-ATR-5 The Pathway Auditors use the Legato NetWorker User Program to retrieve files from the Audit Archive. The Networker User Program runs as a client on the Audit Workstation.

I-ATR-6 The files retrieved from the Audit Archive are placed into a directory from which the Audit Trail Extractor can copy the data for subsequent extraction activities. The ATE is responsible for deleting the files from this directory.

6.1.6 AUDIT TRAIL EXTRACTOR

The high level design for the Audit Trail Extractor is specified in [17].

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 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 available for collection as soon as they have been placed there or (ii) be renamed as specified in [10] to indicate their availability for gathering.
- The sub directories should not be used to hold any other files.
- The name of each Audit Track file shall conform to the naming standard specified in [18] and hence shall identify the Audit Point and Audit Sub Point
- Access to the Audit Track files for gathering shall be via NFS (for Unix systems) or NTFS (for Windows NT systems). Access to the sub directory shall be limited to the application generating the Audit Track and the Audit Track Gatherer. It should be noted that [19] requires the Audit track files to be written in write-append mode.
- The Audit Track Deleter shall delete all of the files from the sub directories. Once the Audit Track files are available for gathering the generating application should not make any assumptions about their continuing availability in the sub directory.
- In order to minimise the need to re-gather files follow 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.

FUJ00098211

- The systems on which the Audit Track files are generated should be able to store a minimum of 3 days worth of Audit Tracks (based on peak loadings).

6.3 DISTRIBUTED APPLICATION SERVICES

See section 6.1

6.4 INFORMATION MANAGEMENT

See section 6.1

6.5 NETWORKING SERVICES

The Audit Server will use the standard Pathway network services.
The Audit Network Requirements are:

1. The network shall support the transfer of up to 120GB (based on the total of one peak day's data plus two typical days data) of (application) data into the Audit Server per campus within an 18 hour period per day.
2. The network link into the Audit Server shall continue to operate across a single point of failure. The LAN interface card(s) in the Audit Server shall be regarded as part of the network for the resilience design

.

6.6 PLATFORMS

See Appendix B for the hardware requirements of the Audit Server and the Audit Workstation. They will be based on the standard secure NT build. Detailed design/implementation work may require modifications to the standard build.

7. SYSTEMS MANAGEMENT

The Audit Server and the Audit Workstation will be managed via Tivoli, and will be based on the standard Tivoli NT build.

8. APPLICATION DEVELOPMENT

There are no additional requirements from the high level design on the application development process.

FUJ00098211

9. SYSTEM QUALITIES

9.1 AVAILABILITY

The requirement to be able to provide a complete audit trail of transactions for agreed timescales (primarily 18 months with some 7 year) within the Horizon system to the auditors (both Pathway Internal and customers') places a requirement to maintain the availability of the Audit Tracks with a high degree of assurance.
The need to ensure that the Audit Server does not fall significantly behind in the collection of Audit Tracks on a single campus places further constraints on the design.

9.1.1 FAILURE MODES

The Horizon system architecture has been designed to provide a high degree of availability at New Release 2. The design of the Audit Data Storage and Retrieval facilities utilises a number of these inherent availability features while providing a number of features its self.
The following list identifies the major failure modes relating to the Audit Data Storage and Recovery design. These are detailed together with the counter measures utilised in [23].

Campus Level Disaster

Loss of Correspondence Server machine

Loss of Other Platforms Generating Audit Tracks

Loss of Audit Server

Availability of Audit Data on the Audit Server and in the Audit Archive

Component Failures During the Operation of the Audit Server

9.1.2 AVAILABILITY OF DATA ON TAPE

There is a need to be able to recover data up to 7 years after it has been written to tape. DLT tapes have an inherently long shelf life, typically quoted at over 20 years, and the archive strategy generates and maintains two copies of all of the archived data (on tape) which are held on different sites.
The actual physical environmental conditions for operation and storage will be dependent upon the actual media used. It is recommended that the media, operational environment and storage environment but will typically be compatible with the following:

FUJ00098211

Operational Environment

Operating Temperature:	10C to 40C
Humidity (Non Condensing):	20% to 80%
Maximum Wet Bulb Temperature	25C

Storage Environment

Ambient Temperature	18C to 26C
Humidity (Non Condensing):	40% to 60%
Stray Magnetic Field	< 4,000G

9.2 USABILITY

Human users of the Audit Server fall into two categories:

- Operations/Administrator staff
- Pathway Internal Audit Staff.

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 using standard facilities provided by COTS software including Legato NetWorker, Windows NT and Tivoli.

Some of the activities require relatively long elapsed times, e.g. recovery of tapes from secure storage and loading them onto the Audit Server, to complete. 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 Trail Extraction facilities is addressed in [17].

9.3 PERFORMANCE

9.3.1 VOLUMETRICS

The peak loading of the Horizon system is expected to generate approximately 58.5GB of audit data in one day, with an average of approximately 30.7 GB per day, see [1].
The sizing of the Audit Server has to be able to cope with a situation where the server has been out of operation for two days. It is a design aim that the Audit Server shall be

FUJ00098211

able to archive those two days worth of data plus the current day's data on the day after the outage.
It is estimated that 3 days worth of such data would peak at approximately 120GB (based on one peak followed by two typical days). The sizing of the Audit Server and other components should be able to cope with this maximum load in a single day.
Each campus has to be able to deal with the maximum load.

9.3.2 CORRESPONDENCE SERVER HARVESTING

Of the peak days audit data approximately 36.6GB is expected to be from the Correspondence Server message store, with approximately 19.6GB on a typical day. Thus a maximum load in a single day for correspondence server audit data would be approximately 75.8GB (based on one peak day followed by two typical days).
This data will have to be read from the correspondence server disks by the Audit Server. Assuming a 2 fold increase for protocol overheads the network will have to handle approximately 150GB of data. Assuming a sustained throughput of $30Mbs^{-1}$ over a $100Mbs^{-1}$ Ethernet would give a minimum elapsed time of approximately 12 hours.

9.3.3 MD5 ALGORITHM PERFORMANCE

The MD5 secure hashing algorithm will place a significant load on the Audit Server. There is evidence that MD5 can be implemented to achieve a throughput (on large files) of better than $2.0MBs^{-1}$ on a single Pentium Pro processor. For the purpose of this HLD a throughput of $2.0MBs^{-1}$will be assumed. At maximum load the estimated time for a single processor to handle this amount of data would be approximately 16 hours. Thus to get the time slot needed for sealing down to below 12 hours on such a day 2 Pentium Pro processor would be needed, assuming a fairly linear scaling on a multiprocessor system this would give a window of 8hrs for sealing.

9.3.4 IMPACT ON OTHER PARTS OF THE HORIZON 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.
There are major impacts on the network and the Correspondence Servers. The requirements the design places on the network are summarised in section 6.5.
The impact on individual Correspondence Servers is outside the scope of this document. However this design assumes that multiple Correspondence Server Clusters can be deployed at any stage during the roll out.
Impacts on other components are outside the scope of this design.

FUJ00098211

9.4 SECURITY

9.4.1 IDENTIFICATION AND AUTHENTICATION

All users (including administrators) of the Audit Workstation and Audit Server shall log onto systems using the Secure Id facilities defined in [21]. Each user shall be uniquely identifiable.

9.4.2 AUDIT

Details of the main file operations on the Audit Server shall be audited as defined in section 0. 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

All Legato NetWorker notifications shall be logged to appropriate log files which shall be archived daily.

9.4.3 DOMAIN STRUCTURE

The Audit Workstations and Audit Servers shall operate in a separate domain. The Horizon NT domain structure will need to support this configuration.

9.4.4 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.

9.4.5 PHYSICAL ACCESS CONTROLS

All Audit Server and Audit Workstation hardware shall be held in physically secure areas where physical access to the systems is controlled.
The tapes used to store the Audit Archive, the Write-Once CD ROM used to store copies of the Audit Track Seal Database and any backups of the Audit Server and Audit Workstation will be stored under similar physical security.

FUJ00098211

9.4.6 ROLES

There shall be separate roles for:

Legato Administration
Audit Server (inc. Audit Workstation) Administration
Pathway Audit Staff
Tape Operator.

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

9.4.7 ACCESS CONTROLS

The access control configuration of the Audit Server shall be compatible with the Access Control Policy [21].
The Pathway Audit Staff role shall have access to the Legato User Program and shall be able to recover any file from the Audit Archive. They shall not normally 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 Pathway Audit Staff can only recover files into a retrieval/extraction working area, it shall not be possible for the Pathway Audit Staff to recover files to other parts of the Audit Server directory structure.
The Pathway 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 NT and Legato NetWorker privileges.

9.4.8 TIVOLI TASKS

Tivoli tasks run with Administrator privilege and hence are highly trusted. All such tasks on the Audit Server should have undergone the strict control mechanisms Pathway is employing to manage their release into the live system.

9.4.9 DATA BACKUP

All non transient data, including any Network Server Bootstrap, on the Audit Server shall be backed up on a regular basis.

FUJ00098211

9.5 POTENTIAL FOR CHANGE

The Audit Architecture, [1], 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 Horizon 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 a third retention period (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

The introduction of New Release 2 involves the migration of the Audit Data Storage and Retrieval facilities from the current Release 1c facilities to the facilities described in this design

10.1 RELEASE 1C AUDIT FACILITIES

The Release 1c audit facilities, [6], provide a standalone Audit Workstation which is used by Pathway internal auditors to view Audit Trails and to provide Audit Trail information to customer auditors as requested.
Audit Tracks are collected from:

- Correspondence Server

  The standalone Audit Workstation is temporarily connected to the campus LAN and a connection made to a Correspondence Server. R-Query is then used to filter the contents of the message store and to extract the requested information to the Audit Workstation.

- Database Audit Trails

  These files are transferred by tape to the Audit Workstation on a periodic basis.

10.2 OVERALL MIGRATION STRATEGY

The strategy for the overall migration of Release 1c to New Release 2, as defined in Release 1c to Release 2 Migration Strategy, [7], and elaborated in Release 1c to Release 2 Migration Details of Implementation, [8], is based on the temporary

FUJ00098211

operation of Release 2 counter PCs in Post Offices against the Release 1c Correspondence Server.

Once all existing Release 1c counter PCs in Post Offices have been replaced with New Release 2 counter PCs the systems in the data centres will be upgraded to New Release 2 components. It is at this time that the New Release 2 Correspondence Server is introduced.

When the New Release 2 Correspondence Server is introduced it will be linked to the Release 1C Correspondence Server and the Riposte message store replication facilities will be used to replicate the contents of the Release 1c Message Store on the New Release 2 Correspondence Server.

As a consequence of this approach the Release 1c message store will have a period of time where it contains New Release 2 messages and the New Release 2 Message Store will have a period of time where it contains Release 1c messages.

In Release 1c the Correspondence Server Message Store never has any messages archived, i.e. removed from the Message Store and stored offline. Consequently the Release 1c Message Store during the migration will contain data which is more than 100[2] days old and does not need to be migrated to New Release 2. In order to ensure that this "out of date" data is not unnecessarily migrated to New Release 2 it will be removed from the Message Store.

10.3 AUDIT MIGRATION STRATEGY

The (proposed) strategy for the migration of audit facilities and data within the overall migration from Release 1c to New Release 2 is based on the porting of the New Release 2 Audit Agent Harvester back onto the Release 1c Correspondence Server. The logical stages of the approach as follows:

1. The New Release 2 Audit Agent Harvester is ported to and validated with the Release 1c Correspondence Server. The Release 1c Correspondence Server archive facilities are validated.

2. Prior to the migration of any of the Post Offices to New Release 2 counter PCs, and any thinning of the Release 1c Message Store a full copy of the message store is taken to tape, as a backup for fallback purposes.

3. The Audit Server is connected to the Release 1c Correspondence Server and the Audit Agent Harvester is run to copy all of the data in the Release 1c Correspondence Server Message Store contents to the Audit Server. The Audit Agent Harvester and the Audit Server then continue to operate

[2] It is understood that although the Riposte archive functions have not been used in Release 1c the default setting for an archiving period has been set to 100 days. It is therefore proposed to use that value.

FUJ00098211

throughout the migration period.

4. The Release 1c message store is thinned, i.e. all messages over 100 days old are removed.

   The Release 1c version of Riposte includes functionality to archive messages which have reached their expiry date. The default expiry date has been set to 100 days from the creation of each message. This facility has not been enabled on the Release 1c implementation.

   The thinning process will be implemented by enabling the archiving and sending the archived messages to a null device, hence removing them from the system. Once the thinning process has been completed the archiving facility will be disabled.

5. The New Release 2 counter PCs are introduced over a period of time interoperating with the Release 1c Correspondence Server.

6. Once all of the New Release 2 counter PCs have been introduced and immediately prior to the linking of the Release 1c Correspondence Server to the New Release 2 Correspondence Server the Release 1c Message Store will be copied to tape, for backup and fall back purposes.

7. The Release 1c Correspondence Server is connected to the New Release 2 Correspondence Server and the Riposte replication facilities are used to transfer the Message Store. All messages replicated to the New Release 2 Correspondence Server will be audited. When the New Release 2 Correspondence Server Message Store has been built the Release 1c Correspondence Server is removed from the system.

This means that there will be Release 1c messages in the New Release 2 Correspondence Server Audit Track and the audit extraction facilities will have to cope with them.

The other central components which operate at Release 1c and will interoperate with the New Release 2 Correspondence Server will be replaced in a single upgrade activity. Prior to that upgrade the Release 1c components will continue to audit as they have previously. As part of the upgrade all of the Audit Tracks from the central Release 1c components will be archived and a copy passed to the Release 1c Audit Workstation.
When the New Release 2 central components are switched on the archiving mechanisms described in this HLD will be deployed.

FUJ00098211

11. SOLUTION IMPLEMENTATION STRATEGY

Just do it.

12. COSTS, RISKS AND TIMESCALES

Potential non compliances with requirements include:

1. Use of Optical Archive Media, as specified in [3], DLT tapes are used instead
2. The architecture of the Audit Server and its HLD support the retrieval times specified in Audit Data Retrieval Requirements, [5], not those specified in Audit Trail Functional Specification, [3].

Although the Audit Server has been sized to cope with pre NR2+ workloads there is a risk that the actual implementation will require a more powerful platform to be introduced prior to NR2+.

FUJ00098211

APPENDIX A POSSIBLE FUTURE CHANGES

1. Building of more comprehensive indices of information within the Audit Archive files held on tape.
2. Use of a tape silo to hold the audit archive
3. Compression of files prior to gathering and their decompression on the Audit Server
4. The collection of large system logs directly from the system they are created on.
5. Specification and provision of a more powerful Audit Server configuration to handle steady state work loads.

FUJ00098211

APPENDIX B HARDWARE SIZING AND CONFIGURATION

Audit Server

There will be one Audit Server at each campus.

Sizing

The Audit Server requires the following partitions at NR2
C – System Software
D – Transient file system & Auditors area
E - Legato

The sizing of the E partition is critical to the running of Legato and has been sized using the following criteria:-

Estimated Files gathered per day (Steady State)	1500
Allowance factor for future release enhancements	4*
Percentage of files retained for 7 yrs	5%
Working days per week	6
Working weeks for 18 month retention	78
Working weeks for 7 yr retention	368
Legato space per index entry	250 bytes

Applying the allowance factor and 7 yr percentage gives

Files per day for 18 month retention = 1500 * 4 * 95% = 5700
Files perday for 7 year retention = 1500 * 4 * 5% = 300

Legato space = No. of files per day * days per week * retention weeks * legato space per entry/1000000 to give result in megabytes

Legato space for 18 month files = 5700 * 6 * 78 * 250/1000000 = 667Mb
Legato space for 7 year files = 300 * 6 * 368 * 250/1000000..= 166Mb
Legato software = 25Mb

Total sizing for E partition = 858Mb
C partition = 2GB
D partition = Total disk space – C - E

The D partition must be sized to adequately cover the following :

FUJ00098211

Audit Track Seal Database Size
Assuming that each entry requires 250 bytes and there is an average of 1.5 entries per file collected gives a total of approximately 192MB disk space required.

Audit Track Temporary Storage
In collecting the Audit Track files the Audit Server will need a substantial disk buffer in order to hold files which are being queued for sealing, hoarding and extraction.
It is recommended that the disk store should be able to hold up to 50% of a peak day's Audit Track collection to buffer the collected files and to hold the same amount of data for extraction purposes. The Audit Architecture, [1], specifies a peak collection of 58.5GB

The Transient File System & Auditors area has a total estimated disk storage requirement is just over 60GB. The configuration should be able to be expanded to 120GB of disk.
The above totals do not include space for RAID overheads which need to be added to the actual configuration.

Number of DLT Tape Drives

A need for separate tape drives to hold the following files has been identified:

1. TMS Audit Track files
2. Non TMS files with an 18month retention period
3. Non TMS files with a 7 year retention period
4. Files to be transferred to the other site (Plan to use a single tape to hold both 18 month and 7 year retention periods)
5. Tape for retrieving files from the Audit Archive
6. Spare

Processing Capacity

It is estimated that at peak load when trying to catch up with a backlog of files following a :

Hoarding will require 1x 200MHz Pentium Pro
Gathering will require 1x 200MHz Pentium Pro
Sealing will require 2x 200MHz Pentium Pro
Recovery and Extraction will require 2x 200MHz Pentium Pro

Giving a total of 6x 200 MHz Pentium Pro processors. This number is considered to be high from an NT scalability point of view and would require a very high end choice of system. Given that the initial work load can be dealt with by a four processor system it is recommended that a 4 (200MHz Pentium Pro or faster) processor configuration is purchased initially and a replacement system with faster CPU is processed at a later date if required. This is considered to be a lower risk approach and may well prove the cheaper route without taking risk into account.

FUJ00098211

Memory Requirements

The system will require a minimum of 256MB, it should be able to support at least 512MB and should preferably be extensible to 1GB memory.

I/O controllers

The system should support a minimum of 2 $100Mbs^{-1}$ Ethernet Controllers
A minimum of 4 disk controllers
A minimum of 2 tape controllers
Provide at least 2 expansion slots

Resilience Features

The disks storage should be implemented in a RAID configuration so that single disk failures can be accommodated. It is desirable that the disks support Hot Pull.
The system should preferably be configured with redundant power supplies and cooling facilities.

Audit Workstation

There will be one Audit Workstation at each campus and one at FEL01.
The Audit Workstations should be expandable especially with respect to memory and disk capacity.

Minimum Configuration:

17" Screen
Mouse and Keyboard
200MHz Pentium
8GB disk
64MB Memory
DLT Tape
Write Once CDROM
Ethernet interface

FUJ00098211

APPENDIX C INTERFACES

FUJ00098211

FUJ00098211

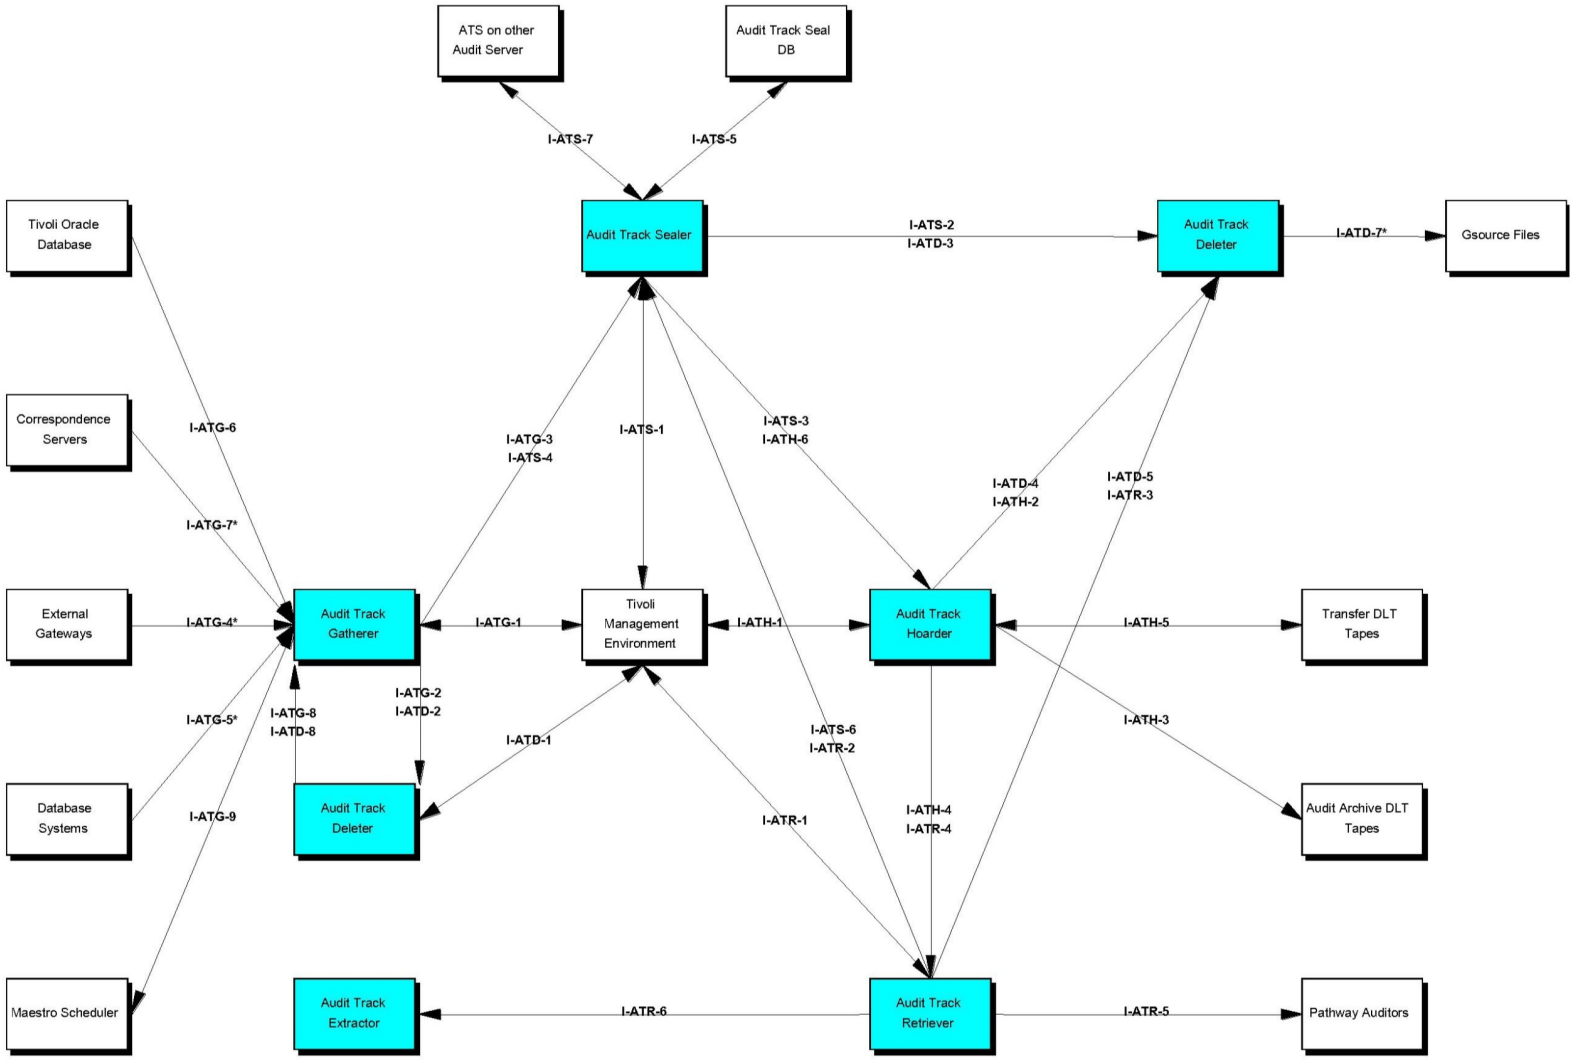

FUJ00098211