FUJ00001742

Document Title: TMS Architecture Specification

Document Type: Architecture Specification

Release: N/A

Abstract: This document provides a description of the architecture of the TMS component of the Horizon solution provided by Pathway at BI3 to meet the requirements of Post Office Ltd.

Document Status: APPROVED

Originator & Dept: Patricia Morris, Application Products Delivery Unit

Contributors: David Hollingsworth, Peter Wiles, James Stinchcombe, Gareth Jenkins, Simon Fawkes, Dave Tanner

Internal Distribution:

External Distribution:

Approval Authorities:

Name	Position	Signature	Date
Tony Drahota	Manager, Architecture and Systems Design		
Bob Booth	Post Office Ltd		

FUJ00001742

0.0 Document control

0.1 Document history

Version	Date	Reason	Associated CP/PinICL No.
0.1	26/03/99	Initial draft	
0.2	08/06/99	Comments from Dave Cooke	
0.3	07/10/99	Incorporating comments from review; removal of BA and BES material.	
0.4	03/11/99	PSTN descriptions removed.	
0.5	12/10/00	Re-structured and refined, incorporating review comments on earlier versions.	
0.6	27/11/00	Incorporating comments on V0.5. Restructured in line with *OPS Architecture Specification.*	
0.7	21/12/00	Cross-references to other documents expanded in line with comment on *OPS Architecture Specification* V0.10.	
0.8	18/05/01	Incorporates comments on V0.7.	
0.9	11/07/02	Version for internal review.	CP3161/ CCN850.
0.10	29/07/02	Version for internal review.	CP3161/ CCN850.
0.11	31/07/02	Version for review. Changes (side-barred against V0.7). Highlighted in yellow: arising from Post Office Ltd comments on V0.7. Highlighted in blue: additions and changes arising from CP3161/CCN850 References to ICL removed References to POCL amended to Post Office Ltd. Section 2.1.1, 2.4.1.4, 3.3.2.1, 3.3.2.3.1, 4.7.2.4, 4.7.3, 5.1.2: references to NBS added Section: 2.4.1, 2.4.3.1: references to CNIM added Section 3.3.1, 4.7.1.1, 5.1.3.2: configuration changes since V0.7 documented Section 2.4.2, 2.4.3.1, 4.5, 5.4: communications changes since V0.7 documented Section 0.7, 2.4, 4.2: references to WebRiposte added Other marked amendments in response to comments.	CP3161/ CCN850.
0.12	03/10/02	Section 0.4: CCS added Section 1.3: formatting corrected Section 2.2: Figure 2-2 corrected. Section 2.1.2.2: reference to appendix corrected. Section 2.4.1.1: reference to copy of message store added. Section 2.4.1.4: superfluous paragraph removed.	CP3161/ CCN850.

FUJ00001742

		Section 2.4.2: 'dial-up' removed to reflect new FRIACO service. Section 2.4.2: 24-hour outlets added. Section 2.4.3.1: Sentence removed Section 3.2.2: third paragraph reworded Section 3.2.2: ISDN dial-out connections added Section 3.3.1: paragraph moved from below to above Figure 3-1. Section 3.3.1: explanatory sentence about LAN connections added Section 3.3.1.2: explanatory sentence about clusters added Section 3.3.1.3: Backup added to sentence about Audit server. Section 3.3.2.1: new subheading for Generic Agent Servers added. Section 3.3.2.4.1: information about NBS authorisation agent corrected and expanded. Section 4.2.1.3: information about message expiry periods added. Section 4.3: sentence added to clarify resilience of single-counter outlets. Section 4.5: information about the Counter Call Scheduler added. Section 4.5.2: information about the replication of priority messages added. Section 4.7.1.1: last sentence third paragraph reworded for clarity. Section 4.7.2.3: information from stock declaration added as a source for recovery Section 4.7.2.4: 'some' removed from on-line services lost during communication outages. Section 4.7.3: 'the' replaces 'any' (in the context of duplicate messages) Section 5.1.2: reference to encryption of PIN numbers and sensitive data added. Section 5.1.3.2: references to digital signatures made explicit. Section 5.1.3.2: Digital signing: 'Simple' removed from 'Standard public key technology'. Smart card acknowledgement added. Section 5.1.3.2: information added on message content encryption and MAC protection Section 5.3: authentication added to list of cryptography usages.	
0.13	08/10/02	Version for Post Office Ltd review. Document History for V0.12 amended to include change in section 5.1.2 Section 2.1.2.1: 'all' removed. Section 3.3.1: Bullet point added re replication. Section 3.3.1: Final bullet point. Sentence added re LAN failure.	CP3161/ CCN850.

FUJ00001742

		Section 5.1.2: Typo corrected (repeated text). Section 5.1.3.2: Audit Data retention period added Section 5.1.3.2: 'digitally signature' removed. Section 5.3: 'four' updated to 'five'. Section 5.3: 'or' replaced by 'both…and'. Section 5.4: section replaced with updated text.	
1.0	10/10/02	Approved version.	CP3161/ CCN850.
1.1	09/01/03	Version for Post Office Ltd review. Section 5.1.3.2: Audit retention period amended from 15 to 7 years.	(Withdrawal of CP3268) CP3375
2.0	10/01/03	Approved version.	

0.2 Review Details

Review Comments by :	
Review Comments to :	

Mandatory Review Authority	Name		
Fujitsu Services	David Hollingsworth	Peter Wiles	David Johns
	James Stinchcombe	Tony Drahota	Glenn Stephens
	Graham Hooper	Gareth Jenkins	Dave Cooke
	Simon Fawkes	Dave Tanner	Will Dawson
	Duncan Macdonald	Chris Bailey	Adrian Goodwin
	Geoffrey Vane	Cliff Wakeman	
Post Office Ltd	*Bob Booth		
Post Office Ltd Contracts and Commercial Team			
Optional Review / Issued for Information			
Fujitsu Services	Martin Whitehead	Allan Hodgkinson	Brian Orzel

(*) = Reviewers that returned comments

FUJ00001742

0.3 Associated documents

Reference	Version	Date	Title	Source
TD/ARC/030			OPS Architecture Specification	Pathway
TD/STD/004			Generalised API for OPS/TMS	Pathway
BP/DES/003			Counter Hardware Design Specification	Pathway
TD/SPE/006			TMS Hardware Design Specification	Pathway
RS/POL/003			Access Control Policy	Pathway
NB/SPE/011			Network Banking RAC Data Flow Model	Pathway

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

0.4 Abbreviations/Definitions

Abbreviation	Definition
API	Application Programming Interface
APS	Automated Payment Service: counter application supported by Horizon.
BCV	Business Continuity Volume
BIOS	Basic Input/Output System
C	A UNIX-derived programming language
C++	Object oriented version of C
CCS	Counter Call Scheduler
CHAP	Challenge Handshake Authentication Protocol; Microsoft protocol used to support authentication in ISDN
Cluster	Group of Correspondence Servers, all handling the same set of Outlets and replicating data between each other for resilience purposes
CNIM	Counter Network Information Monitor; Counter-based ISDN monitoring service
CRC	Cyclic Redundancy Check
DLE	Digital Local Exchange
DRS	Data Reconciliation Service
DSS	Formerly Department of Social Security, now Department of Work and Pensions

FUJ00001742

DWP	Department of Work and Pensions
EMC	ElectroMagnetic Compatibility
EPOSS	Electronic Point Of Sale Service: counter application supported by Horizon.
FAD	Financial Accounting Division
FI	Financial Institution
HSM	Hardware Security Module
IP	Internet Protocol
ISDN	Integrated Services Digital Network
LAN	Local Area Network
LFS	Logistics Feeder Service interface to Post Office Ltd's SAPADS system, used to ensure that outlets have adequate and timely supplies of stock including cash.
LUC	Look Up Cluster; server called by an agent to identify which cluster an outlet belongs to.
MAC	Message Authentication Code
NB	Network Banking
NBE	Network Banking Engine. The system that handles the interface between the Horizon system and the Financial Institutions (FIs) that have reached agreement to provide automated banking services in Post Office outlets.
NBS	Network Banking Service. The application that supports banking functionality within the Horizon architecture.
NDIS	Network Device Interface Specification
NIC	Network Interface Card
NTFS	Windows NT File System
OBCS	Order Book Control Service; counter application supported by Horizon, which supports a similarly named DSS application.
OCX	OLE Custom Control
OLE	Object Linking and Embedding
OPS (1)	Office Platform Service. The provision and support of the hardware and software at outlets including the Desktop environment of the Horizon system.
OPS (2)	Oracle Parallel Server
PMMC	Post Master's Memory Card
POCL	Post Office Counters Ltd (now Post Office Ltd)

FUJ00001742

POLO	Post Office Log On; process used by post office managers when switching on a counter PC
Riposte	Retail Integrated Point Of Sale system in a Transaction Environment: product from Escher that provides both the infrastructure and the Desktop environment of the Horizon system. Now superseded by WebRiposte.
RMS	Riposte Message Server: message storage and replication mechanism of Riposte.
RPC	Remote Procedure Call
SMS	System Management Service
SNMP	Simple Network Management Protocol
SSC	System Support Centre based at Bracknell
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TIP	Transaction Information Processing: Post Office Ltd application that handles transaction data returned from Horizon.
TMS	Transaction Management Service. The hardware and software required for the replication, transmission and management of transactions committed to the Horizon Riposte message store and Pathway Data Centres, or vice versa.
UDP	User Datagram Protocol
VPN	Virtual Private Network
WAN	Wide Area Network
WebRiposte	An enhanced version of Riposte providing additional services based upon web-enabled functionality.

FUJ00001742

0.5 Changes in this Version

Version	Changes
2.0	None.

0.6 Changes Expected

Changes

0.1 Trademarks and Acknowledgements

Riposte and WebRiposte are trademarks of Escher Group Ltd., Cambridge, Massachusetts.

Windows NT is a registered trademark of Microsoft Corporation in the USA and/or other countries.

FUJ00001742

0.2 Table of Content

FUJ00001742

FUJ00001742

1.0 Introduction

1.1 Purpose

This document is an introduction to the architecture of the Transaction Management Service (TMS). It describes the hardware and software components of the service and its relationship to the Office Platform Service (OPS) and other elements of the Pathway solution delivered at BI3. It is produced in conjunction with the documents listed in section 1.3, *Related Documents*.

1.2 Readership

This document is intended for application developers within Pathway or elsewhere. It has been developed to give an overview of the architecture of TMS and to enable developers to plan the development of new applications to operate within the Pathway solution.

1.3 Related Documents

Documents that are referred to in this document, and that should be read in conjunction with it, are as follows:

- Counter Hardware Design Specification

This document describes the hardware that is used by the Office Platform Service.

- TMS Hardware Design Specification

This document describes the hardware that is used by the Transaction Management Service.

- Generalised API for OPS/TMS

This document is designed to facilitate application development in the Pathway environment. It is intended to augment the documentation supplied by Escher by setting the Horizon implementation into context. It gives additional information for application developers about the architecture and facilities of OPS, TMS, and the Electronic Point Of Sale Service (EPOSS).

- OPS Architecture Specification

  This document covers the following topics:

  - The hardware and software components and architecture of OPS.
  - The style and architecture of the OPS Desktop.
  - How the architecture supports changes and extensions.

- Access Control Policy

This document defines the access control policy to be followed across the Pathway solution.

For publication details refer to section 0.3.

FUJ00001742

1.4 Scope

1.4.1 Document Set

This document forms part of the set that defines the environment that supports counter applications within the Horizon system.

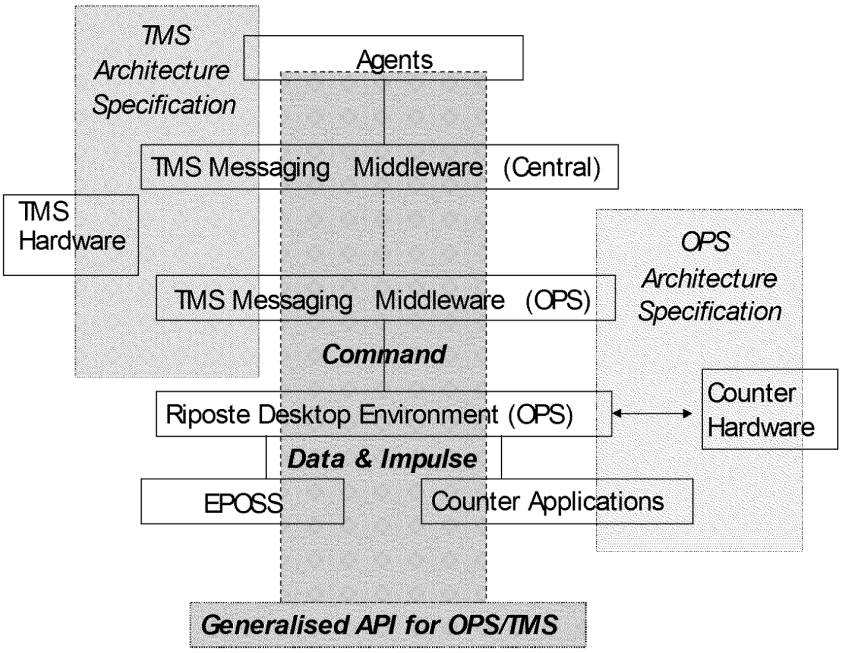

Figure 1-1 Document set

Figure 1-1 relates the components of the system to the documents that describe them, giving an indication of the scope of each document.

The scope of each document is described in section 1.3, *Related Documents*.

Each of the applications that support the Post Office Ltd business requirements has components that operate within the Riposte Desktop architecture and interact with TMS. The components are described in the remaining sections of this document.

How applications interact with other parts of the TMS architecture, and the APIs used by each application, are described in *Generalised API for OPS/TMS*.

FUJ00001742

1.4.2 Contents

This document is organised as follows:

Section	Contents
Section 1	introduces the document and its position within the document set.
Section 2	is an overview of the context within which TMS operates and a brief description of the architecture of TMS in terms of its hardware and software.
Section 3	describes the hardware and networks that support TMS. It also gives details of hardware resilience features.
Section 4	describes the software components of TMS, their organisation and functions. It also gives details of software resilience features.
Section 5	discusses the security considerations associated with TMS.

FUJ00001742

2.0 Overview

This section provides an overview of the architecture of the Transaction Management Service (TMS) and the context in which it operates. It contains the following sections:

- The business context within which TMS and the Horizon system have been developed.
- The technical context within which each application operates.
- The design principles that have been implemented in the TMS architecture described in this document.
- An overview of the architecture of TMS.

Pathway provides Horizon to meet Post Office Ltd's requirements for the automated provision of counter services in Post Offices. Horizon includes the technical infrastructure and supporting services as well as the business applications that deliver the contracted services. Figure 2-1 illustrates the major logical components of the distributed Horizon system.

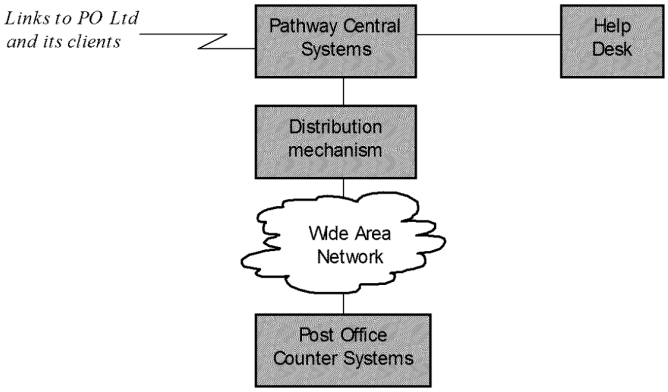

Figure 2-1 Major logical components

FUJ00001742

2.1 Business Context

Horizon provides a number of business services to Post Office Ltd and its Clients, and is designed to be capable of providing more in the future. Post Office Ltd collects payments for a range of Clients: utilities and local councils, for example. Many of these provide Automated Payment systems, such as swipe or smart cards, with the details of the payments being forwarded to the AP Client by Horizon.

Horizon services can be viewed as vertical 'stripes' across a number of horizontal layers, as shown in Figure 2-2.

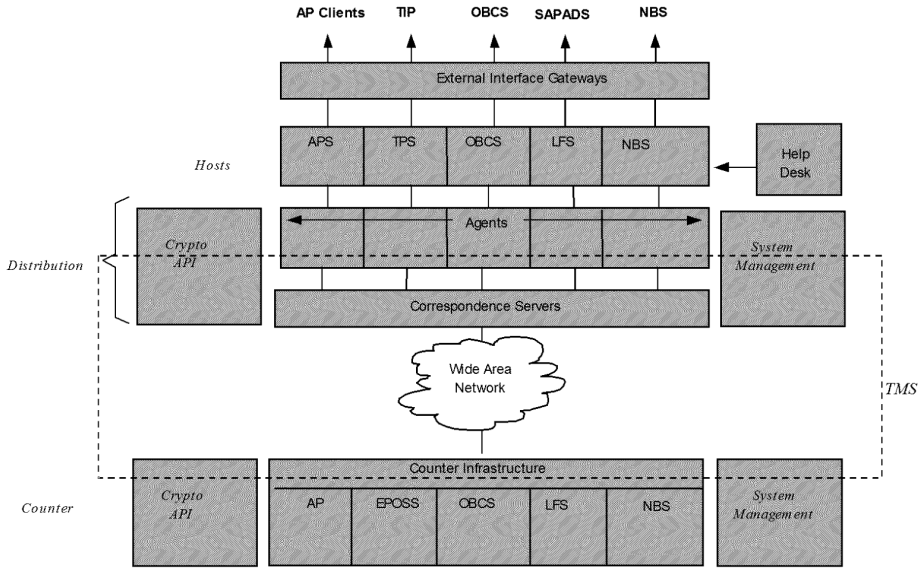

Figure 2-2 Vertical services

2.1.1 Business services provided

The services provided are as follows:

- Electronic Point of Sale Service (EPOSS) to provide support for customers purchasing goods or services, and the production of an electronic Cash Account that reflects such purchases.
- Automated Payment Service (APS) to support AP Clients such as the utility companies and others, who provide the means for their customers to make incremental payments based on cards and other devices.
- Order Book Control Service (OBCS) to check the validity of DSS Order Books and to make payments.
- Logistics Feeder Service (LFS) to ensure that outlets have up-to-date supplies of stock, including cash.

FUJ00001742

- Network Banking Service (NBS) to provide access to external banking authorisation services and support for associated accounting and reconciliation.

Each of the services is treated as separate, except that all transactions from all of these services are also transmitted to Post Office Ltd for their own subsequent processing.

2.1.2 Additional services provided

2.1.2.1 Key Management

Pathway Key Management Service (KMS) supports the business applications by managing the generation, delivery and life cycle of cryptographic key material. KMS provides a context-specific Cryptographic Functions API to applications.

For more information about the Key Management Service, refer to *Generalised API for OPS/TMS,* section 6.2, *Security,* and Appendix B, *Cryptography and Key Management*.

2.1.2.2 System Management

System Management controls the delivery of software to the counter.

For more information about System Management, refer to *Generalised API for OPS/TMS,* section 8, *System Management*. Full details are given in Appendix C, *System Management,* of the same document.

FUJ00001742

2.2 Technical Context

The high level architecture within the Pathway solution incorporates three layers of processing functions:

- Central services layer
- TMS layer
- OPS layer

The delivery of end-to-end services normally involves some processing within each of these three layers.

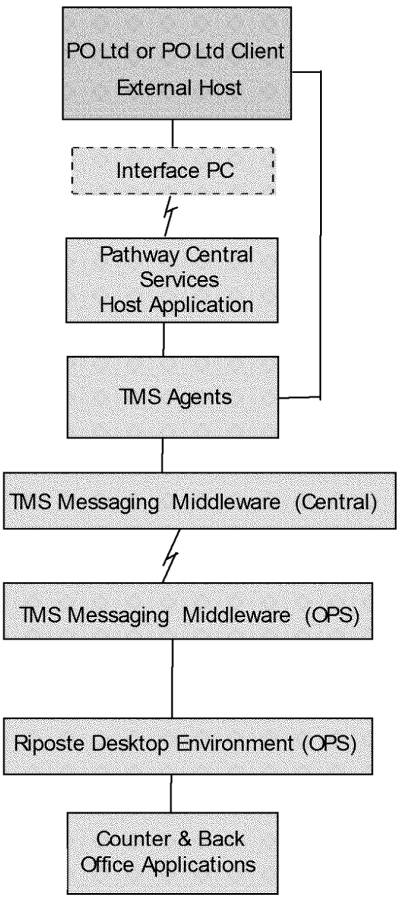

Figure 2-3 TMS in the context of the Pathway solution

Figure 2-3 gives an end-to-end view of the Pathway solution and the position of the Transaction Management Service within it.

FUJ00001742

2.3 Design Principles

The model used in the design stage of the Horizon solution has been based on the idea of two 'super layers':

- Business services
- Technical environment

Each of these has been considered in light of a number of design principles:

- Security

The security of TMS has been considered in terms of authorisation, access control, data integrity and encryption; these are outlined in section 5.0, *Security*.

- Performance

The performance criteria for TMS are that the agreed service levels are met.

- Availability

The resilience, recovery, and reliability of TMS are discussed briefly in sections 3.3 *TMS Hardware Resilience*, and 4.7 *TMS Software Resilience*, in this document.

These topics are covered more fully in *Generalised API for OPS/TMS*.

2.4 Overview of TMS Architecture

TMS can be regarded as the hardware and software that replicates, transmits and manages transactions and other data in the form of messages. TMS manages messages, in both directions, between the counter applications within outlets and the Pathway Data Centres.

TMS hardware consists of:

- Correspondence servers and agent servers within the Data Centres.
- The hardware that supports the Local Area Networks (LANs) in the Data Centres and the outlets, and the Wide Area Networks (WANs) between the Data Centres and the outlets.

Note that some hardware provides services for more than one component of the Pathway system. For example, the LAN in the Data Centre also supports host services, and the WAN also supports the Tivoli system management software. This hardware is described in *TMS Hardware Design Specification*.

The software that comprises TMS messaging middleware is as follows:

- Riposte messaging software.
- Riposte Message Server (RMS), with additional facilities (WebRiposte) enabling future web applications to be supported.
- Riposte APIs used by counter and agent applications.

2.4.1 TMS Layer Components

TMS provides facilities for the transfer and switching of messages between the following processors:

FUJ00001742

- OPS platforms (at outlet level).
- Hosts within the Pathway Central Services layer.
- Other attached computer systems, such as external hosts.

TMS comprises the distributed messaging (RMS) middleware, and includes the underlying communications service. It also includes the APIs that agent tasks use to transfer data for host applications and to the messaging middleware.

2.4.1.1 Messaging middleware

The messaging middleware is realised in the form of a set of distributed message stores provided on a set of central correspondence servers and on each outlet OPS platform. Facilities are provided for the replication of messages between message stores with delivery characteristics according to message priority (high or normal). The same replication mechanism is also used to provide resilience of message storage and message store recovery following failure; additional resilience is provided by a second copy of the encrypted message store, which each counter PC creates each night.

Each counter PC copies its encrypted message store (message store archive) overnight. The message store is encrypted so that security is not compromised, and compressed to optimise the speed of the recovery operation.

2.4.1.2 Message store API

An API is provided to the message store to enable the creation and retrieval of messages on the correspondence servers (for use by agent tasks) and on the OPS platform (for use by counter or back-office applications). APIs are also provided for housekeeping functions: for generating indexes and for archiving.

2.4.1.3 Message ports

In addition to conventional message store read/write operations, this API also supports a ‘message port’ facility through which an application (either TMS agent or OPS counter application) may directly read incoming messages, filtered according to the specified attribute criteria. This is the mechanism used by interactive agents or transaction-orientated counter applications to listen for an online message or a response to an interactive message.

2.4.1.4 TMS agent interface

TMS agents normally interface to a host service within the Pathway Central Services domain (see Figure 2-3), although, where appropriate, they may also interface directly to an external host. Direct interfacing is normally employed where the agent functionality is relatively ‘thin’ and there is no additional processing value provided by the Pathway Central Services. For example, this style of interface is appropriate for certain types of external authorisations.

The on-line Network Banking Service requires a real-time interface between an agent in the Horizon architecture, and the NBE. The interface is used to pass requests and certain classes of confirmations to the NBE and to read authorisations from the NBE and return them to the counter by writing them to the message store. (For further details, refer to *Network Banking RAC Data Flow Model*.)

FUJ00001742

2.4.2 TMS Layer Communications

The TMS layer incorporates the underlying data communications service utilising TCP/IP and UDP/IP, primarily over ISDN. Normally ISDN connections are used, where this cannot be achieved, Satellite and Frame Relay may be used.

Some ISDN-connected outlets require to be permanently on-line during the business day, or in some instances for 24 hours, in order to service NBS transactions; others dial up on demand. In some cases, outlets may be permanently connected during part of the working day and have an on-demand connection during off-off-peak hours. The Counter Network Information Monitor (CNIM) installed in the Gateway PC at each outlet manages the connection to the Data Centre, using either a dial-up connection or permanent connection, according to configuration information particular to the outlet and the time of day.

The CNIM has the following additional functions:

- It reports on the availability of the network to application-level processes that need to know
- It attempts to make regular calls to the Data Centre, following a detected network outage, until service is resumed.
- On request, it gathers network statistics and trace data according to a supplied trace level, and makes the relevant trace and log files available to System Management processes.

Within each multi-counter outlet communications are extended to all OPS counter PCs, via the local Ethernet LAN. This data communications service is used by the TMS messaging middleware (over UDP/IP) and other services such as SMS (over TCP/IP).

2.4.3 TMS Interface with OPS and Counter Applications

The message store, on each PC within an outlet, forms the lower boundary between the Transaction Management Service and the Office Platform Service. The TMS API provides facilities to the applications running on OPS to create and retrieve messages within the local message store. Replication between the local message store and those on the other OPS platforms at the outlet, and to and from the central correspondence servers, occurs automatically as part of the core TMS messaging functionality.

This process is essentially transparent to local applications (APS, for example) running on the OPS. Applications access the message store directly using the TMS API, provided by Riposte, to create, retrieve or parse messages independently of any replication activity.

2.4.3.1 Replication

Replication to and from other OPS nodes within the outlet occurs virtually instantaneously. Replication to and from the central correspondence servers is immediate if a high-priority message is used (an on-line banking request, for example) or if the outlet is at that time permanently connected to the Data Centre.

For other non-priority messages in outlets not connected permanently to the Data Centre, replication occurs in response to one of the following occurrences:

- A replication time interval is exceeded (defined by a configurable system parameter)

FUJ00001742

- A critical volume of messages is reached

In practice, a connection is made at least every hour during the day and at least every hour and a half at night.

2.4.3.2 Synchronisation

Where immediate synchronisation is required between the Data Centre and an outlet, it can be achieved using a priority message in the following ways:

- By a counter application writing a priority message. The priority message initiates immediate replication with the correspondence server central message store. The process may take some time to complete.
- By using a real-time read message port facility, in an agent at the Data Centre, in order to accept and process an incoming message from the local OPS message store.

FUJ00001742

3.0 TMS Hardware

This section describes briefly the two components of the hardware architecture that supports TMS.

- Platforms
- Networking

It includes brief descriptions of the operating systems and network software that are associated with these components. For further information, refer to *TMS Hardware Design Specification*.

3.1 Platforms

Platforms include the following hardware and software components:

- Data storage devices
- Processors
- Network connections
- Terminals
- Printers
- Operating systems and transaction processing software, which manage the hardware in support of the business applications running on the platforms

The platforms relevant to TMS are described in the *TMS Hardware Design Specification* and *Counter Hardware Design Specification*.

3.2 Networking

Networking Services support the distribution of data and applications by providing interconnection and interworking services over both local and wide area networks. These facilities are required to make connections between the following elements of the network:

- The Pathway Data Centres at Bootle and Wigan
- Pathway Data Centres and the Post Office outlets
- Counters within Post Office outlets

The communications infrastructure supports a variety of types of connection and bandwidth, depending on the needs of the services using the link.

Alternative routes are provided where the business needs dictate this.

3.2.1 Links in Data Centres

There are high-speed, static-cabled routes between the two Data Centres. Within the Data Centres, a number of Network Active Devices (hubs, routers and firewalls) are configured for their specific positions and functions in the network. Each Data Centre contains high and low

FUJ00001742

speed Wide and Local Area Networks, and configurations that support the various methods used for connection to the other Data Centre and to Post Office outlets.

3.2.2 Links between Data Centres and Outlets

The Data Centre LAN network is designed to use the IP connection between the outlets and Data Centres to support a Virtual Private Network (VPN).

There is an Access LAN in each Data Centre to which the outlets are connected, using an appropriate router. Correspondence servers do not need to know the link level connection method to an outlet.

Links between Data Centres and outlets include the following connection types:

- ISDN (permanent connections, daytime permanent connections or dial in on demand).
- Satellite
- Frame Relay
- ISDN dial-out connections

3.2.3 Links between Counters in Outlets

Counters in outlets are linked via a LAN using UDP over IP. System management software uses TCP/IP.

3.2.4 Protocols

Riposte mainly uses UDP over IP. The other major protocol used within the Horizon network is TCP/IP. Each platform has one or more IP addresses. Each Gateway PC has two network cards (one for WAN use, one for LAN use). The AutoConfig process allocates every outlet a WAN IP address (ISDN, Frame Relay or satellite) and a LAN IP address. It is necessary to allocate both for every outlet, since an outlet may change from being a single counter outlet to having more than one counter.

Other protocols are used for specialised purposes, including Simple Network Management Protocol (SNMP).

FUJ00001742

3.3 TMS Hardware Resilience

This section contains details of TMS hardware resilience and covers:

- Correspondence server resilience
- Agent server resilience

The resilience, recovery, and reliability of OPS are touched on in *OPS Architecture Specification* sections 3.3 *OPS Hardware Resilience*, and 4.5 *OPS Software Resilience*, and are fully described in *Generalised API for OPS/TMS,* section 6 *Other Functions*.

3.3.1 Correspondence Servers

There are a number of correspondence servers at each Data Centre. Each is part of a cluster of four correspondence servers, which service the same outlets. Note that:

- Although Figure 3-1 shows 1,250 outlets connecting to each correspondence server in cluster 1, in effect the total of 5,000 offices serviced by cluster 1 may connect to any of the correspondence servers.
- Each correspondence server in a cluster is connected to the other three servers in its cluster and all four replicate their message stores to each other.

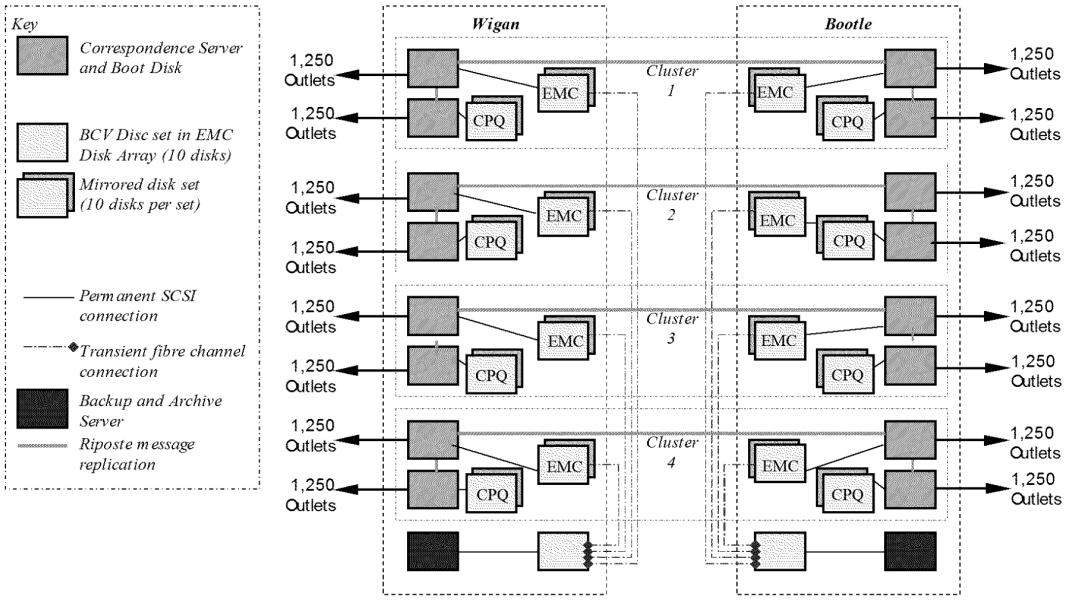

Figure 3-1 Correspondence Server Clusters

The main features of correspondence server platform resilience are as follows:

- Correspondence servers are replicated at each Data Centre.

FUJ00001742

- Riposte message replication is used to keep the correspondence servers in a cluster in step with each other and with the outlets serviced, and to bring a server up to date following its replacement.
- Mirrored disks are used to hold the message stores in case of disk failure.
- Message store archiving is used to restrain the disk space needed.
- Multiple LAN connections are used to cater for LAN or LAN card failure. Each correspondence server is connected to both LANs on its site and uses both to communicate with other correspondence servers. Correspondence servers use only one of the LANs to communicate with outlets. If one LAN fails in a Data Centre, the connections to outlets on the correspondence server on the failed LAN will be picked up and serviced by the correspondence server on the other operational LAN.

3.3.1.1 Processors

Processor resilience is provided by four servers comprising each cluster. The clusters contain one active server in each Data Centre. Windows NT system disks are hardware mirrored within the servers.

3.3.1.2 Disk failure

The Riposte message store data is held within a disk array that provides hardware-enabled disk mirroring.

Two correspondence servers per cluster, one at each Data Centre, use EMC-connected disks; the other two correspondence servers in the cluster are connected to Compaq disk arrays.

Should an EMC disk fail, data from its mirror is automatically copied to a warm standby disk that then becomes the second plex. The failed disk is then replaced, and the replacement becomes the new warm standby disk.

3.3.1.3 Tape backup

To provide for loss of both plexes or a data corruption, regular backups of the Riposte message store on the EMC-connected correspondence server are taken to tape. Backups are scheduled in such a way as to ensure that outlets are always able to connect to one of the correspondence servers in the cluster. To minimise the downtime of the Riposte Message Server when a backup is taken, the RMS is stopped and a copy of the message store is taken using EMC Business Continuity Volume (BCV) technology. Riposte is then restarted. The BCV is copied to tape through the Backup/Audit server.

If it becomes necessary to restore the disk volumes, they are copied from tape and then Riposte message synchronisation is used to bring them up to date.

3.3.1.4 LAN failure

Each correspondence server is connected to both discrete LAN segments and can use either. Riposte is configured so that if it fails to access a neighbouring correspondence server for replication, it will try an alternative route, which is configured to go via the other LAN card (on both systems).

FUJ00001742

3.3.1.5 Riposte Message Server failure

The Riposte Message Server is monitored by Tivoli system management software and restarted if it fails.

3.3.2 Agent Servers

3.3.2.1 Generic agent servers

Agent servers are high performance Pentium III PCs. They run under Windows NT Server, and support both Oracle and Riposte client software. They form the interface between the host central servers and the correspondence servers. The agent applications act as clients to both, by 'pulling' information from one and 'pushing' it to the other. This occurs in both directions.

3.3.2.2 NBS agent servers

Agent servers are high performance Pentium III PCs. They run under Windows NT Server, and support Riposte client software. They provide hardware support for the encryption of data to and from the NBE and form the interface between the NBE and the correspondence servers for real-time transactions.

3.3.2.3 Agent server redundancy

There are a number of agent servers at each Data Centre. Each runs a number of agent services, some interactive (running all the time) and some bulk services which are scheduled by Maestro. There is no inherent relationship between service and server, and each agent server contains the full range of agent service software.

The bulk agent software is designed to distribute the workload across the servers that are available at start-up. Work is divided into chunks, typically in the order of a thousand chunks, which are then distributed among the available agent servers.

If a bulk agent service fails, a sweep-up process at the end detects that some work has not been completed. It schedules it to one of the surviving servers. Any Oracle locks set by the failed agent expire after an average of five minutes.

3.3.2.4 Failure of agent services

Interactive agent services are monitored by Tivoli system management software. If one fails, Tivoli will detect the fact and start another instance of the agent on another agent server. The failure is detected by the inability to access the failed server, and so a LAN or LAN card failure as well as a server failure are all treated as server failures.

Each agent server is configured with two Windows NT services, each in a separate disk partition. Their features are as described for correspondence servers in the subsections of 3.3.1 *Correspondence Servers*.

3.3.2.4.1 NBS authorisation agent

The active NBS authorisation agent is connected to both the correspondence servers within the cluster at one Data Centre, and receives on-line service requests. This request is forwarded via the NBS authorisation agent to the NBE. If there is a failure of the NBS agent, or of both the correspondence servers that it communicates with, another agent instance running as a hot

FUJ00001742

standby at the other Data Centre will take over. There is, however, a period during the change to the standby agent when no service is offered for the outlets serviced by that agent's cluster.

FUJ00001742

4.0 TMS Software

4.1 Software Components

This section describes the software components of the Transaction Management Service, which are as follows:

- Riposte Message Server

Riposte provides the TMS Message Server on each counter PC and on the correspondence servers in the Data Centres.

- Messaging middleware, provided by Riposte

Riposte provides the messaging infrastructure, or middleware, that supports the distribution of messages between counter PCs at the outlets and correspondence servers in the Data Centres. Applications do not include a component at the correspondence server layer.

- Agents

Agents are the only application service components that interface directly with TMS at the Data Centre. They are responsible for transforming the file (or set) view of any host application into the message-based view that is appropriate for the counter application. They also transform Riposte messages from Desktop applications into records for host applications. Agents are discussed only briefly in this document. For details refer to *Generalised API for OPS/TMS,* section 7 *Agent Interfaces*.

The topics covered in this section are as follows:

- Riposte Message Server
- Message replication and synchronisation
- Riposte clusters
- Message transfer
- Transactional integrity
- TMS software resilience

4.2 Riposte Message Server

The Riposte Message Server (RMS) is implemented as a Windows NT service, with configuration information stored in the NT Registry. Error reporting is via the Windows NT application event log. Configuration and system administration facilities are controlled centrally and cannot be changed at an outlet. Riposte Message Server interfaces on counter PCs are shown in Figure 4-1.

WebRiposte provides a set of additional facilities to the Riposte Desktop and Message Server. The existing Riposte APIs, the Message Processor, Message Store and message replication to a remote Correspondence Server continue to support existing 32-bit Riposte desktop applications (APS and EPOSS, for example).

FUJ00001742

New APIs are available for future use by post-BI3 web-based applications, and new web services are provided to support web applications (for example by providing access to reference data).

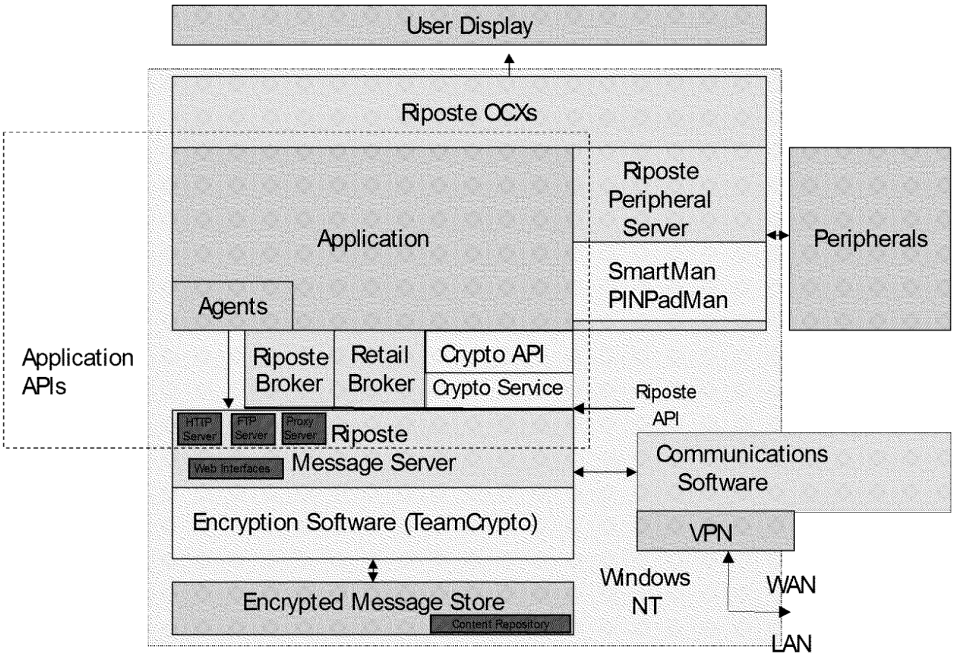

Figure 4-1 Riposte Message Server interfaces on counter PCs

The message store itself is implemented as one or more standard Windows NT File System (NTFS) files. On a counter PC, it is normally one file. On a correspondence server, it is normal practice to map the message store on to one file for each physical disk, thus enabling Riposte to optimise disk performance. Riposte provides facilities for adding new disks; the facility is not automatic and requires manual intervention.

4.2.1 Riposte Messages

Within the message store, all information used by Riposte is held as messages. Messages can be created as a consequence of information being passed down to counter applications, or as a result of customer transactions at the counter. They are also used to record more persistent data such as product price information.

4.2.1.1 Message formats

All messages have a unique message identifier that includes a group number related to the outlet Financial Accounting Division (FAD code), a PC (node) number within the group, and a sequence number within that node.

FUJ00001742

Refer to *Generalised API for OPS/TMS,* section 4.1 *EPOSS* for the use and format of the various types of message, including messages generated by the system and application Reference Data, as well as persistent objects.

4.2.1.2 APIs

The message-handling APIs are provided by Riposte and are implemented via Remote Procedure Call mechanisms. APIs are provided to retrieve messages, to parse messages, and retrieve attribute names and values. Similarly, there are APIs to enable an application to construct a message and to define its attributes. These APIs support natural C or C++ interfaces. Desktop applications that are coded in Visual Basic use a separate Retail Broker interface that has access to Riposte APIs. For more details of message-handling APIs, refer to *Generalised API for OPS/TMS,* section 5.5 *Riposte functions.*

4.2.1.3 Message expiry

All messages, whether messages or temporal persistent objects, are stored within the Riposte Message Server for a period determined by their individual expiry period. Expiry periods are set by the application, depending on the type of message. System-wide default values define the minimum and maximum expiry periods; there are different message expiry periods for correspondence servers and for counter PCs in the outlets. If applications set an expiry period that is outside the system default minimum and maximum, the values are overridden. If an application does not set an expiry period itself, the system default value is used.

4.2.2 Markers

Markers are used by applications to reliably delineate a cut-off point for a set of messages. For example, when it is necessary to balance the outlet at the end of the day, a marker can be used to decide which messages are balanced, and which are held over to the next day. Markers are used also by Riposte to co-ordinate message stores in the event of a failure.

4.2.3 Checkpoints

A checkpoint is a collection of markers for all outlets in the system, and provides a delineation point for the entire system.

4.3 Message Replication and Synchronisation

Each PC within a single outlet is defined as being part of a Riposte group. The group also includes a number of correspondence servers. Each correspondence server will be part of many groups. When a message is created on one member of a group, it is stored locally and then automatically distributed to all other members of the same group. When a new member joins a group, it automatically brings itself into line with the other members of the group. Co-ordination is achieved by the RMS initiating frequent interchange of markers between the local group members.

Within each node in an outlet, separate message runs are maintained for each node in the group.

FUJ00001742

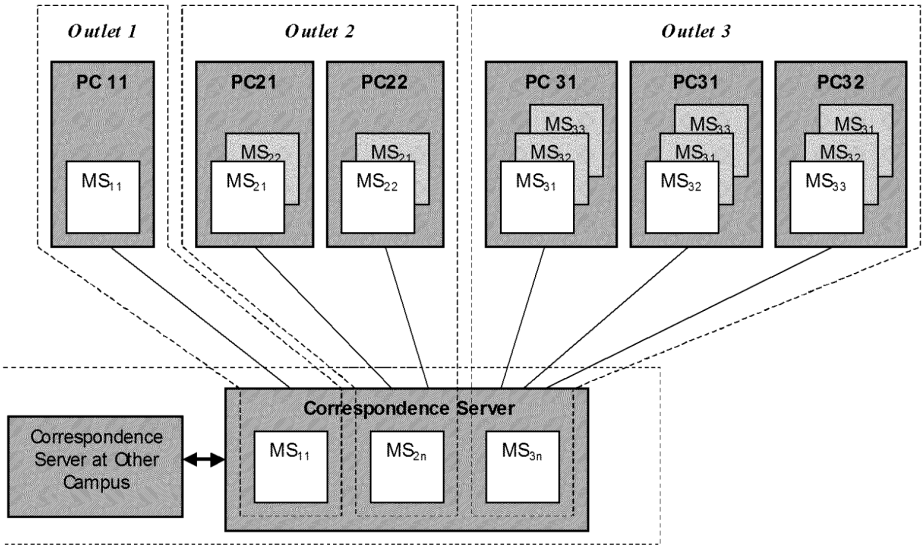

Figure 4-2 Riposte message store replication and synchronisation

Figure 4-2 shows three outlets: with one, two, and three counters respectively. The messages belonging to the three-counter outlet (MS_{31}, MS_{32} and MS_{33}) are replicated within each of the other counter PCs at that outlet, as well as in the correspondence server serving that outlet. This latter replication is not instantaneous, but will happen the next time that the link to or from the outlet is established.

Note that in single-counter outlets, the RiposteMirror service is used to provide a second copy of the message store on a removable hard disk. In this way, all outlets have a minimum of two 'counters' and message runs.

4.4 Riposte Clusters

The total set of correspondence servers is divided into a set of clusters. Clusters reduce the amount of replication that is required, since normal message replication is confined to the members of the cluster. Each correspondence server need only hold the message stores corresponding to the groups in its cluster.

Where an agent or counter application needs to write a Riposte message to another group, it does not know whether this is in the same cluster or not. The supporting multi-Riposte library calls the Look-Up Cluster (LUC) server, which returns the identity of the cluster to which it relates. Multi-Riposte then sends the message to the appropriate group.

4.5 Message Transfer

Except in the case of outlets that have a permanent connection, Riposte only sets up a call when a priority message has to be passed in either direction, or when a normal message has to be passed and a pre-configured time interval expires. A software application, the Counter Call

FUJ00001742

Scheduler (CCS), detects urgent business-level data that needs to be sent to the Data Centre within a specified time but which does not require immediate transfer. The CCS ensures that synchronisation takes place so that urgent data of this sort is transferred within the required timescales.

Once a call has been established, all waiting messages are transferred in both directions. The call is terminated on expiry of an interval of idle time.

There is usually some delay between a message being created at a counter position and that message reaching the Data Centre.

4.5.1 Protocols

Messaging between instances of Riposte uses User Datagram Protocol (UDP) over the Internet Protocol (IP). Any failures are detected by next marker exchange. Riposte requests the retransmission of any missing messages.

4.5.2 Priority Messages

Priority messages are used by applications that need to make an immediate communication with a remote neighbour. An example is an authorisation, where it is necessary to determine at once if the clerk is to be authorised to complete the transaction. A priority message causes the communications link to be opened (unless a connection is already established), and is then sent immediately. It is sent again at its normal position in the message run. An implication of the use of UDP is that even a priority message may sometimes be lost. However, the normal message synchronisation processes will ensure that it is eventually transferred (the time the transfer takes place depends upon the connection interval). The counter sends a single priority message to the Gateway PC in the outlet; the Gateway PC then replicates the message to each of the four correspondence servers. The reason for this is that many of the real-time agents only process messages that are received in real-time (they will not process the message if it has been lost). Since multiple copies of the priority message are sent as multiple UDP datagrams, the chance of all copies of a message being lost is extremely small; the agent detects and ignores duplicate messages.

4.6 Transactional Integrity

Transaction Management capability provided by Riposte ensures that a set of related actions to be carried out in such a way that either they all complete, or they are all considered not to have started. This is achieved by the use of an atomic commit function. The essential feature of this function is that the application is able to determine the start of a transaction sequence, and subsequently commit all the actions carried out within it. If it fails to commit the transaction, then all the actions carried out are rolled back and deemed not to have happened.

A transaction can fail for a number of reasons. For example, the hardware may fail before it completes.

Riposte utilises atomic replication to ensure that transactions are replicated as a complete set of messages; an application cannot read the first message in a transaction until all the component messages are securely replicated.

FUJ00001742

4.7 TMS Software Resilience

This section contains details of TMS resilience and covers:

- Correspondence server message store resilience
- Counter message store resilience
- Agent management of message recovery

4.7.1 Correspondence Server Message Store Resilience

The RMS message store is replicated across all the correspondence servers in a cluster. The RMS message store on each server is mirrored in RAID, so that eight copies of the data are held centrally.

4.7.1.1 Message store recovery on correspondence server failure

When a correspondence server fails, the other correspondence servers will continue servicing the work for the cluster's outlets while the problem with the failed server is corrected (for some failures this is automatic, see below). Tivoli runs a script to move all interactive agent services that have failed as the result of the correspondence server's failure to agent servers utilising another correspondence server. NB and NBS agents are resilient to a single correspondence server failure, in that they are already connected to the other correspondence server on the same site.

During the time that the failed server is being repaired, it will become out of step with the active server. Correspondence servers regularly exchange message markers to synchronise their message stores. When the failed correspondence server is brought back on line, this process is automatically used to bring the new server up to speed. The agents ensure that the correspondence server is not used until this recovery is complete.

Bulk agent instances are distributed between all correspondence servers. If a correspondence server fails, any work being carried out by a bulk agent instance is picked up by another agent instance connected to the alternative correspondence server within the cluster.

4.7.2 Counter Message Store Resilience

Riposte ensures that messages created at a counter are replicated both to the other counter PCs within the outlet, and to a number of correspondence servers in the Data Centres. It provides for automatic recovery of transaction data. Thus if a counter PC fails, its message can be recovered from another counter PC within the outlet, or from the correspondence server at the centre.

4.7.2.1 Resilience in single counter outlets

Counter message store resilience in single counter outlets is fully described in *OPS Architecture Specification* section 3.3.2.2, *Single counter site*.

4.7.2.2 Resilience in multiple counter outlets

Counter message store resilience in multiple counter outlets is fully described in *OPS Architecture Specification* section 3.3.2.3, *Multiple counter site*.

FUJ00001742

4.7.2.3 Loss of a complete outlet

The loss of a complete outlet could be caused by fire, flood, or theft of all counter systems. Any transactions that took place within the outlet, but have not been replicated to the Data Centre, will be lost. Where the equipment rather than the outlet as a whole is destroyed, any outstanding transactions can be reconstituted from stock declaration and tally roll records. When the outlet is reconstituted, Riposte enables it to identify the last transaction number secured at each counter. The remainder can be manually re-keyed from stock declaration and tally roll records.

4.7.2.4 Failure of communication link

Losing an outlet's communication link means that the counters in an outlet will not be able to communicate with their correspondence servers, and vice versa. Riposte message replication will not operate. This is not necessarily an issue in normal working. Messages are stored and only transmitted between the outlet and the Data Centre when the link is established, either on a timer basis or because of the need to send a priority message. When the link is re-established, the Data Centre will be brought up to date.

If the Gateway PC or the link itself fails, work can continue, although on-line services will be unavailable. Certain other operations that would normally make a real-time connection to the Data Centre will time-out and the counter Clerk may be required to make a telephone call to the Horizon System Helpdesk. When the PC or the link is repaired, Riposte message replication brings the outlet and the correspondence servers back into line.

4.7.3 Agent Management of Message Recovery

Agents are responsible for managing recovery and message resynchronisation between the host layer and the counter layer following a failure.

Agents need to co-ordinate the updates that they make to a host system with those they make to the RMS (via RPC calls). Both hosts and counter applications need to be designed to accept that some data may be duplicated but none will be lost.

Riposte contains transaction features that allow multiple messages to be committed as one atomic unit. These ensure that messages are replicated as a unit, and another processor cannot read one message in the message run until all the messages in it have been replicated to that processor.

Equivalent facilities are available at the host layer in such software as Oracle, and should be designed into other application types as appropriate.

The general strategy towards message recovery for both bulk and interactive agents is as follows:

- Applications should be able to detect duplicated messages and ignore any duplicates.
- Agents must be designed so that following a failure, they duplicate data rather than ignoring it.
- On failure of an agent it is restarted automatically by the system.

FUJ00001742

- NBS agents are run on separate platforms because of their particular security and hardware requirements (Hardware Security Module (HSM) cards are needed for PIN translation)

FUJ00001742

5.0 Security

This section describes TMS Security and covers the following topics:

- Security domains and the security features that are relevant to TMS
- The use of access controls at the counter and over the network
- The use of cryptography
- The use of VPN

These topics are covered more fully in *Generalised API for OPS/TMS,* section 6.2 *Security*, and Appendix B *Cryptography and Key Management.*

5.1 Security Domains

5.1.1 Generic security domains

Security architectures revolve around the concept of interworking 'security domains'. A domain, in this sense, is viewed as a collection of one or more platforms and the interconnecting network, which can be considered to be physically secure and which present a common interface to the rest of the world.

A view of domain structures is shown in Figure 5-1.

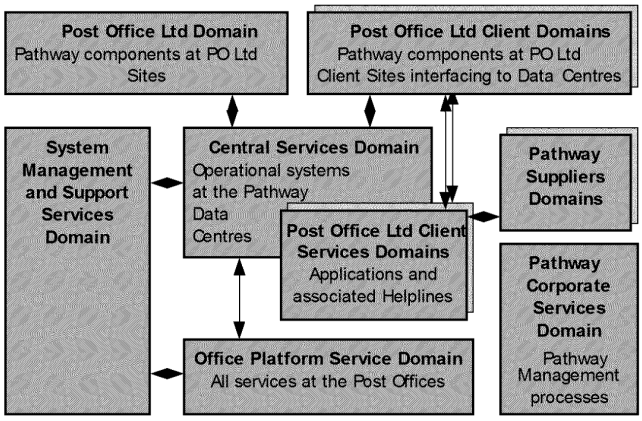

Figure 5-1 Generic view of Pathway domains

These domains represent major areas of responsibility for services. Each has different characteristics that affect the type of access controls that needed. The only security domains that are relevant to security issues relating to TMS are as follows:

- Data storage sub-domain of the Central Services domain.
- Office Platform Service domain.

These topics are described in sections 5.1.3 and 5.1.4 respectively.

FUJ00001742

5.1.2 NBS security domains

The Network Banking Service has particular security requirements and its security domain interfaces are shown in Figure 5-2.

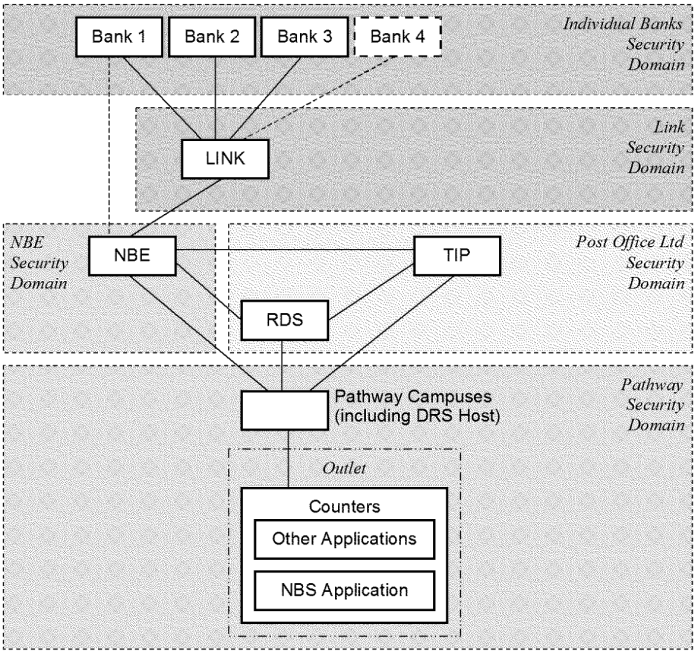

Figure 5-2 NBS security domains

Between the counter and the NBS authorisation agent, NBS messages in transit are protected by the encryption of PIN numbers and sensitive data, and by digital signature. From the NBS authorisation agent to the NBE they are protected by a Message Authorisation Code (MAC).

5.1.3 Data Storage Sub-domain

The data storage sub-domain comprises the main central data repositories owned and managed by Pathway. These include the Riposte message store, and the major host application databases.

5.1.3.1 Domain boundaries

The Riposte message store is also found in, and forms a conduit to and from, the Office Platform Service (OPS) domain.

The data held in the host databases belongs to Post Office Ltd or its clients, and Pathway must ensure adequate separation between the data owned by different organisations.

Data is held for audit purposes within the archive service.

FUJ00001742

5.1.3.2 Security features

The security features that apply in the data storage sub-domain are:

- Cyclic Redundancy Check

  Riposte adds a Cyclic Redundancy Check data field (identified by the preceding attribute name 'CRC') automatically to messages when they are created. This field is used to protect and ensure the integrity of messages.

- Auditing

  To meet the requirements of both Pathway and Post Office Ltd, audit information is logged in response to selected events. The events are one of the following types:

  - Human activity, such as logins.
  - Automatic system events, such as file transfers from Post Office Ltd.
  - Business-related events, such as the generation of a request to issue a benefit book, or user access to a particular data item.

  All new Riposte messages, even those that are purely transient in nature, are audited daily at the correspondence server.

  A complete audit stream is provided that includes all transactions carried out at the counter and messages sent to the counter by host applications. These messages are extracted in the first instance to filestore. The files of the current day's messages are passed to the Archive Service. Audit data gathered but not yet secured is held in a redundant disk configuration so that a single disk failure will not result in loss of the data. Audit information is extracted to an audit server in each Data Centre and written to EMC Centera, a high-capacity, highly resilient device that holds several years' worth of data. Audit data is retained for seven years and can be retrieved as required; audit records are presented in exactly the same way as recent records when retrieved, although they are subject to filters appropriate to the selection and the audit authority for which the selection is being made.

  Each day the live Riposte message store is purged of expired messages, so that messages older than their expiry period are no longer present there. If these messages are required subsequently, they are retrieved in their raw form from the audit stream and loaded to an empty message store on the archive server, in a form that enables the standard Riposte facilities (RQuery) to be used to browse it.

- Message sequencing

  The sequencing of messages and the use of markers and checkpoints ensures that individual messages cannot be lost.

- Encryption

  The Riposte message store and the swap file are encrypted. This is described in *OPS Architecture Specification* section 4.6 *Encryption Software*.

- Digital signing

  Standard public key technology is used for digital signing.

FUJ00001742

Under public key technology, protected messages are digitally signed by a private key and validated using the private key's matching public key.

Transactions are digitally signed at the post office and then verified at the harvesting agent that takes the transactions from the correspondence servers, or vice versa. Complete files of transactions may be signed before transmission to Client hosts.

After a smart card transaction has been processed by the agent, a digitally signed acknowledgement is sent back to the outlet from the correspondence server.

NBS transactions are protected by digitally signing (and verifying) all request, authorisation and confirmation messages exchanged between the counter and the NBS agents. For messages generated in the outlet, the digital signature is applicable to the content of the message, including the PIN block (if applicable) and encrypted data values (if applicable). For messages generated in the Data Centres, the digital signature is verified when it is returned to the counter. Transaction messages that fail the signature verification will be cause a 'decline' to be returned, and an associated application event will be recorded in the System Log files, for onward forwarding to the Security and System management authorities.

- Message content encryption and MAC protection

  For NBS transactions, PIN numbers entered by customers to authorise transactions, and other sensitive data (card details), are encrypted before being digitally signed and passed to the correspondence server.

  The agent then interacts with the HSM for the translation of the encrypted PIN value between the Horizon and NBE domains. In addition, the HSM provides MAC authentication and verification and encryption of sensitive data in messages exchanged with the NBE.

5.1.4 Office Platform Service Domain

Generalised API for OPS/TMS, section 6.2 *Security,* describes aspects of the security of the Office Platform Service.

5.1.4.1 Authentication of counter PCs

There is normally no permanent link between counter PCs and the Data Centres. Where the outlet is connected by ISDN, the Gateway PC in each outlet periodically establishes a link with a Data Centre, during which all outstanding messages in either direction are exchanged. Although this mechanism provides a satisfactory level of authentication for ISDN-connected outlets, it is not available for use if outlets are connected via satellite or permanently connected via Frame Relay.

5.1.4.1.1 Virtual Private Network

Horizon uses a mechanism known as Virtual Private Network, or VPN, which provides an encrypted tunnel between counter PCs within an outlet, and the counter PC and Data Centre. It is described in section 5.4.

FUJ00001742

5.1.4.2 Authentication of new counter PCs

The POLO process described in *OPS Architecture Specification,* section 4.6 *Encryption Software*, assumes that the counter PC is a fully accredited member of the Horizon network, and contains information that may be used to support these authentication processes

The method by which new counter PCs are added to the network is outside the scope of this document.

FUJ00001742

5.2 Access Controls

TMS access controls are provided at the counter by the Riposte Message Server and Windows NT. The access controls provided to OPS by the Riposte Desktop software are described in *OPS Architecture Specification,* section 5.4 *Usability Features*, and *Generalised API for OPS/TMS,* section 6.2.2 *Menus and applications.*

5.2.1 Counter Access Controls

Riposte uses the access control facilities of Windows NT to ensure that it is the only system service able to create messages in the message store. Each message records the identity of the user currently logged onto Riposte at the time the message was written.

5.2.2 Central Access Controls

Standard NT facilities are used to control access to correspondence servers and agent servers, using secure IDs. The BIOS on all Windows NT platforms is configured to prevent booting from the diskette or CD drives. An administration password is used to control future access to these configuration features.

An access token is obtained during user authentication and is passed to the user's first process. Subsequent processes are created by this first process, and the token is passed to them as they are created. When the user tries to access any system resource, his access token is checked against the security attributes of the object. If the two match, the user is permitted to access the object. This process requires the use of NTFS which supports access privileges alongside every file.

5.2.3 Network Access Controls

The Data Centres are contained within a secure physical boundary, and most systems within them are part of the general domain.

There are External Interface Gateways Domains. These are accessed by external bodies, and hence are located on separate LANs that are separated by a firewall from the main Data Centre LAN.

Firewalls are used to protect parts of networks from one another, by only allowing traffic to flow between a defined set of network end points on either side of the firewall.

FUJ00001742

5.3 Cryptography

Cryptography is used for five purposes.

- To authenticate messages.
- To protect data on communications links that pass outside the control of Pathway or its suppliers or customers.
- To protect individual messages from their creation to use.
- To protect the confidentiality of data held in the system.
- To protect sensitive data input via a PIN pad.

Pathway provides a combination of measures to protect both the integrity and the confidentiality of types of data held or transmitted within Horizon.

The following types of encryption are used:

- Symmetric encryption
- Asymmetric encryption, used for sealing messages for integrity, not for data encryption
- Digital signatures, similar to asymmetric encryption, in that a two-part key is used
- One-way encryption, used to encrypt data such that it can never be decrypted, used for verifying secret data such as passwords.

FUJ00001742

5.4 Virtual Private Network

The Horizon Network utilises the Utimaco SG Virtual Private Network (VPN) product which provides encryption over the Wide Area Network (WAN), between outlets and the Data Centre, and on the Local Area Network (LAN) within an outlet.

5.4.1 VPN on Outlet LANs

Where there are more than two counters (including the Gateway PC) at a Post Office, VPN is employed on the LAN. The Utimaco SG VPN product is installed on all counters (where there is more than one counter position) and provides for encryption and decryption for all data transferred between counters, and between counters and the Gateway PC.

5.4.2 VPN between Post Office Outlets and the Data Centres (WAN)

5.4.2.1 Data Centre Inbound traffic

All traffic destined for the Data Centre from any counter position at an outlet is transferred via the Gateway PC. The Utimaco SG VPN product resident on the Gateway PC encrypts the data prior to delivering it to the WAN communications card within the Gateway PC. The WAN communications at an outlet can be either an Eicon ISDN card or an Ethernet card (for satellite-connected outlets). At the Data Centres there are a number of VPN servers which decrypt the inbound data packets prior to forwarding them to the correct destination (correspondence servers and Tivoli servers).

5.4.2.2 Data Centre Outbound traffic

For outbound traffic from the Data Centre, the VPN servers encrypt the data from a correspondence server to an outlet, and the Gateway PC decrypts the data packet. If the data packet is destined for a counter, the data is encrypted again for the local LAN.

5.4.2.3 Outlet Inbound connections

The Gateway PC at each outlet is protected from malicious incoming connections by the use of SG VPN. The SG VPN software rejects any network connections that it cannot successfully decrypt.