FUJ00243143

Document Title:	End to End Application Support Strategy
Document Reference:	SVM/SDM/PRO/0875
Release:	Release Independent
Abstract:	Defines the operational responsibilities of the units involved in the end to end support of the HNGX solution software in relation to each other.
Document Status:	APPROVED
Author & Dept:	Adam Woodley
External Distribution:	None
Security Risk Assessment Confirmed	YES, none identified

Approval Authorities:

Name	Role	Signature	Date
Steve Bansal	Fujitsu Senior Service Delivery Manager	See Dimensions for record	

See HNG-X Reviewers/Approvers Matrix (PGM/DCM/ION/0001) for guidance on who should approve.

FUJ00243143

0 Document Control

0.1 Table of Contents

FUJ00243143

FUJ00243143

FUJ00243143

0.2 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1	03/09/2010	First Draft	N/A
0.2	26/11/2010	Changes as a result of comments Title change to reflect focus on HNGX application support not infrastructure support. Definitions changed to use ITIL® terminology Defined generic 1st – 4th line support qualities Defined 2nd line "virtual team" to explain the lack of a dedicated 2nd line team in RMGA Added definition of Peak response category advice after investigation.	N/A
0.3	18/07/2011	Minor changes as a result of comments Change of document name to reflect content	N/A
0.4	28-Jul-2011	Further change of document name from "Process" to "Strategy"	N/A
1.0	28-Jul-2011	Approval version	N/A
1.1	10-Aug-2020	Revisions as a result of general review by owner.	N/A
1.2	01-Sep-2020	Minor changes as a result of comments	N/A
2.0	23-Sep-2020	Approved Version	

0.3 Review Details

See HNG-X Reviewers/Approvers Matrix (PGM/DCM/ION/0001) for guidance on completing the lists below. You may include additional reviewers if necessary, but you should generally ***not exclude*** *any of the mandatory reviewers shown in the matrix for the document type you are authoring.*

Review Comments by :	
Review Comments to :	Adam Woodley and Post Office Account Document Management

Mandatory Review	
Role	Name
Fujitsu Senior Service Delivery Manager	Steve Bansal
Fujitsu Senior Commercial Manager	Helen Venters; Post Office Account Commercial Mailbox
Service Architect	Phil Boardman
Application Lifecycle Manager	Graham Allen
Optional Review	
Role	Name
Service Architecture Manager	Alex Kemp
Security Architect	Dave Haywood
MAC Team	Sandie Bothick
Online Services SDM and Service Introduction	Sonia Hussain
Information Security Manager	Jason Muir

FUJ00243143

Document Manager	Matthew Lenton
Systems Management Manager	Jerry Acton
POA Delivery Executive	Dan Walton
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	Name

0.4 Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 *(DO NOT REMOVE)*	5.0	03 June 2009	RMGA HNG-X Generic Document Template	Dimensions
SVM/SDM/SD/0004			End to End Application Support Strategy	Dimensions
SVM/SDM/SD/0005			Application Support Service (4th line): Service Description	Dimensions
SVM/SDM/SD/0006			Systems Management Service: Service Description	Dimensions
SVM/SEC/POL/0005			Community Information Security Policy for Horizon & Horizon Online	Dimensions
SVM/SDM/PRO/0018			POA Operations Incident Management Procedure	Dimensions
DES/APP/DPR/0008			Obfuscation of Counter / BAL-OSR data for 4LS	Dimensions
DEV/GEN/TEM/0009			Support Guide Template	Dimensions

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

0.5 Abbreviations

Abbreviation	Definition
AD	Applications Division
AMCS	Application and Multi Cloud Services
BIF	Business impact forum
CET	Counter Eventing Team 1st line team monitoring PO counter events
COTS	Commercial off the shelf. COTS purchases are alternatives to in-house developments.
DPA	Data Protection Act
HDI	Help Desk Interface
ITIL®	Information Technology Infrastructure Library

FUJ00243143

Abbreviation	Definition
KB	Knowledge Base
LST	Live System Test
MAC	Major Account Controller Team 2nd line support team
MO	Model Office
MSS	Management Systems Support
PCI	Payment Card Industry
Post Office ISD	Post Office IT Service Desk 1st line support team
Prescan	Takes place prior to allocating incidents to support team members. Ensures that any incident which can be turned round quickly (e.g. known error, insufficient evidence) does not wait for the attention of a diagnostician who may be working on other duties.
QC	Quality Centre. Incident logging system used by test teams within RMGA
Peak	Fujitsu services incident and release management system
POA	Post Office Account
SRR	Service readiness review
SMC	Systems Management Centre. 1st line team providing the Systems Management Service.
SMG	Systems Management Group
SSC	Software Support Centre. 3rd line application support
TfSNow	Triole for ServiceNow. Incident logging system used by first line and second line support units

0.6 Glossary

Term	Definition
Event storm	A system event that is repeating itself so quickly that it causes performance issues with the server on which it runs or the event recording infrastructure.
Incident	An 'Incident' is any event which is not part of the standard operation of the service and which causes, or may cause, an interruption or a reduction of the quality of the service.
Normal service operation	Service operation within Service Level Agreement (SLA)
Workaround	Method of avoiding an incident or problem, either from a temporary fix or through access to an alternative service. A method to bypass a recognized problem in a system. A workaround is typically a temporary fix that implies that a genuine solution to the problem is needed.
Resolution	Resolution is the action taken to repair the root cause of an incident or problem, or to implement a workaround.

FUJ00243143

0.7 Changes Expected

Changes

0.8 Accuracy

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

0.9 Security Risk Assessment

Security risks have been assessed and it is considered that there are no security risks relating specifically to this document.

FUJ00243143

1 Introduction

This document provides an overview of the support strategy and the interfaces between support units which provide application support for HNG. It also describes qualities, obligations and objectives expected from those units.

This document is effectively a high level design for application support. It does not attempt to define the low level procedures by which each support unit will deliver the qualities described here.

1.1 Exclusions

This document excludes any detail of the management of operational problems and all hardware related incidents. It also excludes details of the business impact forum which is a governance process and not a support process.

This document also excludes any specifics of software and reference data distribution support (SMC, MSS and SMG).

Test teams do not interface formally with the support chain when they are testing a future release. Their interface is with the release management process and directly with the development unit resolving faults in the release. Details of these interfaces are excluded from this document. Live System Test are an exception. When testing the resolution for an outstanding problem with the live release they still fall within the release management process. In some cases they may notice a new issue within the live release when they will generate a new Peak. For the purposes of this document they are treated as a 3rd line unit under these circumstances.

The remainder of this section describes extant situations and expectations which influence application support.

1.2 Support chain

There are normally 4 levels of support within a conventional support organisation:

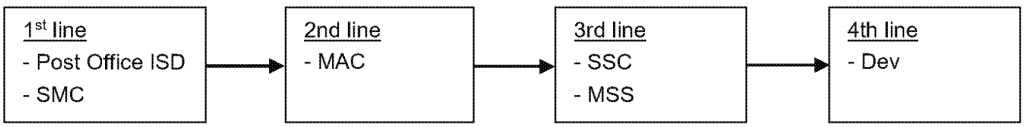

The support strategy expects that incidents will be raised by users and then passed through the chain of support units until a resolution can be supplied to the user. It is important that an incident starts at 1st line and then follows each stage of the chain as appropriate. This ensures:

1. The incident is quickly defined and logged
2. An initial response is given
3. Priority is correctly evaluated
4. The correct skills are applied such that a resolution is supplied quickly
5. The call is correctly recorded, auditable and relevant metrics can be produced.

As incidents move from left to right across the support chain they become:

- More difficult to resolve
- More time consuming to resolve

FUJ00243143

- The training level and cost of the staff resolving the incident rises
- Tooling and supporting infrastructure costs rise

Support costs and timescales for resolution increase as the incident moves to the right. Hence the effort spent "moving support to the left". Ensuring that the incident is resolved as early in the chain as possible reduces the cost and increases customer satisfaction (assuming a first time fix is achieved).

1.3 Logging systems

For historic reasons POA use two logging systems. 1st and 2nd line groups use a system called TfSNow which is primarily an incident management system. All other support, development, test and release management teams within POA use a system called Peak (no it is not an acronym, it's a name!) for incident, problem and release management. The two systems are closely coupled using an interface called the Help Desk Interface (HDI).

An incident is passed over the HDI as it moves between 2nd and 3rd line support groups. When the incident has been "moved" in this way it is no longer being actively progressed in the sending system.

Updates made in either system can be reflected in the other, this allows additional information to be passed between support groups while the call is not active in one of the systems.

The use of two logging systems imposes some restrictions on incident processing:

1) The HDI must be monitored by 2nd and 3rd line units to ensure that it is working correctly.

2) Attached evidence cannot be moved between the two systems. This means that evidence files can only be passed by reference (within the information logged).

3) TfSNow has a limit of 4000 characters within a single update which sometimes results in lost information.

The requirement for detailed evidence sets starts at 3rd line. This is where the interface has to exist between the two logging systems. Using Peak from this level ensures that evidence can be easily passed between 3rd to 4th line support groups and that facilities such as obfuscation and encryption (features of the Peak system) can be applied to the evidence files.

1.4 Restoring normal service

It is incumbent on this support route to restore normal service operation as quickly as possible and minimise the adverse effect on business operations. The restoration of service is considered to be completed once it has been documented and communicated by a support team to the end user who raised the incident. This does not necessarily imply that a change has been made since an incident as perceived by the end user may not be impacting service operation. Restoration of service has targets dependent upon the priority assigned to the incident. These priorities are defined in section 8 of this document.

Once normal service operation has been re-established it may be necessary to resolve the root causes of incidents (Problem management in ITIL terms). Although it is expected to resolve the problem as quickly as possible this strategy is focused on the resolution of the problem rather than the speed of resolution and as such no targets are defined.

1.5 Security restrictions

Access to certain parts of the system or certain support tools must be restricted in order to conform to DPA or PCI data security standards. The application of these standards within the POA support community is defined by the POACS Security team in various documents which are not reproduced here. The following describes restrictions that exist and impact the support community.

FUJ00243143

1.5.1 PO counter access

Direct access to live Post Office counters is currently only available to the SSC for retrieval of log files. It depends on access to the SSN servers and the use of a Secure Shell application called RCA Client. Data visible using this access is subject to DPA restrictions and thus this access cannot be allowed to offshore units.

1.5.2 Obfuscation of logs

Certain log files must be processed to obscure personal details that exist within them (see section 2.6.1 for details) before they can be passed to support teams outside the European Union. As new log files are generated by system enhancements development units need to be aware of the DPA and ensure that information in any new log files is either benign in DPA terms or that appropriate changes are made to the obfuscation tool. The POA CS Security team can advise on DPA issues.

1.5.3 Access to and repair of user data

Access to POA data has been secured using the standard security principle of separation of duties. Separation of duties ensures that an individual can not complete a critical task by themselves. For example: someone who submits a request for reimbursement should not also be able to authorize payment. An applications programmer should not also be the server administrator or the database administrator - these roles and responsibilities must be separated from one another.

In POA the separation of duties principle has been implemented by ensuring:

- Development units cannot have update access to any of the system data.
- Database administration functions are carried out by IS staff.
- Data repair is carried out by SSC staff, following Post Office approval.

The DPA requires that access to personal data remains within the European Union and PCI data security standards mandate physical security restrictions must be applied where update access is allowed to user data. Currently the only units which fulfil all these requirements for data access are the SSC and Unix Operations. The responsibility for data correction is vested with the SSC although Unix Operations sometimes act under SSC authorisation.

1.5.4 Audit servers

The audit servers provide a full audit trail of all information on the HNG system. In order to ensure that this audit trail is irrefutable the teams which have the ability to change data (i.e. SSC) must not also have the ability to change the audit trail. For this reason, Audit server 3rd line support rests with the Audit development team and not the SSC.

2 Qualities and principles

This section describes qualities and principles that apply to all support groups working within POA. It is expected that each support group will define the processes and procedures that implement these principles.

2.1 Accurate call logging and updating

All support groups are expected to receive incidents passed from other support groups and to ensure that any incidents so received are maintained on the relevant call management system:

FUJ00243143

- 1st to 2nd line TfSNow
- 2nd to 4th line Peak

The initial description logged should include a full explanation of the problem, the accurate recording of any references supplied and the actions that were taken which caused the problem to occur.

2.1.1 Updates

When relevant additional information has been made available that information should immediately be added to the incident to ensure it reaches the diagnostician currently working on the call.

The 1st line agent making such updates (in TfSNow) should ensure they are made as a HDI Comment (non-HDI updates are not transmitted over the HDI to the investigating unit in Peak).

The 2nd – 4th line diagnostician processing the incident should ensure that relevant updates are applied on a regular basis appropriate to the priority of the call. These Peak updates should be made using a numbered response category to ensure they go over the HDI and become available to 1st line should the end user request an update. These "HDI updates" are often repeated verbatim to Post Masters when they ring in asking for an update so diagnosticians should ensure that any HDI updates they make are appropriate to the audience who will receive them (i.e. the end user).

2.1.2 Closure

Ensure that every incident reported has a resolution recorded on the logging system and that the resolution is acceptable to the end user. Where the final response has been entered on Peak it will be returned to the 1st line call management system (TfSNow) for communication to the end user and final closure. The description should include a full explanation of the resolution. It is acceptable for any level of support to agree closure but where that agreement has been made outside 1st line support it should be recorded in the incident to ensure that 1st line do not make unnecessary calls to the end user.

2.1.3 Withdrawal

Support units have the facility to withdraw an incident, resulting in its automatic closure on the active incident management system. This should be rarely used and any usage must be communicated in advance to the team currently processing the incident to ensure that they stop investigation.

It is expected that a request to withdraw an incident will originate from the end user who raised it. When this is not the case then governance should be in place at each support unit to ensure the facility is not misused and that the end user is informed.

2.2 Filtration

All support units are expected to resolve as many incidents as possible and only pass on those that are relevant to the next line of support, hence it being called a filtration process.

The first level of support to discuss the incident with the customer should close any incident for which the problem and resolution is already known to the support community. When this does not happen the incident is deemed to be a filtration failure.

When an incident is closed the diagnostician entering the closure also applies a category for the closure. Categories such as, insufficient evidence, published known error and user error indicate that the call should have been filtered rather than sent on to the next level of support. See incident closure categories for a list of closures categories that are considered to be filtration failures.

FUJ00243143

2.3 Review

It is essential that each support group carries out regular reviews of incidents returned as filtration failures and reports on this OLA and the actions being taken to address failures.

2.4 Timely transfer of incidents

All support groups must ensure that any incidents that will require the attention of another level of support are passed in a timely manner. The exact timings are detailed in section 8 of this document.

The timings vary according to the total time allowed for resolution of the incident in the contract between Fujitsu Services POA and the customer. These timings will therefore be dependent on the priority of the incident, with (for example) less time allowed for an "A" priority call than will be permitted for a "D" priority.

These timings should not be used without consideration being given to resolution. Since the requirement is the resolution of a call (not simply its transfer within timescales) then it is acceptable for a support group to retain the call past its normal time if it is confident that it can provide a resolution within the maximum time allowed for the incident.

2.5 Metrics

Each support level needs to produce call metrics to describe the service they supply and the quality of that service. These should be made available to all other support units in a shared location and serve as a guide to the efficiency of all parts of the support chain. These metrics should form the input to regular review meetings within each support unit to improve service.

Metrics produced should include:

2.5.1 Incident counts

1. Count of incidents input to the support unit per day (daily input)
2. Count of incidents closed by the support unit per day (daily output)
3. Incidents remaining in the support unit at the end of the day (daily work in progress)
4. Count of incidents input to the support unit in the last 34 days* (monthly input)
5. Count of incidents closed by the support unit in the last 34 days* (monthly output)

*A figure of 34 days is used to prevent the figures being affected by months which start end and / or end with a weekend. Such months make comparisons unrepresentative.

2.5.2 Filtration percentage

Filtration is an important measure of the efficiency of a support unit. It shows the number of calls transferred that should not have been passed onto the next level of support. This is expressed as a percentage of the total calls the support unit closes.

For example:

100 calls closed by 2nd line in period, of which 25 calls returned from 3rd line in "black mark" categories:

2nd line filtration rate to 3rd line = 75%

This shows that the 3rd line support group had to deal with an additional 25 calls that should have been stopped by previous support units.

FUJ00243143

See incident closure categories for a list of closures categories in use and which constitute filtration failures (AKA black marks).

Whilst 100% filtration is generally expected and strived for a 5% deficit (i.e. 95% filtration) on a rolling three month basis is accepted by all parties.

Filtration metrics are required for:

1) Total number of filtration failure calls - broken down by Peak category

2) List of Peak failure reference numbers

2.5.3 Incident life in team

A measure of the performance against incident closure / transfer time scales:

- The average time to transfer / close an incident by priority per team. Average time an incident is within a team split by final priority on transfer.
- The average age of all open incidents in the team by priority.

2.6 Evidence gathering

All support groups gather and evaluate evidence as part of the problem solving process. It is also a responsibility of all support groups to ensure that relevant evidence is referenced or attached to all incidents that are passed to another support group.

Where relevant evidence has not been supplied it is highly likely that the support group that the call is passed to will just return the call asking for (and detailing) further evidence. Such a response counts as a "black mark" for filtration purposes and will be reflected in the sending support group's filtration figures. Specifically excluded from this measure are instances where:

- Although the evidence was inadequate, no documentation existed describing the relevant evidence required.
- Occasions where the evidence required was unobtainable.

What constitutes "relevant" evidence can be determined by:

1. Examination of the support guide for the area of the system being investigated.
2. A search of existing knowledge entries.
3. Examination of high and low level design documentation.

It is an inconvenient truth that 1st line support groups do not use the same incident logging system as the rest of the support chain. When passing incidents from 2nd to 3rd line, evidence files can only be supplied by quoting a reference to the file. When passing evidence via Peak any relevant evidence files are attached directly to the incident.

2.6.1 Obfuscation

There are requirements within the data protection act for the processing of personal data. It is essential that all support groups are aware of this restriction and do not allow the transfer of any evidence files offshore if they contain personal data. Current offshore units are:

1. 4LS for Counter and BAL-OSR
2. SMC

Areas currently identified as potentially containing personal data (see DES/APP/DPR/0008):

FUJ00243143

- Counter OSR / BAL message log file
- Counter application log file
- Database exports (e.g. CSV exports – Message Journal Exports)
- Screen captures of live system data*
- Audit data extracts (content of message journal)*

**Not handled by obfuscation tool*

In order to allow the use of such information offshore an obfuscation tool has been developed for use on the log files that are known to contain sensitive information before passing to any external support team. The tool has now been integrated into Peak. Information on its use can be found in the FAQ section of Peak.

If anybody suspects that personal data is present in other log files they should raise an incident with the POA CS Security team so that it can be checked before the incident is sent to an offshore support group. It may be necessary to process the information onshore or have the obfuscation tool enhanced.

2.7 Removal of duplication

All support groups should ensure that they do not pass to the right duplicate incidents, i.e. incidents which are repetitions of an incident which has already been passed to the next line of support. They should either retain the duplicate incidents within their own call logging system or close them as duplicates:

a) 1st line units retain duplicates under a "master call" and to ensure that when the resolved incident is received from 2nd line, the end user is contacted and duplicated calls incidents closed within TfSNow.

b) 2nd – 4th line support units normally immediately close the incidents as duplicates because they add no value to the support process at these levels. This results in the incidents being returned to 1st line (TfSNow)

Duplicate incidents are only acceptable where the symptoms reported by the customer did not match the symptoms recorded in the original incident, and which therefore could not reasonably have been identified as a duplicate.

Failures will be reflected in filtration figures where the incidents are closed in the "duplicate incident" category in Peak by subsequent support units.

2.8 Training

Training on new facilities added to the system will be provided by the architects and development units that design and implement those facilities. A checkpoint for this will be included in the SRR for each new release. This training takes the form of one or more of:

- Classroom training (workshops)
- One-to-one training
- Training collateral
- Updated support guides

It is expected that the training provided is then relayed from right to left through the support chain. The form and content of the training is likely to change as it is modified to ensure it is suitable for the new audience at a different level of support.

FUJ00243143

It is the responsibility of the unit receiving the training to ensure that it is completed and fit for purpose. Inadequate training on new facilities will be reflected in reduced filtration figures for the receiving unit so it is in their best interest to ensure that the training material is timely and appropriate for its audience.

2.9 Knowledge base maintenance

All support units are expected to search, enter, update and maintain the information in the support knowledge base. Criteria detailed in section 11.

Because of the time constraints applied to 1st line units there should be no requirement for them to raise a KB for every new incident. It is a requirement that 2nd line ensure that a KB is generated for any new incident encountered.

2.10 Support tools

Support tools normally result from a requirement placed on development for a new release or are generated by 3rd and 4th line support groups when necessity demands. In either case these tools must have the following qualities:

1. They must be tested on a test rig (normally LST) before they are deployed in the live environment.
2. It must be possible to restrict their use on an individual basis. This is to ensure that DPA / PCI rules can be enforced dependent upon the person (not just the team) using the tool.
3. They must be self documenting or have suitable documentation written in the form of a work instruction.

Such support tools will always be preferable to various diagnosticians generating manually crafted scripts or SQL which need individual testing, are rarely documented and often represent duplication of effort.

Support tools generated outside development units are not expected to be subject to the CP process since this represents an unacceptable overhead to this type of tool. It is expected that any support unit generating support tools will put in place a process or work instruction to ensure:

1. No duplication of effort takes place
2. The effort required to write the tool does not exceed any benefit gained from it.

3 1st line support

1st line support units comprise Post Office IT Service Desk (ISD), and SMC for system monitoring. From the point of view of 2nd line support they can be treated as one unit since incidents arriving at 2nd line support will always have passed through one of these units first.

1st line support log incidents by directly interacting with the user or from monitoring systems. They clearly document incident symptoms based on customer perception or observed alerting information. The Post Office ISD are trained to the same level as the user they should resolve all issues where the cause is user training or environment. 1st line resolve incidents by the identification of knowledge base entries and the application of defined scripts (flowcharts). 1st line support provide the “touch point” with the end user (e.g. Post Masters or in some cases internal units such as Unix Operations) and are responsible for ensuring the end user is kept informed of progress of their incidents and taking any escalations from that end user. It is incumbent upon the subsequent lines of support to provide the information and escalation routes necessary for 1st line to do this.

FUJ00243143

3.1 Hardware calls

1st line support should filter all hardware calls and ensure they are routed to the correct supplier for resolution.

3.2 System monitoring

1st line support are also responsible for monitoring the live estate and taking corrective actions for all critical events seen. This role is currently fulfilled by:

1. SMC: Data centre event and schedule monitoring

For each event seen 1st line must check for the event in documentation (knowledge base and support guides) for the relevant system subsection and to take the documented action (if the required support tools are available). A new incident should be raised for each critical event that is not already documented and passed 2nd line support teams for action.

1st line support will also ensure that "event storms" are correctly handled. Where an event storm is being produced by a server, or by a subsystem in the POA solution AND where the cause of that event has already been documented, then first line should, where possible, take appropriate action to prevent the system becoming "swamped" with repeat events.

3.3 1st line obligations to 2nd line support

The section describes obligations inherent in the interface between 1st and 2nd line support units. These are in addition to the general obligations defined in section 2.

1) Filter all hardware calls and route them to the appropriate unit for hardware support. No hardware calls should be passed to 2nd line except in circumstances where Fujitsu Services POA need to be made aware of:

1. Hardware behaviour that can be influenced by application design changes

2) Ensure that any incident which requires investigation by 2nd line support is passed onto the call management system used by Fujitsu Services POA and assigned to the correct 2nd line support team.

3) Addition reports to the metrics defined in section 2.3:

1. Number of incidents passed to 2nd line by priority.
2. Time taken between receipt of the incident and transfer to 2nd line by priority.
3. Time between receipt of the incident by 2nd line and the resolution passed back to 1st line by priority.

4) Monitoring: Ensure that no critical events on the live system go unnoticed and ensure that appropriate action is taken for each event. In this context, a critical event is one which is logged as critical, and appears in red on the Tivoli screens and with a red marker in NT event logs.

No event calls should be passed to 2nd line without documentation having been consulted (Support guides and knowledge base).

5) Monitoring: Ensure that for any incident which has been resolved and passed back to TfSNow, the end user has been contacted and made aware of the closure.

4 2nd line support

2nd line support staff can be described as expert users of the system. They use the symptoms documented by 1st line to understand the error and then gather additional information / logs from

FUJ00243143

standard sources within the system to provide evidence of the environment and other factors present at the time of the error.

From their understanding of the system and the use of reference equipment that simulates what is available to the end user, 2nd line can devise procedural workarounds (authorised knowledge base entries) to alleviate incidents.

2nd line produce knowledge entries for all incidents using the symptoms collected by first line and where possible add a definition of the root cause problem and any solution or workaround found. They also generate scripts and procedures for 1st line, produce explanatory documentation and MI reports. 2nd line have a level of access to the system that allows the use of simple GUI support tools.

Some of the 2nd line responsibilities are being partially fulfilled by 1st line or 3rd line units providing a supportive 2nd line function as follows:

Initial production of knowledge entries	SMC, SSC
Evidence gathering	SSC
Operational changes	SSC
Procedural workarounds	SSC
KB reference applied to Peak	SSC
Priority assessment and change	SSC

4.1 2nd line obligations to 1st line support

The section describes obligations inherent in the interface between 2nd and 1st line support units. These are in addition to the general obligations defined in section 2.

1) HDI monitoring: Where updates are made to the calls which are of relevance to 1st line then the second line support unit will ensure that these updates reach the first line call logging system (currently TfSNow).

2) To ensure that any resolutions or workarounds that are passed back to 1st line have been tested. The exception to this rule is the case where resolutions are being passed specifically to be downloaded to the POA test rigs for testing.

4.2 2nd line obligations to 3rd line support

The section describes obligations inherent in the interface between 2nd and 3rd line support units. These are in addition to the general obligations defined in section 2.

1) To ensure that the priority of any incident is assessed and recorded correctly. No calls should be passed to 3rd line support whose priority does not conform to the specification defined in section 8.

2) No incidents should be passed to 3rd line without a valid KB reference recorded on TfSNow as a HDI comment, or recorded in the Peak references.

5 3rd line support

3rd line support groups within POA include:

SSC – 3rd line support for POA written application code.

MSS – 3rd line support for software distribution and event management

FUJ00243143

3rd line support staff apply analytical skills to the symptoms and evidence gathered by 1st and 2nd line and undertake in-depth investigation into incidents. They have detailed knowledge of the system based on documentation and source code inspection.

Trained on operating systems, COTS packages that underlie the application and the coding languages used within the application, they are also expected to self train by examination of support guides, design documentation written for the components of the end user application. They will also have access to development and package management tools to allow the production of specialised diagnostic code, scripts and support tools.

It is incumbent upon the 3rd line support unit to produce a work around and on 4th line to produce the final code solution to any software problem. This does not preclude the production of a workaround by other units or negate the requirement for 4th line to provide assistance in the generation of a workaround.

The SSC are responsible for the implementation of any workarounds that require data changes to the live system. They are the only unit who can be granted with authorisation and sufficient physical security controls to perform this function.

5.1 3rd line obligations to 2nd line support

The section describes obligations inherent in the interface between 3rd and 2nd line support units. These are in addition to the general obligations defined in section 2.

1) Where updates are made to the calls which are of relevance to 1st line or 2nd line then the 3rd line support unit will ensure that these updates reach TfSNow.

2) To ensure that any resolutions or workarounds have been tested and have been correctly authorised via a TfSNow Change Task. The exception to this rule is the case where workarounds or resolutions are being passed specifically to be downloaded to the POA test rigs for testing.

5.2 3rd line obligations to 4th line support

1) To ensure that any incident which requires investigation by 4th line support is assigned to the correct team dependent on the specific product in which the incident has occurred.

2) To ensure that the priority of any incident is assessed and recorded correctly.

3) To ensure that for any incident passed to 4th line support, the exact area of the problem has been identified, and wherever possible a workaround already produced.

4) To ensure that for any code error a probable solution is indicated prior to passing to 4th line support, and wherever possible, the possible solution has undergone limited testing.

6 4th line support

Have intimate knowledge of narrow areas of the system and are ultimately responsible for the production of permanent fixes to repair the root cause of an incident or problem in the live application. Trained in development languages and coding techniques there is often overlap between 4th line support and development roles. 4th line assist 3rd line with workarounds and resolution of incidents, produce test scripts for testing of code fixes and unit test those fixes.

4th line support within POA is supplied by various onshore and offshore units. Since 4th line own the interface with development (and the functions are often vested in the same people) they are also tasked with ensuring that various development obligations to the support groups are met, which includes ensuring that evidence is obfuscated where required.

FUJ00243143

6.1 4th line obligations to 3rd line support

The section describes obligations inherent in the interface between 4th and 3rd line support units. These are in addition to the general obligations defined in section 2.

1) To ensure that the incident reported is correctly resolved and the resolution recorded on the Peak system and the incident and resolution passed back to 3rd line. Where appropriate this should also contain the method of recreation of the problem. Add JIRA's related to the resolution as Peak references.

2) To ensure that the incident is resolved within the total time allowed by the contract between the customer and Fujitsu Services POA Account. Specific targets for timescales are documented in section 8 of this document. However, in most cases the provision of the fix is at the discretion of the Release Management Forum, and the target for the provision of any fix therefore is as specified by that forum.

3) To ensure that any resolutions or workarounds that are passed to 3rd line have been tested. Where they are being recommended for application to the live system they must have been correctly authorised via the TfSNow Change process.

4) To ensure that when a resolution is produced the baseline reference is added to the relevant KB entry that describes the problem.

5) To ensure that 3rd line is supplied with training and documentation relating to new releases of the POA Account solution in sufficient time to enable 3rd line staff to become familiar with the product prior to its release, and in sufficient time to enable 3rd line to adequately train other support staff.

Preferably this knowledge transfer should be a continual process during the course of development but an adequate timescale is 6 weeks prior to data centre release or model office for counter releases.

6) Ensure that support guides have been written or updated for all new facilities in HNG. An outline description of the contents of a support guide is given in section 10 of this document

7) To ensure that 3rd line support groups are supplied with read access to all source code developed within POA Account development prior to the release to live of that component.

8) In addition to the metrics defined in section 2.5: to ensure that the following figures are available to other support units on demand.

1. Total number of calls where resolution has been deferred to a future release
2. Counts of deferred calls by future release.

7 Definition of incident priorities

The definition of an incident priority changes depending upon the stage of the life cycle. Full details of POA incident management process life cycle can be found in SVM/SDM/PRO/0018.

7.1 Support priorities

Support priorities are used from the point when the customer initially logs a call until a work around has been achieved.

Priority A Business stopped	A Post Office unable to trade, unable to process any business. 2nd – 4th line application support units do not provide cover for Post Offices down outside normal working hours (These hours are defined in Section 2.2 of SVM/SDM/SD/0004). A central system failure which will result in a number of Post Offices being unable to process work.

FUJ00243143

	Causes significant financial loss (as agreed between POL and POA Customer Services) Results in data corruption or unrecoverable data loss. Outage of key infrastructure
Priority B Business restricted	A Post Office restricted in its ability to transact business e.g. 50% of counters unable to trade or trading with restricted business capability. Has an adverse impact on the delivery of service to a number of end users. Causes a financial loss that impacts POL and/or POA reputation (as agreed between POL and POA Customer Services) If a PCI Major Incident process is invoked.
Priority C Non critical	A Post Office working normally but with a known disability, e.g. an interim solution (workaround) has been provided. Has a minor adverse impact upon the delivery of service to a small number of end users
Priority D Non urgent	Insignificant and usually cosmetic error, either a trivial documentation error or spelling error on the system. Single-user affecting incidents on non key functionality Non user affecting incidents NOTE: This is the default priority if 1st line do not provide an "POA severity" in the TfS incident.
Priority E Internal incidents	Internal incidents that do not affect a Post Office.

7.2 Development priorities

Is it envisioned that development units may want to redefine these incident priorities once a satisfactory work around has been agreed and documented with 3rd line support and POA customer service. This is the stage where:

1. A code or documentation fix is required.
2. The release management process cuts in (defining target timescales).
3. Support definitions of priority become meaningless.

Within development groups the sequence of addressing peaks tends to be based on the requirement for understanding and resolving peaks (returning to standard service) and then fixing peaks:

1st priority - analysis of new peaks

2nd priority - fixing targeted peaks

3rd priority - analysis of peak backlog or deferred peaks

4th priority - preparing fixes for untargeted peaks

8 Definition of incident timescales

FUJ00243143

It is expected that although calls may enter the support chain at a high priority, in the majority of cases support will produce a resolution for the incident, and at that stage the priority of the call will be reduced in order to provide 3rd and 4th line support with sufficient time to allow a root cause analysis and possible code fix.

Since it is incumbent upon the 3rd line to produce an incident resolution and on 4th line to produce the final code solution to any software problem, for the majority of its "life" any incident should be with one of those two units.

Once a resolution has been generated for an incident the root cause fix may be deferred to a later release of the software and the targets specified below no longer apply. The POA release management process takes over at this point.

8.1 Full life times

Target times to resolve software incidents are as follows:

A Priority	2 working days
B Priority	4 working days
C Priority	7 working days
D Priority	28 working days

Note that these are targets and that no formal SLA or penalties apply to these timescales

8.2 Transfer times

The target times within each line of support are show below. They represent the maximum time each team should retain an incident before transfer to the next level of support. These times are measured from the **time the incident was logged**.

	1st to 2nd	**2nd to 3rd**	**3rd to 4th**	**4th resolution**
A priority	30 mins	2 hours	1 day	2 days
B priority	1 hour	1 day	2 days	4 days
C priority	1 hour	2 days	4 days	7 days
D priority	1 hour	7 days	14 days	28 days

9 Peak closure categories

The KB column is used to indicate whether a KB should be raised (or amended). "Opt," indicates that it is left to the discretion of the person closing the call.

The Fail column has a "Yes" if use of this category is counted as a filtration failure (AKA black mark) against previous support groups.

Code	KEL	Fail	Meaning/Usage
60	Yes	No	**9.1.1 S/W fix released to call logger**

FUJ00243143

			Code fix has been tested and can be (or has been) released into the live estate.
61	Yes	No	**9.1.2 Build fix released to call logger** Build fix – i.e. configuration, registry edit, incorrect DLL loaded etc - has been tested and can be (or has been) released into the live estate.
62	Opt.	Yes	**9.1.3 No fault in product.** Indicates that the product is working to specification. No changes are required in software code, scripts, hardware, documentation, work instructions or training plans. Really indicates that previous lines of support have completely mis-diagnosed the problem. (See also 66, 70, 94, and 98).
63	Yes	No	**9.1.4 Programme Approved. No fix required** Rarely used. Covers the case where there IS a fault in the product and this is acknowledged by both Fujitsu and POL, but the fault is there as a result of an agreed design specification, and Fujitsu would require POL to fund any correction. MUST NOT be used without approval from HNG programme manager or authorized representative.
64	Yes	Yes	**9.1.5 Published Known Error** Should only be used when the resolution of the problem is documented in a KB and does not require the call to be passed to this support group. When the KB is raised after the call is logged the call should be closed as Unpublished known error. To be used when there was already a KB in existence when the call was passed on. And the KB fully described the problem. And no changes have been made to the KB as a result of diagnosing this or other later calls. (c.f. cat 65). KB reference MUST be quoted in response text.
65	Yes	No	**9.1.6 Unpublished known error** To be used instead of 64 to close those calls where this support group now knows the problem (and optionally solution), but where no KB was visible at the time the call was passed on. New KB to be raised or old KB to be updated with every use of category 65. KB reference MUST be quoted in response text.
66	Yes	No	**9.1.7 Enhancement request** To be used when the error complained of is not a fault against the original specification, but Post Office agrees that a change is needed to avoid the error in future. CP number or other equivalent reference should be quoted in the response.
67	Yes	No	**9.1.8 Solicited Known Error**

FUJ00243143

			To be used to cover those cases where a call has been sent to the support group in response to a specific request in a published KB. Ideally it would only be used if it is clear in the call text that the previous support group have indeed spotted the KB and have sent the call in as a result. If there is no such indication, then category 64 should be used instead.
68	No	No	**9.1.9 Administrative response** Only to be used for closing calls which cannot be closed in a legitimate category for "administrative" reasons – e.g. incident incorrect changed by the system (Peak, TfSNow or the HDI); Test calls; Miss-routes; Double escalates; Unintended escalates etc. Not to be used as a catch-all for "unable to decide which category to use". See also 200 – "Withdrawn by user".
70	Opt.	No	**9.1.10 Avoidance Action Supplied** To be used when there IS a fault in the product (usually a one-off), but for whatever reason, there is no time or justification for fixing it in the current (or any future) release. Typically this will be used for migration or build problems or when the facility is to be withdrawn at a future release. Should include a KB reference in the response if there is even the slightest chance of a recurrence. MUST include a clear avoidance action that can be taken by the user / support on this or a subsequent occurrence.
72	No	Yes	**9.1.11 Duplicate call** To be used when two calls are discovered to relate to the same incident. (E.g. when both Post Office ISD and SMC event management report an error message). Not recommended for use when 2 calls for separate incidents can be traced to a single root cause (e.g. code error). In that case use a combination of appropriate code for the first incident and then published or unpublished known error for the second and subsequent incidents.
74	Yes	No	**9.1.12 Fixed at Future release** Genuine error in the product but cannot be closed in 60 or 61 because the fix will not be available until a later release. Applies to manuals / support guides etc as well. Release initially targeted to contain the fix should be quoted in the response text as well as the "target release" field. (See also enhancement request for use when target release is unknown).
90	No	No	**9.1.13 Reconciliation – resolved** Special category for closing reconciliation calls, Only used between SSC and Security Operations incident management.
92	Yes	Yes	**9.1.14 Suspected hardware fault** Use when the problem was caused by a fault in a piece of hardware or the network. Record the symptoms and cure in a KB to aid future diagnosis of similar problems.

FUJ00243143

94	Yes	Yes	**9.1.15 Advice and Guidance given** This code should be used as an alternative to "No fault" or "User Error" in those cases where it should be obvious to everyone that a product is working to specification – e.g. documented feature or symptoms described in a support guide. Can also be used to highlight cases where the end customer or someone in the support chain could benefit from further training in this area of the product.
95	Yes	No	**9.1.16 Advice after investigation** Similar to "Advice and guidance given" but used in the situations where it was not obvious that the system was working to specification. i.e. User documentation or support guides do not contain the information the advice given is based on.
96	Yes	Yes	**9.1.17 Insufficient evidence** Use when it is crystal clear from the KB or other published sources what evidence is required to accompany a call of this class, but that information has not been supplied (Check that the call post-dates the latest update of the relevant KB). The incident is returned to TfSNow for the 2nd / 1st line support teams to provide the evidence. See also special case use of 62 - "No fault"
97	Opt.	No	**9.1.18 Unspecified insufficient evidence** To be used when the problem cannot be resolved from the evidence supplied and more evidence is being requested, but there was nothing to tell the other support groups that this particular piece of data is required to investigate this class of call. KB should be raised or amended to give future guidance, if this class of evidence may be needed in future. A support guide update may also be appropriate. The incident is returned to TfSNow for the 2nd / 1st line support teams to provide the evidence.
98	No	Yes	**9.1.19 User Error** Close in this category when it is clear from the evidence that the end user has caused the problem by doing something incorrectly. The documentation or training available to the user concerned should specifically cover this incident – if it does not then consider using 94 instead. No KB is required – documentation should already cover the case.
100	No	No	**9.1.20 Route call to TfS** Special code which "closes" the call on Peak and causes it to be routed to the system used by AMCS NT / Unix Operations or Networks, without it registering as "closed" or "finished with" on their systems.
200	No	No	**9.1.21 Call withdrawn by user** Specific case of an Administrative Closure. To be used when the user (e.g. PM) or other support group specifically request that a call is returned to them without the

FUJ00243143

			currently assigned team doing any further work on it.

10 Outline contents of a support guide

Support guides are generally written by a combination of the architect / designer for the product and the developers of that product and are based on the document management template reference DEV/GEN/TEM/0009.

Although intended for the support community the support guide should be produced in time for the initial test cycles of a new facility. This will help testers understand what they are testing and also validate the contents of the support guide.

10.1 Overview of the facility

Often cribbed from relevant design HLDs (for which cross-references should also be provided)

- Purpose and functional overview
- The way in which it performs the function

10.2 Documentation

A list of document references, with titles, associated with the facility. This is intended give the support community a list of reading matter that can be used to self train.

- Table of document references and titles. This will certainly include POA written HLDs, LLDs and design notes.
- If the facility includes COTS products then references to any externally produced support or user guides (normally these will have been pre-registered in Dimensions).

10.3 Server definition

Define the servers and workstations which support will need to access in order to support the product.

- Which servers and / or workstations are involved in the delivery of the facility
- List any COTS products and version that should be installed that the facility depends on.
- Define the location of configuration files relevant to the facility or COTS product.

10.4 Diagnostics and logging

Define what log files are produced by the facility and where the support groups can find them. Can the level of information logged be changed, if so how? How are they interpreted?

- Location, configuration and layout of log files
- What other diagnostics are written
- How can the diagnostics be configured
- Supply samples of diagnostics

FUJ00243143

10.5 Errors and messages

This is intended to provide support with guidance on common messages or errors produced by the software. It is not helpful to simply copy and paste a list of error texts into the guide without any further detail!

It is expected that this section can be enhanced with:

- Any messages encountered by testers that they did not understand – support will probably see the same thing in live!
- Any errors that result due to miss-configuration of the test rigs

The information will probably take the form of a table including:

- The exact text of the error / message
- What events / activities cause the message to be produced
- What avoidance action should be taken (KB reference where applicable)
- What impact will it have on the service the facility provides

10.6 Code base and APIs

In general support groups would expect that all code written is self documenting with comments included to aid diagnostics:

- Outline of the code standards and languages used for the various parts of the facility.
- Where can support find the source code, which source repository. Define location and type (CVS, Subversion etc).
- Define administrator for source repository (i.e. where can a user / password be obtained to access).
- Details of external APIs provided (may be a cross-reference to appropriate LLDs)

10.7 Support route

Definition and contact list. Who supports the product at what levels. At 1st to 4th line this is likely to be just a definition of which POA support groups are involved.

In development, define who the lead developers and architects are.

If any interface touches units external to POA define:

- Their role in the support process
- How to contact them
- What hours they work
- What level of service is offered
- Expected clearance timescales for problems of different severities

NOTE: This may already be provided in an interface specification document in which case a cross reference to that document is adequate.

11 Knowledge base maintenance

FUJ00243143

One knowledge base is maintained within POA

11.1 KB

The KB is the master knowledge base for all incidents being progressed through the support chain. It was developed by the SSC who also maintain the servers it runs on. Because of this history, the SSC have also become the arbiters of the information within the KB. The KB was previously referred to as Known Error Log (KEL) but was renamed to clarify that not all entries relate to Errors.

It is also the SSC's responsibility to support the KB system and to allow access to this register to all other support units so that they can enter details within their area. Support for the KB system is only provided during normal SSC working hours.

11.1.1 New KB generation

All support units have access to the KB system and are able to generate new knowledge entries. A new KB should be generated for each **new** incident that is raised on the live system. Before any new KB is generated it is **essential that an extensive search** of the KB is done to ensure that a duplicate is not being created.

It is not necessary to fill in all fields when generating a new KB. This may happen when a 2nd line diagnostician can define the symptoms but has been unable to determine problem and solution. It can then be passed as a skeleton KB to the next level of support who should update with further details when known. In these circumstances, the KB reference **MUST** be added as a reference on the Peak.

A minimum KB would consist of:

1. Title
2. Summary
3. Type (Knowledge Base Information / Knowledge Base Action / Knowledge Base Fault)
4. Peak reference
5. Symptoms

This should be considered to be a minimum standard.

11.1.2 KB authorisation

When a KB is created or updated is has to be authorised before it can be seen by all users of the KB. This is an SSC function. The KB web site will send an automated email to the SSC duty manager to request authorisation.

It is possible to amend the authoriser in specific cases if required.

It is expected that the SSC duty manager will complete the authorisation within one working day. If this does not happen please follow the usual escalation process.

11.1.3 KB rejection

A new or updated KB may be rejected if:

- The update adds no value
- Information on the KB is incorrect
- Minimum information has not been specified (see 11.1.1)

FUJ00243143

If a KB is rejected a reason is given by the authoriser and an email is sent to the originator. It is then the originators responsibility to either:

- Correct the rejected KB based on feedback from the authoriser.
- Respond to the authoriser if they do not agree with the rejection.
- Request that the SSC delete the rejected KB.

It is not expected that a KB will be rejected where a minor change can be made by the authoriser that will result in a satisfactory KB.

NOTE: If a rejection is not responded to within 2 weeks the new KB details are deleted.

11.1.4 KB updates

Updates to existing KBs can be made by anyone with access to the KB system.

- There is no requirement to update a KB with a new Peak number simply because another incident has been seen for an unresolved issue. Counting this type of incident is the function of Post Office ISD master call processes.
- KBs should be updated with workaround or resolution details as soon as they are available.
- KBs should be updated with the details of a baseline when generated by 4th line
- KBs should be updated with details of when a resolution was delivered to the live estate. KBs can then be deactivated under these circumstances since it is possible that code regression may occur later (they should not be deleted).
- KBs have version numbers, it is possible to view a deprecated version and highlight the differences between versions. The KB system previously only retained the last 3 versions of an entry, but has since been changed to retain all versions.

When a KB is updated it will need to be authorised again. The old version of the KB is still visible until the authorisation has taken place.

11.1.5 KB deactivation

KBs should be deactivated when:

- A resolution has been delivered to the live system (and evidence obtained to demonstrate that the problem no longer exists, where possible).

Deactivated KBs are retained within the KB system but are not visible when searching, unless the 'Include deactivated KBs' flag is selected.

11.1.6 KB deletion

KBs should be deleted when:

- The function they refer to no longer exists in the live system
- It is a duplicate of an existing KB with the same problem and resolution.
- The KB contains misleading information and no update is appropriate.

KBs can only be deleted by members of the SSC. If you think a KB should be deleted then update the KB and replace the KB summary with the text "KB SHOULD BE DELETED: Reason" and specify "Reason". The SSC diagnostician you direct the update authorisation to will then check and delete the KEL if they agree on the reason.