Appendix 10: Report by Professor Sommer

Al-Sweady Inquiry:

Digital Forensic Report on

Liverpool Server

Peter Sommer

Summary

I am asked to carry out a digital forensic examination of a computer system informally known as the Liverpool Server. The server is said to have been used by the Princess of Wales Royal Regiment (1 PWRR) during May 2004, the period with which the Al-Sweady Inquiry is concerned.

Prior to my involvement the Inquiry's investigators had carried out a series of examinations based largely on the search in electronic documents for keywords thought to be significant; they had also used a series of standard techniques for recovering deleted material. These examinations based on keyword searching are continuing. I am asked to carry out a more detailed review and in particular to look for unusual patterns of deletion, testing a hypothesis that certain records may have been deliberately deleted so as to thwart the activities of any subsequent inquiry. It is not part of my remit to duplicate the keyword searching exercises.

The conclusions I have been able to reach are rather more limited than I would have liked. Much of the detail of this Report is aimed at describing the various tests and procedures I have attempted to use.

The main reasons for the limited scope of my findings are:

1. Forensic examinations of computers produce their best results when it has been possible to seize and properly preserve the contents of the computers very

shortly after the events of interest. Although data on computer storage media is surprisingly persistent, the longer the period between the events and the examination the greater the chance that important data has been over-written. I include some estimates about the overall capacity of the hard disk storage capacity and the extent to which it has been used and over-written.

2. Although the Inquiry is concerned with events in May 2004, 1 PWRR continued to use the computer system until October 2004. I have attempted to reconstruct its subsequent history. In October 2004 a handover took place to the Welsh Guards. It seems very likely that standard handover procedures involved the deletion of user accounts and files so that those relating to 1 PWRR were deleted. The Welsh Guards used the computer system throughout their six-month deployment until April 2005. Thereafter there were a further three handovers to units who each had six-month deployments and who each appear to have carried out a deletion of material relating to the previous units. The computer system was apparently decommissioned in September 2006 when the Iraq theatre closed. It appears that in mid-2010 hard disks relating to the Liverpool Server and very many others were collected and sent to the MoD Defence Archive System, the aim being to add to MoD's "collective memory" from which a wide variety of lessons and other research could be drawn. The disks were sent to PJHQ at Northwood Hills. As the aim was research and not criminal investigation no precautions or procedures to preserve the integrity of the disks in their then state were taken; some modification and over-writing of data took place. There were also other periods during which the computer was started up and disks were viewed. It was not until 2011 that the Royal Military Police carried out proper preservation via forensic disk imaging to produce the material that is available for me to examine.

3. What is left is a potentially compromised possible "crime scene" from which nevertheless some limited conclusions are possible. The methods used to carry out the deletions between unit tours have left some records of the names of files and folders, though often not their content. There are a series of ghosts of user accounts and files, references to their previous existence rather than anything more substantial. These "ghosts" are a function of how the data was deleted and

the way in which the computer's internal directory file records operate. However, even where there are substantive files there may be doubts about when they were created, modified, or viewed.

4. The "Liverpool Server" is in fact two machines, Liverpool 1 and Liverpool 2, both configured to run a Microsoft product called Exchange. This consists of a "server" operating system (Windows Server), that is, one designed primarily to serve the needs of several individual users each with their own computer so that they have accounts, can share and distribute data and have centralised storage as the storage on their own machines. Another function is to provide links to the outside world including, if required, the Internet. Sitting on top of the server operating system is an additional product called Exchange Server which provides sophisticated email and other facilities. The combined product is in very wide use in businesses and organisations world-wide and can be configured in many different ways to suit local requirements. Many servers can be placed in what is called a "forest" of servers within a hierarchy where all are able to communicate with each other, following specific rules to do so. A particular feature is the ability to replicate data between individual attached computers and servers and servers between themselves.

5. I am asked merely to examine one pair of such servers. To my knowledge so far none of the individual computers/workstations used by 1 PWRR, almost certainly either HP/Compaq towers or laptops, that were "served" by Liverpool 1 and Liverpool 2 have been recovered or are available for examination. To my knowledge also, although it seems very likely that some of the data originated in Liverpool 1 and 2 will have been replicated to other servers and in particular to servers located at PJHQ Northwood Hills, none of this material has come to my attention. I make a recommendation that searches are carried out by the MoD to see if tower workstations, laptops and replicated material or back-ups thereof for the period of interest. I note that, according to a *INET System Management Manual* of September 2005 (section 12) a database of every asset in the INET system is supposed to have been held on the Unicenter ServicePlus Service Desk application on a server called Hertfordshire.

6. I have come across indications that there were routines to back-up the Liverpool Servers. However it appears that even in the best of circumstances data backed up in May 2004 would have been over-written by May/June 2005. The aim of the back-up routines was to recover from a computer disaster rather than to create a formal archive of activity. It also appears that the back-up equipment did not function well in the desert environment of Iraq. As a result no back-up material has come to light.

7. I have been asked to operate under "secret" conditions. Although no restraint has been placed on my access to the contents of the Liverpool Servers I have had to treat files on them as secret and then ask Al-Sweady Inquiry staff to negotiate with MoD to selectively downgrade them so that I could refer to them in this report. There have also been restrictions on the types of examination I could carry out. In normal circumstances where I am asked to report on the contents of the computer I am provided with a "forensic" image copy of the disk which I can then examine on my own equipment using such analytic tools as I deem most suitable; more-over I can conduct the work in my own time at my own regular premises. In this instance I was asked to examine the forensic disk images using equipment at the premises of the Iraq Historic Allegations Team (IHAT) at Trenchard Lines, Upavon, Wiltshire. The equipment was set up to use only one type of computer forensic analysis software, AccessData's Forensic Tool Kit. The hardware was configured so that it was not possible to introduce further software; neither was it possible to download exhibits without going through an elaborate procedure. It is not for me to question MoD judgements about the need to keep secret details of its computer and communications infrastructure or of personnel and operations but I need to record the impact of these limitations on the extent of my investigation. The costs of examination have also been greater and the time taken to complete this Report much longer. Although I cannot say that there are obvious lines of investigation that I wished to pursue and was unable to do so, nevertheless it is possible that some-one with unrestricted access to the forensic disk images and no limitation on the range of forensic analysis tools deployed may reach more extensive findings than I have been able to.

But I have concluded that there is at this point and on the evidence available no obvious indication of deliberate deleting of key documents, emails and other files originated in May 2004, either at that time or subsequently.

Instructions

1. I am asked to carry out some investigations on a computer known as "the Liverpool Server". In particular my instructions from the Al-Sweady Inquiry are:

    a. Provide us with a chronology of deletions and overwriting made:

        · Before 14th May 2004 (the date of the battle with which the Inquiry is concerned) to establish what had been normal practice in terms of deleting/overwriting;

        · Between 14th May 2004 and October 2004 (the period during which one the groups of soldiers we are investigating, the PWRR, was stationed at Camp Abu Naji) to establish whether normal practice changed; and

    b. From October 2004 to August 2006 (when unrelated groups of soldiers were stationed at the Camp) to compare the practice of subsequent groups.

    Identification of, and an informed view, on any abnormal deletion or overwriting from 14th May 2004 onwards on any items which might have been expected to have been identified by the Inquiry's searches.

2. These relate to issues 73 and 74 of the List of Issues identified by the Inquiry.

3. Many other computers located by the MoD and by the Inquiry are of interest to the Inquiry; these are the subject of what is called the "Sensitive Case". I am aware, at an anecdotal level, of some of the issues raised by the Sensitive Case. However, my instructions are limited to the "Liverpool Server"

Qualifications

4. I have been providing advice and expert evidence in relation to computers and computer-derived evidence since 1985. My instructions for criminal matters have included, among others, cases involving terrorism, large-scale software

piracy, large scale international intrusions into computer systems, murder, fraud, obscene material, immigration offences, and the acquisition and distribution of collections of pictures of child sexual abuse. My instructions in civil proceedings have included claims of defamation, breach of contract, breach of confidence, passing off and Family Court matters. I have been instructed in South Africa in a matter alleging state corruption and in Australia in a dispute about the consequences of state regulatory action; I have also appeared before the Solicitors Regulatory Tribunal. I have recently been instructed by prosecutors in the Special Tribunal on the Lebanon sitting in the Hague and by the International Criminal Court in respect of charges against Uhuru Kenyatta.

5. I am currently a Visiting Professor at the Cyber Security Centre at de Montfort University and am also a Visiting Reader at the Open University Department of Computing and Mathematics where I am "course consultant", that is the main author, for a MSc course module on Forensic Computing and Investigations. For 17 years I taught and researched information system security at the London School of Economics ending up as a Visiting Professor. Between 2002 and 2006 I was first External Evaluator and then External Examiner for the MSc course run by the Centre for Forensic Computing at the Defence Academy , Shrivenham (Cranfield University). I was Joint Lead Expert for the computing speciality within the scheme run by the Home Office-sponsored Council for the Registration of Forensic Practitioners (CRFP) until the Council's closedown in March 2009. Since 2008 I have been on the Digital Forensics Specialist Group which advises the Forensic Science Regulator. I am on the editorial board of the journal *Digital Investigation.* (http://www.journals.elsevier.com/digital-investigation/).

6. I am the author of, among other things, *Directors and Corporate Advisors' Guide to Digital Investigations* published by the Information Assurance Advisory Council and now in its third edition (www.iaac.org.uk) and co-author of the OECD study *Reducing Systemic Cyber Security Risk.*

7. A full CV is attached as PS-1.

Material Considered

8. I have been supplied with forensic disk images, described in more detail below, corresponding to the Liverpool 1 and Liverpool 2 servers. They had already been installed on equipment provided by IHAT and were viewed via the computer forensic analysis software suite AccessData Forensic Tool Kit, versions 3.4 and 4 ("FTK").

9. I have had the benefit of extended conversations with Michael Moore, an investigator employed by the Al-Sweady Inquiry, and Jim Priddin, Royal Military Police, who was responsible for overseeing the forensic imaging of the original disks and carrying out some initial inquiries. I have compared what they have said with what I was able to observe from the disk images themselves. In some instances some of the information I have used comes into the category of hearsay, particularly in relation to the development of chronologies of events; I have included these elements where I believe them to be consistent with what I am able to determine for myself and where I hope they add clarity to Report; I have sought to flag these instances in this Report.

10. I have been provided with a series of interim reports created by Mike Moore and colleagues which describe their own examinations of the Liverpool Servers.

11. A few substantive documents, identified in more detail below, found on the servers appear to refer to how the servers and other relevant equipment were being managed and configured during 2004.

Arrangement of Report

12. As much of this report is concerned with a series of detailed examinations the form and nature of which are likely to be unfamiliar to most readers I have had to design a report format which provides necessary background explanations.

Some of these will appear in the main narrative while others have been placed into Appendices.

13. I begin with an attempted reconstruction of how the two servers are likely to have been functioning during May 2004 and afterwards

14. I follow this up with another reconstruction, a history of what appears to have happened to the servers since 2004 and up to the point at which their contents were properly preserved so as to become the precise form in which they are available for examination today.

15. In Appendix 1 I describe the range of techniques available for examining and recovering data from computers, together with the terminology used. In Appendix 2 I describe in relevant outline what I have been able to gather about Army's computer services architecture based on Microsoft Exchange.

16. I report on what I have been able to gather about the history of Liverpool 1 and Liverpool 2 and in particular events after May 2004 up to the point where a safe forensic copy was made of them.

17. I then report on the tests I have run to establish patterns of usage such as might indicate deliberate deletion during May 2004 and immediately afterwards.

18. Finally I have some recommendations for further inquiry.

The role of Liverpool 1 and Liverpool 2.

19. I have not had access to central Ministry of Defence and Army explanations of how computers were used for communications and data storage in 2004 and afterwards but it is possible to infer a great deal for my purpose by looking at the configuration and contents of the disk images presented to me for examination. As referred to above at paragraph 11 a number of extant documents were found on the servers which are operational manuals, relate to training or are handover notes. I can add to these from my own knowledge of the core Microsoft products and the ways in which they are likely to have been deployed.

20. The Army based a great deal of what later came to be known as the Defence Information Infrastructure (DII) on a product called Microsoft Exchange. At the relevant time, when this part of the MoD system appears to have been called INET, it consisted of a hierarchy of mid-sized machines known as servers which, at the very local level, provided facilities to "serve" the needs of individual within individual units in terms of communications, messaging, emails, storage of documents and so on by acting as a hub. But these local servers were also able to communicate upwards and with each other. In more technical jargon, each server acted as a local domain controller but the servers together were operated as a single domain within a flat IP network. Users could access the overall system with the same username/password credentials irrespective of location, The actual facilities on each server consisted of an underlying operating system called Windows Server (at various times Window 2000 Server and Windows Server 2003) and on top of that an application called Exchange Server (also at various times in "2000" and "2003" versions). A system administrator would need to make some settings in the overall configuration in the underlying operating system and some in Exchange Server itself. For the purposes of this Report, where I refer to "Exchange Server" I mean the entire system. Appendix 2 gives more detail of the capabilities of the Exchange product and the many ways in which it can be configured.

21. Connections between various forward-based servers, other such servers and servers at headquarters are likely to have relatively limited and slow. In a normal commercial environment, speed and quality of connections are seldom a problem as they can take place over physical lines and, with the protection of encryption, via the Internet. From the documentation I have seen, most of the connections would have been via satellite but with restricted speed capability.

22. At Camp Abu Naji it appears there were two such servers, referred to as Liverpool 1 and Liverpool 2. According to a INET Guide Document from 2003, this was a normal arrangement. They served the needs of between 25 and 35 workstations, towers, laptops and other devices.

23. According to a INET Guide Document from 2004, laptop and work stations would have used the operating system Windows 2000, Microsoft Office 2000, an Internet browser and a small number of utilities. Each laptop or work station was set up using a facility called "roaming profiles", so that each authorised user could log in and see their own "desktop" on any connected terminal or laptop. Once they connected to the server users would also have seen what appeared to be additional disk drives: Drive U was for personal files, Drive S was for shared files, for all official work, while Drive P was a public file store, accessible to all users. Drive R was labelled "Registry" and was designed to hold all completed and finalised documents and important emails - it was managed by central Registry staff. (Note: this should not be confused with a hidden feature of modern Microsoft operating systems also called the Registry which holds system configuration details – ordinary users never see this). Users also had access to an Intranet (web-based service but only accessible within the military domain and not available to the world-wide Internet audience). The intranet apparently contained the latest internal bulletins, a notice board and other types of administrative information.

24. On the servers according to the forensic disk image I have been examining, in addition to folders one might associate with central system administration and operating system functions, there are major folder areas (a "folder" is Microsoft-speak for "directory"; each major folder holds a hierarchy of sub-

folders). "Public Files" (presumably files for general use with all those with accounts on the computer), "Shared Files" (possibly for material shared more widely with other servers), and "User Profiles" (where one can see the files and configuration details of those with accounts on the server) and "Liverpool Registry", presumably completed and finalised documents. The "Shared Files" contain some details of "TELIC 7" and more files associated with "TELIC 8". These are references to the last two units that used the server. Operation TELIC was the codename under which all of the UK's military operations were conducted in Iraq. Every six months a new battlegroup consisting of many different regiments was deployed to the Iraq theatre. Each battlegroup was assigned the next "TELIC" number. For the purposes of the use of the Liverpool servers, TELIC 7 refers to the Scots Dragoon Guards and TELIC 8 to the Queen's Royal Hussars; I understand from my instructions as well as date and time stamps on the computer that the latter used the server between April and September 2006. TELIC 8 is the most complete and presumably gives the clearest idea of what might have existed for each unit while it was active. The unit in place at the time with which the Inquiry is concerned, 1 PWRR, was part of TELIC 4 and is referred to as such on the Liverpool servers.

25. Each "User Profile" has an associated hierarchy of files. But much of the user data has been deleted so that what remains simply indicates that substantive files once existed. (Readers are referred to Appendix I). Of the various User folders, including those where there is only a "ghost" left recoverable; approximately 80 appear to be based on the names of people. There are just under 20 accounts referring to "DivHq". The vast majority, in excess of 450 consist of "ukbg" followed by a four-digit number. It is beyond my current remit to seek to identify who any of these account holders were but I could assist if so asked.

26. Further user accounts appear in a folder on the server called "Documents and Settings". The most likely explanation is that these are people and entities with accounts directly on the server and hence able to take some action on the server (Windows Server 2000 or 2003) as opposed to the individuals in the "User

Profiles" who would have interacted only via the own laptops or towers. The roles of some of the users with direct accounts on the server can be inferred from their names: "Administrator.Liverpool1" is the main system administrator, "Adminstrator.PJHQUK" presumably is the controller of the entire domain of local unit-based servers, of which Liverpool is but one. The other accounts seem to be in the names of individuals.

27. **Back up and Data Retention Policies** According to a "DII/C Deployed and INET System Administration Course" dated February 2005 (section 3.3) and to be found on the Liverpool Server, a number of back-up arrangements should have been in place. These included " images" of the main system - - which enabled rapid restoration of the complete state of the main operating system facilities; and back-ups of the data files. The aim was to be able to recover from a computer-related disaster. Also on the Liverpool Server were a series of instructions for specific items of back-up hardware and software. I will return to the implications of the back-up tapes later on, but the longest period for which back-up was supposed to be retained was 343 days (just under a year) so that even in the best of circumstances events recorded to tape in May 2004 would have been over-written in May/June 2005.

28. **Replication** An important feature of Exchange Server is the ability to replicate folders (directories) between different servers. Often such processes take place in the background. One reason is to provide an additional form of back-up but a more important one is to enable users on other servers to have instant access on their local machine to important data created elsewhere within the organisation. Please see Appendix 2 for more detail.

29. **Relationship between Liverpool 1 and Liverpool 2** I have not so far been able to find any documentation which would enable me to reach a definitive conclusion why it was necessary to have two separate servers at the same location. As we will see the two servers may have had overlapping functions. There is hierarchy of folders on Liverpool 1 called "Liverpool 2 Backup" and containing "Home Directories". But these do not appear to be full back-ups

carefully and routinely collected so that they could be part of a disaster recovery program.

Subsequent History of Liverpool 1 and Liverpool 2

30. The following is my current understanding of the way in which the hardware comprising the Liverpool Servers were deployed. I stress that this is anecdotal information, included here for convenience but not within my personal direct knowledge. The army unit which has fallen under suspicion – the Princess of Wales's Royal Regiment – appears on the servers as TELIC 4. Their use of the servers is said to have ended on 16 October 2004. They were replaced by the Welsh Guards, who appeared as TELIC 5, who used the server until April 2005 at which point they were replaced by the second Battalion Prince of Wales Regiment, referred to as TELIC 6, until October 2005. Thereafter the Scots Dragoon Guards were referred to as TELIC 7 between October 2005 and April 2006. Finally the Queen's Royal Hussars used the server between April and September 2006 are referred to as TELIC 8 at which point they were decommissioned from active service. However, as we will see, the server hardware was intact and usable for some time after that.

31. The Iraqi theatre closed in April 2009.

32. I understand on an anecdotal basis that in mid-2010 a decision was made to transfer large parts of the overall electronic data created during the Iraqi war to the Defence Archive System which, as the name implies, is an electronic archive of all important documents. The aim was to support the MoD's "corporate memory". I understand that the person in charge of this operation was called [redacted] I further understand that the methods used to create the archive did not deploy formal forensic methods to preserve the entire contents of hard disks but rather concentrated on creating copies of the contents of files thought to be of future value. In the course of creating these copies the

contents of hard disks (including those of interest to the inquiry) would have been and in fact were altered. I was also told that it is also possible that further user accounts were created in order to facilitate access to the substantive files.

33. Later I will describe my own examination of date and time stamps which show that the servers were indeed intact and in use at least until November 2008. (see paragraphs 71 ff below).

34. At some point after that it appears that the servers were decommissioned and the hardware split up.

35. Towards the end of 2010 the Royal Military Police (RMP) apparently decided that a series of then on-going investigations needed to be informed by what could be found in computer records and they decided to evaluate what was available. The technical team was led by Jim Priddin who arranged a meeting with [redacted] By that time I understand, and still on the basis of anecdotal evidence, large numbers of individual hard disks had been acquired from theatre and placed in storage.

36. In the end some 1728 hard disks were collected by RMP[1] so that their contents could be properly preserved and subject to forensic scrutiny. It is not clear to me how far they were able to rely on the Asset Database which, according a *INET System Management Manual* of September 2005 (section 12) should have contained "a comprehensive record of every component on the system including network items, and items in stores, out for repair, or scrapped. It provides such detail as Model, Manufacturer, Serial number, Location, Service Status, Installation date etc. and a log of all activities."

37. Jim Priddin has said to me that once his team was in place proper "ACPO" guidelines for evidence handling and preservation were followed. Appendix I describes the data preservation procedures. From what I have been able to see for myself there is no reason to doubt that from February 2011 onwards proper handling and preservation procedures were indeed followed. Many if not all of

[1] This number may have increased as a result of activities by investigators since this report was researched

these hard disks have by now been installed for examination within a large forensic analysis system based at IHAT.

38. A particular problem was that a number of the hard disks had originated from servers such as Liverpool 1 and Liverpool 2 within which they had been made to work together in a RAID array. RAID arrays are used for two main reasons: to extend the storage capacity of a computer where no single physical hard disk would be sufficiently large; to increase performance speed. The RAIDs needed to be reconstructed. In a typical situation, and Liverpool 1 and Liverpool 2 fall into this category, each server would have associated with it 6 hard disks which needed to be arranged in the correct order. It is in fact obvious when the correct arrangement has been achieved as until that point the overall machine would be both unusable in practice and unviewable through specialist forensic analysis software[2]. I understand that it took RMP some man 8000 hours during February and March 2011 to reconstruct the various RAIDs

39. RMP created a series of work records of forensic activity as they proceeded in documents called Blue Books. I have had an opportunity to see some of these records as they apply to JRY/68 and JRY/69, which are the evidence numbers given to Liverpool 1 and Liverpool 2 respectively. The records include photographs of the original hard disk and cover the procedures under which the individual hard disks were re-assembled back into their RAID configurations so that the "Liverpool Server" can be forensically analysed. The records indicate to me that the two servers were properly preserved in the state they were in February 2011. Based on my experience of ACPO handling and preservation guidelines, I can say that we have good continuity of evidence from that date. The resulting forensic disk images are what I have been examining for the purposes of this Report.

[2] In fact each disk in a RAID array contains "header" information which can be used to reconstruct the whole.

Handover Procedures

40. I have attempted to identify what was supposed to happen as each army unit – regiment - was replaced by a new one. I have not so far located any procedure manual which describes in detail how handover was supposed to take place or what arrangements were supposed to be made to back up folders and files. I have found some informal notes passed from unit to the next but these seem to be limited to referring to specific unresolved problems with particular items of hardware and software. The explanation below is based partly on anecdote and partly on my examination of the servers.

41. As represented on the server, each "TELIC" had specific hierarchies of folders of files covering its areas of activities as well as accounts for its users. Each "user" area also had within it a hierarchy of folders and files; some of these would have applied to any user of any Exchange system, others were specific to their roles.

42. Here, for example, are screengrabs of the folders associated with TELIC8, which was the last to use Liverpool1. The second screengrab is of the opened "G1" folder which appears in the first screengrab.

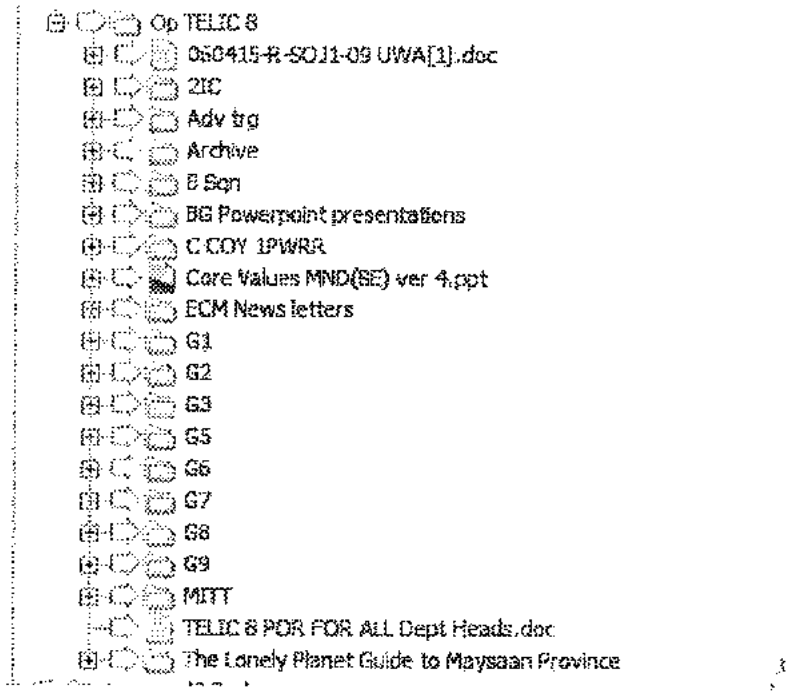

[3] TELIC8 from LiverpooC-top.jpg

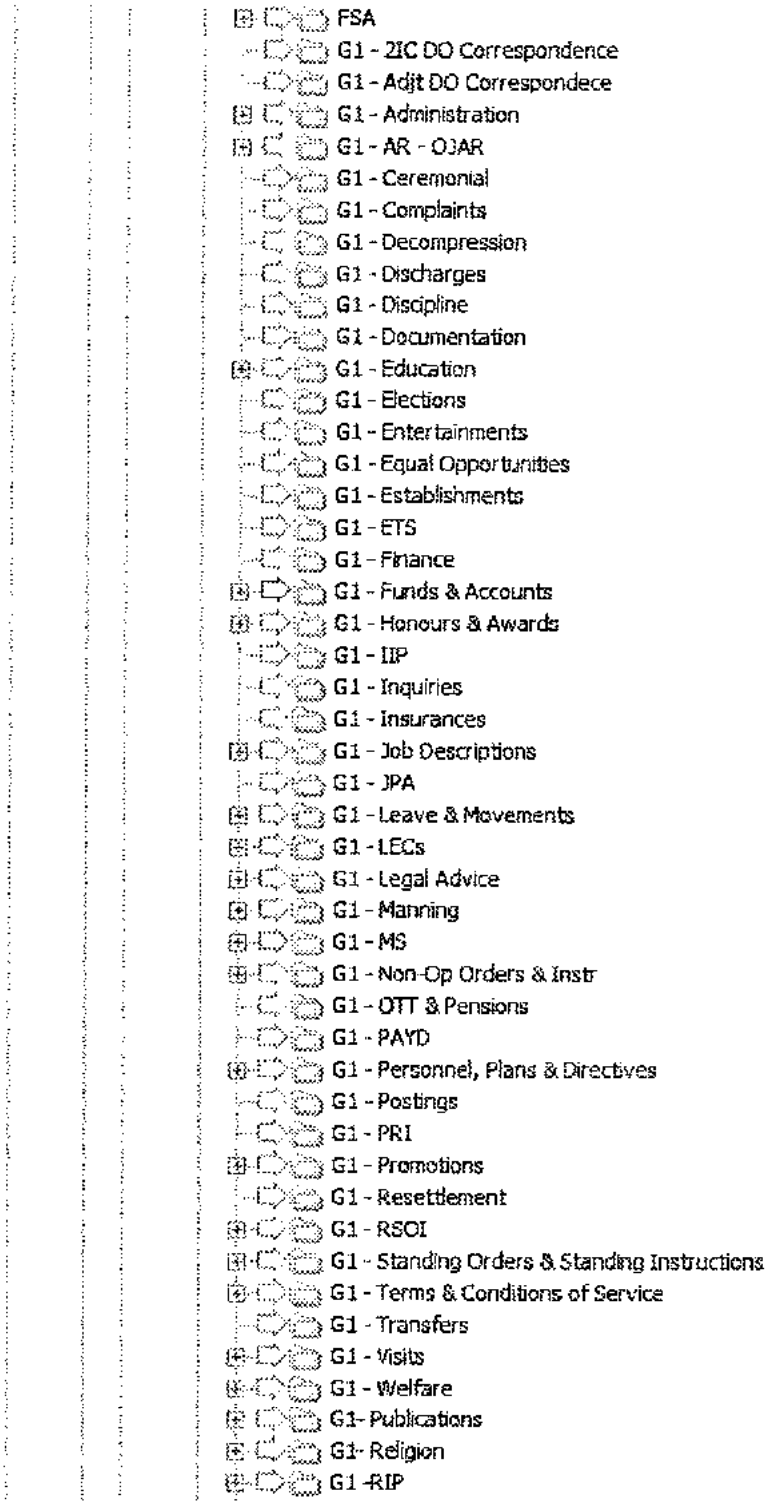

[4]

43. The files above occupied approximately 45 GB of hard disk space, some 331188 individual files in all.

[4] TELIC8 from LiverpoolC_G1.jpg

44. The practice seems to have been that when a TELIC unit left, its hierarchy of files was deleted, using simple delete commands, no attempt being made to use secure deleting methods of the type that would thwart subsequent recovery. (The reader is referred at this point to the more detailed explanations in Appendix 1). Many older files would have been eventually lost through over-writing. However, because of the operation of the internal Master File Table (MFT) which keeps track of the location of files, although the substantive contents of the files may have been lost, there are often entries in the MFT which point to their having existed. It also seems likely that there was little rigour in the deletion exercise; the aim was less the desire to protect sensitive material from the eyes of the subsequent unit than to create space in which the new unit could mount its files. Indeed, I am not able to say, from the material I have seen, whether the deletion was carried out by a unit when it ended its tour or was carried out by the following unit at the beginning of its tour. One reason for my uncertainty is that Microsoft operating systems do not create specific records for when a deletion takes place.

45. As a result in the form in which the Liverpool Server is available for examination, we have reasonably full files for the activity of TELIC8 (Queen's Royal Hussars) but "ghosts" in the form of index entries that files existed for previous TELICs; the further back in history, one goes, the less complete these MFT index entries.

46. As it appears that deletion at the end of each TELIC's tour of duty was routine it is plainly going to be difficult to look at the "ghost" entries and determine if they were deliberately deleted as part of some cover up process as opposed to routinely deleted as part of a handover when 1PWRR were replaced by the Welsh Guards in October 2004. However because of the lack of rigour in the deletion process, some files from TELIC 4's tour do remain.

47. Deletion activity did not, from what I can tell, extend to email activity, a matter which I take up later at paragraph 79ff.

Examinations

48. I now turn to describing my own examinations.

Circumstances of Examination

49. When I was first instructed I was shown a copy of an agreement between the Inquiry and the MoD. I have been asked to operate under "secret" conditions. I was asked to examine the forensic disk images using equipment at the premises of the Iraq Historic Allegations Team (IHAT) at Trenchard Lines, Upavon, Wiltshire. The equipment was set up to use only one type of computer forensic analysis software, AccessData's Forensic Tool Kit. The hardware was set up so that it was not possible to introduce further software; neither was it possible to download exhibits without going through an elaborate procedure. In normal circumstances where I am asked to report on the contents of the computer I am provided with a "forensic" image copy of the disk which I can then examine on my own equipment using such analytic tools as I deem most suitable; moreover I can conduct the work in my own time at my own regular premises.

50. AccessData's Forensic Tool Kit is a comprehensive suite of tools for the analysis of disk images; I was already familiar with earlier versions of it. However, because of the rate of change in computer operating systems, applications and hardware, coupled with their great complexity, no forensic computing examiner is content simply to rely on just one product. Rival products include Encase and X-Ways; in addition there are many specialised single-purpose tools.

51. Although no restraint has been place on my access to the contents of the Liverpool Servers I have had to treat files on them as secret and then ask Al-Sweady Inquiry staff to negotiate with MoD to selectively downgrade them so that I could refer to them in this report. It is not for me to question MoD judgements about the need to keep secret details of its computer and communications infrastructure or of personnel but I need to record the impact

of the limitations on my extent of investigation. The costs of examination have also been greater.

52. In all I attended Upavon over 9 days, 5 separate visits, on each occasion copying screenshots, copies of a limited number of substantive files and "exports" of various tables showing file names and associated date/time stamps and other data. This material was submitted to MoD for clearance.

53. Although I cannot say that there are obvious lines of investigation that I wished to pursue and was unable to do so, nevertheless it is possible that some-one with unrestricted access to the forensic disk images and no limitation on the range of forensic analysis tools deployed may reach more extensive findings than I have been able to.

Hardware, Storage Capacity

54. There are two machines, identified as JRY/68 and JRY/69 - Liverpool 1 and Liverpool 2. Each of the original machines, when reconstructed, consisted of some hardware containing five hard disks arranged in what is known as a RAID configuration. (In fact there were six had disks associated with each server but only 5 were necessary to make the reconstruction; presumably the sixth had disk was a spare). This arrangement of hard disks is quite commonly deployed either to increase the overall storage capacity, in order to speed up access to data or a combination of these. Back in 2002/3, when the servers were being installed, readily-available hard disk capacities were limited to 30-50GB. By 2013 such has been the improvement in disk technology that single hard disks of 2-3 TB (2000-3000GB) are available for under £100.

55. The five physical disks would almost certainly have appeared to the regular user viewing them via "Explorer" or "File Manager" as two disks, C: (by convention the first hard disk) and D: (the second hard disk). Each physical disk in the server had a capacity of 36.4GB, three were used to create the "C" drive and the remaining two became the ""D" drive. This arrangement would conform to usual recommendations for setting up Microsoft Exchange. There

is constant background activity within that product, with different parts of the program calling on each other; the two-disk approach gives performance benefits.

56. The individual hard disks making up the array had been given arbitrary exhibit numbers – JRY/68 A, B, C, D, E, F but this is how they resolved themselves after reconstruction:

- Liverpool 1, Drive C: JRY/68/BEA
- Liverpool 1, Drive D: JRY68//CF
- Liverpool 2, Drive C: JRY/69/BED
- Liverpool 2, Drive D: JRY69/CF

57. The overall storage capacity of each server is thus 5 x 36.4 GB, around 180 GB. This is significant for the hope to be able to recover deleted data. (Please see the "Data Recovery" section in Appendix 1 for a fuller explanation). At various times between 2004 and 2006, the servers appear to have been operating quite close to full capacity. As a result, once data was formally deleted, the period for which it remained recoverable was quite short, as the sectors would have been over-written with new data.

58. The AccessData Forensic Tool Kit, when asked to establish overall hard disk storage capacities cannot easily to do so in terms of historic circumstances. Because it is able to make use and present information about files that were on the disk even if they are no longer present, its reports of the amount of data stored on a hard disk can be misleading. This is what it reports:

Server	Drive (as it appears to the user)	Physical capacity	As Reported by FTK
Liverpool 1	Drive C: JRY/68/BEA	109 GB	117.7 GB 1184607 files
Liverpool 1	Drive D: JRY68//CF	73 GB	71.21 GB 167704 files
Liverpool 2	Drive C: JRY/69/BED	109 GB	161 GB

			772538 files
Liverpool 2	Drive D: JRY69/CF	73 GB	55.86 GB 222.863 files

59. The main lesson to draw from these figures is that the hard disks had been substantially used since 2004 and up to the point at which they were forensically imaged in 2011. It is thus not surprising that very little in the way of the contents of substantive files extant in May 2004 but subsequently deleted can now be recovered.

Software Configuration

60. Both servers have a similar configuration and conform to Microsoft recommendations for setting up Exchange Server on top of Windows Server. That is to say, the RAID disk which appears as "C" holds mostly Windows Server files and application program files also and contains the user accounts of ordinary users within a folder

61. The RAID disk which appears as "D" contains a relatively small number of user accounts for those who will have responsibilities for system administration; these include, from their names, remote identities for various divisional headquarters – "divhq" plus an identifying number – and administrators based back in the United Kingdom.

62. Liverpool 1 Drive C (JRY/68/BEA): Aside from standard folders one would expect to see in any Exchange set-up, the main folders are

a. Liverpool Public Files: This consists largely of deleted folders which have no content, for example "Camp Clearance map", "FPS, and references to various Ops. (These are "ghosts")

b. Liverpool Registry: This contains a folder for Easy CD Creator software (used for burning CDs)

c. Liverpool Shared Files: This contains some extant folders of operational status, including references to TELIC 7 and TELIC 8. There are also some MX reports apparently related to TELIC 6. I understand that these extant files have already been viewed for content as part of the keyword search exercise and I have not carried out any further inquiries beyond noting their existence.

d. There are a series of what I am assuming are "operational" folders with names such as "Mission Critical". A hierarchy of folders for "Op TELIC 7" occupies some 1218 MB and contains some TELIC 6 documents including MX reports . I understand that these have been reviewed for content separately as part of the "keyword search" exercises. The earliest appears to date from December 2004, which is too late for this inquiry. There is also a hierarchy of folders for " Op TELIC8" .

e. Liverpool 2 Backup: this contains deleted "Home" folders and as mentioned above at paragraph 29 above, does not seem to have been part of any regular disciplined back-up exercise.

f. User Profiles

g. In addition there are a number of folders holding source program files, presumably so that they could be installed or re-installed if needed

63. Liverpool 1 D (68CF): Aside from Server files and program files, approx. 30 users, presumably for Server, as opposed to Exchange. Event Logs for March-August 2006; Files with numbers.

64. Liverpool 2 C: (69BED): Aside from standard Exchange folders:

a. Event logs for Liverpool 1 and Liverpool 2 March -- August 2006

b. Home Directories: 4905 MB

c. Deleted "Archive" file

65. Liverpool 2 D: (69CF): Similar but not identical to the contents of the D: disk on Liverpool 1.

66. Both servers contain copies of source files to enable the installation or re-installation of key items of software, including items that might be needed on

the laptops and workstations "served" by Liverpool 1 and Liverpool 2. There is nothing *prima facie* sinister about this; if we recall that that Liverpool 1 and Liverpool 2 were in forward locations to which communications might be difficult, such local storage of software code would mean that repairs and changes could be effected quickly without waiting for the provision from the UK of installation disks, or using the long download times to fetch the files over a data communications link.

67. I also noted the presence of software to support back-up routines. These matched some of the documents found on the servers and which prescribed back-up procedures.

68. I did not find any software that might be used for "anti-forensic" purposes such a secure deletion and the removal of file logs of activity.

Patterns of Usage

69. At the heart of my instructions is the requirement to look for abnormal patterns of usage and in particular abnormal deletions. As noted in paragraph 44 above, it appears that the routine practice at handover between successive units, TELICs, was to delete the user material of the previous unit. It is only because simple deletion methods were used, as opposed to more rigorous "secure" techniques, that it is possible to see any evidence of the activity of TELIC4 in 2004.

70. In addition, Microsoft Windows Server and Microsoft Exchange do not specifically record acts of deletion[5], though it may be possible from surrounding circumstances to infer when deletion took place. Again, the reader is referred to the "data recovery" section of Appendix I.

[5] There is an exception to this, files which are sent to the "recycle bin" will carry dates when they were deleted; but the recycle bin is of finite size and older files are constantly being discarded

71. I decided to examine a series of date/time stamps on various folders and files to see if I could identify any anomalies. This is what I found:

72. The last date I could associate with TELIC operational activity was 16/08/2006.

73. There were some "File Created" activities on, among others, 20/06/2007; 17/07/2007; 23/01/2008; 30/05/2008; and between 25 and 28/11/2008

74. There were some "File Modified" activities, on, among others, : 08/01/2007; 06/03/2007 13/06/2007; 20/06/2007; 17/07/2007; 12/11/2007; 06/12/2007; 23/01/2008; 22/05/2008; 30/05/2008; 17/06/2008; 20/08/2008; and between 25 and 28/11/2008.

75. Activity on all of these dates implies that at this stage both Liverpool Servers were intact and functioning. At the very least on the dates above, one or other of the machines was started up and material viewed.

76. The most extensive activity took place between 25 and 28/11/2008 and further examination shows that much of it is attributable to some-one with the user account "[redacted]gbr"; a search of email traffic carried out with Al-Sweady investigator Mike Moore shows that the owner of this account was a [redacted] who was probably an engineer based at PJHQ. His activities seemed to involve the downloading of material, possibly source code files, from other servers. It is also possible that various files were deleted during this period as well, though the files I have identified as being possibly deleted seem to relate to files which had originally arrived on the server in March 2006, outside the period of interest to the Inquiry.

77. I stress that so far I have found that none of this activity after August 2006 gives ground for immediate suspicion. However there is clearly a gap in the narrative of what happened to the Liverpool Servers between that date and the point in February when RMP began evidence preservation.

78. I recommend, as first step, that Al Sweady investigators attempt to trace Ken Sutherland and ask him for explanations. These could, if necessary, be

reconciled with extant email evidence. Should it prove necessary I am willing to test [redacted]'s account of his activities against evidence on the server.

Records of Email activity

79. The software tool used for forensic analysis has been able to recover daily records of email activity on a monthly basis. The following screengrab shows activity, as reported by the forensic analysis tool, for May 2004:

80. But we can compare this with activity for adjacent periods:

[6] Emailactivity_52004.jpg

April 2004

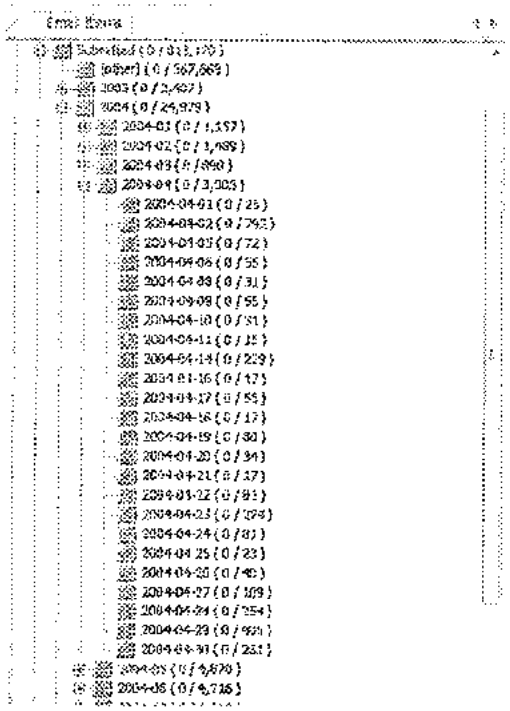

June 2004

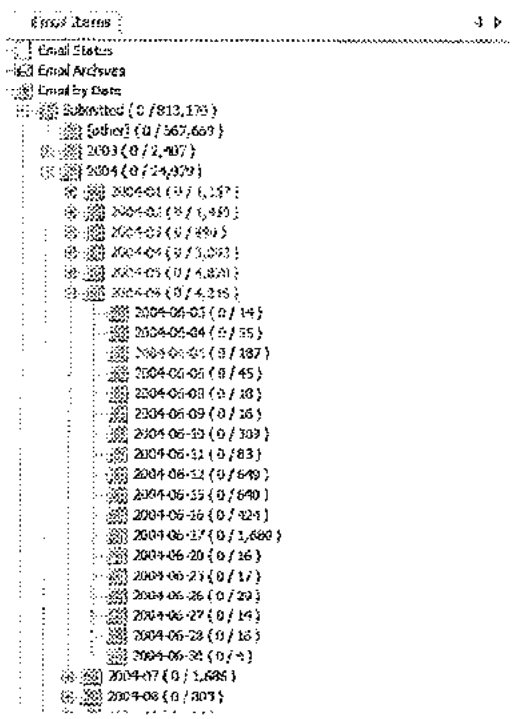

7

[7] Emailactivity_42004.jpg; Emailactivity_62004.jpg

81. And also for how TELIC5 behaved at the end of 2004:

November 2004

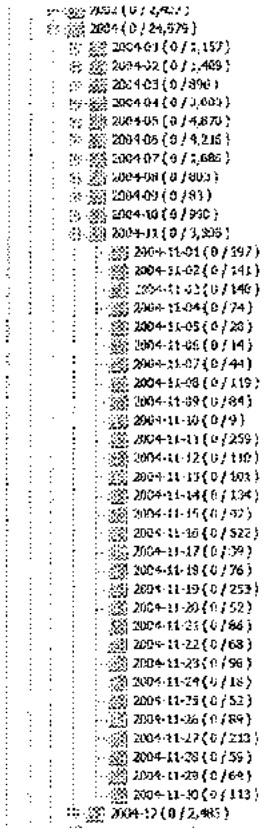

December 2004

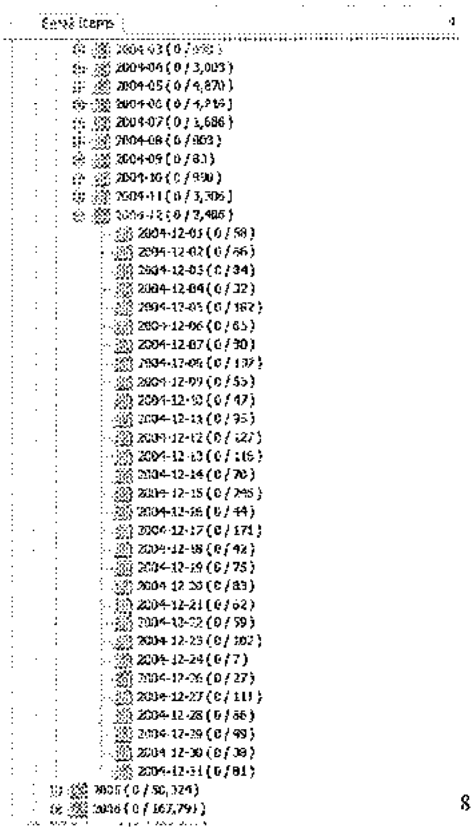

82. It will be seen that although email activity for 14 May 2004 is reported as low, when looked at overall patterns of email activity this lack is not obviously anomalous. I have captured further instances of monthly email activity which I can produce if required.

83. It may be that other investigators, from their wider knowledge of the circumstances, are able to provide explanations of patterns of monthly email activity which can vary from several hundred a day to figures in the low tens.

[8] Emailactivity112004.jpg; Emailactivity122004.jpg

Event Logs

84. An important feature of Windows Server is the collection of logs of various system events. In a running system they are viewed via a facility called Event Viewer and the most important relate to the System, Security and Applications. Potentially these also can identify anomalous activity as well as crashes and conflicts. However, upon examination I was only able to identify event logs in any form for after March 2006.

Conclusions

85. I now return to the remit I have been asked to address:

a. Provide us with a chronology of deletions and overwriting made:

- Before 14th May 2004 (the date of the battle with which the Inquiry is concerned) to establish what had been normal practice in terms of deleting/overwriting;

- Between 14th May 2004 and October 2004 (the period during which one the groups of soldiers we are investigating, the PWRR, was stationed at Camp Abu Naji) to establish whether normal practice changed; and

b. From October 2004 to August 2006 (when unrelated groups of soldiers were stationed at the Camp) to compare the practice of subsequent groups.

Identification of, and an informed view, on any abnormal deletion or overwriting from 14th May 2004 onwards on any items which might have been expected to have been identified by the Inquiry's searches.

86. There is no longer much evidence to make it possible to provide a chronology of deletions and overwritings because:

- The contents of the hard disks which made up the servers were not properly preserved until February 2011; however they had been in continuous use until September 2006 and in occasional use between September 2006 and February 2011
- Recovery of deleted material relies on the fact that the contents of a deleted file remain on the disk until such time as routine re-use of the sectors occupied by that file become over-written by newer files. The longer the period between when a file was deleted and the point at which recovery is attempted, the lower the chances of success. Once a file is deleted some entries in the computer's file index – Master File Table – may remain but not their content

- For most practical purposes computers with Microsoft operating systems do not retain for any lengthy period a record of the date when a file is deleted; it is sometimes possible to infer the date of deletion from surrounding circumstances but again the longer the period between deletion and the attempt to reconstruct events the lower the chance of success,
- It was the policy on handover that the previous TELIC unit's files were deleted; 1PWRR was TELIC 4; the last was TELIC8 so that there were four routine deletions of unit files before the computers were properly preserved
- The total storage capacity of each server was 180 GB; each sector on the hard disk formerly occupied by 1PWRR files would have been over-written with new material simply as part of the routine use of the computers.
- Such files as remain from May 2004 on the servers are likely to be the result of deliberate decision or an incomplete global deletion action.

87. Documents found on the server show that there was a back-up policy, designed to enable recovery from a disaster. Documents also suggest that back-up hardware was likely to have been available at Camp Abu Naji. The longest period for which the back-up was held was 343 days. There is software on the server to facilitate the back-up. I am not aware that any back-up material exists

88. Facilities on the server show that there were user accounts for remote, top-level administrators. It seems high likely that copies of key files would have been replicated in several distant locations including PJHQ. Where emails and files were sent from TELIC4 during May 2004 it is reasonable to inquire whether copies exist at the locations of and in the computer records of, likely recipients. But such inquiries are outside my remit.

89. Liverpool 1 and Liverpool 2 were designed to "serve" a series of laptops and workstations at Camp Abu Naji. It would be worth investigating how many of these still exist and remain relatively unaltered and unused since May 2004;

forensic examination of any discovered may shed additional light on the electronic records of the events of 14 May 2004.

Recommendations

1. Attempt to locate the MoD's Asset Management System database and records for May 2004 as they applied to computer equipment deployed in activities in the Iraq theatre.
2. Conduct, in so far as this has not already been done, an investigation of what has happened to the laptops, towers and workstations in use at Camp Abu Naji in May 2004.
3. Attempt to identify and locate system administrators for the Liverpool Servers in May 2004; obtain statements about their activities and ask them to verify or contradict analyses made in this Report, particularly as they relate to document retention and inter-unit handover as this may shed more light on the ways in which the Servers were in fact used.
4. Attempt to identify and locate system administrators, presumably located at PJHQ, for the overall INET/DII system in May 2004; obtain statements about their activities and ask them to verify or contradict analyses made in this Report, particularly as they relate to document retention and inter-unit handover as this may shed more light on the ways in which the Servers were in fact used.
5. Conduct, in so far as this has not already been done, a search for copies of files and emails originated at Camp Abu Naji in May 2004 and which were replicated by the normal operation of Exchange Server to servers and computer resources located at PJHQ Northwood Hills.
6. Conduct, in so far as this has not already been done, a search for copies of files and emails originated at Camp Abu Naji in May 2004 and which were replicated by the normal operation of Exchange Server to other servers and computer resources which were part of INET/DII.
7. Obtain a statement from [redacted] who was apparently responsible for siphoning off information from hard disks in 2010 for the MoDs' Archive in order to have a fuller narrative of the processes involved and the state of the equipment at the time.

8. Obtain a statement from [redacted] about his activities in November 2008 in order discover his remit and to have a fuller narrative of the processes involved and the state of the equipment at the time.

Statement of truth

1. This statement is true to the best of my knowledge and belief and I make it knowing that, if it is tendered in evidence, I shall be liable to prosecution if I have willfully stated anything which I know to be false or do not believe to be true.

2. I confirm that I have made clear which facts and matters referred to in this statement are within my own knowledge and which are not. Those that are within my own knowledge I confirm to be true. The opinions I have expressed represent my true and complete professional opinions on the matters to which they refer.

3. I understand that my primary duty is to the Public Injury when preparing written statements and giving evidence and I have complied with that duty.

4. I believe my statement to be accurate; it covers the issues raised by my instructions and reflects my views as an independent expert.

5. Where relevant my statement includes any information of which I have knowledge, or of which I have been made aware, that might adversely affect the validity of my conclusions.

6. I have indicated any sources of information upon which I have relied on in my statement.

7. Those instructing me will be informed immediately, with written confirmation, if my existing report requires correction or qualification.

8. I understand that my statement, subject to any corrections before swearing as to its veracity, will form the evidence to be given under oath.

9. I understand that an expert may assist any cross-examination on my report.

10. I confirm that I have not entered any arrangement whereby the amount or payment of my fees is in any way dependent on the outcome of the case.

Signed

[redacted]

Peter Sommer

21 June 2013

Appendix 1: Primer on Computer Forensics and Data Recovery Methods and Terminology.

Basics

1. Computers do not by default create reliable audit trails of all activity on them. Some applications, for example those used in finance, may do so. But computer operating systems do create records of some activity, for example by placing time and date stamps on files, by logging certain events and by having ever-changing configuration files which reflect occasions when programs were added and modified. Computer forensics relies on the collective research of many examiners who have observed how operating systems and applications function and create records – and have been able to derive rules for, say, particular Microsoft products. The observations can then be incorporated into specialist analytic tools. But the results may be imperfect; this is why the findings of a computer forensic examination may be ambiguous.

2. Data in and around computers is highly volatile. Start a computer up and data is written to disk; during close down more data is written to disk. Open up a file directory (folder) and some date-and-time stamps may change. If there are picture files and you ask to see "thumbnails" in the directory – new files are created. Connect an external USB drive to a computer and open up its file directory - some date-and-time stamps will change; if it is the first time that USB device is being connected to that computer, alterations will occur in a normally-hidden part of the operating system called the registry. Open up a file itself and not only will date-and-time stamps change but it is entirely possible that in the background temporary files are being created and still further changes occurring on the hard disk.

3. Because data is constantly changing it is necessary to use special techniques to freeze the scene, to capture a snapshot of the state of a hard disk or file at a specific identified point in time. The process is called forensic disk imaging and usually requires specialist hardware to prevent the disk medium being written to during the process of acquisition. Software captures every sector of a

disk or other storage device, including those that appear to be unused or empty. This ensures that hidden files as well as the remnants of deleted files are available for examination. The result of the deployment of the specialist hardware and imaging software is an image file which can then be subject to analysis; copies can be made so that many can work on the same material simultaneously. It is an important feature of forensic imaging that the technician carry it out produces a detailed note of his actions, including precautions taken and any difficulties encountered.

4. Basic examination techniques that can be used include:

   - Keyword searching. This can include searching the entire contents of disk whether or not the words are located within an extant file. But words which are within databases, spreadsheets and specialist formats such as Adobe Acrobat may not always be located
   - Chronology and time-line building. One can list all file date and time stamps, or a sub-set, in order to show a sequence of events. But not all computer activity may be identified as dates may refer only to the most recent change, not to previous changes

Date and Time Stamps

5. Date and Time stamps are of immense value in any examination of computers and files originated from computers. But there are also many traps for the unwary and hasty.

6. The main uses are:

   - to develop a chronology of events in and around a computer
   - to help identify the author or last user of a file or transaction
   - to produce alibi or absence-of alibi evidence

7. Date and time stamps are frequently created within computers as they go about their business and are also recorded by devices external to a computer that is being examined.

8. The main stamps associated with each file found on a computer are:

   - **File created** the date and time at which the file was first created on this medium (it may have been created earlier on another disk and transferred to this disk by means of floppy, USB solid state disk or via download)
   - **Last written** the date and time at which the file was last modified. This can give rise to some apparent anomalies if a file was originated on one computer and then copied on to another. Thus: supposing I have on my computer a file I finished editing on 11 January 2001. If I now copy this file on to your computer on 24 November 2005, your computer will show for this file a "file created" date of 24 November 2005 but a "last written" date of 11 January 2001.
   - **Last accessed** the date at which the file was last "touched" by an application on this computer but not altered. "Touched" may mean the same as "viewed but not altered" but a program such as Windows Explorer and some antivirus programs will "touch" the file to the extent that the last accessed date is altered even though the file has not been viewed by anyone.
   - **Entry Modified** column, pertinent to NTFS (Windows NT, Windows 2000, and Windows XP) and Linux file-system files, refers to the pointer for the file-entry and the information that that pointer contains, such as the size of the file. If a file was changed but its size not altered, then the entry modified column would NOT change. However, if the file *size* has changed (from eight sectors to ten sectors, for example), then this column would change. The Entry Modified column is of relatively limited value when developing typical chronologies of events

9. The operating system does not record when a file was *deleted* unless the file is in the *Recycle Bin*

10. Some applications programs, including word-processors, may generate additional information. This is known as *metadata* and in Microsoft Word can be accessed via the "Properties" screens.

11. Log files and other files found on computers may use a variety of non-obvious notations for time – Unix time. The examiner must realise which is being used and then deploy a piece of software to convert into more familiar and useful notation.

Data Recovery

12. When a user wants to locate a file, he will do so via a program called "My Computer" or "Windows Explorer" or "File Manager" or something similar – these programs get their information from the part of the hard-disk which acts as the index. Files that can be located in this way are described as "normally visible". Some operating systems, including those from Microsoft, by default hide away certain critical files so that the user does not inadvertently alter them. But options available in Windows Explorer can render these files visible to the user.

13. As is reasonably well known, when the user of a computer "deletes" a file on a hard-disk, or deleting occurs because of the routine operation of a particular program, the data which makes up the file is not in fact immediately deleted but changes are made in a special part of the hard-disk which maintains an index of the physical locations of all the files on the disk. Under the operating systems of interest in this case it is called the Master File Table (MFT) The programs "My Computer" or "Windows Explorer" or "File Manager" and similar get their information from the part of the hard-disk which acts as the index (the MFT).

14. A computer file in fact goes through several stages of "deletion" on a computer which uses the Windows family of operating systems: In the first instance the data remains physically on the hard-disk, but the entry for it disappears from the folder in which it had been and re-appears in another folder called the Recycle Bin. From here it can readily be recovered to its original location. When recovery takes place the date-and-time stamp information which Windows automatically records is also recovered. The Recycle Bin in essence a safety device against careless discarding of a file.

15. The Recycle Bin is of finite size and depending on circumstances, eventually removes files either because they are "old" or because the overall capacity is exceeded. The Recycle Bin keeps track of its contents by means of a normally hidden database file called INFO2. The INFO2 file can be recovered and analysed forensically.

16. Thereafter the data that makes up the deleted files is still on the disk and significant amounts of information about the file are still retained in the "index" part of the hard-disk. By using a software utility readily available from several manufacturers and sold on the High Street, it is possible to recover all the data, including the name) and associated date-and-time stamp information. Deleted files are normally found in the folders in which they were originally stored. But this is only possible provided that the disk sectors in which the data resides have not been over-written by subsequent activity. The period for which the deleted data is retained depends on the overall capacity of the hard disk and the space being taken up by the newer "live" files. Eventually the deleted files will become partially or wholly over-written and thus beyond recovery.

17. Sometimes the index, the MFT, may retain information about a file such as its name and the date data associated with it, but because some or all of the sectors in which the file was held have become overwritten, recovery of the contents of the file may be only partial, or not possible at all. This phenomenon can often be seen while using forensic analysis tools such as EnCase and AccessData FTK.

18. There is a further stage, when the "index" no longer retains any information about the physical location of the file. At the same time, information held in the directory about the time a file was created, or modified, or also viewed, is also lost. The data that is still on the disk is referred to as being in "unallocated space" or in "unallocated clusters" Recovery may still be possible but a different more specialist types of software utility are required.

19. They fall into three categories. The first is called "data carving". This simply looks at everything on the hard-disk and seeks out either the unique file signature associated with a particular sort of file or simply looks for character sequences that are thought to be relatively unique to a document.

20. However when files are recovered by this method it is often not possible to recover its name, any of the associated date-and-time stamp information, or the original folder. Sometimes only fragments of files can be found. Most data carving software simply gives each file in a specific format an ascending number – the first file is called 1.gif, or 1.doc, 1.html, etc, the second becomes 2.gif etc. The absence of context may limit the number of inferences that can be made about a file, particularly if several people have had access to the computer.

21. A second method is to look for records of deleted folders and other deleted features of the MFT. Again, these can be located by looking for their characteristic signatures. But, having found the folders it may be possible to inspect their contents to see if there is information about files that were within those folders. This information may point to those files so that they can be recovered. If such recovery is possible then the *names* of the files and also *date and time* information may also be recoverable. This technique can be referred to informally as "recovering deleted folders". However, as with the simpler data-carving techniques there is also likely to be a loss of *context* – it is not possible to recover the name of the folder or any dates associated with it, though one may be able to do so for the *contents* of the folders. The technique may also recover folders from an earlier installation of an operating system as opposed to the current one. As before, the absence of context may limit the

number of inferences that can be made about a file, particularly if several people have had access to the computer.

22. The software utilities to carry out the recovery of deleted folders technique have to be quite sophisticated. In the first instance they have to look for deleted records of folder and then interpret their contents. But the contents may point to locations on the hard-disk which are now occupied by quite different files, the original data having been over-written. Different utilities (and indeed subsequent versions of the same utility) may come to different conclusions, so that in all instances the program's immediate suggestions (for that is what they are) have to be tested by inspection of the recovered file which has to be manually assessed for completeness and consistency.

23. The third method relies on what is usually called "keyword searching" or "string searching". A program scans the entirety of a hard disk or other data storage medium looking for instances of words or strings which are relatively unique and might point to matters of interest. At each "hit", the *surrounding* material is examined manually and then captured as a potential exhibit. This method finds file fragments. However it only works when the file holds text, or recognisable words. Some file formats, for example PDFs, databases, and email archives do not hold words as straight-forward text.

24. When a disk is reformatted or a new operating system placed over an existing one, data associated with the previous installation is not deleted. Some of the disk sectors will be over-written by the new installation but others will remain until, during normal use, they too become over-written. Until then, the data in these sectors is not accessible by ordinary means to the regular computer user. That is because the MFT has been overwritten. But it can be searched for using the data-carving and folder recovery methods described above. This data too is referred to as being in "unallocated space".

25. The only point at which there is reliable information about when a file was deleted is when the file is still within the Recycle Bin, as the information is held within the INFO2 database file referred to above. It is sometimes possible to recover earlier versions of an INFO2 file and inspect it for details of files that

were in it – and this might include times of deletion. But apart from this, the date of deletion can only be guessed at: it will be some indeterminate time *after* the "last accessed" date.

Appendix 2: Microsoft Exchange: an Overview

1. Microsoft Exchange is a very widely used and complex product designed to support the functioning of organisations of all kinds and sizes. In particular it allows for complex messaging both between individuals within a specific configuration and, using the Internet email complete with attachments to the outside world. It also allows for extensive archiving of the activities of the organisation.

2. Individual personal computers (which may include laptops and other devices) are known as "clients" while the central computing facility is known as a "server". A particular feature is that in a large set up there can be several servers arranged in what is sometimes referred to as a "forest"; there is often a hierarchy of such servers.

3. Where there are many such servers it is possible to arrange things so that a top-level server can carry out administrative tasks, including adding new programs and managing users and their specific levels of access, centrally. It is also possible to configure matters so that a server controlled remotely – a high-level user can sign in to a distant server and use it as though he was physically adjacent.

4. The product has gone through a number of versions; the ones likely to be of significance in this instance are Exchange Server 2000 and Exchange Server 2003. The full range of services is in fact delivered by a combination of two products Windows Server 2000 and the actual package called Exchange Server 2000. Windows Server 2000 is the fundamental operating system, the equivalent if you will of Windows XT or Windows 7 as used in personal computing, but with the difference the essential technical infrastructure for

large numbers of individual computers to "sign in" and have accounts on it. "Exchange" provides additional sophistication including full scale email, web-serving, and conferencing.

5. Other functions include instant messaging and conferencing plus the ability to create considerable archives of documents which can either be stored for the sole use of their originator or can be widely shared.

6. A key feature is the ability to replicate data as between a personal computer and archives on the local server and also as between servers within the overall hierarchy. In one of the most popular protocols for email all emails created on an individual's personal computer are kept in a synchronised copy on the local server.

7. One consequence of this is that, depending on the precise set up, if a user decides to delete an email on their personal computer and then goes to the "deleted" folder into which the deleted email will first go and then deletes the copy of the "deleted" email, that email will for all practical purposes also disappear from the respective archive on the server. (At that point it may be possible to use forensic techniques to recover the deleted email either on the laptop or on the server; but either of these would be a non-trivial exercise and would only be possible for a relatively short periods after the deletion took place).

8. Replication has many other uses; whole folders can be updated and synchronised between servers. One advantage is that replication provides back-up in the event that a server suffers a failure, but a more important one is that it enables individuals dispersed about an organisation with multiple servers to have instant local access to information which may have been created elsewhere on another, distant server. Such arrangements can be particularly valuable, as in military networks, where communication links may not be available or "up" all of the time.