IRMA
INFORMATION RISK MANAGEMENT & AUDIT
◆ A SPECIALIST GROUP OF THE BCS ◆

JOURNAL
volume 14 number 5 winter 2004 ISSN 1741-4229

Programme for members' meetings 2004 – 2005

Tuesday 7 September 2004 Late afternoon	**Computer Audit Basics 2: Auditing the Infrastructure and Operations**	16:00 for 16:30 KPMG
Thursday 7 October 2004 Full day	**Regulatory issues affecting IT in the Financial Industry**	10:00 to 16:00 Old Sessions House
Tuesday 16 November 2004 Full day	**Networks Attacks – quantifying and dealing with future threats**	10:00 to 16:00 Chartered Accountants Hall
Tuesday 18 January 2005 Late afternoon	**Database Security**	16:00 for 16:30 KPMG
Tuesday 15 March 2005 Full day	**IT Governance**	10:00 to 16:00 BCS – The Davidson Building, 5 Southampton Street, London WC2 7HA
Tuesday 17 May 2005 Late afternoon AGM precedes the meeting	**Computer Audit Basics 3: CAATS** ***Preceded by IRMA AGM***	16:00 for 16:30 KPMG

Please note that these are provisional details and are subject to change.

The late afternoon meetings are free of charge to members.

For full day briefings a modest, very competitive charge is made to cover both lunch and a full printed delegate's pack.

For venue maps see back cover.

Contents of the Journal

GUIDELINES FOR POTENTIAL AUTHORS

The *Journal* publishes various types of article.

Refereed articles are academic in nature and reflect the Group's links with the BCS, which is a learned institute governed by the rules of the Privy Council. Articles of this nature will be reviewed by our academic editor prior to publication and may undergo several iterations before publication. Lengthy dissertations may be serialised.

Technical articles on any IS audit, security, or control issue are welcome. Articles of this nature will be reviewed by the editor and will usually receive minimal suggestions for change prior to publication. News and comment articles, dealing with areas of topical interest, will generally be accepted as provided, with the proviso of being edited for brevity. Book and product reviews should be discussed with the appropriate member of the editorial panel prior to submission. All submissions should be by e-mail and in Microsoft Word, Word-Pro, or ASCII format. Electronic submission is preferred.

Submissions should be accompanied by a short biography of the author(s) and a good quality electronic digital image.

Submission Deadlines

Spring Edition	7th February	Autumn Edition	7th August
Summer Edition	7th May	Winter Edition	7th November

The ***Journal*** is the official publication of the Information Risk Management & Audit Specialist Group of the British Computer Society. It is published quarterly and is free to members.

Letters to the editor are welcome as are any other contributions. Please contact the appropriate person on the editorial panel.

Editorial address:
47 Grangewood,
Potters Bar
Herts, EN6 1SL
Email: john@lhscontrol.com

Designed and set by Carliam Artwork, Potters Bar, Herts
Printed in Great Britain by PostScript, Tring, Herts.

Editorial

A letter from the cable company ntl arrived this week. The opening paragraph read "You may have noticed that since the beginning of September there have been some delays getting through to our Customer Service teams. This is because we've started a programme of service improvements to help us give you a better level of customer service next year". Yeah, right. They degrade the service while improving it! When will these clowns learn that any service improvement programme should not degrade what is already a very poor service. They also forgot to mention that they were half way through their improvement programme and had only succeeded in making things worse. True successors to BT!

I wanted to renew my subscription to PC World so I called their "telephone hotline" only to be told that they were doing "maintenance" on their database and could I call back in a couple of days? Whoever heard of doing maintenance during the working day? I am feeling like one of the grumpy old men in that TV programme, so I will just finish up with one of my perennial favourites at this time of the year, the J Sainsbury's so called "store locator". Every year I want to check the Christmas opening times of my local large store and every year I use the so called Sainsbury store locator to find it. Only I never can, because it doesn't turn up if I enter my post code and because they don't use fuzzy matching, unless you enter that period after the "St" of "St. Albans", which is the nearest town, you can whistle for one of their flagship stores. I have complained about this for years, but nothing happens and now I get messages that they are receiving a large number of emails so tough luck on getting a response. No wonder they have slipped down the rankings. No-one can find their stores. Do you remember the debacle of their Nectar loyalty card registration? You still can't register your card on-line, but they do provide a telephone number.

I decided to sign up for the new eVat service. They wanted five pieces of information from me, one of which was my exact original VAT registration date. Now I vaguely remember that it must have been sometime in the last quarter of 1989 when I set up my business, but the exact date escapes me. Tough luck, as all five pieces of information have to be correct and their help desk is closed at weekends. I also need separate PINs to change my VAT registration details and to submit my return. Do I see an auditor's hand in preventing me from registering for this service?

Compare all of the above woes with a mail order firm in the USA. I placed my order over the web, got immediate confirmation that it was an in-stock item, an immediate email that my order had been taken, a next day confirmation that it had left the warehouse, a tracking number in case I wanted to know where in the world it was and the item arrived three days later. So it is possible to get it right. I suspect you have a few stories on these issues. Drop me a line and air your grievances, or plaudits, in our letters column. No-one has written for ages and I'm feeling a bit lonely. Perhaps that is why I am feeling particularly grumpy, although in truth it is the content of the next paragraph that has brought on my mood.

Now for some good news and some very bad news. Jeremy Jaynes a US spammer has just been sentenced to nine years in jail for his activities. Jaynes was ranked as the eighth most prolific spammer in the US, sending out some 100,000 messages each month. What amazed me was that he has reportedly made £13 million from saps who responded to his various offerings. The very, very bad news? Marta Andreason, the brave EU whistle blower, was finally dismissed by the outgoing Commission after being suspended for over two years. They didn't have the guts to do it while they were still in their jobs, but waited until the last possible minute to avoid any further mud being thrown at them for a disgraceful dereliction of duty at a time when the European Court of Auditors reports that it cannot give assurance over the spending of 93% of the EU's £70 billion budget. So much for corporate governance and the (un)ethical behaviour of our unelected appointees. It makes me really angry that these people are appointed by governments who do not meet even their own economic rules for continuing membership of the euro. A depressing conclusion to the year.

On to more positive things, the content of this issue. Celeste Rush concludes her paper on computer forensics, while Bob Ashton our correspondent from the antipodes updates us on business continuity planning in Australia. Mark Smith has negotiated some excellent member only benefits and Greg Krehel provides some useful advice on identifying those really important things to prove your case.

The festive season is now upon us and I send you the compliments of the season and a productive new year.

John Mitchell

The Down Under Column

Bob Ashton – IRMA Oceania Correspondent

Australian Business Continuity Management Standard

The Australian Prudential Regulatory Authority (APRA) is the prudential regulator for the Australian financial services industry. APRA has recently released a draft Prudential Standard on Business Continuity Management (BCM) which will be mandatory for the entities it regulates.

The new Standard will replace the previous arrangements, which only required organisations to have in place a disaster recovery plan for information technology, and covers all significant aspects of BCM. As such it provides an excellent high level check list for any auditor contemplating the review of BCM.

The Standard is based on a whole of business approach and is prescriptive in nature. It will impose a statutory requirement on Internal Audit, or an external expert, to periodically review the Business Continuity Plan and provide an assurance to the Board that:

- The BCP is in accordance with the Business Continuity Management Policy
- Addresses the risks it is designed to control
- Testing procedures are adequate and have been performed satisfactorily

Organisations are required to report on compliance with the Standard in the Risk Management Declaration which they are required to make annually. Details of non compliance must be reported along with an action plan for rectification.

The following mandatory requirements of the Standard can be used by auditors anywhere as a standard of best practice:

1. Risk Assessment
   Organisations must identify plausible disruption scenarios that may lead to short, medium and long term disruptions to critical business functions and assess the likelihood of these scenarios occurring.

2. Business Impact Analysis
   A BIA identifying all critical business functions, resources and infrastructure and assessing the impact of a disruption on these must be created.

3. Recovery Strategy
   An appropriate recovery strategy, subject to a cost/benefit analysis, must be developed

4. Business Continuity Plan
   A Board approved BCP must be maintained. The BCP is defined as the documented procedures and information which enable the organisation to respond to a disruption, recover and resume critical business functions and return to normal operations in an orderly manner.

   The BCP must cover all business units, critical business functions, resources and infrastructure, including operations located interstate, subsidiary companies providing specialist services to the organisation and arrangements with service providers, to ensure a whole of business coverage. Minimum requirements:

   a. The procedures to be followed in response to a material disruption to normal business operations. These procedures should enable the organisation to manage the initial crisis and recover and resume the critical business functions, resources and infrastructure outlined in the BCP in the nominated timeframe.

   b. Detailed procedures for restoring normal business operations. These should include an orderly entry of all business transactions and records into the relevant IT systems and the completion of all verification and reconciliation procedures.

   c. A list of all resources needed to run operations in the event the primary site is unavailable. This would include computer hardware and software, printers, faxes, telephones, and stationery. Additional resources include suitably trained staff and relevant documentation such as insurance policies and contracts.

   d. A communication plan for notifying key internal and external stakeholders if the organisation's BCP is invoked.

   e. Consideration of business continuity as part of any proposed material outsourcing agreement with a third party service provider.

   f. Relevant information about the organisation's alternative site for the recovery of business and/or IT operations.

5. Testing
   Organisations must thoroughly test BCP on a regular basis (but at least annually or more frequently if there are material changes to business operations) to ensure that the BCP is capable of meeting its objectives.

6. Communication Plan
   A BCP should contain a communication plan for notifying key internal and external stakeholders if the BCP is invoked. This would include staff, regulators, customers, counterparties, service providers, market authorities and media. The communication plan should clearly identify the staff authorised to deal with the media if the BCP is invoked.

7. Outsourcing
   Business continuity should be considered as part of any proposed material outsourcing agreement with a third party service provider. Proper due diligence should be conducted in this regard as part of the decision making process when assessing service providers. The contract with the service provider should include a requirement that the service provider have a BCP and testing program in place and provide for regular reporting to the regulator.

8. Alternative Sites
   Alternative sites should be located at sufficient distance from the primary site to minimise the risk of both sites being unavailable simultaneously. In particular, alternative sites should not be on the same power grid or telecommunications network as the primary operational site. Where primary operations are located in the central business district of a major capital city, the alternative site would be expected to be situated outside that CBD in order to minimise the risk of both sites being impacted by a wide area disruption. A review and assessment of the capacity of the alternative sites should be undertaken at least annually.

ISACA IS AUDITING GUIDELINE

The Information Systems Auditing and Control Association has recently published an exposure draft titled IS Auditing Guideline Business Continuity Plan (BCP) Review from IT Perspective.

This complements the above document in that it provides in dot point form a more detailed list of items to be considered in reviewing some of the areas covered by business continuity management.

FOR MORE INFORMATION
www.apra.gov.au
www.isaca.org

Member Benefits Discounts

Mark Smith

We have negotiated a range of discounts for IRMA members see below:

Software

Product	*Discount Negotiated*	*Supplier*
Caseware Examiner for IDEA (mines security log files for Windows 2000, NT, XP)	15%	Auditware Systems (www.auditware.co.uk)
IDEA (Interactive Data Extraction and Analysis)	15%	Auditware Systems (www.auditware.co.uk)
Wizrule (data auditing and cleansing application)	20%	Wizsoft (www.wizsoft.com)
Wizwhy (data mining tool)	20%	Wizsoft (www.wizsoft.com)

Events

Event	*Discount Negotiated*	*Contact*
Computer and Internet Crime 2005 (www.cic-exhibition.com)	15%	Paul Webster paul@panpres.co.uk
All Unicom events (www.unicom.co.uk)	20%	Julie Valentine julie@unicom.co.uk

We are looking to extend this range of discounts to include additional events, training courses, computer software or other products that our members may find beneficial. If you have any suggestions for products we could add to the list, please contact Mark Smith (mark.smith@lhp.nhs.uk), our Members' Benefits Officer, and he will be happy to approach suppliers.

Creating & Using Issue Analysis Memos

Greg Krehel

This is the fourth in a series of articles which deal with best practice in compiling information required for civil or criminal litigation. The best practices described are equally relevant to the audit process. In this article Greg explains the basics of identifying the principle and supporting characters in a case – Ed

Introduction

Of the hundreds of hours you invest in a case, the handful needed to work up an issue analysis memo could easily be the most important. Make this simple case analysis tool a standard for every matter and you'll always have a tight grasp on issues and arguments and ultimately on the case itself.

Your ownership of case issues will permeate other critical activities — taking depositions, drafting briefs, evaluating the facts, reviewing documents, and so on. Preparing an issue analysis memo can also result in numerous less obvious benefits, e.g., it can be used to dramatically improve your demonstrative evidence.

Details on how to create an issue analysis memo and the many advantages of doing so follow below. Before proceeding, let me make the following disclosure: The ideas presented in this article do not require our CaseMap case analysis tool to implement, and the article itself makes no mention of CaseMap. However, CaseMap's issue spreadsheet does make it easy to create issue analysis memos, and, as such, I'm hardly a neutral party. That said, I hope and believe you will find that my issue analysis recommendations are built on solid reasoning and that it was this logic that shaped CaseMap's design, not vice versa.

Creating Issue Analysis Memos

Here are twelve pointers for conducting effective issue and argument analysis and developing an analysis work product.

A Complaint is Not the Answer

Over the better part of two decades, my partner, Bob Wiss, and I managed a trial consulting firm. In each of the 500+ matters on which we were engaged to conduct mock trials, our efforts began by requesting background materials that would educate us about the case and its issues. A handful of times, we received a work product the trial team had specifically prepared to summarize issues and arguments. In the remaining instances, the closest approximation of an issue summary was generally the latest amended Complaint and Answer.

These pleadings certainly do present the claims, counterclaims, and cross-claims that belong in an issue analysis memo. However, while a recitation of claims is necessary, it is by no means sufficient. There are numerous reasons a Complaint and Answer should not be substituted for an issue analysis memo.

First off, the Complaint and Answer focus almost exclusively on the top-level legal issues. They rarely specify the elements required to prove each claim.

A second problem with using pleadings in lieu of an issue analysis memo is that, once the cycle of amending the Complaint and Answer ends, pleadings become frozen in time. They're not working documents that can be used to capture evolving thinking over the many months or years leading to trial.

The third and most important shortcoming of relying on pleadings for a case issue synopsis results from the fact that pleadings will by definition fall into enemy hands. As such, we're not going to use them to display our thinking regarding best arguments and hardest-hitting evidence – knowledge that's essential to a proper issue analysis. In the majority of Answers, the defense's position regarding each claim is pithily presented as "Denied." This tactic makes complete sense given the Answer's true purpose, but also renders the Answer an absolutely worthless stand-in for an analysis report intended to clarify our arguments.

The bottom line is that any case worth filing or fighting deserves a purpose-built issue analysis memo that makes thinking regarding issues and arguments explicit.

Begin Before the Beginning

The filing of a Complaint is the gunshot that starts a case. But since plaintiff counsel authors this document, they've obviously been thinking about the matter and its claims for an extended period beforehand. As often as not, defense counsel is also stewing on the potential case long before a Complaint is filed, as the defendant is normally well aware of the dispute that may result in litigation.

You should start an issue analysis memo for each new case as soon as your noodling on the matter begins. It only takes a few minutes to jot down your initial impressions of case issues.

Be sure to trap all possible claims, counterclaims, and cross-claims, as well as any arguments that you're already aware could be made about them. In other words, get down the issues all parties are likely to introduce, not just your own.

Use early drafts of the memo to frame the Complaint or Answer, but then keep this analysis document hard at work until the case is resolved by settlement or trial.

The World of Issues and Arguments Isn't Flat

The best way to organize your issue analysis memo is as an outline, not a flat list.

An outline is the proper way to deal with the hierarchical relationship among claims and their elements. In an outline, elements are nested below the claims to which they relate. Visual presentation maps legal reality.

In contrast, a flat list masks the relationship between parent claim and child elements. Consider a Fraud claim. Proving Fraud requires a showing of these elements: Intent, Reliance, and Loss. In a flat list of issues, Fraud, Intent, Reliance, and Loss are

displayed on equal footing. This is a counterintuitive presentation of the issues for those who understand the parent-child relationship between claims and elements. And it's an extremely misleading one for clients and others not versed in the law.

An outline structure also provides the best way to organize thinking regarding themes and arguments. Here again, an outline handles the real world correctly and a simple listing falls, well, flat. Arguments are typically marshaled in support of our position on a claim or one of its elements. Using an outline, arguments are easily binned under the claim or element to which they relate. Conversely, if we added arguments to a list that already contained claims and their elements, we would just be compounding the mess caused by parent and child dimensions appearing on the same level.

A final problem with using flat lists to organize issue thinking is that they quickly grow into ungainly monsters. As argument ideas are added, a list just gets longer and longer and longer.

An outline is a far more elegant method for dealing with the increase of issue ideas over time. It can be viewed fully expanded or collapsed so that it hides all child nodes below a chosen depth. For example, an outline can be set to present only the handful of top level issues or just the top two levels.

Don't Let the Pendulum Swing Too Far

A multi-level outline beats the pants off a one-level list. Does it follow that an issue outline with many levels is better than one with just a few? No, definitely not. There's no reason to expect an issue outline that's six levels deep to be twice as good as one that's three levels deep. In fact, a six-level issue outline is likely to be counterproductive overkill.

In the first months of working up a case, it's best to keep the outline quite simple — two levels in most areas, maybe three levels in a few. In the early going, it's rare for the evidence to be clear or for the trial team's thinking to have gelled. As such, there's little basis to create a sophisticated issue hierarchy at this stage. Doing so would result in needless reworking, and would likely spook others on the trial team.

As the case proceeds towards trial, an issue outline can and should gain depth. Nonetheless, in even the most complicated case, it's rare to need an outline that's over four levels deep. Three levels is sufficient to capture claims as the top level, the elements of each claim as the second, and arguments that can be made in support of each element as the third.

Pardon Our Permanent Dust

An issue analysis memo is, by design, an unending work in progress. If everyone on the trial team understands this fact, you'll be free to use the outline to foster communication and thinking.

On the other hand, if you fail to set expectations properly, you'll get a fraction of the possible benefit from this analysis tool. You'll tinker with the outline, but won't share it with clients and other trial team members for fear of their negative reactions to something rough hewn. Or, you'll limit the issues listed to the most obvious and the least controversial. Definitely a case where the best is the enemy of the good.

In addition to making verbal efforts to set expectations, why not make an introduction that serves this purpose a part of every issue analysis memo? Here's a draft:

"Pardon Our Dust! Please note that the following issue analysis memo is a draft. We expect it to evolve substantially over the many months leading to trial. We use this document to capture even the roughest ideas so they may be evaluated, shared, and improved. You'll find that different portions of the outline are at varying stages of refinement. Again, please understand that this memo isn't intended to be a polished report, but rather a tool that helps us think and that provides a way to receive your valuable input. Thank you."

Include Key Factual Disputes

Most cases involve dozens, if not hundreds, of disputed facts. When you get right down to it, isn't a disputed fact in essence an issue? We claim the fact is true. The opposition claims it isn't. Or vice versa. All parties can present evidence in an attempt to persuade the factfinder to see things the "right" way.

Does this mean every disputed fact belongs in your issue analysis memo? No. The vast majority aren't critical to the way the factfinder is likely to decide the case. They should simply be entered in your fact chronology and flagged as disputed by one party or another.

While the majority of disputed facts can be dealt with in a fact chronology, the handful that emerge as case lynchpins should be treated as issues and added to your outline.

… And Extralegal Issues, Too

The bulk of every issue analysis memo will be devoted to legal claims and the elements and arguments related to these claims. But your outline should also trap thinking regarding the extralegal dimensions that may influence the way jurors and even the judge respond to the case.

Such extralegal issues are typically tied to the emotional reactions evoked by the plaintiff, the prosecution, and/or the defendant. Has the plaintiff sustained such grievous injuries that jurors' cognitive processes could be swamped by sympathy? Does a corporate defendant have a stained reputation in the community? If so, your issue outline deserves an issue on the topic.

Extralegal issues don't have to be of the "Elephant in the Room" class. Assume, for example, that you represent the defendant in a toxic tort matter where the plaintiff's damage demands seem excessive. In such a case, there might be good cause to include a Plaintiff Greed issue in your outline.

Greed obviously isn't a legal issue in the case. And it certainly isn't one you're likely to argue at trial. However, by including a Greed issue in the analysis memo, you're in a position to consider what facts, if any, would prompt jurors to see the plaintiff as motivated by avarice.

Naming Issues

One challenge when working up an issue outline is how to phrase or name each issue. The first instinct of many new issue analysts is to use a descriptive statement as the issue's name, e.g., "Third National Bank Breached its Fiduciary Duty to Hawkins." There's nothing wrong with this approach per se, but why not skip the formality and adopt whatever name trial team members would find most natural to employ as they discuss the case?

Use my Conversation Test to evaluate name candidates. Plug each issue name you develop into the following sentence: "Did

you learn anything important about Issue X at the Lang depo today?"

Let's try the Conversation Test using the hypothetical issue discussed above. "Did you learn anything important about Third National Bank Breached its Fiduciary Duty to Hawkins at the Lang deposition today?" That flunks. How about "Did you learn anything about Fiduciary Duty at the Lang depo today?" Much better.

The move to simpler issue names is made at the expense of having the name itself indicate such details as who allegedly wronged whom. However, as explained further in the following topic, this knowledge is better captured in ways that won't result in a monster moniker.

Get Some Meat on the Bones

To make the most out of an issue analysis memo, it should be more than a skeleton outline of issue names. Flesh out your thinking by capturing a detailed description of each issue and argument. Explain legal jargon and, more importantly, provide a summary of the key evidence regarding the issue.

Even if you're the only attorney working on a case, it's still worth spending the few minutes necessary to generate these issue descriptions. They'll be of great value to your client and to expert witnesses. And the very process of writing up a description helps clarify your thinking.

Creating descriptions also provides a way to test the issues you've defined. If you struggle to pen a good description, it may well mean that the issue needs to be recast.

The Good, the Bad & the Ugly

In addition to trapping a description of each issue, get down a paragraph or two explaining how you feel the evidence on each issue cuts. Is our position on the issue weak or strong and why do you feel this way? While making such an evaluation isn't essential in the early months of a case, the sooner you move from simply describing the issues to also evaluating them, the better.

Sometimes we get queasy about making our evaluations explicit. Once our opinions are out it the open, others get to play the critic and black-hat our analysis.

I'd argue that any client worth having will appreciate the fact you're pushing the analysis process forward. You'll get kudos for sparking a debate that can either be held before trial or as part of a postmortem following a courtroom disaster.

Here's a strategy to employ if you agree that evaluating case issues is critical, but still don't want to be the first up to the tee. Why not conduct a brainstorming session devoted to evaluating the strengths and weaknesses of our position on each case issue?

By holding such a session, you get everyone to lay down their cards at the same time. An issue analysis memo provides both a structure to guide the meeting and a container for the group's thinking.

Off-Off-Broadway

While thinking about case issues, you'll surely come up with new twists that don't fit neatly into the existing issue hierarchy. Provide a home for these brainchilds by adding a top-level item named "Ideas," or something analogous, at the bottom of your issue outline. Visit the Ideas node frequently to see if the issues in it can be matured and eventually promoted into the primary outline structure.

Please don't misconstrue this recommendation as implying that only fully-formed, "perfect" issues deserve to appear elsewhere in the outline. As hopefully made clear earlier in this piece, your issue analysis memo should be considered an Off-Broadway production. Think of your Ideas node as being Off-Off-Broadway.

This Ideas node is another big plus made possible by using an outline to organize your issue thinking. Your extra-rough ideas are easily dumped into an area that can be collapsed and hidden from view when not in use.

Employed at Will

Since becoming an advocate for case analysis a decade or so back, I've been asked to critique a good number of issue analysis memos. A common problem I've observed is that once an issue or argument has been added to the outline, it's rarely, if ever, removed.

Even though an outline makes it possible to deal gracefully with a large number of issue ideas, there's no reason to clutter the memo with issues that aren't pulling their weight.

Barring the judge's granting of some portion of a Motion for Summary Judgment, the issues related to claims and elements of claims have a guarantee of lifetime employment in your outline. But that's not true for the arguments you list under each claim or for the extralegal dimensions you may add as well. These items are employed at will and can be terminated at any point. Give them a semi-annual review and make sure they should be retained.

If giving issues and arguments the pink slip makes you uncomfortable, you can always transfer them out of your primary outline and demote them below the Idea category described in the "Off-Off-Broadway" topic above.

Putting Issue Analysis Memos to Work

Let's now shift gears from crafting an effective issue analysis memo to using it. Here are some of the ways an issue and argument outline can be put to work on your behalf.

Thinking Clarified

What's the most important reason to make issue analysis memos a standard practice?

Creating this document crystallizes your personal thinking about a case.

Our minds are incredible thinking machines. But try to consider more than a handful of items at once and that amazing mind of yours is sure to be overwhelmed. Getting thinking out of your head and into an issue analysis memo allows you to deal with case issues in mind-sized bites. Your memo makes it possible to back up from your thinking so you can reflect on and improve it.

Trial Team Educated

Want a great way to give clients, expert witnesses, and new trial team members a snapshot of the case?

Use your issue analysis memo as an instructional aid — by sending it along to be read independently or by employing it as a prop that structures a verbal case overview. If the case warrants, why not turn its issue outline into a PowerPoint presentation?

Consensus Built

Interested in a tool that helps the trial team achieve a common understanding regarding issues and arguments — both what they are and how they cut?

An issue analysis memo acts as a central repository for the team's thinking and makes areas of agreement and disagreement readily apparent. Once it's clear where thinking diverges, the task of reaching consensus becomes far easier.

New Case Closed

Want a great addition for your practice development kit?

When you meet with prospective clients, take along a sampling of the issue analysis memos you've prepared for completed matters. Why not also hand out the very roughest beginning of a memo for the prospective case? We're talking five minutes of work that can yield $500,000+ in revenue.

Demonstratives Refined

Interested in making the most of your courtroom visuals while also keeping the cost of these aids under control?

Use your issue analysis memo to validate demonstrative evidence ideas before they're produced for use in court. How? This benefit isn't as obvious as the ones covered to this point, so let me explain in some detail.

Before giving the green light to start production of courtroom graphics, print out a copy of your issue analysis memo and complete the following steps:

1. Obtain a list of all demonstrative evidence ideas the team plans to have produced as final graphics.

2. For each item on the list of planned demonstratives, review your issue analysis memo and determine which issues the visual will help communicate. Jot the name of these issues down next to the name of the visual.

3. When you bump into a visual that doesn't seem to support any issue, ask what purpose that graphic is going to serve. If you don't come up with a darn good answer, strike the idea. Smile — you've just saved $500 or more in graphics costs.

4. When you finish the issue evaluation of all demonstrative ideas, tally the number of visuals that relate to each issue. Then review these counts to see if they make sense. It's a good bet you'll find that some issues have too many visuals devoted to them and that others are naked of demonstrative support. If that's true, eliminate the weakest graphics planned for issues where things went overboard and cook up some new demonstratives for those issue areas starving for visual attention.

I'm not suggesting that every issue must have demonstratives that help communicate our position on it. Nor am I suggesting that visuals should be equally distributed across issues. However, by making conscious choices about how demonstratives are distributed across case issues, you'll end up with a particularly persuasive set of courtroom graphics.

New Associates Productive

Could you use a method for leveraging the value of new associates, while also providing them with superb training?

Turn your issue analysis memos into a baton that passes analysis responsibility along to your new hires. Here's the process:

1. Ask new associates to read a handful of issue analysis memos from prior cases (and maybe this article as well). Answer their questions regarding them.

2. Give these associates the Complaint, Answer, and any other appropriate paper for an ongoing matter. However, do not give them the current issue analysis memo for this case. Have them draft their own issue outline for the matter. Take a red pen to their efforts, and then sit with them to explain the mark-up. Finally, give them the actual issue analysis memo for the case so they may compare it to their draft.

3. Repeat Step 2 once or twice more.

4. Assign your apprentice issue analysts to a brand new case that's relatively simple and for which no issue analysis memo yet exists. Have the associates create the first draft. Critique what they've done and have them update. Make them responsible for pushing the issue analysis memo forward over the life of the matter.

Chronology Fleshed Out

Interested in a way to fill in gaps in your fact chronology?

The following issue-driven approach to brainstorming on case facts makes it easy to develop a comprehensive chronology. And, as explained below, if this method is employed early in a case, it can produce results that are invaluable during discovery.

Here's how to use your issue analysis memos to drive fact brainstorming:

1. Open a memo and focus on the first issue.

2. Search your mind for any and all facts that support your position on this issue. Enter each of them into your fact chronology. Assuming discovery is still open when you conduct this brainstorming exercise, don't limit your thinking to facts with solid sources. By getting down prospective facts, i.e., those with undetermined sources, you're developing a shopping list that can guide your discovery efforts.

3. When you've exhausted the facts that support your position on the issue, step into the opposition's shoes and conduct the exercise from their point of view. Again, push the envelope.

4. Repeat steps 2 and 3 for each issue in your outline.

Please give this fact brainstorming tactic a spin as soon as you draft your first issue analysis memo. I believe you'll be very impressed by the results.

Facts and Documents Organized

How else can an issue analysis memo assist in other case analysis efforts?

It can improve the fact chronologies and document indexes that (I sure hope) you're creating for each case.

Include a Linked Issues column in your fact chronology spreadsheet. Use it to list the names of the issues on which each fact bears. Add an equivalent column to your document index, and capture your assessment of the issues to which each document relates.

Assuming your fact chronology lives in a database program, once facts have been linked to issues, you'll be able to filter the fact display down from all facts to just those relevant to a particular issue. Ditto for the document index.

When you enhance a fact chronology and a document index as described above, you're simultaneously improving issue analysis by making it easy to identify the specific evidence related to each claim.

Warning: I've seen a fair number of trial teams issue-code facts and documents without first having carefully defined the case issues themselves. Talk about a recipe for disaster! Be sure to work up a solid issue analysis memo before someone starts issue coding facts and documents willy-nilly.

Future Cases Leveraged

Could you use an analysis tool that becomes all the more valuable over time?

If you make issue outlines standard practice, you'll have one. Sure issue analysis memos have a dramatic impact on the cases for which they're originally created. But they also become a fantastic resource you'll turn to again and again in future matters. Work on the next analogous case gets off to a faster start as you have a library of argument ideas to fuel the issue analysis process.

Conclusion

As I hope you can tell by this point, I'm in favor of thinking hard about our cases. However, let me wrap up by arguing against thinking just this once: Please don't spend another minute thinking about the issue analysis memo concept. Act!

Please pick a case and hack out the first draft of an issue analysis memo for it today.

About the Author

Greg Krehel is CEO of DecisionQuest's CaseSoft division (www.casesoft.com). CaseSoft is the developer of litigation software tools including CaseMap and TimeMap. He can be contaced at gkrehel@casesoft.com.

Computer Forensics Science – part II

Celeste Rush

This is the second of a two part series dealing with this very interesting subject – Ed

4 Storage media

Storage media usually refers to the permanent storage of data in a non-volatile way; that is retains the data without the use of electrical power. There are many different types of media that store data in a more or less permanent way. These include:

- Floppy disks and hard disks,
- Compact discs (CDs),
- Digital versatile/video discs (DVDs),
- Tape,
- Flash memory,
- Zip and Jaz disks,
- Microdrives,
- Magnetic optical

Many of these media work using the same principles. An analysis of both floppy and hard disks would help to understand these principles which then may be applied to the other various media forms.

1. Floppy Disk Systems

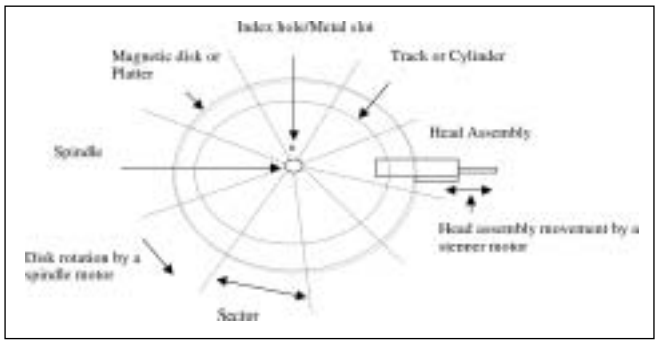

Figure 18 - Model of a Floppy Disk Drive Unit

The construction of the floppy disk unit is quite simple in structure and concept. A circular flexible (hence the term 'floppy') plastic disk or *platter* is coated on both sides with a magnetic material and is encased in a protective envelope or hard plastic covering. When inserted into the drive the disk is locked into place on a *spindle* and is rotated at a constant speed by a *spindle motor*. The *head assembly* consists of two read/write heads which both make contact with the disk on opposite sides. These heads are moved in discrete steps across the disk by means of a *stepper motor*. As the heads are in direct contact, the speed of the rotation must be limited to protect the lifespan of the disk and disk drive unit. The standard 3.5-inch, 1.44-MB floppy rotates at a maximum of 360 revolutions per minute.

The magnetic material that is passed under the upper and lower heads during the disk rotation is seen as very narrow circular strips called *tracks*. The two tracks together are called a *cylinder*.

Data is electromagnetically written to the disk on a given track by one of the two heads for the duration of one rotation. The number of discrete steps available to the stepper motor will determine how many cylinders are possible. To read the data, the appropriate head is moved to the correct step position and reads the disk surface electromagnetically for one rotation.

The *index hole*, used in older floppy disks, marks the beginning and end of a track. The 3.5 inch 1.44 MB floppy provides a metal slot to be mated with a pin on the drive spindle, which registers the disk/spindle motor relationship and, through timing pulses, determine the start and end of tracks. (See figure 18)

Normally, a track is too long large for storing data so the disk is divided into a number of equal sized *sectors*. Older 5.25-inch floppy disks use 9 sectors per track (spt) and later ones use 15. The standard 3.5-inch 1.44-MB floppy has 18 sectors per track.

Timing pulses can determine the position of each sector. In this way, any information unit may be uniquely identified and accessed. The addressable elements are identified through a *physical CHS address*: C for the cylinder number (from 0 – normally outermost), H for the head number (from 0 – normally topmost), and S for the sector number (from 1).

1.2 Formats and FM – MFM Encoding

The *format* of the information recorded is important in helping to identify sector addresses. The format will determine *what* sequences of *bytes* are used to represent the required data structures. An *encoding method* determines *how* bits are encoded on the magnetic surface.

Bits of 1's and 0's are written to a magnetized disk through a flux change in the electromagnetic frequency, usually a reversal. A serial sequence of bits are written and interpreted as 8 bit bytes on a track. When the magnetic head reads a magnetized area it will distinguish between the two bits through the flux changes in the magnetization which will produce a signal. However, if there happened to be a string pattern of only 1's or 0's then the 8-bit string of would not generate any signals to be 'read'. To remedy this situation, an encoding format is used to ensure the correct reading of the number of flux changes in the magnetization regardless of the bit pattern.

There are two encoding formats used by floppy disks: Frequency Modulation (FM) and Modified Frequency Modulation (MFM).

In Frequency Modulation, information is stored in a *bit cell*, which consists of a *clock* bit and a *data bit*. Each flux change through a reversal of magnetization, or transition, will set either a clock bit or data bit so that 1 bit is encoded as two transitions and a 0 bit is encoded as one transition. Therefore, there is a clock bit present for every bit cell and a data bit present only where there is a 1 in the data stream but not where the data stream contains a 0.

For example:

The hexadecimal number 42h, representing the ASCII letter 'B', is written as 0100 0010. This is shown in figure 19 with two 0 bits added on each of the sides.

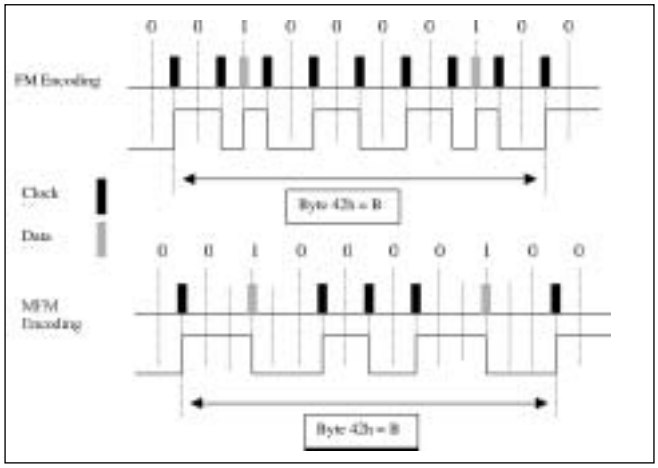

Figure 19 - FM and MFM Encoding

In FM encoding the magnetic flux reversals are shown with the clock bits for each bit cell and two set data bits for the byte 42h. However, many of the clock bits are redundant and as a result less information density is possible.

The Modified Frequency Modulation encoding, also shown in figure 19, offers twice the amount of information density than FM. MFM was established from two rules:

1. A data bit is always written if it is equal to 1.
2. A clock bit is only written if the proceeding bit cell as well as the current bit cell does not have a data bit set.

[Sammes et al 2000, pp 94]

MFM encoding was used on the earliest hard disks and is still the standard for floppy disks today. As MFM doubles the capacity of storage than FM, these disks became known as *'double density'*. Although hard disk technology has moved on, the standard 3.5-inch 1.44-MB floppy disk has remained basically unchanged for over a decade. This is largely due to the fact that the standard floppy disk has become and indispensable and inexpensive media to save and transport small files to use on virtually any PC regardless of the platform.

For each FM and MFM encoding methods a slightly different low-level formatting is specified. However, it is still a requirement to have for each track on the disk, a *start of track* sequence of bytes, a *sector format* sequence, which includes the CHS address and the data block for each of the sectors, and an *end of track* sequence of bytes. *[Sammes et al 2000, pp 94-95]*

2. Hard Disk Systems

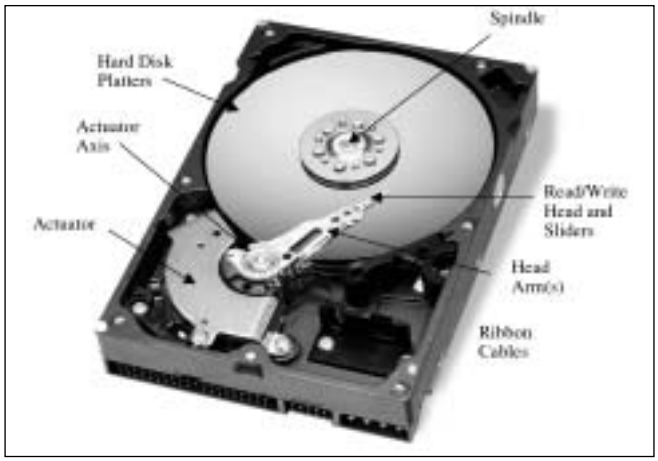

Figure 20 – Inside of a Hard Drive

Hard disks are usually the primary permanent storage media in a PC. Invented in the 1950s, hard disk, originally called "fixed disks" or "Winchesters", started as large disks up to 20 inches in diameter, but could only hold a few megabytes of data. Modern hard disks are now capable of holding hundreds of gigabytes at the fraction of the costs of those first created. Although technology has made vast improvements in data storage the design of the hard disk has remained basically the same.

There are several components in a hard drive. These include platters, spindle motor, actuator, head arm(s), read/write heads, logic board and ribbon cables. What they are and how they work will be considered in turn.

2.1 Platters

Hard disks comprise of one or several flat, mirror smooth, round disks called *platters*. The size of the platters in a hard disk will give the overall physical dimensions known as the *form factor*. Every platter in a specific hard disk will have the same diameter and are arranged stacked one on top of another, held together by a *spindle* that runs through a hole in the middle of each platter. Older disks were 5.25". The most common platter dimension now is the 3.5". Other sizes include 2.5" or 3.5" for laptops, 1.8 or 1.3 for PC cards, and 1.0 for the Compact Flash.

Similar to the floppy disk, both sides of each platter are generally used for information to be written to within concentric circles called *tracks* and the set of parallel tracks on each of the platters is called a cylinder. At any location of the head positioning arm, all tracks under all heads are one *cylinder*. The cylinder number is one of the three address components required to find a specific address. The other two are the head number and sector number (CHS addressing). A sector is a smaller unit of track whose size is determined by formatting but normally holds 512 bytes. A combination of two or more sectors on a single track is called a *cluster* – the basic storage unit of a disk. [Casey 2001]

When used as an address component, the sector and location refer to the sequence number of the sector around the track. Typically, one sector stores one user record of data. The number sectors per track used is dependent on the system type, the *controller* (a component in a PC that controls mass storage devices) capabilities, the drive encoding method and interface used (discussed below). [Seagate 2003]

The *capacity* of a hard disk is the amount of memory, which can be stored on a disk (or on a tape drive's data cartridge) and is measured in megabytes (MB) or gigabytes (GB). This is usually given as *formatted* capacity. The *unformatted* drive byte capacity is equal to the product of the bits per track (bpi), number of heads, and number of cylinders. [Seagate 2003]

2.2 Platter Substrate

The underlying material that supports the media layer containing the recorded data is known as the *substrate*. This material must be rigid but lightweight, easy to work with, inexpensive, and magnetically inert. Traditionally an aluminum alloy has met these criteria. However, as the speed of the platters increase and the *fly height* (also known as the *head gap* or *floating height*) of the *read/write head* gets lower, it has been necessary to look for alternatives to aluminum that will provide a smoother surface to prevent *head crashes*. (See 'Heads' below) These alternatives include glass, glass composites, and magnesium alloys.

Glass substrates, although more fragile, provide a greater rigidity

and uniformity of the magnetic surface and can be made much smoother and thinner than aluminum, allowing faster spindle speeds. However, the magnetic layer is harder to deposit on glass, making it more difficult to achieve the same read densities.

To address this problem, IBM have recently developed the use of antiferromagnetically coupled (AFC) media, which contains a second magnetic layer oriented opposite the primary layer to reinforce the magnetic orientation. This added layer enables a greater amount of data that can be stored in a given area on the hard disk platter measured as the *areal density*.

2.3 Areal density

Track density is a measure of how many tracks can be put down within an inch of radius on the platters. The bit density, also known as *linear* or *recording density*, measures how closely the bits are packed within a length of track. [Kozierok 2001]

The areal density is the product of the bit density, or bits per inch (BPI), and the track density, or tracks per inch (TPI). As the surface of a disk is two-dimensional this is usually expressed in bits per square inch (BPSI). *[HardDiskInfo (undated)]*

Areal density is illustrated in figure 21.

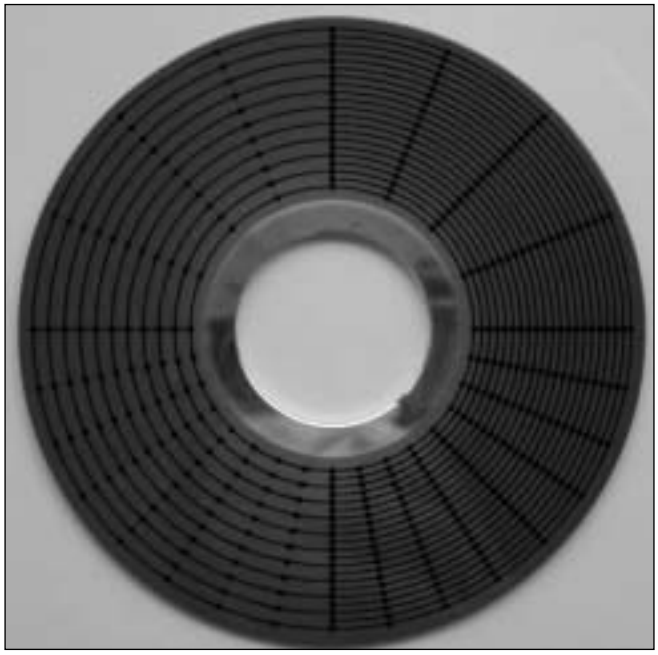

Figure 21 – Areal density illustration

[The PC Guide – Charles Kozierok 2001]

The left half of the circle illustrates low track density and the right half shows high track density. The upper half shows low bit density and the bottom half shows high bit density.

Put together, the upper left quadrant has the lowest areal density, the upper right and lower left have the greater density and the bottom right has the highest density. [Kozierok 2001]

2.4 Spindle

The *spindle motor*, or *spindle shaft*, is the rotating hub structure to which the disks are attached. It must provide stable, reliable and consistent turning power for thousands of hours of continuous use, in order for the hard disk to function properly. Many drive failures are actually failures with the spindle motor, not the data storage systems.

For over 15 years the maximum spindle speed was 3,600 revolutions per minute (rpm). Since then the speeds have increased up to 4500 rpm, 7200 rpm, 10000 rpm and most recently up to 15000 rpm. This increase of speed has presented an engineering challenge. Problems encountered include the read/write head going off track, causing a mis-read and a *rotational miss* which will then require another full rotation before the read can be tried again. An off-track head during writing can cause noise or even the overwriting of an adjacent track.

2.5 Actuator and head arms

The *actuator* is a device used to position the *head arms* synchronously to different cylinders on the platter. It is a very important part of the hard disk and is the only operation on the hard disk that requires active movement in changing heads electronically and changing sectors. Speed of change is of utmost importance in any physical motion as it is typically 1000 times slower than any electronic action. [Kozierok 2001]

There are two basic types of actuators: stepper motor actuator used in floppy disks and older hard disks, and voice coil actuator, used in modern hard disks.

A stepper motor actuator stops only at predefined "steps" as it rotates either clockwise or counterclockwise. The head assembly arms will move in or out one position accordingly. Each position defines a certain track on the surface of the disk. This is known as an absolute positioning system.

In older disks, the use of a stepper motor had several drawbacks. If a disk expanded or contracted, there was no way to compensate for the heads coming out of synchronisation with the tracks and data loss would result. This also meant that they had to be low-level formatted regularly to ensure proper head/track alignment.

In contrast, the voice coil actuator is known as a relative positioning system. The voice coil dynamically moves the head assembly in and out over the surface of the platters very accurately to any particular cylinder position in a closed-loop feedback system called a servo system. This is done through electromagnetic attraction and repulsion. A coil is mounted within an assembly containing a strong permanent magnet. When current is applied, an electromagnetic field is generated that will move the heads in a particular direction depending on the attraction or repulsion of the magnet. The actuator changes position by rotating on an axis.

As well as better accuracy, the seek time, or time it takes for the read/write heads to travel from one cylinder to another including the settling time, is much faster in a voice coil system than the stepper motor system.

Track density is very low in disks with a stepper motor actuator in compensation for the wide track width needed to ensure the head would be able to locate them even if slightly misaligned. A 1.44-MB standard floppy disk, which still uses the stepper motor actuator, has a track density of 135 tracks per inch (TPI), whereas modern disks using the voice coil actuator are achieving an average of over 67000 TPI or 35.3 Gigabytes/inch2. However, due to new developments within the area of magnetic media, that figure is continually rising. [IBM Research News (undated)] (See 'Platters' above)

2.6 Head(s)

One of the main differences between a floppy drive and a

hard drive is that the read/write heads do not touch the surface of the disk except when the drive is switched off. As the spindle motor rotates the disk, the heads, all attached to the same *actuator*, fly freely above the surface supported by aerodynamic pressure. When the disk is powered down the heads are 'parked' automatically on an unused track. [Sammes et al 2000, pp 99]

There are several types of head: *monolithic, composite, thin-film, magnetoresistive* (MR), and *Giant MR* (GMR). The monolithic head was one of the first types manufactured, which is made of a single block of ferrite, a ceramic material. The composite head was an improvement on the monolithic and consists primarily of non-magnetic material with a small ferrite core.

Inductive thin-film heads read and write data using a conductive coil wrapped around a very small magnetic stylus and is produced through a process similar to microchip fabrication. [MediaTek 1998/1999/2000]

Magnetoresistive (MR) technology, designed to support very high track densities, uses a special material in which the electrical resistance changes in the presence of a magnetic field. As it passes over the magnetic patterns on the disk it determines the strength of the magnetic field is determined an electrical pulse is created in response to flux reversals. However, MR technology cannot be used for writing and so is mounted alongside a conventional inductive thin film write element. [*Ibid*.]

The Giant Magnetoresistive (GMR) read element, first introduced by IBM in 1997, is now used in virtually all disk drives. It differs in MR technology in that it uses a 'spin valve' to create a much stronger electrical signal enabling much more information to be stored on the hard disk [IBM undated]

See Internet <http://www.research.ibm.com/research/gmr.html > for a detailed description on how GMR technology works.

If the heads contact the surface of the disk while it is at operational speed, a loss of data and damage to the surface of both head and disk may result. This event is known as a *head crash*. Although a lower head *fly height* (also known as the *floating height* or *head gap*) is needed to read a higher areal density on the disk, head crashes are avoided through a better hard disk enclosure to eliminate contamination and special mounting techniques to eliminate vibration and shock.

IBM provides a description of the chain of events following a write command:

"When a command is made to store some data on a disk, the following chain of events occurs:

- *The data flows into a cache where it is encoded using special mathematically derived formulae, ensuring that any subsequent errors caused by noise can be detected and corrected.*
- *Free sectors on the disk are selected and the actuator moves the heads over those sectors just prior to writing. (The time it takes the actuator to move to the selected data track is called the "seek" time.)*
- *Once over the data track, the heads must not write the data until the selected free sectors on that track pass beneath the head. This time is related to the rotation speed of the disk: the faster the speed, the shorter this "latency" period.*
- *When it's time to write, a pattern of electrical pulses representing the data pass through a coil in the writing element of the recording head, producing a related pattern of magnetic fields at a gap in the head nearest the disk. These magnetic fields alter the magnetic orientations of bit regions on the disk itself, so the bits now represent the data."*

[IBM (1996) 'Meet your hard Drive'
Copyright 1996 IBM Corporation, available at Internet
http://www.research.ibm.com/research/gmr/basics.html]

3. Logic board or IDE/ATA

Figure 22 – The back and front of a Maxtor-DiamondMax Plus 9 ATA-133 80GB 7200 rpm Hard Drive

All modern hard disks are made with an intelligent circuit board integrated into the hard disk unit to communicate with the rest of the computer via the data bus. This is known most commonly as the Integrated Drive Electronics (IDE). However, as other systems also integrate the controller electronics onto the drive, such as the Small Computer Systems Interface (SCSI), this term is rather misleading. The term *Advanced Technology Attachment* (ATA) is more appropriate and defines the standard interface between the hard disk IDE system and the AT-style PC to which it is connected. [Sammes et al 2000]

The original ANSI standard first defined in 1991 called ATA-1 defined the physical, electrical, transport, and command protocols for computer storage devices. Revised standards were subsequently drawn up to take into account rapid improvements in disk technology. The ATA-3 standard is a minor revision of the ATA-2 standard (Enhanced Integrated Drive Electronics or EIDE) and which introduces the *Self-Monitoring and Reporting Technology* (SMART). Ultra-ATA, another standard (ATA-4), refers to the use of a higher speed *Direct Memory Access* (DMA) transfer mode running at 33.3 Mbytes/second. Ultra-ATA is also called *Ultra-DMA* or *Ultra ATA-33*. Since the ATA-33, there have been further standard updates: Ultra ATA/66 (ATA-5), Ultra ATA/100 (ATA-6) and most recently ATA/133. [RITECC undated]

Another ANSI standard is the ATA Packet Interface (ATAPI) which allows CD-ROM and tape drives to be plugged into the standard IDE interface to be configured as master or slave, as one would do in a hard disk.

The logic board or IDE is normally mounted on the bottom of the base casting, exposed to the outside (See figure 22). It is separated from the base casting using foam or other cushioning material. A flexible ribbon cable links the read/write heads to the logic board (See figure 20).

The entire hard disk, collectively called the *head-disk assembly*, is mounted into a physical enclosure designed to protect the heads and platters from dust and contaminants present in the outside environment. Small *breather holes* are used in many drives to allow air to pass between the inside of the drive and the outside environment. This enables equilibrium of air pressure should the drive be moved to a different altitude or environmental temperature. A permanent *breather filter* covers the holes to prevent dirt or dust from entering the internal chamber. The airflow within the sealed chamber is created from the rotation of the platters and an added *recirculating filter* is used to catch any minute bits of debris should it somehow make it into the chamber. If the assembly is opened, the drive will become contaminated and unusable.

4. Encoding methods and formats

Hard disks use the same approach to encoding and low-level formats as the floppy disk. However, due to the development of the speed and capacity of the hard disk new encoding methods were developed. There are three encoding methods used:

- Modified Frequency Modulation (MFM)
- Run Length Limited (RLL)
- Advanced Run Length Limited (ARLL)

Again, slightly different low-level formats are used, including differences between hard disk manufacturers.

With RLL encoding, there are no clock bits in the recorded stream. Instead the run length of any sequence that can occur without a 1 bit is limited to something that is short enough for the controller to maintain synchronization. RLL 2,7 is the most widely used method where there are at least two and at most seven 0 bits between any two 1 bits in the recorded stream; thus 2,7. There are other RLL methods, which include RLL 1,7 and RLL 3.9 and are referred to as Advanced RLL or ARLL.

In RLL encoding there is no clock bit, and so there is a flux density saving in the region of 3 to 1 over MFM. However, because the bit pattern size is also increased, there is a loss of 2 to 1. So, the overall net gain is about 1.5 to 1 [Sammes et al 2000, pp 101]

One problem with the RLL method is that if synchronization is lost, the result could be a burst error of up to five bits (RLL 2,7). For this reason RLL controllers will use *error correcting codes* (ECC) in their formats rather than the simpler cyclic redundancy checks (CRC) used in the MFM format.

4.1 Hard Disk Low-Level Format

As with the floppy disk, the first hard disk low-level format requires, for each track on the disk, a start of track sequence of bytes, a sector format sequence including the CHS address and the data block for each of the sectors, and an end of track sequence of bytes. They are recorded sequentially along the track in 8 serial bits.

The format itself is controller dependent and format performed by one controller cannot be utilised by another. [Sammes et al 2000, pp 101]

Part of the encoding process involves a byte used as a *sector flag*, which enables the controller to perform *bad sector mapping*.

4.2 Bad Sector Mapping

At the time of manufacture, micro-defects are introduced on the magnetic surface of the disk. There are two types of defects: hard, which usually relates to a surface problem, and soft, which is normally due to a magnetic anomaly. The hard defects can be easily determined by low-level formatting software. However, the soft defects are only found through repeated testing with specialist test equipment. A *manufacturer's defect list* is produced and is often listed on the casing of the drive as a *defect label*.

Each bad sector is specified by its cylinder and head number, which define the track, and the *bytes from index* (BFI) or *byte count after index* (BCAI) value. This specifies the number of bytes from the start of track index mark to where the bad area exists. This index mark is considered the absolute point of reference for the BFI value. [Sammes et al 2000, pp 103]

Because of these defects, it is necessary for the controller to perform a *bad sector mapping* during the low-level formatting process and remove the effects of the defects. This is done in different ways.

If the controller identifies a defect in only a single sector it will shift the sector formatting slightly so the defect occurs in one of the sector or track gaps where it will have no effect. This is called *sector slipping*. If this is not possible then the sector flag 0 bit will be set at 1 to signify that the sector concerned is bad and no access to this sector will be allowed. Sometimes it is possible to format an additional spare or replacement sector on the track to avoid loss of formatted disk capacity. [Sammes et al 2000, pp 104] (See figure 23)

If the controller identifies defects in more than one sector in a track it can mark the entire track as bad by setting bit 1 of the sector flag to 1 in all the sectors of the track so no access will be allowed to any of them. This would effectively create a "hole" in the disk address space where the bad track resides.

Another way to deal with this scenario is for bit 2 to be set to 1 in all sector flags and assigning an alternative track starting from the highest cylinder number. The CHS address values of the ID fields in these bad sectors will now point to the alternative track and so will leave no "hole" in the disk address space even though there is still a loss of formatted disk capacity. (*Ibid.*)

The sector flag details and corresponding bits are illustrated in figure 23.

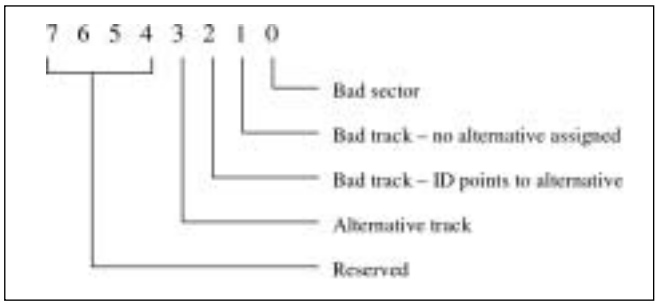

Figure 23 – Sector Flag details [Sammes et al 2000, pp 103]

The first cylinder is usually the one reserved for bad sector mapping and may be assigned an artificial cylinder number '-1'. This area would be used to store the defect list, which would be only accessible by the controller through specific internal commands. There are often two parts to the list: the original manufacturer's list and a *grown defect list*, which may be added to by the subsequent use of vendor specific bad sector mapping commands.

4.3 Zoned bit recording (ZBR)

In earlier disks every track had the same number of sectors. As the tracks are arranged in concentric circles, this arrangement would mean that the tracks on the outside of the platter were much larger than those on the inside. The area constraint of the inner tracks meant that they were packed as tightly as possible and the outer tracks were set to use the same number of sectors by reducing their bit density thus greatly wasting available space given the same bit density limitations.

To utilise this wasted space, modern disks employ a technique called *zoned bit recording* (ZBR), also known as *multiple zone recording* or *zone recording*. This technique involves grouping tracks into zones based on their distance from the centre of the disk. Each zone is assigned a number of sectors per track which increases as one moves outward from the centre.

However, since the angular velocity of the disk is constant and there are more sectors on the outer tracks than the inner, then the data transfer rates for the outer tracks must be higher than those for the inner tracks. [Sammes et al 2000, pp 109]

A side effect of this is that the benchmark tests run on disks when new and then used for sometime will show the disk is getting slower. What has really happened is that the later tests are being run on tracks that are closer to the centre than when new as the file system begins from the outermost cylinder inwards. Table 3 shows the 15 zones of a 3.8 Gbyte Quantum Fireball hard disk with the different data transfer rates for the various zones to compare.

Zone	Tracks in Zone	Sectors per Track (spt)	Data Transfer Rate (Mbits/s)
0	454	232	92.9
1	454	229	91.7
2	454	225	90.4
3	454	225	89.2
4	454	214	85.8
5	454	205	82.1
6	454	195	77.9
7	454	185	74.4
8	454	180	71.4
9	454	170	68.2
10	454	162	65.2
11	454	153	61.7
12	454	142	57.4
13	454	135	53.7
14	454	122	49.5

(From Quantum Fireball TM Product Manual, © 1996 Quantum Corporation.)
Table 3 – Zoned bit recording

The same number of zones and tracks per zone is not a requirement and the number will differ between manufacturers.

Although the CHS address assumes a constant sectors per track value for the whole disk, the CHS remains unaffected by zoning as the controller will translate the address into its own internal zoned address. This means that knowledge of the zoning structure within the disk system is not necessary to address it as zoning is part of the low-level formatting data structure and is hidden from view.

4.4 Head and Cylinder Skewing

Head and cylinder *skewing* occurs in hard disks with more than one read head. The problem it solves is the time delay between changing heads and cylinders and continuing to read sectors consecutively, a gap between which would require a full revolution before that sector would be accessible again. The time factor is greater for cylinders and so the skew will be bigger than heads.

For example, the sequence for accessing consecutive sectors would be:
Cylinder 0 Head 0 Sectors 1 to 17,
Cylinder 0 Head 1 Sectors 1 to 17,
Cylinder 1 Head 0 Sectors 1 to 17,
Cylinder 1 Head 1 Sectors 1 to 17, etc.

Skewing arranges sectors further around the track so the sequence might read:
Cylinder 0 Head 0 Sectors 1 to 17
Cylinder 0 Head 1 Sectors 16 to 15, Head skew = 2
Cylinder 1 Head 0 Sectors 8 to 7, Cylinder skew = 8

This is illustrated in table 4

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
16	17	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
8	9	10	11	12	13	14	15	16	17	1	2	3	4	5	6	7

Table 4 – Head and cylinder skew
[Sammes et al 2000, pp 110]

Sector 1 of head 1 is directly under head 0 sector 3, which will give a two sector gap before head 1 reads sector 1; thus, a head skew of 2. Likewise with cylinder 1 where sector 1 is directly under sector 9 of cylinder 0 head 1 giving an eight sector gap before sector 1 is read again; thus, a cylinder skew of 8. The skew should allow sufficient time to move the head and head assembly without loss of efficiency.

4.5 Interleave Factor and Memory Cache

More relevant to older disks, interleaving is used to overcome the problem of reading two successive sectors without sufficient time to transfer the buffer across from the first into the second by the time it is rotated under the read head. This problem may be due to a slow interface or the fact that error-correcting codes are being used and a whole sector needs to be available in the *buffer* before it can be checked. [Sammes et al 2000, pp 106]

The solution is to physically place the second sector further on. The interleave factor would be the number of sectors in between the consecutive sectors so they read in a numerical sequence. For example: interleave factor 1:1 indicates no interleaving, 2:1 indicates there is one sector in between, 3:1 indicates there are two sectors in between and so on. The interleave factor is set at the low-level formatting process and may be specified to obtain optimum efficiency.

This problem is overcome in modern disks through the use of a memory *cache* that may hold an entire track in the buffer at one time. A *cache buffer* is an integrated solid-state memory inside the disk drive which stores the results of recent reads from the disk and will anticipate the request of information by including sectors before and after the first request. The use of a cache improves performance of the hard disk by reducing the number of physical accesses to the media on repeated reads.

[Kozierok 2001]

5 Formatting process

One of the major differences between a hard disk and floppy disk is the formatting process. While the floppy disk only has two stages, low-level and high-level formatting, hard disks have three stages: Low-level formatting, partitioning and high-level formatting. Low-level formatting is to establish the tracks and sectors with all the address marks, IDs and gaps. Partitioning enables more than one logical volume to be associated with a disk. High-level formatting establishes the file system for the particular operating system.

1. Floppy disk formatting

The standard DOS program FORMAT performs both the low and high-level formatting function. The low-level format function is performed first by writing, to each track, the gaps, sectors, address marks and CRCs and with all sector data areas filled with a pattern of f6h bytes as shown in figure 24. Each track is read and checked for any bad sectors.

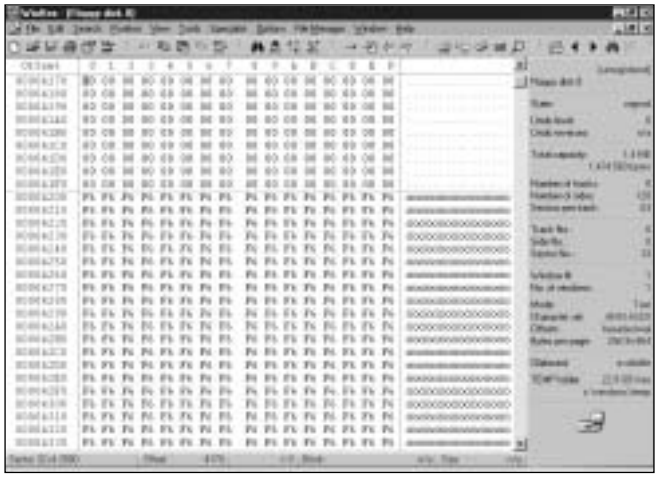

Figure 24 – Pattern of F6h bytes in floppy disk low-level formatting

The high-level formatting is performed next by where the MS-DOS file system will be set up on the disk. This involves writing two *file allocation tables* (FATs) and the *root directory*. Any bad sectors will be marked appropriately.

The disk can be made *bootable* if when initiating the FORMAT program the /S option is selected. This will copy the system files IO.SYS and MSDOS.SYS onto the floppy disk. [Sammes et al 2000]

The last step is the prompt to add a volume label before terminating.

2. Hard Disk Formatting

Unlike the floppy disk, a low-level format cannot be done with a FORMAT program on a hard drive but it is still used in establishing the MS-DOS file system during high-level formatting. Low-level formatting is performed at the manufacturer where the factory written defect list (both hard and soft) can be produced as well as any skewing values or optimised interleave.

More recent disks can be in one of two modes: physical and translation (or non-physical geometry). Disks that use the zoned bit recording, using an internal address translation, will be in translation mode so a low format will not have any effect on the defect mapping files or skewing values as the structures will not be rewritten. However, this kind of formatting will be destructive to any data on the disk. This is referred to as *intermediate* or *mid-level formatting*. Reasons to do this level of formatting may be due to a contracted virus or increasing bad sectors as well as the removal of an operating system.

Some manufacturers supply a special program for low-level formatting or there may be utilities built into the BIOS of the PC or hard disk and accessed using the DOS DEBUG command. Only those with a high level of expertise are expected to use these programs.

Partitioning the drive involves dividing the disk into logical pieces or volumes. This is done even if there is only one partition. Also, some operating systems can combine several partitions of several disks as a single volume (identified by a single drive letter) to achieve enhanced performance or fault tolerance. The standard formatting program is FDISK but others programs are available such as Partition Magic from PowerQuest.

The high-level formatting is performed to establish the *file system* for the particular operating system. Similar to the floppy disk, two copies of the *file allocation tables* (FATs) are written as well as the *boot sector* and *root directory*. However, there is no need to mark the FATs with any bad sectors as this mapping would have already been done by the controller and would appear perfect to the FORMAT program.

The same options as those for the floppy disk are available to make the hard disk partition bootable under certain conditions and a volume label may also be set to the partition(s).

3. File Allocation Table (FAT) and Clusters

The *file system* of an operating system defines the structure for organising and locating data on the disk. Data is stored in *clusters*, a unit measured in bytes. A file may be stored in one or across many clusters. These need not be contiguous and the file data may be *fragmented* across clusters in different parts of the disk if consecutive clusters are not available.

Most file systems store data in a hierarchical tree structure where the top level of the hierarchy is called the *root* or *root directory*. Below this are *directories* or *folders* (and sub-directories) to contain files and to manage the files.

The starting cluster reference number of a file is listed in the directory area of the storage media, e.g. floppy diskette or logical partition of a hard disk drive. The *file allocation table* (FAT) is a central record keeper that links the various clusters assigned to one file and keeps tracks where file data is stored.

The file system for defining the organisation and location of data on the disk will differ, depending on the operating system used. The most familiar are those used by the Microsoft operating systems:

- FAT12 - developed for the DOS operating system, using 12-digit binary number (12 bits) for cluster information. Used in small hard disks (under 16MB), it is also used to format floppy diskettes.
- FAT16 - for 16-bit allocation table entries on disks larger than 16MB and is supported by all Microsoft operating systems as well as OS/2 and Linux.
- VFAT - a virtual FAT system driver introduced in Windows for Workgroups 3.11 and is supported by Windows 95. It provides a protection mode and the capability for using long file names over the 8.3 limitation of the original FAT16.
- FAT32 – uses a 32-bit allocation table to support FAT32

partitions of up to 32GB. Supported by OSR 2 version of Windows 95 (95b), it is incompatible with MS-DOS, Windows 3.x, Windows 95a, and Windows NT.

- NTFS - Windows NT native file system designed to be more robust and secure than any other MS file systems. It supports a feature called *hot fixing*, whereby the operating system will detect and automatically relocate a bad sector on the disk to a good sector and marks the bad sector as unusable. In theory it supports up to 16 exabytes (equal to one thousand petabytes, one million terabytes, or a billion gigabytes) and allows volumes to span two or more partitions. A *master file table* (MFT) contains information about all the files and data stored on disk.
- EFS – Encrypting File System included in the Windows 2000 version of NTFS and above, which supports file encryption based on public key cryptography and digital certificates.

Other file systems include:

- CDFS – the Compact Disc file system
- HPFS – the high-performance native file system of OS/2
- Ext2fs, VFS, and Journaling file systems for Linux
- Macintosh Hierarchical File System (HFS)
- Network file systems

[Shinder 2002, pp 196]

4. Cluster sizes and slack space

A cluster is the minimum amount of space that can be assigned to any file. A file does not use only a part of a cluster under the FAT file system. Instead, the amount of space used by the file is rounded up to an integer multiple of the cluster size. So, if a file is only one byte, it will still use up the space of the whole cluster. (See figure 25 below). The file will be allowed to expand in the cluster until the maximum cluster size is reached. *Any* amount over the maximum will take up the entire space in the next available cluster, and so on. [Kozierok 2001]

The space left at the end of a file is known as the *slack* or *slack space*. The amount of slack space varies between FAT versions and is affected by the size of the volume partition. The different FAT cluster size ranges are:

- Fat12 - 0.5 KB to 4 KB
- Fat16 - 2 KB to 32 KB
- FAT32 - 4 KB to 32 KB

[Kozierok 2001]

To illustrate, figure 25 is a screenshot of the file properties of the word 'Hello' saved as a text file in Notepad. The image on the left was saved on a FAT32 system with only one logical partition. The image on the right was saved on a NTFS system with three logical partitions.

The size of the file is equal at 5 bytes for each character in the file. However, the size of each file on the disk differs considerably where the one on the left takes up a full 32 KB of disk space in contrast to only 4 KB on the right. It must be pointed out, however, that although FAT32 indicates a 32 KB cluster size, in reality only 28 of the 32 are used as 4 are 'reserved' (or one sector per cluster) as a means to control the size of the FAT itself. [Sammes et al 2000, pp 159]

As the cluster size reduces, the more clusters are needed, all being equal. This has two immediate effects. First, the amount of slack space is reduced and a more efficient use of disk storage is achieved. Secondly, this will require that a files larger than the cluster size are split into potentially many clusters that may not be sequential or contiguous and so add to the level of *fragmentation* on the disk and reduce file access efficiency. The latter effect can be minimised by running a *defragment* program, which will reassign clusters, as far as possible, so they are sequential and contiguous.

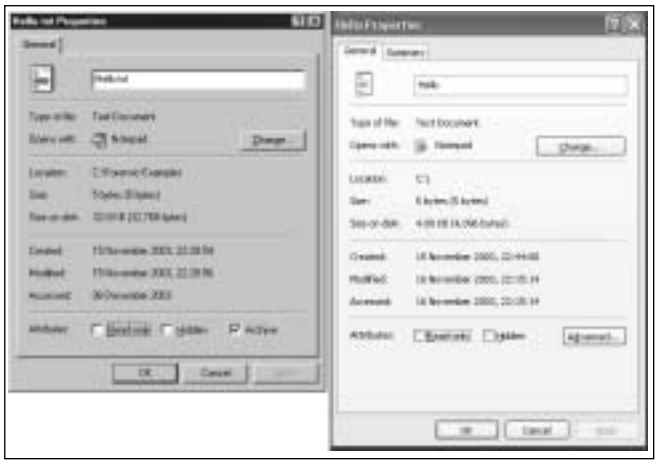

Figure 25 - File properties of a file saved FAT32 and NTFS

5. POST/Boot Sequence

The *boot sector* is the first physical sector (cylinder 0, head 0, sector 1) on a floppy or hard disk and contains the *boot record* or, a *master boot record* in the case of the hard disk. The boot record stores information about a disk's logical and physical makeup. This includes the number of:

- bytes per sector
- sectors in a cluster
- sectors on the disk
- sectors per track
- sides on the disk.

Also included is the *media descriptor byte*, which distinguishes between a hard disk and the various types of floppy disks, so might indicate, for example, that the disk is a 3.5" floppy disk with a 1.44-megabyte capacity. A media descriptor is also stored as the first byte of the FAT.

The *bootstrap loader* is a program included in the boot record of bootable disks that loads the operating system during the POST sequence. The bootstrap loader is itself initially loaded into memory by a small program located in the system BIOS. The bootstrap loader in the *master boot record* of a hard disk with more than one partition indicates which partition contains the hidden system files used to start the operating system (e.g. IO.SYS and MSDOS.SYS).

5.1 Power On Self Test (POST)

The POST sequence software, held in the Read Only Memory (ROM) Basic Input Output System (BIOS), initialises as necessary and then carries out diagnostic tests on each of the various hardware components of the system. Before entering each step in the sequence, the BIOS writes a one byte identifying code, usually to I/0 port 80h, which signals a successful completion of the previous step. This code is commonly referred to as a POST *code*.

As this happens before the display system has been activated, a special POST *code reader* may display the POST code values and used as a diagnostic tool in the event of a

hardware failure, which is often signalled by a system lockup. By viewing the POST codes, the last valid action may be identified and thus track down the device that has failed.

The system speaker is used to generate *beep codes* to identify an error sent from the BIOS. The beep codes vary from same length beeps, low and high tones, a series of long and short beeps and beeps and pauses. A successful completion of all the POST diagnostic tests is indicated by one short beep.

[Sammes et al 2000, pp 139]

Figure 26 shows an image of a section of motherboard containing the BIOS, CMOS and lithium backup battery.

Figure 26 BIOS on a VIA KT133 Chipset motherboard

5.2 CMOS RAM

The battery-backed Complimentary Metal Oxide Semiconductor (CMOS) RAM chip provides some of the information to establish required BIOS variables during the POST routines. This information will typically include the date and time settings, the hard drive types and disk geometry parameters, the memory configuration, and any power management or password protection settings.

The date and time settings are updated by a Real Time Clock (RTC), which is read during the booting process and explains why system clocks are sometimes out of synchronisation. Rebooting will cause the RTC to be reread and synchronisation may be restored. The clock, CMOS RAM and battery are usually all integrated into a single chip. [PC Tech Guide 2003]

The BIOS allows access to a SETUP program so the user may alter the CMOS settings. This is entered during the boot sequence by pressing F1 or F2, CRTL-ALT-ESC, CRTL-ESC, DEL or ESC, among others. [Sammes et al 2000, pp 141]

6. Master Boot Record and Partitions

There are four 16 byte partition entries held in the master boot record (MBR) partition table. Each of these four may refer to a *primary* partition or one that would contain a bootstrap loader in the first sector and associated operating system code elsewhere in the partition. Any primary partition may be flagged as 'active' or 'bootable' but only one may be flagged at any one time. This would allow each of the four partitions to contain a different operating system, with only one active at any one time, the selection of which can be made by the use of special software. Figure 27 is a screenshot of a hexadecimal listing of a 40 GB hard disk in the hexadecimal editor, WinHex, for an analysis of the partition table.

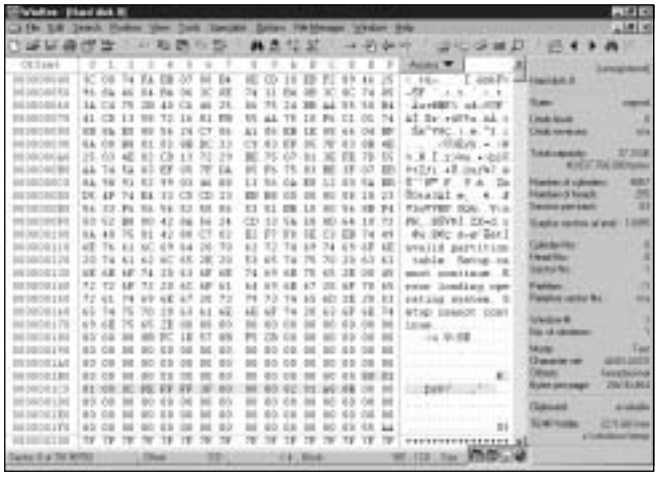

Figure 27 – Hard disk partition table in WinHex

The partition table will always start at address 1BEh in the master boot record sector. There can only be four entries in the MBR partition table and each entry is 16 bytes long. The first entry for this example is highlighted in figure. Address 1FDh marks the end of the four entries. Immediately following this are he final two bytes of the partition table, which are always the hex values 55 and AA, shown at addresses 1FEh and 1FFh respectively.

The hex value 80 at address 1BEh identifies this as the active partition (the value being 00 if not). This is known as the *boot flag*. The next three bytes, beginning at address 1BFh, represent the starting CHS address of the active partition. These values are 01, 01 and 00, which translates to start of head 1, start of sector 1, and start of cylinder 0.

1C2h gives the partition type as 0C, which indicates a FAT 32 partition Logical Block Address (LBA) using INT 13 extensions. [Sammes et al 2000, pp 249] The next three bytes from 1C3h-1C5h give the CHS address at the end of the partition. These values are FE, FF and FF, giving end of heads 254, end of sectors 63, and end of cylinders 1023 respectively. Although the last two figures do not seem correct at first glance this is because the sectors only use the first six bits of byte FF and the cylinder will use the last two bits added with all eight bits of the next byte FF totalling ten bits in all.

The four bytes starting at address 1C6h is the Logical Block Address (LBA) or the sectors preceding partition 1. The values 3F 00 00 00 are held in little endian format and so need to be reordered to 00 00 00 3F and calculated as LBA sector 63. The next four bytes are similarly in little endian format. These start at address 1CAh and give the partition size as 82 91 A8 04. Reordered this becomes 04 A8 91 82, or 78,156,162 sectors.

As this hard drive has only one partition table entry the remaining addresses from 1CEh to 1FDh all have the value 00 followed by the end of partition table signature 55 AA.

To overcome the limit of four partitions on a disk, one of the entries in the MBR partition table can be set to be an *extended* partition entry (type values 05h or 0Fh). [Sammes et al 2000, pp 148] This will allow one partition to be extended into one or more *logical* or *secondary* partitions.

Although there are no formal rules about how partition tables work, most versions of FDISK appear to conform to a certain set of rules:

- In the master boot record there can only be up to four *primary* partition entries *or* up to three *primary* partition entries and up to one *extended* partition entry.
- In an extended partition there can be up to one *secondary* partition entry and up to one *extended* partition entry.
- Only primary partitions can be marked as "active" (bootable) and only one of them can be marked as such at any one time.
- It is usual for a partition table to start a head 0, sector 1 of a cylinder and the boot record to start at head1, sector 1 of a cylinder.
- The slots in the partition table can be used in any order and an unused slot can occur in the middle of the table.
- Some operating systems may indicate that a partition spans or starts beyond cylinder 1024 by setting the starting and ending CHS values all to FFh.

[Sammes et al 2000, pp 147]

7. Virtual memory

Virtual memory is a technique used to increase the amount of memory available to programs. Virtual memory uses space on the hard disk to simulate Random Access Memory (RAM). This space is known as the *swap* or *paging file*.

Programs must be loaded from the disk to RAM before they can be run. The more RAM that is available, the more programs can be run simultaneously, and the larger the programs can be. Windows can "swap" or "page" program instructions and data between RAM and the hard disk as needed. By swapping unneeded data from RAM to disk, the operating system can free up RAM, and run more and larger programs than the computer's RAM would otherwise allow.

The space reserved on the disk for virtual memory can shrink and grow dynamically to meet the changing requirements of the system, so the potential size of the virtual memory can be larger than the current size of the swap or paging file.

The Windows 98 and ME swap file is named *WIN386.SWP*. In Windows NT, 2000 and XP, it is known as a paging file and is named *PAGEFILE.SYS*.

The term 'virtual memory' is often used synonymously with "swap file" or "paging file". However, the term has also been used to refer to the total working memory, including RAM plus available disk space, used by the operating system to run programs.

6 Hide and Seek

Having investigated the disk geometry down to a low level in the previous sections, it may be easier to understand how information may be intentionally or unintentionally stored and where a forensic analyst may look for it. The following are some of the places information may be found and retrieved. Although forensic evidence may be discovered in many other places, the ones discussed below are those most relevant to the main thread of this project.

1. Bad sectors

A possible place to hide information would be in the sectors that have been marked bad by the controller. By using specific controller commands the marking of bad sectors may be deliberately done even if the sectors are perfectly good. For instance, in the case of a floppy disk, the WRITE DELETED SECTOR command will permit the hiding of information and the same can be read by the READ DELETED SECTOR command. No one would be aware of this information or be able to access it without reassigning the track.

The commands used in IDE disks are vendor specific but with this technical knowledge it can be done here as well. A telltale sign that this may be the case is if there appears to be one formatted tracks too few for the known physical disk geometry. [Sammes et al 2000]

2. Hidden partitions

With specialist software such as PartitionMagic, it is possible to mark a partition as *hidden* so the operating system will not access it so any information contained therein would be invisible. Considering the possibility for extended partitions this may create problems for the analyst. Sometimes when two hard drives are used and both have extended partitions the drive letter assignments to the extended partitions may change if one of the drives were removed. Analysts often mount hard drives in removable trays where this could be a significant disadvantage. [Sammes et al 2000, pp 151]

Unused sectors at the beginning of each partition between the partition table sector and the boot sector are also places where information could be safely hidden without detection by normal use of the file system. These are the places where viruses often hide. The free space in the sector following an extended partition table is also a good place to hide information.

Sometimes a partition may be deliberately or accidentally modified using a disk editor, for example, so that one or more of the working partitions are no longer recorded. By comparing the physical disk size with the sum of all the partition sizes, one could detect for this condition.

Other potential opportunities for hiding and discovering information:

- The translated disk geometry may not permit access to the entire physical disk.
- Logical partitions may not fill the whole of the first extended partition.
- The usable files area may not extend to the end of the defined partition because of the cluster size. The last few sectors may be checked for hidden information.
- Clusters may have been marked as bad in the FAT in order to hide information.

Deleted files and information that has been unwittingly saved in temporary files or buffers may also be recovered.

3. Slack

Once a file has been deleted, it is not erased or removed from the media. Only the first character of the deleted file in the File Allocation Table is changed to hexadecimal 'E5'. This character will inform the operating system that the clusters associated with that file entry is again available for reuse. [Wolfe 2003]

By using programs such as UNDELETE.EXE and UNERASE.EXE one may automate the recovery of deleted files. However, for more difficult cases Norton Disk Editor would be more effective. [Sammes et al 2001, pp 162]

The deleted information may remain latent on the disk for

some time before the clusters are written to again. However, once that happens the new file may only partially fill the last, or only, cluster assigned to it. The remaining space may still hold information from the deleted file or even files written and deleted prior to this. This is known as *cluster slack space*.

Likewise, where the last part of a file does not completely fill the standard sector buffer, there may be information from a previous file in the remaining space. This is known as *buffer slack space*.

Unallocated clusters may also contain information from previous files.

4. Temporary files

In order to provide a more *friendly* user interface to its' users, Microsoft Windows have compromised confidentiality and privacy in different ways.

For example, if when typing a document file the computer *crashes* for any number of reasons it is typical that the user will be able to retrieve their work even if they had not saved it to the hard disk before the crash. This is because the software will periodically save the file automatically as a temporary file (TMP), give it a name, for example, ~WRL2680.tmp, and store it somewhere on the hard disk. This TMP file may be invisible to normal searches and will continue to exist on the disk even if all traces of the original file are removed. [Caloyannides 2002]

5. Encrypted files

A temporary file is particularly relevant when attempting to retrieve encrypted files. For example: Before a file is encrypted it exists in an unencrypted form or *plaintext*. When a file is encrypted using the Windows Encrypting File System (EFS) a copy of the plaintext is made as a backup in case something goes wrong with the encryption process and may later be recovered by a forensic examiner. Also, the plaintext might be stored as a temporary paging file (pagefile.sys) before encrypted.

Through the process of decrypting and re-encrypting a file may leave a plaintext copy stored temporarily on the disk. For instance, if the encrypted file is printed and the System32\Spool\Printers folder is not, spool files will contain unencrypted copies of the encrypted files. [Casey 2002]

5.1 Steganography

Another way files may be encrypted is by the use of *steganography*, or literally translated as 'covered writing'. A pure steganographic system inserts information in a cover – a concept known as "security through obscurity". Information is inserted into digital objects such as images, audio files, videos and text, which then becomes a *stego object*. The highest possible similarity of the cover and stego object is preserved to prevent a casual viewer from noticing additional information.

Steganography is becoming a significant issue in computer forensics, as important information may be embedded in seemingly innocuous data and therefore making its extraction all the more difficult. [Mohay et al 2003]

It is possible to detect even minute changes in the stego object if the original cover is found. This may be found by searching through the file system, unallocated space or even the Web browser cache to find and download the relevant media so a comparison can be made.

6. Swap/paging files

As well as temporary files, Windows operating system automatically creates its own swap file (or paging file) of 20 megabytes or more. It is not normally viewable by Windows Explorer and is hidden by the operating system. Almost anything could be placed in a swap file by Windows; including passwords, encryption keys or other useful data. [Wolfe 2003]

7. File information

Along with the automatic save process, metadata, or data about data, is also saved in a different location on the disk. This metadata may include the following information:

- Date and time of creation
- Date and time of modification
- Last access date
- Serial number of the particular copy of the software that created it
- The author of the file

Any discrepancy of the system time and date with the actual time based on external sources of time and date may be calculated through a *dynamic time and date stamp analysis*. Dynamic time and date stamp analysis is a method for correlating times and dates contained within a file to the modified, accessed and created/change (MAC) of status times of the file. This method can be used to standardise the apparent file MAC times to the actual time and allow the examiner to derive the approximate actual system time. [Weil 2002]

This analysis is particularly useful to an examiner when trying to determine if a system was involved in a specific intrusion into another system.

8. RAM data recovery

Random Access Memory (RAM) is used to enter all data. RAM is allocated as required and functions as another hard disk or RAM-disk. However, when the computer is switched off this information is lost; or so it would appear.

Recovering data from RAM and magnetic disk media is difficult but not impossible. In fact, truly deleting data from magnetic media is very difficult. This is due to the fact that when data is written to the medium, the write head sets the polarity of most, but not all, of the magnetic domains. This is because of the inability of the writing device to write in exactly the same place each time, and partially due to the variations in media sensitivity and field strength over time.

When data is written, the media will record either a one or a zero depending on the polarity (see encoding). If a zero *and* a one are both overwritten with a one, the disk circuitry will read both new values as one. However, deviations in the head position from the original track may leave significant portions of the previous data along the track edge relatively untouched. Figure is a simplified illustration of how small differences in head positions may leave the images on the magnetic medium. Each track contains an image of everything written to it. [Gutmann 1996]

The dependence of media *coercivity* on temperature can affect the overwrite capability. Coercivity is the intensity of the magnetic field needed to reduce the magnetisation of a ferromagnetic material to zero after it has reached saturation. [Komag Inc (undated)]

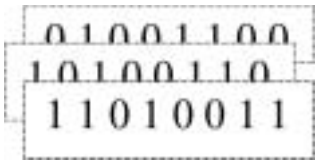

Figure 28 Layers of magnetised patterns

For example, if data was initially written at a temperature where the coercivity was low (thus penetrating deep into the media) and the data was overwritten at a temperature when the coercivity was high, the overwrite performance can be affected. Also, long-term ageing can have a marked effect on the erasability of magnetic media. This relates to the amount of time it has been stored on the media and not on the age of the media itself. [Gutmann 1996]

By using specialist circuitry it is possible to work out what the previous layers 'contained'. Access to specialised equipment is needed to analyse the surface of the magnetic medium. A technique known as magnetic force microscopy (MFM) is used for imaging magnetisation patterns and is derived from scanning probe microscopy (SPM). This process is discuss in detail in Peter Gutmann's paper, 'Secure Deletion of Data from Magnetic and Solid-State Memory', and companion paper, 'Data Remanence in Semiconductor Devices'.

Both papers are available at Internet <http://www.cs.auckland.ac.nz/~pgut001/>.

In the former paper, Gutmann outlines a 35-pass overwrite technique to 'scrub' the latent images left layered on the track. Alternative methods, requiring specialised equipment, involve degaussing, a process whereby the recording media is returned to its initial state, and physical destruction.

Gutmann concludes:

> "Data overwritten once or twice may be recovered by subtracting what is expected to be read from a storage location from what is actually read. Data which is overwritten an arbitrarily large number of times can still be recovered provided that the new data isn't [sic] written to the same location as the original data (for magnetic media), or that the recovery attempt is carried out fairly soon after the new data was written (for RAM). For this reason it is effectively impossible to sanitise storage locations by simple overwriting them, no matter how many overwrite passes are made or what data patterns are written. However by using the relatively simple methods presented in this paper the task of an attacker can be made significantly more difficult, if not prohibitively expensive."
>
> *[—Peter Gutmann (1996), 'Secure Deletion of Data from Magnetic and Solid-State Memory - Conclusion']*

Conclusion

An in-depth knowledge of the materials a computer forensic analyst may encounter is vital to a successful analysis. Different sets of interpretative rules need to be applied to data stored on computer media in order to extract meaningful information.

The intention of this project was to go to the deepest levels of computer media to understand where and how information is stored. Using this knowledge, an analysis was possible on the data stored, which produced a set of meaningful results.

A background on the science of computer forensics was presented as well as principles and guidelines authored by standards committees and working groups.

The number system used by computers was investigated in its various forms and representations. Using the information obtained, an analysis on file types was possible with the aid of the hexadecimal editor, WinHex, by Stephan Fleischmann.

A low-level analysis of floppy disks and hard disks enabled an understanding of the components of magnetic media, its properties and functions. How data is written to a disk and how it is organised was examined including the various features, which enable information to be hidden and discovered.

Problems encountered and future research

The scope of this project included an investigation of the activities of the expert witness and that of the defence council. Due to unforeseen problems with physical mobility, it was not possible to take this to its logical conclusion and therefore had to be omitted from this report. Interviews did take place with expert witness, Dr. John Race, as well as defence barrister, Charles Bott, both of whom were very helpful and informative. It is intended that these areas of investigation will be continued in future research.

Witnessing a forensic examination was also part of the original scope; however, due to the ongoing investigations as part of *operation ore* in potential forensic labs for an observation to take place, this was not legally possible. However, forensic specialist, David Watson, kindly allowed access to his forensic laboratory where many of the tools and equipment used for forensic analysis was available for inspection. Again, due to the aforementioned, further investigations were not possible but will be continued in future research.

Although the original project scope was somewhat altered, the intent of the project was satisfied by the in-depth analysis of computer media and the first-hand analysis of information contained on a hard disk, thus gaining a valuable understanding the forensic analysis process that can be built upon in future research.

Appendix A

Hexadecimal Tables

From hexadecimal editor, WinHex (by Stephan Fleischmann)

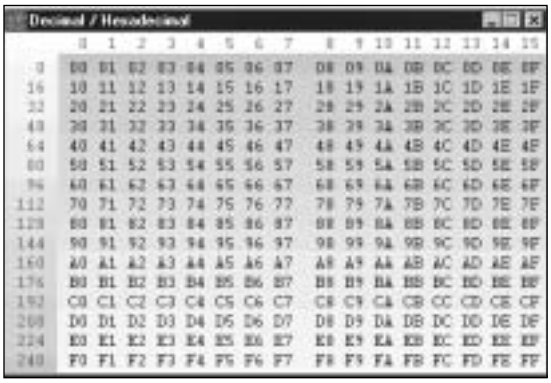

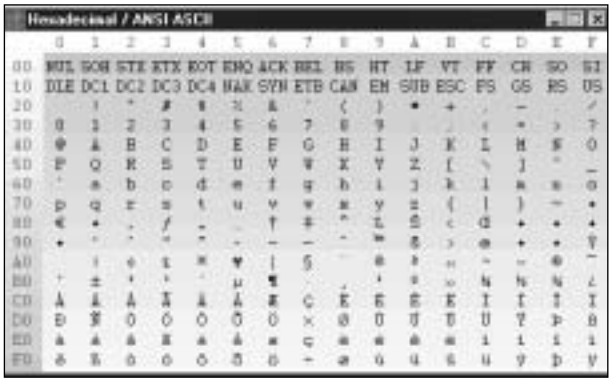

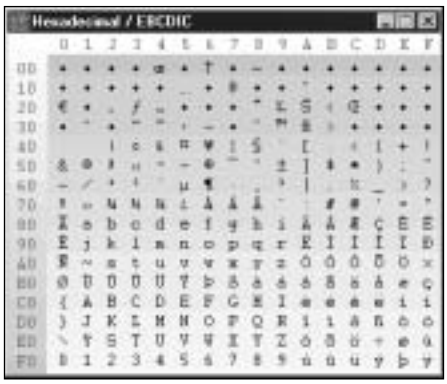

Appendix B

Principles of computer-based evidence

The Principles of Computer-Based Evidence

Four principles from the *Good Practice Guide for Computer Based Electronic Evidence:*

Principle 1: No action taken by law enforcement agencies or their agents should change data held on a computer or storage media which may subsequently be relied upon in court.

Principle 2: In exceptional circumstances, where a person finds it necessary to access original data held on a computer or on storage media, that person must be competent to do so and be able to give evidence explaining the relevance and the implications of their actions.

Principle 3: An audit trail or other record of all processes applied to computer based electronic evidence should be created and preserved. An independent third party should be able to examine those processes and achieve the same result.

Principle 4: The person in charge of the investigation (the case officer) has overall responsibility for ensuring that the law and these principles are adhered to.

[ACPO 2003]

General Principles Applying to the Recovery of Digital Evidence, IOCE Best Practice Guide v1.0

1. The general principles, that have been adopted as G8 recommendations relating to digital evidence, that should be followed by forensic laboratories are as follows:
    1. The general rules of evidence should be applied to all digital evidence.
    2. Upon seizing digital evidence, actions taken should not change that evidence.
    3. When it is necessary for a person to access original digital evidence that person should be suitably trained for the purpose.
    4. All activity relating to the seizure, access, storage or transfer of digital evidence must be fully documented, preserved and available for review.
    5. An individual is responsible for all actions taken with respect to digital evidence whilst the digital evidence is in their possession.
2. All activity relating to the seizure, examination process, and presentation of evidence access, storage or transfer of digital evidence must be documented, preserved and available for review.
3. Responsibility for maintaining evidential value and provenance is a personal, not corporate issue. If an individual has acknowledged responsibility for an item by signing an access log they are responsible for all actions taken in respect of that item until such time as it is returned to store or formally transferred to another individual.

[IOCE 2002]

Bibliography

Books

Caloyannides, Michael A. (2002) *Desktop Witness: The do's and don'ts of personal computer security*, John Wiley & Sons, LTD, GB

Casey, Eoghan (2001) *Digital Evidence and Computer Crime*, Academic Press, UK

Mohay, G., Anderson, A., Collie, B., De Vel, O., McKemmish, R. (2003) *Computer and Intrusion Forensics*, Artech House Computer Security Series, MA, USA

Sammes, Tony, Jenkinson, Brian (2000) 2nd printing (2001), *Forensic Computing: A Practitioner's Guide,* Springer-Verlag London Limited, G.B.

Shinder, Debra, L. (2002) *Scene of the Cybercrime: Computer Forensics Handbook,* Syngress Publishing, Inc., USA

Vacca, John R. (2002) *Computer Forensics: Computer Crime Scene Investigation,* Charles River Media, Inc., USA

Web addresses

ACPO (2003) *Good Practice Guide for Computer Based Electronic Evidence V3.0*, The Association of Chief Police Officers (ACPO) available at Internet

< http://cryptome.org/acpo-guide.htm > last accessed 13 November 2003

BCS (2003) 'BCS Welcomes New ACPO Guidelines on Rules of Evidence for Computer Based Crime', Press Release 28 August 2003, available at Internet <http://www1.bcs.org.uk/DocsRepository/05400/5449/acpo.htm > last accessed 13 November 2003

Brezinski, D., Killalea, T. (2002) 'RFC3227 Guidelines for Evidence Collection and Archiving', Copyright (C) The Internet Society (2002), available at Internet< http://www.ietf.org/rfc/rfc3227.txt?number=3227 > last accessed 14 November 2003

Casey, Eoghan (2002) 'Practical Approaches to Recovering Encrypted Digital Evidence', *International Journal of Digital Evidence*, Fall 2002, Volume 1, Issue 3, available at Internet < http://www.ijde.org/archives/docs/02_fall_art4.pdf > last accessed 15 November 2003

DXing.com (2003) 'UTC/GMT Conversion', Copyright 1999-2003 by Universal Radio Research available at Internet < http://www.dxing.com/utcgmt.htm > last accessed 2 December 2003

Gutmann, Peter (1996) 'Secure Deletion of Data from Magnetic and Solid-State Memory', first published in the *Sixth USENIX Security Symposium Proceedings,* San Jose, California, 22-25 July 1996, available at Internet <http://www.cs.auckland.ac.nz/~pgut001/pubs/secure_del.html > last accessed 15 November 2003

HardDiskInfo (undated) 'Glossary', HardDiskInfo.com available at Internet <http://www.harddiskinfo.com/Sections/Glossary/Home.asp > last accessed 15 November 2003

IBM (undated) 'GMR: A Giant leap for IBM Research', IBM Corporation, available at Internet <http://www.research.ibm.com/research/gmr.html > last accessed 15 November 2003

IBM (1996) 'Meet your hard Drive', Copyright 1996 IBM Corporation, available at Internet < http://www.research.ibm.com/research/gmr/basics.html > last accessed 15 November 2003

IBM Research News (undated) 'IBM's New Magnetic Hard-Disk Drive Media Delays Superparamagnetic Effects', IBM Corporation, available at Internet <http://www.research.ibm.com/resources/news/20010518_whitepaper.shtml > last accessed 15 November 2003

IOCE (2002) 'Guidelines for Best Practice in the Forensic Examination of Digital Technology' draft v1.0, International Organization on Computer Evidence, IOCE 2002 Digital Evidence Standards Working Group available at Internet <http://www.ioce.org/2002/ioce_bp_exam_digit_tech.html > last accessed 13 November 2003

Komag Inc. (undated) 'Glossary', Komag Incorporated, available at Internet <http://www.komag.com/company/glossary.html > last accessed 21 November 2003

Kozierok, Charles M. (2001) 'The PC Guide', Site Version: 2.2.0 -Version Date 17 April 2001, available at Internet < http://www.PCGuide.com > last accessed 15 November 2003

MediaTek (1998/1999/2000) 'HardDriveBasics3', Copyright MediaTek, available at Internet < http://209.35.237.117/system/hdd/hddbasics3.htm > last accessed 15 November 2003

Noblett, Michael G., Pollitt, Mark M., Presley, Lawrence A. (2000) 'Recovering and Examining Computer Forensic Evidence', *Forensic Science Communications,* October 2000, Volume 2, Number 4, Available at Internet <http://www.fbi.gov/hq/lab/fsc/backissu/oct2000/computer.htm> last accessed 15 November 2003

PC Tech Guide (2003) 'Components/Motherboard', PC Tech Guide, available at Internet <http://www.pctechguide.com/01mboards.htm > last accessed 2 December 2003

RFC Draft, v1.2 1996/11/20, available at Internet < http://www.catb.org/~esr/magic-numbers/rfc-draft > last accessed 23 November 2003

RITECC – (undated) 'ISA and PCI Bus Slots' Regional Information Technology and E-Commerce Committee, available at Internet < http://www.ritecc9.org/isa.php > last accessed 15 November 2003

Seagate (2003) 'Seagate Glossary', Seagate Technology LLC, available at Internet <http://www.seagate.com/support/glossary/ > last accessed 15 November 2003

Smith, Richard M. (2003) 'Microsoft Word bytes Tony Blair in the butt', available at Internet <http://www.computerbytesman.com/privacy/blair.htm >) last accessed 28 November 2003

The Unicode Consortium (2003) 'FAQ Unicode and ISO 10646' Unicode, Inc. available at Internet <http://www.unicode.org/faq/unicode_iso.html > last accessed 22 November 2003

The Unicode Consortium (2003a) 'Technical Introduction' Unicode, Inc. available at Internet <http://www.unicode.org/standard/principles.html > last accessed 22 November 2003

The Unicode Consortium (2003b) 'What Is Unicode?' Unicode, Inc. available at Internet <http://www.unicode.org/standard/WhatIsUnicode.html > last accessed 22 November 2003

Verts, William T. (1996) 'An Essay on Endian Order', 19 April 1996, available at Internet <http://www.cs.umass.edu/~verts/cs32/endian.html > last accessed 2 December 2003

Weil, Michael C. (2002) 'Dynamic Time & Date Stamp Analysis', *International Journal of Digital Evidence*, Summer 2002, Volume 1, Issue 2, available at Internet <http://www.ijde.org/archives/docs/02_summer_art4.pdf > last accessed 13 November 2003

Whitcomb, Carrie Morgan (2002) 'An Historical Perspective of Digital Evidence', *International Journal of Digital Evidence*, Spring 2002, Volume 1, Issue 1, available at Internet <http://www.ijde.org/archives/docs/historical_perspective.pdf > last accessed 13 November 2003

Papers and Journals

ACPO (1998) 'Good Practice Guide for Computer Based Evidence' V1.0, The Association of Chief Police Officers (ACPO) Crime Committee, 25 March 1998

Pollitt, Mark M. (undated) 'Computer Forensics: An Approach to Evidence in Cyberspace', Federal Bureau of Investigation, Baltimore, MD, USA

Pollitt, Mark M. (1995) 'Principles, practices, and procedures: an approach to standards in computer forensics', *Second International Conference on Computer Evidence*, Baltimore, MD, 10-15 April 1995, Federal Bureau of Investigation, Baltimore, MD.

Wolfe, Henry (2003) 'Evidence Analysis', *Computers & Security*, Volume 22, Issue 4, 2003, pp. 289-291 © 2003 Elsevier Science Ltd. All rights reserved. PII: S0167-4048(03)00404-8

Celeste has recently completed an MSc in Information Security at Westminster University. She can be contacted on RushLSE97@aol.com

IRMA
INFORMATION RISK MANAGEMENT & AUDIT
◆ A SPECIALIST GROUP OF THE BCS ◆

BCS
THE BRITISH COMPUTER SOCIETY

Membership Application

(Membership runs from July to the following June each year)

I wish to APPLY FOR membership of the Group in the following category and enclose the appropriate subscription.

CORPORATE MEMBERSHIP (Up to 5 members)* *Corporate members may nominate up to 4 additional recipients for direct mailing of the Journal *(see over)*	£75
INDIVIDUAL MEMBERSHIP *(NOT a member of the BCS)*	£25
INDIVIDUAL MEMBERSHIP *(A members of the BCS)* BCS membership number: _____	£15
STUDENT MEMBERSHIP (Full-time only and must be supported by a letter from the educational establishment). Educational Establishment: _____	£10

Please circle the appropriate subscription amount and complete the details below.

INDIVIDUAL NAME: (Title/Initials/Surname)
POSITION:
ORGANISATION:
ADDRESS: POST CODE:
TELEPHONE: (STD Code/Number/Extension)
E-mail:
PROFESSIONAL CATEGORY: (Please circle) 1 = Internal Audit — 4 = Academic 2 = External Audit — 5 = Full-Time Student 3 = Data Processor — 6 = Other (please specify)
SIGNATURE: DATE:

PLEASE MAKE CHEQUES PAYABLE TO "BCS IRMA"AND RETURN WITH THIS FORM TO

Janet Cardell-Williams, IRMA Administrator, 49 Grangewood, Potters Bar, Herts EN6 1SL. Fax: 01707 646275

ADDITIONAL CORPORATE MEMBERS

INDIVIDUAL NAME:
(Title/Initials/Surname)

POSITION:

ORGANISATION:

ADDRESS:

POST CODE:

TELEPHONE: (STD Code/Number/Extension)

E-mail:

PROFESSIONAL CATEGORY:

1 = Internal Audit	4 = Academic
2 = External Audit	5 = Full-Time Student
3 = Data Processor	6 = Other (please specify)

INDIVIDUAL NAME:
(Title/Initials/Surname)

POSITION:

ORGANISATION:

ADDRESS:

POST CODE:

TELEPHONE: (STD Code/Number/Extension)

E-mail:

PROFESSIONAL CATEGORY:

1 = Internal Audit	4 = Academic
2 = External Audit	5 = Full-Time Student
3 = Data Processor	6 = Other (please specify)

INDIVIDUAL NAME:
(Title/Initials/Surname)

POSITION:

ORGANISATION:

ADDRESS:

POST CODE:

TELEPHONE: (STD Code/Number/Extension)

E-mail:

PROFESSIONAL CATEGORY:

1 = Internal Audit	4 = Academic
2 = External Audit	5 = Full-Time Student
3 = Data Processor	6 = Other (please specify)

INDIVIDUAL NAME:
(Title/Initials/Surname)

POSITION:

ORGANISATION:

ADDRESS:

POST CODE:

TELEPHONE: (STD Code/Number/Extension)

E-mail:

PROFESSIONAL CATEGORY:

1 = Internal Audit	4 = Academic
2 = External Audit	5 = Full-Time Student
3 = Data Processor	6 = Other (please specify)

IRMA
INFORMATION RISK MANAGEMENT & AUDIT
◆ A SPECIALIST GROUP OF THE BCS ◆

BCS
THE BRITISH COMPUTER SOCIETY

Management Committee

CHAIRMAN	Alex Brewer	alex.brewer@morganstanley.com
SECRETARY	Siobhan Tracey	siobhan.tracey@booker.co.uk
TREASURER	Jean Morgan	jean@wilhen.co.uk
MEMBERSHIP	Celeste Rush	rushlse97@aol.com
JOURNAL EDITOR & SECURITY PANEL LIAISON	John Mitchell	john@lhscontrol.com
WEBMASTER	Allan Boardman	allan@internetworking4u.co.uk
EVENTS PROGRAMME CONSULTANT	Raghu Iyer	raguriyer@aol.com
LIAISON - IIA & NHS	Mark Smith	mark.smith@lhp.nhs.uk
LIAISON - LOCAL AUTHORITY	Peter Murray	cass@peterm.demon.co.uk
LIAISON - ISACA	Ross Palmer	ross.palmer@hrplc.co.uk
MARKETING	Wal Robertson	williamr@bdq.com
ACADEMIC RELATIONS	David Chadwick	d.r.chadwick@greenwich.ac.uk
	David Lilburn Watson	dlwatson@bcrm.co.uk

SUPPORT SERVICES

ADMINISTRATION	Janet Cardell-Williams t: 01707 852384 f: 01707 646275	admin@bcs-irma.org
OR VISIT OUR WEBSITE AT	**www.bcs-irma.org**	Members' area Userid = irmamembers Password = irma2004

BCS IRMA SPECIALIST GROUP ADVERTISING RATES

Reach the top professionals in the field of EDP Audit, Control and Security by advertising in the BCS IRMA SG Journal. Our advertising policy allows advertising for any security and control related products, service or jobs.

For more information, contact John Mitchell on 01707 851454, fax 01707 851455 email john@lhscontrol.com.

There are three ways of advertising with the BCS IRMA Specialist Group:

The Journal is the Group's award winning quarterly magazine with a very defined target audience of 350 information systems audit, risk management and security professionals.

Display Advertisements (Monochrome Only) Rates:

- Inside Front Cover £400
- Inside Back Cover £400
- Full Page £350 (£375 for right facing page)
- Half page £200 (£225 for right facing page)
- Quarter Page £125 (£150 for right facing page)
- Layout & artwork charged @ £30 per hour

Inserts can be included with the Journal for varying advertising purposes, for example: job vacancies, new products, software.

Insertion Rates:

For inserts weighing less than 60grams a flat fee of £300 will be charged. Weight in excess of this will incur additional charges:

- 60-100grams: 14p per insert
- 101-150g: 25p per insert
- 151-300g: 60p per insert
- 301-400g 85p per insert
- 401-500 105p per insert

Thus for an insert weighing 250g it would cost the standard £300 plus weight supplement of £210 (350 x 60pence) totalling £510.

Discounts:

Orders for Insert distribution in four or more consecutive editions of the Journal, if accompanied by advance payment, will attract a 25% discount on quoted prices.

Direct mailing

We can undertake direct mailing to our members on your behalf at any time outside our normal distribution timetable as a 'special mailing'. Items for distribution MUST be received at the office at least 5 WORKING DAYS before the distribution is required. Prices are based upon an access charge to our members plus a handling charge.
Access Charge £350. Please note photocopies will be charged at 21p per A4 side.

Personalised letters:

We can provide a service to personalise letters sent to our members on your behalf. This service can only be provided for standard A4 letters, (i.e. we cannot personalise calendars, pens etc.). The headed stationery that you wish us to use must be received at the Office at least ten working days before the distribution is required. Please confirm quantities with our advertising manager before dispatch. If you require this service please add £315 to the Direct mailing rates quoted above.

Discounts: Orders for six or more direct mailings will attract a discount of 25% on the quoted rates if accompanied by advance payment

Contacts

Administration

Janet Cardell-Williams,
49 Grangewood, Potters Bar, Hertfordshire EN6 1SL
Email: admin@bcs-irma.org
Website : www.bcs-irma.org

BCS IRMA Specialist Group Advertising Manager

Eva Nash Tel: 01707 852384
Email: admin@bcs-irma.org

Venue for Full Day Briefings

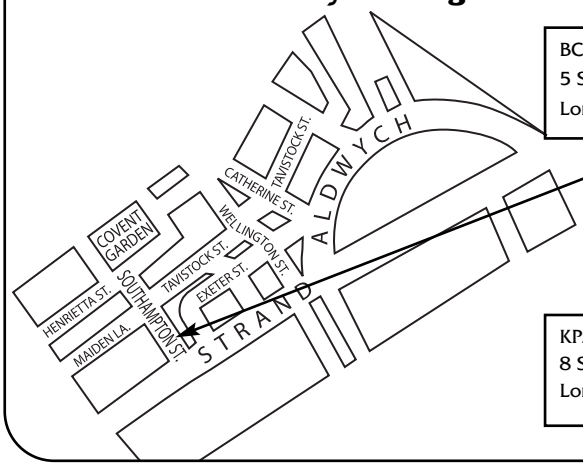

BCS, The Davidson Building,
5 Southampton Street,
London WC2 7HA

KPMG
8 Salisbury Square
London EC4

Venue for Late Afternoon Meetings