Insights into the ISO/IEC 27001 Annex A

By Dr. David Brewer FBCS, Dr. Michael Nash FBCS

December 2010

Abstract

ISO/IEC 27001 is a specification for an Information Security Management System (ISMS). It contains an annex, Annex A, which catalogues a wide range of controls and other measures relevant to information security. At first view, it appears that all an organisation has to do is select the controls that it believes that it needs from this catalogue. However, there is a requirement to carry out a risk assessment. The purpose of this is to identify the controls that are actually required. Over the years arguments have raged between the users of ISO/IEC 27001as to the relative importance and relationship between these two requirements.

This paper reports on research carried out by Gamma Secure Systems Limited (Gamma) over the period January 2007 to December 2010 to investigate the relationship between these two requirements. We discover that if an organisation wishes merely to ensure coverage of the Annex A controls then the scope of the risk assessment is highly constrained. Indeed, we discover that it is possible to generate a small set of templates that once completed will fulfil the risk assessment requirements of the standard and guarantee coverage of the Annex A controls, whilst not necessarily providing a risk assessment that adequately addresses the organisation's real exposure.

Introduction

This paper reports on research carried out by Gamma Secure Systems Limited (Gamma) into the relationship between the ISO/IEC 27001 [1] requirements for the Statement of Applicability (SOA) and its requirements for risk assessment/risk treatment. This research was carried out over the period January 2007 to December 2010.

We begin by introducing these requirements and their relationship as intended by the standard. The paper continues by explaining the trigger for this research and why it raised in our minds the question of whether it is possible to specify the risk assessment to generate requirements for controls exactly matching the 133 controls in Annex A to ISO/IEC 27001:2005. We next explain how we set about answering this question and our discovery of three cardinal triggering events, which together demonstrate that the answer to this question is a resounding “yes”. We then show that by further analysis of the Annex A controls we were able to derive the assets, threats and vulnerabilities to which these controls refer, and perhaps more significantly *template risk treatment specifications*, which in actuality represent the *minimum work required* to satisfy the SOA and risk assessment/ risk treatment requirements.

In parallel with this research we extended the risk assessment method first proposed by Brewer and List [2] to facilitate the quantitative determination of residual risk. This extension necessitates the estimate or measurement of the effectiveness of controls as determined by the way that they modify the frequency or likelihood of the occurrence of an incident and the severity of its consequences. We found that we were unable to do this using the vast majority of the Annex A controls, but could do so if we grouped them together in an appropriate manner. We conclude the presentation of our results by describing this work.

Finally we present our overall observations and conclusions.

The requirements and their intended relationship

ISO/IEC 27001 is a specification for an Information Security Management System (ISMS). It contains an annex, Annex A, which catalogues a wide range of controls and other measures relevant to information security. For some people, Annex A is the most important component of the standard, as they regard it as a set of controls that they must apply unless they have a very good reason not to. Such exclusions do occur in the real world: for example, if an organisation does not indulge in electronic commerce then the controls concerning electronic commerce are clearly irrelevant. There is a requirement to document a disposition of the Annex A controls stating whether each applies or not, and giving reasons in both cases. The disposition is called the Statement of Applicability (SOA).

However, there is another explicit requirement in ISO/IEC 27001, which is to select and implement controls to meet the requirements for controls identified by risk assessment and risk treatment. Some people see this as the most important component of the standard, and the primary way that necessary controls are identified. They regard the SOA merely as an exercise to validate the completeness of their risk assessment. Indeed there is a note to requirement 4.2.1 (j) (3) to that effect. Thus, it is argued that if properly carried out, the risk assessment/risk treatment process will identify all the controls that an organisation requires. The process of cross-checking is illustrated in Figure 1. In producing the SOA it would be expected that a large proportion of applicable controls will have been identified by the risk assessment (Area I). However, the organisation may decide that other controls (Area II) are also applicable even though they were not identified in the risk assessment. The existence of these indicates a deficiency in the risk assessment which can then be corrected. Note also the presence of controls identified by the risk assessment which are not in Annex A (Area III), which the organisation needs but would have omitted if it had solely relied on selecting controls from Annex A.

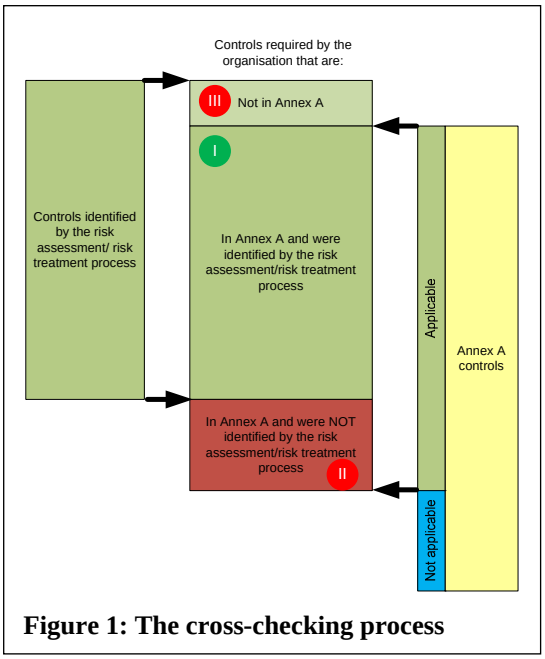

Figure 1: The cross-checking process

The dichotomy of which is the most important requirement – risk assessment or the SOA – is discussed further in Appendix A.

Research trigger

An ISO/IEC 27001 risk assessment requirement is to identify a method that is suitable to the organisation. With the publication of BS 7799-2:2002, the forerunner of ISO/IEC 27001, our first outing as consultants with this new standard was with a UK logistics company. Traditional doctrine in performing a risk assessment is to start by identifying the organisation's assets, threats and vulnerabilities. Indeed there is an ISO/IEC 27001 requirement (4.2.1 (d)) to do this. However, we had such difficulty in explaining what assets and threats were to the Managing Director of the client organisation, that in fear of being shown the door we dared not to enquire about vulnerabilities but to adopt a different line of questioning. We therefore asked the Managing Director about his information security *concerns*, which with reassuring enthusiasm he was most capable of answering. In reviewing his answers, we determined that his concerns were a mixture of *events* and *impact*, where an *event* is a happening that if left unchecked leads to the occurrence of an *impact*. This discovery led to the development of an event-impact driven risk assessment methodology [2], which is summarised in Appendix B.

In this first assignment we identified eight events, which we referred to as:

- S1 – Theft;
- S2 – Acts of God, vandals and terrorists;
- S3 – Fraud;
- S4 – IT failure;
- S5 – Hacking;

- S6 – Denial of service;
- S7 – Disclosure;
- S8 – Law.

We found these events appropriate for other organisations, and soon began referring to them as the *standard eight*. Our SOAs justified the inclusion of controls by reference to an *event* or to a *policy*. At the time, we were unconcerned with the need to reference policies as an additional source of requirements for controls, as this had been standard practice in the world of the Common Criteria [4] for a long time. However, in 2006 we had the opportunity to extend our researches and identified three additional events, which promised to reduce the number of references to policies. We called these events:

- B1 – Inappropriate deployment of people;
- B2 – Failure to maintain proper records;
- B3 – Issuance of wrong documents.

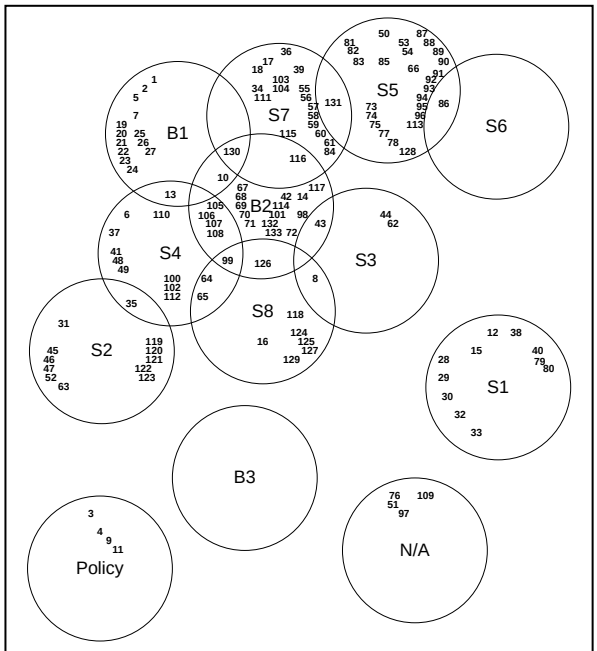

Figure 2: Relationship between events and the Annex A controls

As an exercise, we plotted the relationship between the SOA and these eleven events. Our findings are reproduced in Figure 2. The numbers 1, 2, 3 … 133 refer to the Annex A controls in the order they are presented in the standard, the B1, B2, B3, S1, … S8 to the events as listed above, 'policy' refers to those controls still referenced to policies rather than events and 'N/A' identifies those controls that are not applicable.

The introduction of the three additional events certainly reduced our dependence on policy as a means to justify the inclusion of Annex A controls. However, B3 is empty. There are actually a lot of controls in B3, but none of them are present in Annex A. This is the classic case of Area III in Figure 1. S6 is effectively empty as it includes only one Annex A control, and that is shared with S5. S3 is pretty empty and, with the exception of S1, there are a lot of overlaps.

With a desire to improve efficiency in our consultancy processes, we raised the question "what is the minimum number of events that would justify inclusion of all the Annex A controls". Whatever the number, we anticipated that the corresponding Venn diagram would have no empty sets, no set called 'policy' and few, if any, overlapping sets.

Cardinal triggering events

Sequencing the Annex A controls

A characteristic of the risk treatment plans (RTPs) described in [2] is that the identification of risks, their treatment, selection of controls (where appropriate) and their effects, are presented in the form of a story. We therefore asked if it was possible to sequence the Annex A controls in the form of a story, such that we would consume all but reuse none.

The stories that we had often used in creating the RTPs associated with the standard eight events and those which we had just created for the three new ones gave us a head start, and equipped with this experience the task proved relatively straightforward.

The answer is that the controls can be sequenced in such a manner, albeit that one of them (A.9.2.1 Equipment siting and protection) has to be split into two.

Our results are reproduced in full in Appendix C.

Events

The answer to the original question ("what is the minimum number of events that would guarantee coverage of the Annex A controls?") is not one, however. This is because there is no single discernable event in the sequencing story. What we needed to do was to break that story up somehow. We proceeded by trial and error and soon drew the conclusion that there were three or four *cardinal* events. The four candidates were:

- Inappropriate deployment of people;
- Vulnerability exploitation;
- IT failure;
- Dispossession (of an information container).

On closer investigation, we concluded that the controls associated with the first of these candidate events was really a subset of those required for *vulnerability exploitation* and to a lesser extent, *dispossession* also.

We therefore concluded that the answer to our original question is three:

- EI1 – Vulnerability exploitation;
- EI2 – IT failure;
- EI3 – Dispossession.

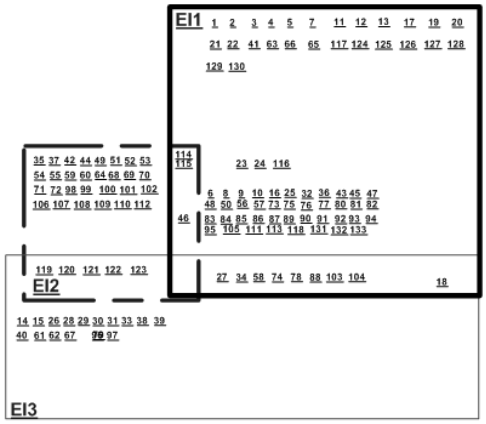

Figure 3: Relationship between the cardinal events and the Annex A controls

Sample definitions of these three cardinal events are given in Appendix D and the resulting Venn diagram relating them to the Annex A controls is presented in Figure 3. The diagram possesses all the anticipated characteristics.

Assets, threats and vulnerabilities

The Brewer-List approach to risk assessment, as presented in Appendix B, does not conform to the requirements of ISO/IEC 27001 because the method *does not require* the identification of assets, threats and vulnerabilities. To attain conformance, these missing elements need to be derived. As an *event* is in actuality an asset-threat-vulnerability triple, derivation is straightforward, albeit superfluous to the risk assessment process. We usually perform this derivation on-the-fly as an integral component of considering the risk treatment options.

Consider three events being respectively the *loss*, *theft* and *misappropriation* of a laptop. These three events are subordinate to the cardinal event of dispossession. We should all be able to visualise instances of these three events, and indeed some of us may well have personal experience of them. In each case, there are two assets. The first is the laptop and its associated operating system and application software. The second is the sensitive business information, and perhaps personal information, that it may contain. The threats are respectively, ourselves (the employee), the thief and the disgruntled employee. There are three vulnerabilities. Firstly and associated with all three events is the size, weight and portability of the laptop, making it easy to loose, easy to steal and easy to forget that someone has it. Secondly, as we can get the data out of the laptop, so in the absence of any controls (such as access control and encryption) can anyone else. This vulnerability is pertinent to *loss* and *theft*, but not to *misappropriation*. This is because the person misappropriating the laptop always had access to the information anyway. The third vulnerability is the attractiveness of the laptop as an item to be possessed or sold on the Illegitimate Market. This vulnerability is associated with *theft* and *misappropriation*, but not *loss* as reasons for losing it have nothing to do with its attractiveness, solely size, weight and portability.

The foregoing indicates the ease by which assets, threats and vulnerabilities can be derived by a non-thought-taxing consideration of the circumstances surrounding the event, and it is for this reason that we refer to these subordinate events as *event-circumstances*. The foregoing may also illustrate the inefficiencies of adopting an asset-threat-vulnerability to risk assessment, as with two assets, three threats and three vulnerabilities, 18 calculations are required as opposed to just 3 in the Brewer-List approach. Realising that not all combinations of assets, threats and vulnerabilities are possible (a conclusion, we may add, which does require some thought process) will reduce the number 18 to 14, but that is a number still significantly greater than 3.

Notwithstanding this observation, we asked if it was possible to derive the assets, threats and vulnerabilities from a consideration solely of the Annex A controls. The answer is "yes" and in Appendix E we present the assets, threats and vulnerabilities that we found and their relation to the three cardinal events.

Template risk treatment specifications

As previously mentioned, a characteristic of the risk treatment plans (RTPs) described in [2] is presentation in the form of a story. Given that we have discovered three cardinal events that together guarantee coverage of the Annex A controls, and that we have an overall story sequence for all 133 controls (see Appendix C), it was natural to ask if it was possible to specify, at least in outline form, the story line for each of the three cardinal events. The answer to this question was again "yes".

We present in Figure 4 an image of the template we created for the vulnerability exploitation event. The template was created using Adobe Dreamweaver. We created three such templates, one for each cardinal event, as part of a tool-based automated risk assessment process now owned by *IMS-Smart* Limited[1]. These cardinal event templates have been used by us in creating four ISMS, all of which have been certified as being conformant with ISO/IEC 27001.

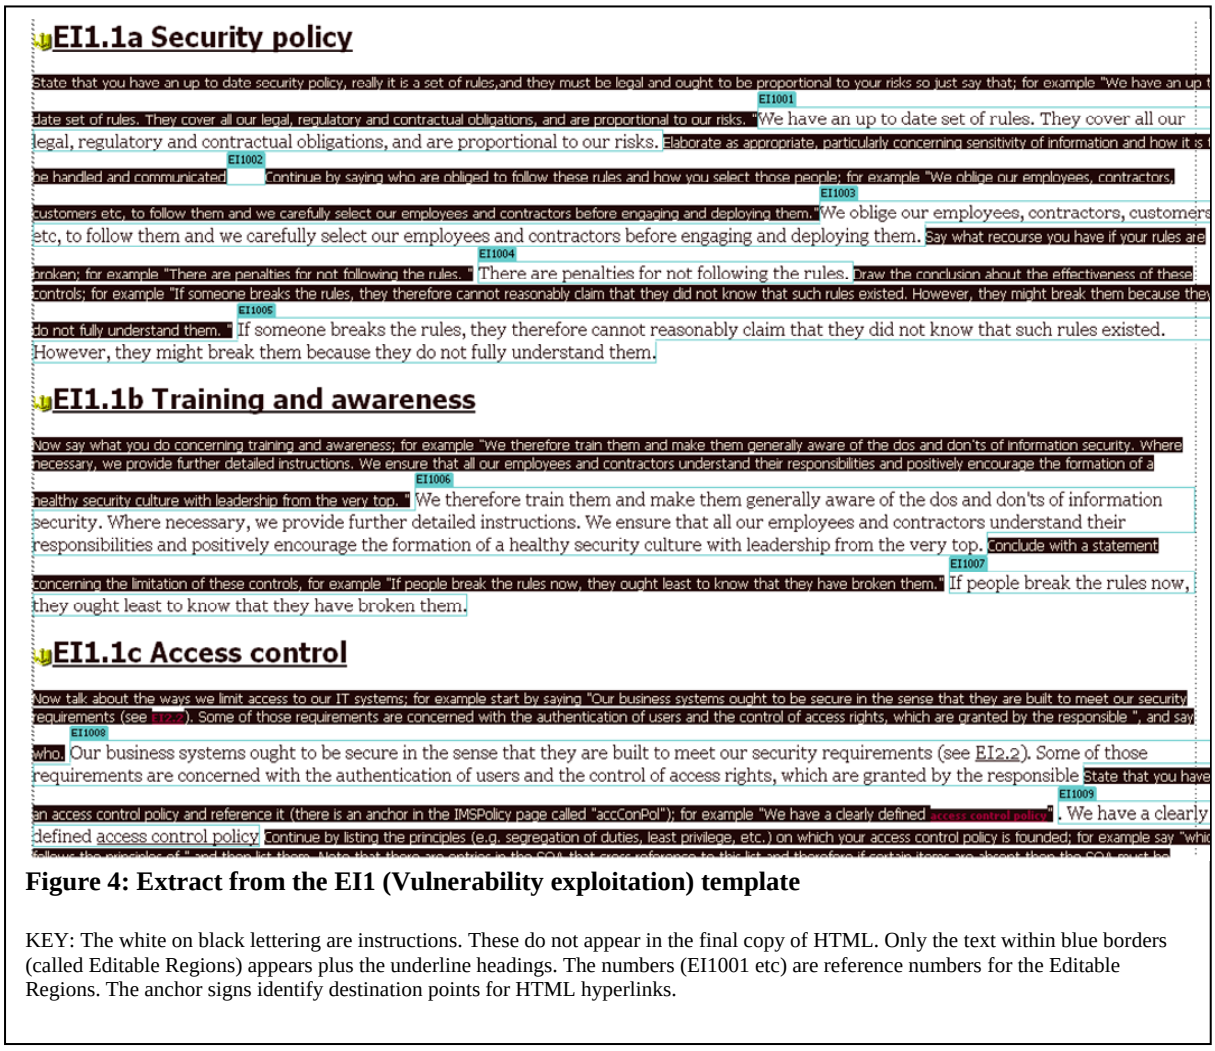

Figure 4: Extract from the EI1 (Vulnerability exploitation) template

KEY: The white on black lettering are instructions. These do not appear in the final copy of HTML. Only the text within blue borders (called Editable Regions) appears plus the underline headings. The numbers (EI1001 etc) are reference numbers for the Editable Regions. The anchor signs identify destination points for HTML hyperlinks.

[1] *IMS-Smart* Limited is a British company that specialises in know-how and technology for integrated management systems. Its web site URL is www.ims-smart.com

Real controls

The final step in the Brewer-List approach, as presented in Appendix B, is to calculate the residual risk. As shown in that appendix, the method expresses risk as a vector with two components: frequency or likelihood of occurrence (FoL) and severity of consequence (Sev). Thus the Brewer-List approach requires the identification of controls that modify either FoL and/or Sev. The vast majority of controls in Annex A do not possess this property.

Consider, for example, the first control: A.5.1.1 - Information security policy document. In what way does any document – on its own – modify either FoL or Sev? The answer is that it cannot. However, when combined with other Annex A controls, the resulting combination, which we refer to as a *real control*, did have the required property of modifying either FoL or Sev.

We have performed this analysis twice, each time for a real ISMS, in fact two of the four referred to in the section above on "Template risk treatment specifications". In each case, we used the story sequencing in the three cardinal event templates, as instantiated for the ISMS in question, to do this. The sequencing in the templates is not identical to that in Appendix C, and the process of instantiating the risk fragments (see Figure 4) can itself change the relationship between the risk fragments and the Annex A controls. Thus no two instantiations will be exactly the same. Nevertheless, in Appendix F we present a sanitised combination of our results, indicating a possible relation between such *real controls* and Annex A controls. In Appendix G we present the reverse mapping for the table in Appendix F.

Notwithstanding our observations on this matter, a new ISO standard was published about a year ago. Called ISO 31000 [5] it provides guidance on risk management in *general*, i.e., not specific to information security but also pertinent to all other disciplines, such as quality, environmental protection, food safety etc, etc. ISO/IEC JTC 1 SC 27 WG1[2], the people responsible for developing ISO/IEC 27001, have resolved that the next revision of ISO/IEC 27001 will conform to ISO 31000. That standard invokes the ISO Guide 73 definition of a *control*, namely *a measure that modifies risk*. Thus it transpires that our *real controls* conform precisely to the ISO Guide 73 definition, whereas the Annex A controls do not – they are mostly just measures. To be fair, some are indeed controls but others are in fact combinations of controls and other measures. In that sense, Annex A is actually quite a mess. Nevertheless as a safety net, or as we have described it in previous publications [6] as an AIL (an Alternative Ideas List) it remains in our opinion extremely useful.

Interestingly, the Brewer-List approach conforms precisely to the requirements of ISO 31000, a standard that does not require the identification of assets, threats and vulnerabilities.

Observations and conclusions

Our first and most important conclusion is that there is a philosophical problem with the current version of ISO/IEC 27001, and that this needs to be sorted out in the next revision. Either controls are selected to meet the requirements of risk assessment, and Annex A is a confidence building double-checking exercise, or controls are selected from Annex A and risk assessment must be satisfied using those controls, plus any custom controls that may be necessary.

This is important for two reasons. Firstly, many existing tools and methodologies propose actual security controls which do not elegantly map onto Annex A, either in terminology or coverage – our paper includes a practical example of this. Secondly, ISO/IEC 27001 is now widely used in conjunction with other management system standards that assume the primacy of risk assessment as per ISO Guide 73 and ISO 31000. Either ISO/IEC 27001 is special and different, or it is not. That answer is important, but what is even more important is that the answer is clear, accepted by everyone, and uniformly applied.

We have shown that risk assessment methodologies *can* generate requirements corresponding to all 133 Annex A controls, and that such approaches can be used in practice. We have also shown that risk assessment philosophies based on events and impacts, as widely used outside the IT arena, *can* be adapted to meet the classical asset, threat and vulnerability approach generally found in IT security. We conclude that greater commonality of terminology and approach is both possible and viable. However, our analysis of "real controls" shows that many current Annex A controls are poorly specified for such convergence.

[2] Joint Technical Committee 1, Sub-Committee 27 Working Group1

Finally, we have demonstrated an alternative, sequence-based approach to structuring the existing Annex A controls that many people may find attractive and perhaps easier to understand.

ISO/IEC 27001 is an important standard, and the consensus approach underpinning international standardisation makes it important that we convince people that the standard has problems concerning Annex A and how it relates to risk assessment, and that these need to be fixed.

References

[1] "*Information technology – Security techniques – Information security management systems – Requirements*", ISO/IEC 27001:2005

[2] "*Measuring the effectiveness of an internal control system*", Brewer, D.F.C., List, W., March 2004, www.gammassl.co.uk/topics/time

[3] "*Information technology – Security techniques –Code of practice for information security management*", ISO/IEC 27002:2005

[4] "*Information technology – Security techniques –Evaluation criteria IT security*", ISO/IEC 15408

[5] "*Risk management – Principles and guidelines*" ISO 31000:2009

[6] "*Exploiting an Integrated Management System*" Brewer, D.F.C., Nash, M. L., List, W., February 2006, www.gammassl.co.uk/topics/ics/MSExploitation.pdf

About the Authors

Dr. David Brewer, FBCS, CITP

Dr. David Brewer has been involved in information security since he left university, and is an internationally recognised consultant in that subject. He was part of the team who created the ITSEC and the Common Criteria, and has worked for a wide range of government departments and commercial organisations both at home and abroad. He was one of the driving forces behind the international ISMS standards, and has assisted many clients to build ISMSs since 1998 in Europe, East Africa, the Middle East and the Far East.

Dr. Michael Nash, FBCS, CITP

Dr. Michael Nash also has a long background in information security. His first involvement came in 1985, working initially within NATO using the US TCSEC "Orange Book", and then setting up and managing the first UK security evaluation facility. He has been involved in international standardisation for more than twenty years, most recently as the Project Editor for ISO/IEC 27010, the ISMS standard targeted specifically at sharing sensitive information between organisations.

Appendix A: A brief history of Annex A

The dichotomy of which is the most important requirement – risk assessment or the SOA – and the entrenched positions that survive to this day are best understood through an examination of the history of ISO/IEC 27001 and, in particular, Annex A.

ISO/IEC 27001 and ISO/IEC 27002

ISO/IEC 27002 [3] is an expansion of the controls summarised in Annex A. It is an *informative*[3] standard that provides *guidance* on how each of these controls might be implemented together with other useful information. It first saw the light of day as a British Standard, BS 7799:1995, and its technical content was produced by a group of ten or so highly experienced information security officers drawn from a wide variety of large public and private sector organisations operating in the UK. It is thus full of practical experience of dealing with a wide range of security issues. Following various revisions to keep pace with advances in technology it became an ISO standard (ISO/IEC 17799:2000) at the turn of the Millennium, and following further revision and renaming was republished as ISO/IEC 27002:2005. It is currently under revision. Once again, technology has moved on and now commonplace IT security products, such as personal firewalls, are not mentioned in the current version, but hopefully will be in the next.

ISO/IEC 27001 was also first published as British Standard, BS 7799-2:1998. Its purpose was to act as a bridge between BS 7799 and certification and from the outset was cast as a *normative* standard. It summarised the BS 7799 controls in the main body of the text, in a section called "Detailed controls"[4] and introduced the requirement for an SOA, recognising that not all controls applied to all organisations. However, it was hurriedly put together and subsequently recognised as a *weak* standard. It was therefore substantially revised and republished as BS 7799-2:2002, and with a modicum of change republished again as ISO/IEC 27001:2005. It is also currently under revision. It is being updated to take account of practical experience gained in applying the standard, really since its major revision in 2002, and to conform with the new risk assessment standard, ISO 31000 [5]. It is also being restructured to align with the new requirements of the ISO Technical Management Board (TMB), for management system standards, as proposed by the Joint Technical Coordination Group on Management System Standards.

Annex A

BS 7799-2:1998 introduced the concept of risk assessment, and also required organisations to justify their selection of controls, with reasons recorded in a Statement of Applicability (SOA). However, there was no explicit requirement to link selection of controls to the results of the risk assessment, and thus it was perfectly legitimate for an organisation to claim that its selection was based upon best practice 7799-1 controls as recorded in the section on detailed controls, or even based upon personal preference!

If the reasons recorded in the SOA did not have to link to the risk assessment, and could not be challenged, those who believed in the primacy of risk assessment for the selection of controls considered that the SOA had little value. During the 2002 revision, an argument was therefore proposed to drop the SOA requirement, together with the "detailed controls". However, this proposal met with resistance from those who regarded this list of "detailed controls", now moved into an annex (Annex A) for structural reasons, as a convenient checklist for certification auditors. A compromise was reached by requiring the selection of applicable controls to be justified by reference back to the risk assessment. In addition, the requirement to record those controls from Annex A that were not selected was retained.

ISO/IEC 27001:2005 extended the SOA to identify controls that were already implemented. It also clarified that controls could be selected for reasons other than meeting requirements identified by the risk assessment. Finally, it recommended Annex A as a starting point for control selection.

Although ISO/IEC 27001:2005 requires that controls are selected to meet the requirements identified by the risk assessment, it also recommends use of Annex A to ensure no important control options are overlooked. This can, and has been, interpreted to mean that all controls from Annex A should be included, except where there is

[3] *Informative* standards are guidance documents and are cast in the language of *should* not *shall*. By contrast, *normative* standards are requirements documents and are cast in the language of *shall* not *should*.

[4] Section 3 of BS 7799-1:1998 was called "detailed controls" and summarised all the controls and control objectives found in BS7799-1:1995 and expressed them in normative form. In the 1999 version, the "detailed controls" were revised to accord with BS 7799-1:1999 and moved to Section 4. They were moved to Annex A in the 2002 edition.

a clear justification for their exclusion. This is clearly undesirable, as unnecessary controls waste resources and can themselves introduce new risks. It also positions Annex A as the primary source for controls, where many practitioners would prefer to derive controls from other sources, and use Annex A purely as a completeness checklist.

The inclusion of Annex A has therefore once again come under scrutiny during the current ISO/IEC 27001 revision process, but in October 2010 at the ISO/IEC JTC 1 SC 27 WG 1 meetings in Berlin a resolution was passed to retain it. The principal argument in favour of removing it was that WG 1 should make its mind up: is ISO/IEC 27001 supposed to be a risk-based standard or a control-based standard, and if the former the Annex is unnecessary. The arguments for keeping it were unchanged from 2002: those who just like it because they regard it as a useful guide for organisations lacking in security experience and those who place great reliance on risk assessment but nevertheless regard Annex A as a useful checking mechanism. Nevertheless, we expect a change of wording in ISO/IEC 27001 to stress that Annex A is not the only possible starting point for selection of controls.

Certification

Accredited certification is always against a *normative* standard. Thus, ISMS certification was in the first instance against BS 7799-2:2002 and for the past five years it has been against ISO/IEC 27001:2005 and will continue to be so until the next version of ISO/IEC 27001 is published, and so on. Certification has never been against BS 7799-1, ISO/IEC 17799 or ISO/IEC 27002 although there continues to this day a misguided belief .

Appendix B: The Brewer-List approach to risk assessment

Introduction

The Brewer-List approach to risk assessment was first proposed in 2002 in response to difficulty in communicating with the managing director of a logistics company. The problem was that the director could not comprehend the concepts of assets, threats and vulnerabilities. He was, however, perfectly capable of articulating his concerns.

The method that resulted has now been used in a variety of certified management systems in the domains of information security, quality and business continuity in Kuwait, India, Mauritius, Saudi Arabia, the UK and the US.

It was first published in 2004 by Brewer and List in their paper "Measuring the effectiveness of an internal control system" [2]. We present below an updated version of the approach. It consists of eight steps.

Step 1

The organisation starts by producing a list of *concerns*.

In the original work that led to the publication of the Brewer-List paper, the organisation – a logistics company - produced a list which included: acts of God, denial of service, disclosure, duplicate delivery, export violations, failure to deliver, hacking, IT failure, premature release [of exam papers], prize fraud and theft.

Step 2

The organisation analyses its list of concerns, refining and separating them into *events* and *impacts*. An *event* is a happening that if left unchecked leads to the occurrence of an *impact*.

For example, *IT failure* (an event) could lead to the *inability of the organisation to carry out its business* (an impact).

Step 3

The organisation estimates the *frequency or likelihood* (FoL) of the occurrence of each identified event. The dimensions of FoL are reciprocal time.

Using a logarithmic scale ,suitable units are: 1 = once a decade; 2 = once a year; 3 ≅ once a month; 4 ≅ every few days and 5 ≅ several times a day. Using such a scale allows risk to be presented in a graphical form, see for example Figure A-1. However, care should be taken by the organisation not to limit its estimate of FoL to say a 1 – 5 scale. The value -2, for example, corresponds to once every 10,000 years and the value 7.8 corresponds to once a second. In the IT world, attack frequencies can be very high, e.g. every few milliseconds.

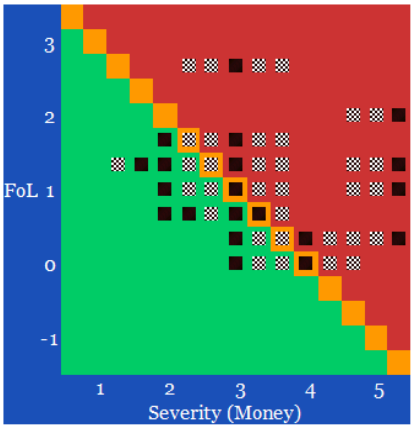

Figure A1: Example of an organisation's estimates of inherent risk

Step 4

The organisation estimates the *severity of the impact*[5] (Sev) should the impact occur. In doing this, the organisation will need to choose an appropriate scale. Money is an excellent choice and surprisingly easy to use, particularly by business people.

Again it is appropriate to use a logarithmic scale. Using US$, 1 could represent $100, implying that 5 represents $1,000,000 etc.

Step 5

The organisation then combines these estimates, pairing the events with the impacts:

Note that an event may lead to several impacts. As an example consider the theft of a laptop. Such an event could lead to the disclosure of sensitive information. That is impact number one. Not having the laptop implies that its owner will at least temporarily be unable to work. That is impact number two, and impact number three is that the owner is now out of pocket to the tune of the cost of replacing the laptop.

In the figure, the grey squares represent areas of uncertainty. If a number of people are asked to estimate FoL and Sev, it is unlikely that they will all give the same answer. Thus there will be a spread of results.

Step 6

The organisation applies its risk criteria to identify those risks that require treatment.

For example, in Figure A1, each black square represents a risk. By this organisation's risk criteria, those in the green area are inherently acceptable – no treatment is required. Those in the red area are unacceptable – treatment is definitely required. Those in the orange area are borderline, and management will decide if treatment is required. Note that for other organisations, the number and position of these areas may be different. They decided by the organisation in accordance with risk criteria of its own choosing.

Step 7

The organisation considers its risk treatment options and decides what it wishes to do.

---

[5] ISO 31000 [2] refers to this concept as the *consequence*. Note that this is not synonymous with the term *consequential impact*, which is a term that refers to risks that are secondary, tertiary, … For example, the inability of an organisation to carry out its business could lead to loss of revenue, which in turn could lead to redundancies and ultimately bankruptcy.

If a risk is to be modified, then the risk treatment plan will associate controls[6] with that risk. Preventive[7] and detective[8] controls modify FoL and reactive[9] controls modify Sev.

Step 8

The final step is to calculate the residual risk. To do this the organisation will need an estimate or measure of the effectiveness of the controls it uses (or will use).

The calculations are straightforward but depend on the way in which the control modifies risk[10, 11, 12].

Figure A2 shows the residual risks that correspond to the inherent risks of Figure 1 following the organisation's choice of risk treatment options.

Note that the results are quantitative and meaningful to Top Management, as both inherent risk and residual risk are expressed in units of $/time.

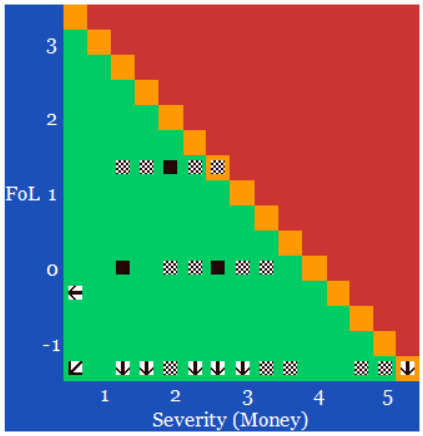

Figure A2: Residual risks corresponding to the example of inherent risks given in Figure 1 organisation's estimates of inherent risk

Note also that some risks (identified by the arrows) are off-scale. This is one of the benefits of not restricting risk evaluations to a 1 to 5 (or any other number) scale. In the real world these risks merit closer attention as they may indicate over control. In such cases, controls could be removed or altered to increase risk whilst still satisfying the organisation's risk criteria. The effect of such alterations could lead to greater business efficiencies and cost savings. Thus the method provides a valuable insight to the economics of internal control, which was one of the objectives of the original Brewer-List paper.

---

[6] In accordance with ISO Guide 73, a control is a measure that modifies risk. Note that most of the controls in Annex A are not controls by this definition – they are merely *measures*.

[7] A *preventive* control acts to prevent the occurrence of the event, or otherwise arrests it as it occurs, thereby preventing the occurrence of the impact.

[8] A *detective* control detects the occurrence of the event, and then acts to prevent the occurrence of the impact.

[9] A *reactive* control detects the occurrence of the impact, and then acts to reduce its severity.

[10] *Example 1*: A control may act to reduce FoL or Sev to zero or other limiting constant.

[11] *Example 2*: If a preventive control relies on some mechanism which would permit a 1 in N chance of being defeated, then it will reduce FoL by the factor N.

[12] *Example 3*: A detective control may lack the capacity to deal with multiple events. In this case, the control may be overwhelmed when the event FoL exceeds some threshold. Similarly a reactive control may be overwhelmed when the event FoL exceeds some threshold, or may otherwise have limited effect if the severity of the impact that the control has to deal with exceeds some (other) threshold.

Appendix C: Sequencing the Annex A controls

Introduction

As part of our research, conducted in 2007, we conjectured whether it was possible to string all 133 controls together in the form of a story. The answer turns out to be "yes" with every control being used once only save one which had to be split into two parts, one part being used in early on in the story and the remainder somewhat later.

The story is related below in nine Acts and is presented in a series of tables with intervening text. The research was originally presented as a lecture.

The nine Acts are:

- Act 1 – Deployment: reducing the likelihood of staff/contractors from causing a security breach;
- Act 2 – A secure work environment: restricting physical access to information in the workplace;
- Act 3 – Outside work: taking care when sending or using (non-IT) information outside the workplace;
- Act 4 – Open (computer) access: controlling access to computers that an attacker can physically access in the workplace;
- Act 5 – Action at a distance: protecting our computers from cyber attack;
- Act 6 – Applications: making sure that our applications are secure;
- Act 7 – Operating conditions: making sure our computer hardware works;
- Act 8 – Does it work? checking that our security controls are working before we are attacked;
- Act 9 – When things go wrong: taking action when there is an incident.

Act 1 – Deployment

Purpose: reducing the likelihood of staff/contractors from causing a security breach.

Internal control is about marshalling our resources to achieve our objectives. We wish to deploy people to do that and we want them to follow our rules. Taking the story up from here as the starting point, we allocate 'controls' as shown in the following table.

Story fragment	Annex A control
So what are our rules?	A.5.1.1 Information security policy document
They are, of course, legal and above board.	A.15.1.1 Identification of applicable legislation A.15.1.2 Intellectual property rights A.15.1.3 Protection of organisations records A.15.1.4 Data protection and privacy of personal information A.15.1.5 Prevention of misuse of processing facilities A.15.1.6 Regulation of cryptographic controls
Now we have some rules, do we have the means to enforce them? Yes, contracts of employment or similar. Let's make sure our rules are in them first. Note that this control covers contractors, etc as well.	A.8.1.1 Roles and responsibilities
Staying with the employees, let's try to make sure we don't hire the bad apples …	A.8.1.2 Screening
And when they join, let's sign the contract to say they and we agree …	A.8.1.3 Terms and conditions of employment
Of course, there are penalties if our staff do not follow our rules (otherwise what do we do if they don't follow them?)	A.8.2.3 Disciplinary process

Story fragment	Annex A control
And we can do something similar with our suppliers and customers …	A.6.2.2 Addressing security when dealing with customers A.6.2.3 Addressing security in third party agreements
Assuming, of course that we know what the risks are.	A.6.2.1 Identification of risks related to external parties

Table C1: Map of story fragments to Annex A controls

So where does that get us?

- We have some rules;
- They are legal;
- They are in all the contracts;
- The contracts are signed;
- Our employees and contactors have been screened.

The likelihood of one of them *deliberately* breaking our rules is now much lower. If someone does, however, they say "*sorry, I didn't know that it meant* that" and people also make mistakes.

Continuing …

Story fragment	Annex A control
We tackle the first of these by training them and making them aware.	A.8.2.2 Information security awareness, education and training
And by ensuring that everyone knows what their responsibilities are, and cooperates	A.6.1.3 Allocation of information security responsibilities A.6.1.2 Information security co-ordination
And we can always write down more detailed instructions where appropriate (note that there are other controls like this)	A.10.1.1 Documented operating procedures
We tackle the second (at least in the first instance) through leadership …	A.6.1.1 Management commitment to information security
Through supervision …	A.8.2.1 Management responsibilities A.15.2.1 Compliance with security policies and standards
And by making it difficult for people to cheat …	A.10.1.3 Segregation of duties
If we need to do something with the people, the mechanisms are in A.8.2.2/3, but if we need to change the rules …	A.5.1.2 Review of the information security policy
And we should learn from others as well as ourselves.	A.6.1.7 Contact with special interest groups A.13.2.2 Learning from information security incidents

Table C2: Map of story fragments to Annex A controls (continued)

So what does this achieve? We have now done our best, using the Annex A 'controls' to counter the inappropriate deployment of people. The residual risks are now:

- People might still knowingly and deliberately break the rules – but they know the consequences if they get caught;
- People will still make mistakes, perhaps through ignorance.

Act 2 – A secure work environment

Purpose: restricting physical access to information in the workplace.

Let us now look at the work environment. We will not worry about fire, flood etc as we will deal with that later. We will, however, worry about the people who are not included in the set of good people who, in Act 1, have been selected and obligated, and who are now trained, aware and competent. Expressed as a Venn diagram (Figure B1), in this and subsequent acts we concentrate on the red area.

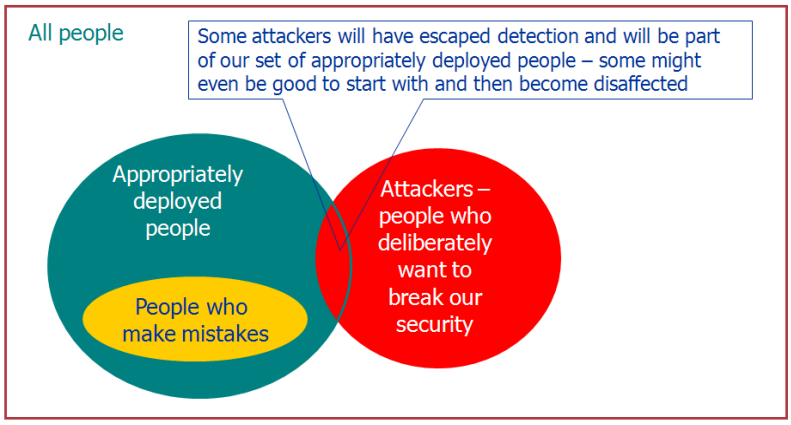

Figure B1: Partitioning of the threat space

Continuing our story …

Story fragment	Annex A control
Let's start by securing the work area, so that all people (especially outsiders) can't go just where they want to…	A.9.1.1 Physical security perimeter A.9.1.2 Physical entry controls A.9.1.3 Securing offices, rooms and facilities A.9.1.6 Public access, delivery and loading areas A.9.2.1 Equipment siting and protection[13]
OK, but what about cleaners and visitors? They may have need to access the work area, but not the information that is stored within it, so let's lock it away, or ensure that it is otherwise safe when we are not there:	A.9.1.5 Working in secure areas A.11.3.3 Clear desk and clear screen policy A.11.3.2 Unattended user equipment

Table C3: Map of story fragments to Annex A controls (continued)

The residual risks are now:

- People might overcome the physical controls (if there is a danger of that, strengthen them – better locks, CCTV, guards);
- Information may have to leave the workplace for all sorts of good business reasons;
- Computers.

Act 3 – Outside work

Purpose: taking care when sending or using (non-IT) information outside the workplace.

Let us now look at what can happen outside the workplace. In this Act, however, we will ignore IT.

[13] This is the only 'control' that is used twice. However, we actually split it into two. In this instance we use the second half: "Equipment shall be sited or protected to reduce the opportunity for unauthorised access.

We start by asking "can our information leave the workplace?" The answer is, of course, "yes" for several reasons:

- We may post it to an organisation that we are doing business with;
- We may take it to a meeting;
- We may talk about it to a business colleague on the telephone
- We may talk about it in a public place;
- We dispose of items containing information.

Continuing our story …

Story fragment	Annex A control
Let us start by making sure that we know about the removal of anything physical	A.9.2.7 Removal of property
Let's next deal with confidentiality. We take precautions depending upon the sensitivity of the information …	A.7.2.1 Classification guidelines A.7.2.2 Information labelling and handling
If we do this, we might want to maintain an inventory of what we have got...	A.7.1.1 Inventory of assets
And make people responsible for looking after them ...	A.7.1.2 Ownership of assets
These are rules, and therefore become part of our security policy. We also want people to sign up to them (then they can't complain if they break the rules and we find out and penalise them for it).	A.6.1.5 Confidentiality agreements
There are all sorts of things these rules should cover as well …	A.9.2.6 Secure disposal or re-use of equipment A.9.2.5 Security of equipment off-premises A.10.7.1 Management of removable media A.10.7.2 Disposal of media A.10.8.1 Information exchange policies and procedures A.10.8.2 Exchange agreements A.10.8.3 Physical media in transit A.10.7.3 Information handling procedures

Table C4: Map of story fragments to Annex A controls (continued)

It would be appropriate next to deal with integrity and availability. Integrity is broken if someone can intercept the information and change it. However, there are no 'controls' in Annex A that deal with this (apart from dealing with electronic messaging, which is IT). Regarding availability, if you are expecting something by post or are in the middle of a telephone conversation and it does dead, you know, but again there no 'controls' in Annex A (apart from IT 'controls') that deal with the case when the loss of availability is not quite so obvious.

Nevertheless, we may continue by taking a slightly different tack...

Story fragment	Annex A control
What happens when someone leaves? We need something to trigger our knowledge of this (other things may have to be done later as well)	A.8.3.1 Termination responsibilities
And then get back any assets we have loaned them	A.8.3.2 Return of assets

Table C5: Map of story fragments to Annex A controls (continued)

Act 4 – Open (computer) access

Purpose: controlling access to computers that an attacker can physically access in the workplace.

Story fragment	Annex A control
At the moment we don't have any computer access control, but if a computer is that sensitive, it ought not be connected to anything, with the tightest ever physical security in place …	A.11.6.2 Sensitive system isolation
For everything else, the first step is to define the rules…	A.11.1.1 Access control policy
Now let's give our people (the ones we entrusted in the first place) an account …	A.11.2.1 User registration
And it is particularly important when someone leaves…	A.8.3.3 Removal of access rights
Make sure they use it and protect it…	A.11.2.3 User password management A.11.3.1 Password use
Ideally (particularly concerning access to security enforcing stuff, like operating systems) …	A.11.5.1 Secure logon procedures A.11.5.2 User identification and authentication A.11.5.3 Password management system
And some access rights…	A.11.2.2 Privilege management
And, of course periodically check that they are correct	A.11.2.4 Review of user access rights
Maybe enforce some restrictions…	A.11.6.1 Information access restriction

Table C6: Map of story fragments to Annex A controls (continued)

The residual risks are now that the 'controls' may be imperfect, e.g. access rights might not be absolutely correct, or can be circumvented, passwords may be broken, but overall control is much tighter. Nevertheless, the risks that now remain will increase with greater connectivity – our networks and/or the Internet. We deal with this next.

Act 5 – Action at a distance

Purpose: protecting our computers from cyber attack.

Story fragment	Annex A control
The first step could be to partition the networks, just like we did with different areas of the working environment… (Note that this control includes firewalls)	A.11.4.5 Segregation in networks
We then need to ensure that users only have access to those parts of the network that we want them to have access to…	A.11.4.6 Network connection control
And the routers do what we want…	A.11.4.7 Network routing control
And they can only connect to the services that we want them to have access to…	A.11.4.1 Policy on use of network services
And use them for only certain purposes…	A.7.1.3 Acceptable use of assets
Now if we can connect to computers at the far ends of our networks, other people might be able to connect to us. Who are they?	A.11.4.2 User authentication for external connections
What is being connected?	A.11.4.3 Equipment identification in networks
Is it possible for an attacker to hijack a session?	A.11.5.5 Session time-out A.11.5.6 Limitation of connection time
Or gain access through a cable?	A.9.2.3 Cabling security
Or by any other means?	A.12.5.4 Information leakage A.12.3.1 Policy on the use of cryptographic controls A.12.3.2 Key management
And if vendors (who know about security) connect to us, let's be particularly careful…	A.11.4.4 Remote diagnostic and configuration port protection
Despite these controls, an attacker might to able to circumvent them for mount a denial of service attack by exploiting some technical vulnerability…	A.12.6.1 Control of technical vulnerabilities
Or plant a virus…	A.10.4.1 Controls against malicious code

Story fragment	Annex A control
And, of course, we need to be able to manage all of this…	A.10.6.1 Network controls A.10.6.2 Security of network services
And if we use mobile code to help us we need to make sure no one else can…	A.10.4.2 Controls against mobile code
Finally if we allow computing on the move, or teleworking we need all of this, with greater security in the IT, (a) because the physical environment is outside our scope of control (b) it is still likely to be connected to us…	A.11.7.1 Mobile computing and communications A.11.7.2 Teleworking

Table C7: Map of story fragments to Annex A controls (continued)

We next turn our attention to our software applications.

Act 6 – Applications

Purpose: making sure that our applications are secure.

Story fragment	Annex A control
What should they do?	A.10.8.5 Business information systems
Which could mean…	A.10.9.1 Electronic commerce A.10.9.2 On-line transactions A.10.9.3 Publicly available information A.10.8.4 Electronic messaging
Whoever builds our applications, we ought to specify what we want in terms of security…	A.12.1.1 Security requirements analysis and specification
And that we have sufficient capacity…	A.10.3.1 Capacity management
Typical requirements that we need to ensure that the user specifies for application security…	A.12.2.1 Input data validation A.12.2.2 Control of internal processing A.12.2.3 Message integrity A.12.2.4 Output data validation
If we outsource development…	A.12.5.5 Outsourced software development
If we do it ourselves, we must ensure that we don't confuse the development environment with the live environment…	A.10.1.4 Separation of development, test and operational facilities
In all cases, only the developers should have access to the source code and the test data…	A.12.4.3 Access control to program source code A.12.4.2 Protection of system test data A.10.7.4 Security of system documentation
The systems must formally be accepted before being put into use…	A.10.3.2 System acceptance
Thereafter, changes must be approved and properly carried out…	A.10.1.2 Change management A.12.5.3 Restrictions on changes to software packages A.12.5.1 Change control procedures
But other things might change, like operating systems. We must ensure that these do not have a bad affect on our applications…	A.12.5.2 Technical review of applications after operating system changes
But application software is easy to get hold of these days, could anyone just install something against our wishes etc…	A.12.4.1 Control of operational software A.11.5.4 Use of system utilities
Or use their own facilities…	A.6.1.4 Authorisation process for information processing facilities
Rather than run the applications ourselves, we could outsource that as well, perhaps in the form of 'software as a service' or part of some larger and more significant outsourcing contract Either way it needs to be controlled in a similar fashion…	A.10.2.1 Service delivery A.10.2.2 Monitoring and review of third party services A.10.2.3 Managing changes to third party services

Table C8: Map of story fragments to Annex A controls (continued)

We have dealt with all the technical IT controls. Nevertheless we still need to pay attention to our computer hardware.

Act 7 – Operating conditions

Purpose: making sure our computer hardware works.

Story fragment	Annex A control
Our IT needs power and appropriate operating conditions	A.9.1.4 Protecting against external and environmental threats A.9.2.1 Equipment siting and protection[14] A.9.2.2 Supporting utilities
And needs to be in a good state of repair…	A.9.2.4 Equipment maintenance

Table C9: Map of story fragments to Annex A controls (continued)

We have now dealt with 112 out of the 133 controls. These are all to do with prevention. In the final two chapters we deal with the remaining 21 controls which are detective and reactive controls.

Act 8 – Does it work?

Purpose: checking that our security controls are working before we are attacked.

Story fragment	Annex A control
Rather than waiting for something to happen, how do we know if they will work? Let's audit, making sure that that does not interfere with the business …	A.15.3.1 Information system audit controls
Let's do some technical checks …	A.15.2.2 Technical compliance checking
And even invite someone else to do that for us…	A.6.1.8 Independent review of information security

Table C10: Map of story fragments to Annex A controls (continued)

Act 9 – When things go wrong

Purpose: taking action when there is an incident.

Story fragment	Annex A control
What happens when things go wrong? But first, how quickly can we find out? We could simply watch..	A.10.10.2 Monitoring system use
People can report things…	A.13.1.1 Reporting information security events A.13.1.2 Reporting security weaknesses
We can log things (all of this is also useful in investigating what happened afterwards as well)…	A.10.10.1 Audit logging A.10.10.4 Administrator and operator logs A.10.10.5 Fault logging
We have the audit data and tools…	A.15.3.2 Protection of system audit tools
We need to protect this information, particularly if it is going to be used in evidence (and remember to preserve that chain of evidence)…	A.10.10.3 Protection of log information A.13.2.3 Collection of evidence A.10.10.6 Clock synchronisation
And liaise with the authorities…	A.6.1.6 Contact with authorities
When there is an incident, we need to know who is doing to do what…	A.13.2.1 Responsibilities and procedures
Recovery might be as simple as restoring a back-up…	A.10.5.1 Information back-up
Or it might require us to deploy our disaster recovery plan, already well thought out and tested…	A.14.1.1 Including information security in the business continuity management process A.14.1.2 Business continuity and risk assessment A.14.1.3 Developing and implementing continuity plans including information security A.14.1.5 Testing, maintaining and re-assessing business continuity plans

Table C11: Map of story fragments to Annex A controls (continued)

[14] This is the second 'use' of this control. This time we used the part that concerns environmental protection, i.e. "Equipment shall be sited or protected to reduce the risks from environmental threats and hazards".

Appendix D: Sample definitions for the three cardinal events

EI1 – Vulnerability exploitation

An attacker exploits a security vulnerability to cause the undesirable disclosure of information, fraud or denial of service. The attacker could be an authorised user of our information, whereby they abuse their authority and do things that they should not. If the attacker is not an authorised user of our information they might attempt to masquerade as an authorised user. Perhaps the attacker can gain access to our information by some other means, or simply eavesdrop.

EI2 – IT failure

Our IT fails because of a hardware or software malfunction. The malfunction can be brought about in a variety of ways, such as lack of power, loss of Internet connectivity, adverse operating conditions (fire, flood etc.), unreliability and specification/design/implementation errors.

EI3 – Dispossession

A physical container of information is dispossessed. Typical containers are documents, envelopes, briefcases, laptops, desktops, servers, PDAs, mobile phones, cameras, magnetic tapes, CDs, DVDs, USB sticks etc. Dispossession could be because the container is lost, stolen, damaged or destroyed, misappropriated (e.g. lost in the post). Dispossession might also be because the container is disposed of or reused (by someone else). In all cases the owner, or rightful user, of the container no longer has the container in their possession.

Appendix E: Assets, threats and vulnerabilities

Introduction

The following three tables present the assets, threats and vulnerabilities identified during the analysis of the Annex A controls. Their relation to those controls is presented in Appendix D.

Assets

Asset	Description	Event association
Data back-ups	The means to recover data (and software) to a previous good state.	EI1, EI3
Electricity	IT requires electricity to work.	EI2
Environmental conditioning	IT will only operate within a given window of temperature extremes, humidity, electric/magnetic fields and so on.	EI2
Forensic computer evidence	The evidence required to determine beyond a reasonable doubt that who did what to commit a particular crime. The evidence must be obtained in a proper manner, there being no possibility that the method of acquisition has in anyway contaminated the evidence, and the chain of evidence must be preserved. Other conditions may apply for it to be admissible in a court of law.	EI1
Information containers	Anything that contains or may contain information, e.g. documents, envelopes, briefcases, laptops, desktops, servers, PDAs. mobile phones, cameras, magnetic tapes, CDs, DVDs, USB sticks etc.	EI3
Information security software/appliances	Software and/or hardware that has a particular role to play in information security, e.g. a firewall, anti virus software, an authentication server etc.	EI1
Reliable communications	Communications between computers is dependent on the reliability of the interconnecting communications systems. There are many forms of communication links: cable, radio, satellite, local area, wide area, fixed and mobile; and they will have variety of characteristics such as band width.	EI2
Reliable hardware	Given electricity and proper operating conditions, the computer hardware still needs to do what it is supposed to do, consistently and on demand.	EI2
Sensitive information	The definition of this asset is organisational specific.	EI1, EI3
Software that does what it is supposed to do	Given reliable computer hardware, the software must do what it is supposed to do, consistently and on demand.	EI1, EI2

Threats

Threat	Description	Event association
Contractors	People who work for the organisation but are not employees. They pose a threat because they might act *inadvertently* against the interests of the organisation. In contrast to employees, despite having a duty of care to the organisation, their loyalty will lie with their own employer.	EI1, EI3
Customers/Clients	Third party organisations that commission and buy an organisation's products. The term customer is used in ISO 9001 to describe the supply chain thus: supplier → organisation → customer.	EI1
Disaffected workers	An employee or contractor who may act *willfully*against the interests of the organisation.	EI1, EI3
Employees	A person employed by an organisation for wages or salary. They pose a threat because they might act *inadvertently* against the interests of the organisation. In contrast to contractors their loyalty lies with the organisation.	EI1, EI3
Fire, water and adverse operating conditions	These pose a threat as at one extreme they may prevent the IT from working properly and at the other destroy it completely, along with personnel and the place of work etc. Adverse operating conditions includes cyclone, dust storms, earthquakes etc. which can lead to destruction through building collapse.	EI2, EI3
Fraudsters	People who commit (or attempt to commit) wrongful or criminal deception intended to result in financial or personal gain.	EI1
Hackers and cyber-criminals	People who use a computer to gain unauthorised access to data.	EI1
New possessor	A person who finds or otherwise comes into the possession of an information container after it has been dispossessed by its owner or custodian	EI3
Spies, competitors and investigative journalists	People whose job is to discover non publicly available information.	EI1
Software developers	People whose job is to develop the software used by the organisation.	EI2
Thieves	People who steal or wish to steal the organisation's property.	EI3
Third party service provider	The provider of services to the organisation. Such services include, water, electricity, communications, software applications etc.	EI1, EI2, EI3
Use	Equipment, such as computer hardware and other electronic (or mechanical) devices, may fail because of being used - or for that matter, not being used.	EI2
Vandals, terrorists, rioters and war	People or circumstances who may damage or destroy the organisation IT and facilities. The organisation may be the target of such actions, or it may be the victim of collateral damage or indiscriminate actions.	EI2, EI3

Vulnerabilities

Vulnerability	Description	Event association
Addressing information can be forged	Addressing information in communication header, for example, can be changed to make it appear that a communication can from somewhere other than it was really sent.	EI1
Arbitrary program code can be executed	A common software vulnerability (e.g. a buffer overflow) that allows the an attacker to cause the victim's computer to execute a short program of the attacker's creation.	EI1
Attractiveness of our information	The very nature of information may make it the target of attack.	EI1
Authentication credentials can be forged	No matter the authentication method, an enterprising attacker can invariably forge the authentication credentials	EI1
Communication lines can be tapped	Devices can be attached to communication lines to intercept the information as it passes.	EI1
Communications systems are outside our scope of control	Often, communication services are provided by third parties. Messages can be intercepted, or data stored on servers read and potentially modified. Unless, the information is end-to-end encrypted, protection is at the mercy of the service provider, and may or may not be good enough.	EI1
Component failure	Electrical components (resistors, capacitors, integrated circuits etc) may fail after use.	EI2

Vulnerability	Description	Event association
Electromagnetic radiation	All electrical devices create electric fields when they are operated. It is possible to detect these and determine for example, what information is being displayed on a computer screen. The strength of the signals is a function of the strength of the source of radiation, the distance between it and the receiver and the nature of any intervening obstacles.	EI1
Ease by which software can be changed	The whole value of software is the ease by which it can be changed. The idea is, of course, to change the software for the better, but it could also be changed for the worse and therein lies the vulnerability.	EI2
Frequency of changes required to keep pace with business needs	Each time that software is changed presents an opportunity for it to be changed for the worse.	EI2
Information containers can be damaged/destroyed	Being physical devices, information containers are not impervious to damage or destruction.	EI3
Information containers are not particularly heavy and can be easily moved	Such is the nature of modern day computing that a great deal of information can be packed into a very small device. Being easily transportable, a great deal of information can be lost.	EI3
Information containers may be valuable	Particular types of information containers (laptops, mobile phones) are valuable and can be resold on the black market.	EI3
Information is extractable from its container	The utility of the container is information, once deposited into the container can be extracted. There is be an intended way to do this, but there may be other ways as well. For example, a hard disc could be removed from the computer that uses it and put into another in order to read it.	EI3
Ignorance, misunderstanding and human fallibility	People, however well intentioned, may act out of ignorance of what is expected of them to maintain security. They may misunderstand what to do and they may make mistakes.	EI1
Network connectivity	The greater the connectivity between computers, the greater the number of people who could potentially attack one of ours.	EI1
People are gullible	People may wish to attack us just because we are who we are.	EI1
Protocols can be manipulated	Design/implementation errors in communication protocols may make them prone to misuse and abuse. There are a large number of different types of known attacks of this nature, e.g. ping flood, smurf, ping of death, teardrop, LAND, ...	EI1
Need for power, connectivity and favourable operating conditions	Without power, connectivity (e.g. Internet) and favourable operating conditions computers will either not function at all, or not function effectively or reliably.	EI2
Reliance of third parties for software development and/or service provision	Reliance is in two forms. First the service is expected to be there when it is required. Secondly it is expected to do what it is supposed to do. Neither are within our scope of control.	EI2
Software can be installed/removed	Because software can be installed/removed, it is possible to install bad programs and remove good ones. This what many viruses attempt to do.	EI1
Software complexity	Modern day software is extremely complex. It is difficult to prove program correctness even in very small programs. It is therefore highly likely that our software contains errors and unintended features.	EI2
Software does not always do what it is supposed to do and often does what it is not supposed to do	Software integrity is the property the software does what it is supposed to do and does not do what it is not supposed to do. We regard the first of these properties as an asset; the second as a vulnerability.	EI1
Uncertainty of user/security requirements	Despite the best intentions of users and software developers it is possible to get the requirements wrong, both from the perspective of business functionality and the perspective of security.	EI2
Wireless telegraphy can be intercepted	Wireless communication works by transmitting in the radio or microwave frequency bands. It is possible to intercept the radiation and decode the signals.	EI1

Appendix F: Definition of Real controls and mapping to Annex A controls

Number	Description	Associated Annex A controls
Human resources security		
1	The selection of people and external parties that the organisation wishes to have access to its information, their obligation to adhere to the policies and procedures of the organisation, the measures that the organisation may have against the person or external party should that person or party fail to fulfil their obligations, and the recourse that the organisation may take to successfully recover its losses, and defend itself accordingly	A.5.1.1, A.5.1.2, A.6.1.5, A.6.1.6, A.7.1.3, A.8.1.1, A.8.1.2, A.8.1.3, A.8.2.3, A.10.8.1, A.10.8.2, A.10.10.4, A.10.10.6, A.11.7.1, A.11.7.2, A.12.3.1, A.13.2.3, A.15.1.1, A.15.1.2, A.15.1.3, A.15.1.4, A.15.1.5, A.15.1.6
2	The education of those people and external parties to ensure that they are information security aware, understand the policies and procedures of the organisation, are aware of their obligations and are competent to fulfil their responsibilities, specifically in the areas of:	A.6.1.7, A.7.2.2, A.8.2.2, A.10.1.1, A.13.2.2
2a	Good password practice	A.8.2.2, A.11.3.1
2b	Social-engineering	A.8.2.2
2c	Discrete behaviour, especially in public	A.8.2.2
3	The routine checking and supervision of people and external parties	A.6.1.1, A.6.1.2, A.6.1.3, A.7.1.2, A.8.1.1, A.8.2.1, A.8.3.1, A.13.1.2, A.13.2.1
4	Reporting of security incidents	A.13.1.1
5	Periodic audit sampling	A.6.1.8, A.15.2.1, A.15.2.2, A.15.3.1
Physical, utility and environmental security		
6	The regulation of physical access to the organisation's premises and assets	A.7.2.2, A.9.1.1, A.9.1.2, A.9.1.3, A.9.1.4, A.9.1.5, A.9.1.6, A.9.2.1, A.9.2.5, A.10.7.3, A.11.3.2, A.11.3.3, A.11.7.1, A.11.7.2
7	Protection against fire, water, temperature and other environmental hazards	A.9.1.4
8	Protection against hardware failure	A.9.2.4
9	Protection against power failure	A.9.2.2
10	Protection against communications failure	A.9.2.2
11	The reliability and management of IT service providers	A.10.2.1, A.10.2.2, A.10.2.3
12	Recording and mustering of assets	A.7.1.1, A.8.3.2, A.9.2.7
13	Protection against loss or damage of assets during transit	A.7.2.2, A.10.7.1, A.10.7.3, A.10.8.3
14	The sanitisation of electronic media before reuse or disposal	A.9.2.6, A.10.7.2
15	The secure destruction of all media	A.9.2.6, A.10.7.2
Technical security		
16	User authentication	A.11.2.1, A.11.2.3, A.11.4.2, A.11.5.1, A.11.5.2, A.11.5.4
17	The regulation of electronic access within and outside the organisation to its IT systems and networks, in accordance with the organisation's access control policies and requirements for the prevention of fraud	A.7.2.2, A.8.3.3, A.10.1.3, A.10.6.1, A.10.6.2, A.10.7.3, A.11.1.1, A.11.2.2, A.11.2.4, A.11.4.1, A.11.4.3, A.11.4.4, A.11.4.5, A.11.4.6, A.11.4.7, A.11.6.1, A.11.6.2, A.11.7.1, A.11.7.2, A.12.4.3
18	The prevention of hijacking of electronic communication sessions, both wired and wireless	A.10.8.4, A.10.8.5, A.10.9.1, A.10.9.2, A.10.9.3, A.11.5.6, A.11.5.7
19	The prevention of active and passive eavesdropping on electronic communications	A.7.2.2, A.9.2.3, A.10.7.3, A.10.8.4, A.10.8.5, A.10.9.1, A.10.9.2, A.10.9.3, A.12.3.1, A.12.3.2

Number	Description	Associated Annex A controls
20	The prevention of masquerading and 'man-in-the-middle' attacks	A.10.8.4, A.10.8.5, A.10.9.1, A.10.9.2, A.10.9.3, A.12.3.1, A.12.3.2
21	The prevention, detection and removal of malware	A.10.4.1, A.10.4.2, A.12.5.4
22	Hardening of operating systems and maintaining patch levels	A.12.6.1
23	Hard disc encryption	A.12.3.1, A.12.3.2
24	The detection of anomalous/abnormal network usage and penetration	A.10.3.1, A.10.10.1, A.10.10.2, A.10.10.5, A.10.10.6
25	The taking and restoration of software and data backups	A.10.5.1
System security		
26	Prevention of the use of unauthorised or illegal system components	A.6.1.4, A.10.3.2, A.12.4.1
27	The use of reputable commercial-off-the-shelf system components	A.12.5.4
28a	The use of reliable system engineering techniques, ensuring that: a. security requirements are fully understood, particularly those concerning the intended business application	A.6.2.1, A.6.2.2, A.6.2.3, A.10.3.1, A.10.8.4, A.10.8.5, A.10.9.1, A.10.9.2, A.10.9.3, A.12.1.1, A.12.2.2, A.12.2.3, A.12.5.4, A.12.5.5
28b	b. the system components is resilient against operator error	A.12.2.1, A.12.2.4, A.12.5.5
28c	c. the system components have been reliably tested before acceptance and use	A.10.3.2, A.12.5.4, A.12.5.5
29	Change management	A.10.1.2, A.12.5.1, A.12.5.2, A.12.5.3
Business continuity		
30	The information security aspects of the organisation's business continuity plans	A.14.1.1, A.14.1.2, A.14.1.3, A.14.1.4, A.14.1.5
Risk avoidance		
31	Not putting information in places where there is little or no protection	A.10.1.4, A.10.7.4, A.10.9.3, A.10.10.3, A.11.5.4, A.12.4.2, A.15.3.2

Appendix G: Mapping of Annex A controls to Real controls

Annex A reference	Name of Annex A control	Real control number (see Appendix F)
A.5.1.1	Information security policy document	1
A.5.1.2	Review of information security policy	1
A.6.1.1	Management commitment to information security	3
A.6.1.2	Information security coordination	3
A.6.1.3	Allocation of information security responsibilities	3
A.6.1.4	Authorization process for information processing facilities	26
A.6.1.5	Confidentiality agreements	1
A.6.1.6	Contact with authorities	1
A.6.1.7	Contact with special interest groups	2
A.6.1.8	Independent review of information security	5
A.6.2.1	Identification of risks related to external parties	28a
A.6.2.2	Addressing security when dealing with customers	28a
A.6.2.3	Addressing security in third party agreements	28a
A.7.1.1	Inventory of assets	12
A.7.1.2	Ownership of assets	3
A.7.1.3	Acceptable use of assets	1
A.7.2.1	Classification guidelines	2
A.7.2.2	Information labelling and handling	6,13, 17, 19
A.8.1.1	Roles and responsibilities	1, 3
A.8.1.2	Screening	1
A.8.1.3	Terms and conditions of employment	1
A.8.2.1	Management responsibilities	3
A.8.2.2	Information security awareness, education and training	2, 2a, 2b, 2c
A.8.2.3	Disciplinary process	1
A.8.3.1	Termination responsibilities	3
A.8.3.2	Return of assets	12
A.8.3.3	Removal of access rights	17
A.9.1.1	Physical security perimeter	6
A.9.1.2	Physical entry controls	6
A.9.1.3	Securing offices, rooms and facilities	6
A.9.1.4	Protecting against external and environmental threats	6. 7
A.9.1.5	Working in secure areas	6
A.9.1.6	Public access, delivery and loading areas	6
A.9.2.1	Equipment siting and protection	6
A.9.2.2	Supporting utilities	9, 10
A.9.2.3	Cabling security	19
A.9.2.4	Equipment maintenance	8
A.9.2.5	Security of equipment off-premises	6
A.9.2.6	Secure disposal or re-use of equipment	14, 15
A.9.2.7	Removal of property	12

Annex A reference	Name of Annex A control	Real control number (see Appendix F)
A.10.1.1	Documented operating procedures	2
A.10.1.2	Change management	29
A.10.1.3	Segregation of duties	17
A.10.1.4	Separation of development, test and operational facilities	31
A.10.2.1	Service delivery	11
A.10.2.2	Monitoring and review of third party services	11
A.10.2.3	Managing changes to third party services	11
A.10.3.1	Capacity management	24, 28a
A.10.3.2	System acceptance	26, 28c
A.10.4.1	Controls against malicious code	21
A.10.4.2	Controls against mobile code	21
A.10.5.1	Information back-up	25
A.10.6.1	Network controls	17
A.10.6.2	Security of network services	17
A.10.7.1	Management of removable media	13
A.10.7.2	Disposal of media	14, 15
A.10.7.3	Information handling procedures	6, 13, 17, 18, 19
A.10.7.4	Security of system documentation	31
A.10.8.1	Information exchange policies and procedures	1
A.10.8.2	Exchange agreements	1
A.10.8.3	Physical media in transit	13
A.10.8.4	Electronic messaging	18, 19, 20, 28a
A.10.8.5	Business information systems	18, 19, 20, 28a
A.10.9.1	Electronic commerce	18, 19, 20, 28a
A.10.9.2	On-line transactions	18, 19, 20, 28a
A.10.9.3	Publicly available information	18, 19, 20, 28a, 31
A.10.10.1	Audit logging	24
A.10.10.2	Monitoring system use	24
A.10.10.3	Protection of log information	31
A.10.10.4	Administrator and operator logs	1
A.10.10.5	Fault logging	24
A.10.10.6	Clock synchronization	1, 24
A.11.1.1	Access control policy	17
A.11.2.1	User registration	16
A.11.2.2	Privilege management	17
A.11.2.3	User password management	16
A.11.2.4	Review of user access rights	17
A.11.3.1	Password use	2a
A.11.3.2	Unattended user equipment	6
A.11.3.3	Clear desk and clear screen policy	6
A.11.4.1	Policy on use of network services	17
A.11.4.2	User authentication for external connections	16
A.11.4.3	Equipment identification in networks	17

Annex A reference	Name of Annex A control	Real control number (see Appendix F)
A.11.4.4	Remote diagnostics and configuration port protection	17
A.11.4.5	Segregation in networks	17
A.11.4.6	Network connection control	17
A.11.4.7	Network routing control	17
A.11.5.1	Secure log-on procedures	16
A.11.5.2	User identification and authentication	16
A.11.5.3	Password management system	16
A.11.5.4	Use of system utilities	31
A.11.5.5	Session time-out	18
A.11.5.6	Limitation of connection time	18
A.11.6.1	Information access restriction	17
A.11.6.2	Sensitive system isolation	17
A.11.7.1	Mobile computing and communications	1, 6, 17
A.11.7.2	Teleworking	1, 6, 17
A.12.1.1	Security requirements analysis and specification	28a
A.12.2.1	Input data validation	28b
A.12.2.2	Control of internal processing	28a
A.12.2.3	Message integrity	28a
A.12.2.4	Output data validation	28b
A.12.3.1	Policy in the use of cryptographic controls	1, 19, 20, 23
A.12.3.2	Key management	19, 20, 23
A.12.4.1	Control of operational software	26
A.12.4.2	Protection of system test data	31
A.12.4.3	Access control to program source code	17
A.12.5.1	Change control procedures	29
A.12.5.2	Technical review of applications after operating system changes	29
A.12.5.3	Restrictions on changes to software packages	29
A.12.5.4	Information leakage	21, 27, 28a, 28c
A.12.5.5	Outsourced software development	28a, 28b, 28c
A.12.6.1	Control of technical vulnerabilities	22
A.13.1.1	Reporting information security events	4
A.13.1.2	Reporting security weaknesses	3
A.13.2.1	Responsibilities and procedures	3
A.13.2.2	Learning from information security incidents	2
A.13.2.3	Collection of evidence	1
A.14.1.1	Including information security in the business continuity management process	30
A.14.1.2	Business continuity and risk assessment	30
A.14.1.3	Developing and implementing continuity plans including information security	30
A.14.1.4	Business continuity planning framework	30
A.14.1.5	Testing, maintaining and re-assessing business continuity plans	30
A.15.1.1	Identification of applicable legislation	1
A.15.1.2	Intellectual property rights (IPR)	1

Annex A reference	Name of Annex A control	Real control number (see Appendix F)
A.15.1.3	Protection of organizational records	1
A.15.1.4	Data protection and privacy of personal information	1
A.15.1.5	Prevention of misuse of information processing facilities	1
A.15.1.6	Regulation of cryptographic controls	1
A.15.2.1	Compliance with security policies and standards	5
A.15.2.2	Technical compliance checking	5
A.15.3.1	Information systems audit controls	5
A.15.3.2	Protection of information systems audit tools	31