Information **S**ecurity **M**anagement **S**ystem

Using attributes for better information security controls

Summary

This paper extends the concept of 'control attributes' introduced in ISO/IEC 27002:2022, discussing a wider variety of factors potentially worth bearing in mind when considering, selecting, designing or reviewing information security controls intended to mitigate unacceptable information risks. It includes pragmatic suggestions on how to make use of control attributes in the business context.

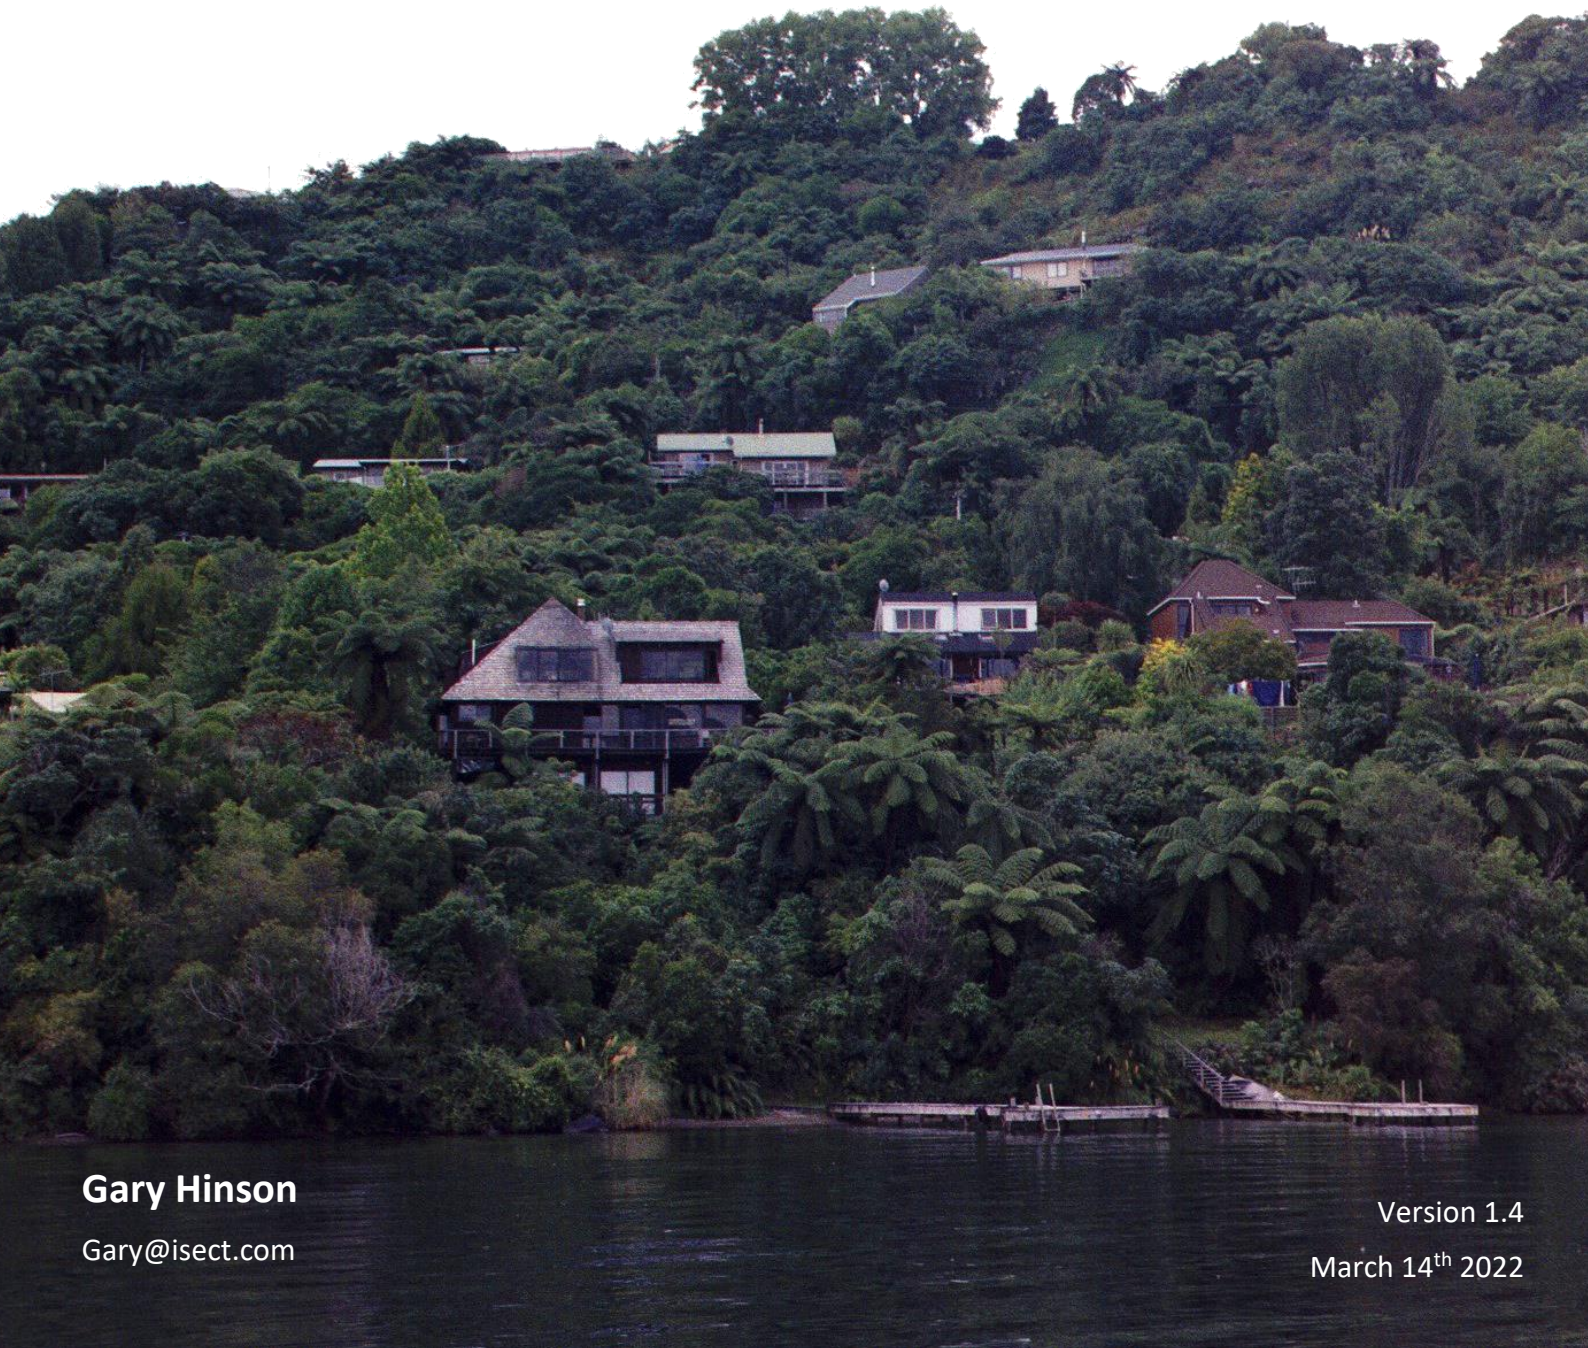

Gary Hinson

Gary@isect.com

Version 1.4

March 14th 2022

Contents

1 Introduction 2

2 Definitions 3

3 Additional control attributes, beyond those in ISO/IEC 27002 3

3.1 Assurance 4
3.2 Complexity 5
3.3 Contextual-fit 6
3.4 Control systems 6
3.5 Failure modes 7
3.6 Formality 7
3.7 Incident scenarios 8
3.8 Integration 9
3.9 Maturity 9
3.10 Measurability 10
3.11 Multi-functionality 10
3.12 Origin 11
3.13 Regulation 11
3.14 Strength 12
3.15 Targets 12
3.16 Transparency 13
3.17 Value 13
3.18 Other controls, other attributes 13

4 Using control attributes in practice 14

4.1 Determining information security control requirements 14
4.2 Discovering relevant controls potentially worth adopting 15
4.3 Selecting or rejecting controls 15
4.4 Strengthening information security (improving control) 16
4.5 Achieving a ‘balanced’ suite of controls 16
4.6 Maximising the value of existing and proposed controls 17
4.7 Reviewing, assessing and auditing controls 17
4.8 Audience views 17

5 Conclusion 18

6 References 18

Appendix: examples 19

Preface

This paper was inspired by the 2022 third edition of ISO/IEC 27002, specifically its use of themes and attributes to structure 90-odd information security controls – a radically different approach to previous versions. Work has commenced within ISO/IEC JTC 1/SC 27 on drafting a new ISO27k standard – ISO/IEC 27028 – covering control attributes. I plan to contribute content from this paper to the committee in order to flesh-out the initial working draft standard and hopefully speed-up its development and release.

I have been reminded to follow the committee's formal process, of course allowing due consideration and improvement by representatives of all the national standards bodies that participate in SC 27. To be clear, my purpose in circulating this paper and inviting public comment is to generate and offer the committee high quality inputs that support and enable - rather than undermine or constrain - the hard-working editors, other experts and the national bodies in any way. *I mean no harm! I come in peace!*

Your feedback comments and (especially) improvement suggestions are *very* welcome: please email Gary@isect.com and/or by all means discuss this online and route your comments to SC 27 through your national standards bodies.

ISO/IEC 27028 will gradually take shape over the next year or three. We have the opportunity to make substantive, creative contributions now, while it is still reasonably fluid and amenable to change.

1 Introduction

In accordance with the information risk management process laid out in ISO/IEC 27005, having identified, assessed and evaluated information risks, the next stage is to select or design the corresponding information security controls if it is deemed necessary to mitigate the risks.

ISO/IEC 27002:2022 provides copious advice on commonplace information security controls. The controls are each assigned to one of four broad categories known as 'themes' (organizational; people; physical; technological) and further tagged with applicable attributes and attribute values from this list:

- **Control types**: preventive; detective; corrective. These attribute values reflect the time relative to an incident at which the controls are most effective (before, during or after occurrence).
- **Information security properties**: confidentiality; integrity; availability. Information security involves protecting these characteristics of information. Individual controls may support one or more properties.
- **Cybersecurity concepts**: identify; protect; detect; respond; recover – another breakdown of cybersecurity activities before, during and after incidents, as explained in ISO/IEC TS 27110 and NIST's Framework for Improving Critical Infrastructure Cybersecurity.
- **Operational capabilities**: governance; asset management; information protection; human resource security; physical security; system and network security; application security; secure configuration; identity and access management; threat and vulnerability management; continuity; supplier relationships security; legal and compliance; information security event

management; information security assurance. These are generally recognised areas or specialisms within the field of information security.

- **Security domains**: governance and ecosystem; protection; defence; resilience. Described as 'information security fields, expertise, services and products', this is just another way to categorise information security controls.

Notice that the attributes are overlapping, not alternatives. Controls can be categorised using the applicable attribute values from all attributes. Furthermore the attribute values themselves are not always distinct, hence controls may possess more than one or a range of values for any attribute, having the characteristics of several attribute values. There is a many-to-many relationship between information security controls, attributes and attribute values.

2 Definitions

- **Attribute** is a type of characteristic of a control[1].
- **Attribute value** is a particular form or quantity, or a category or range of quantities, for a given attribute. Some attribute values are discrete, while others fall across a continuum or spectrum.
- **Control,** or information security control, is a means or mechanism intended to mitigate information risk. Whether it actually does so in reality is another matter.
- **Incident** is the result of a threat acting on or exploiting a vulnerability, causing an impact.
- **Impact** is an adverse consequence on the owners of information and other stakeholders, most notably an organisation that is proactively managing its information risks.
- **Information risk** is risk pertaining to or involving information.
- **Information security** is a body of practices that seeks to protect valuable information, while also facilitating its legitimate exploitation.
- **Risk** is a combination of the probability and impact of adverse consequences arising from incidents.
- **Risk mitigation** means reducing unacceptable risks, usually, but can also mean stabilising or preventing risks from increasing.
- **Risk treatment** means avoiding, mitigating, sharing and/or accepting risks. A given risk may require one or more forms of treatment, perhaps mitigating part and accepting the remainder.
- **Threat** is an external factor that impinges directly or indirectly on information, potentially leading to an incident.
- **Vulnerability** is an inherent weakness or failing that may be exploited by a threat, potentially leading to an incident.

3 Additional control attributes, beyond those in ISO/IEC 27002

Clause 4.2 of ISO/IEC 27002:2022 says, in part: "*The organization can use attributes to create different views which are different categorizations of controls as seen from a different perspective*

[1] In other contexts, 'attributes' may refer to parameters, arguments or settings for software controls (such as the sizes, colours and actions triggered when users click on-screen buttons and selectors) or the control configuration details (such as the rights and permissions granted/denied using access controls).

to the themes". This section suggests several additional attributes potentially worth taking into consideration, as appropriate.

The control attributes outlined in ISO/IEC 27002, plus others described below, are applicable in various circumstances. It is unlikely that an organisation would want to use all of them at once.

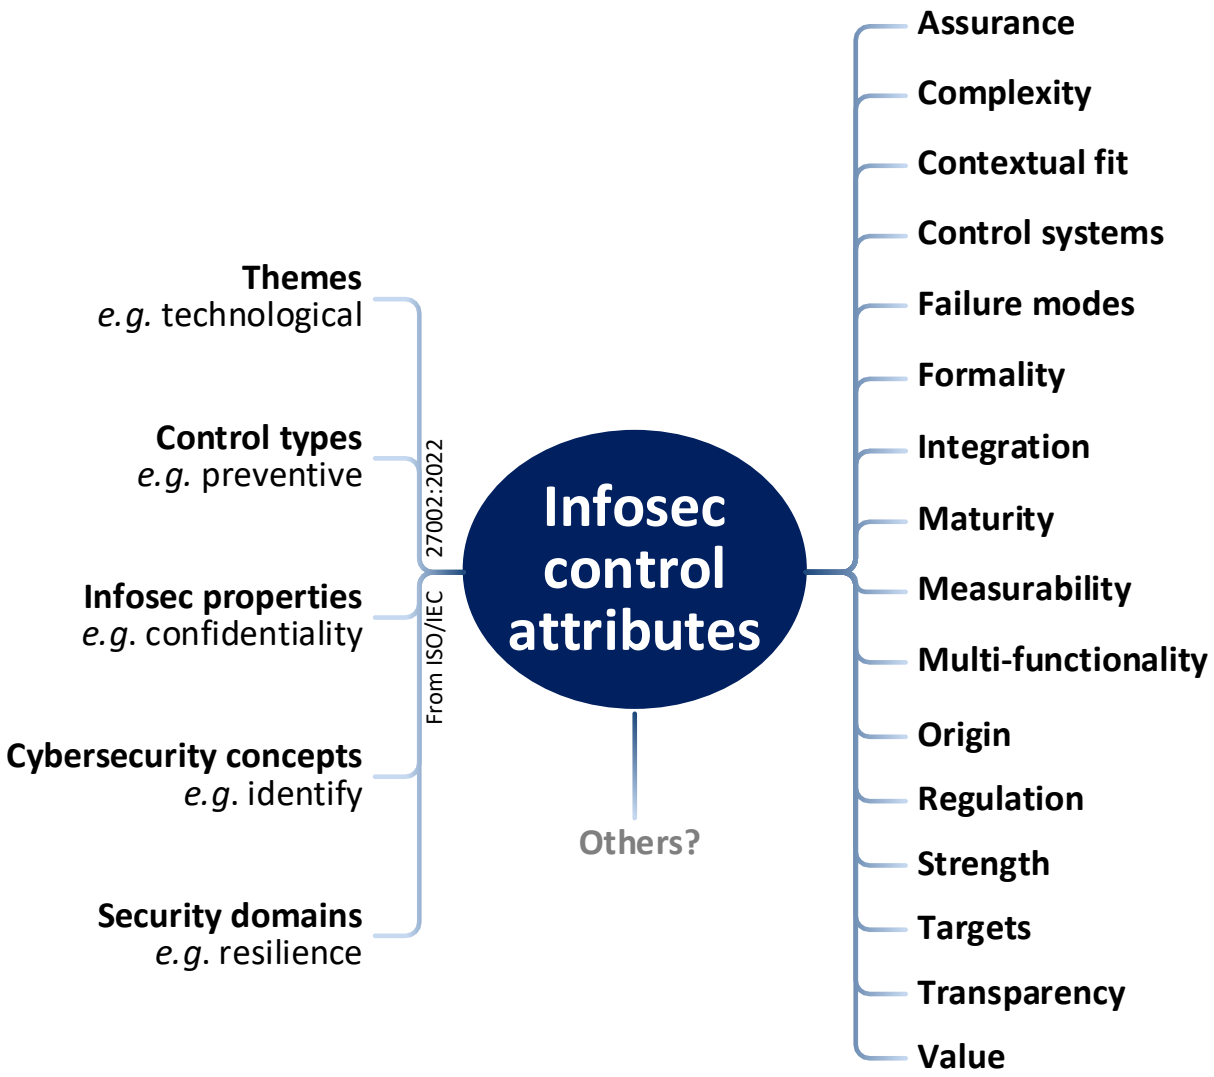

3.1 Assurance

The 'assurance' or 'reliability' attribute concerns *the extent to which a control can be relied upon* to mitigate information risks as intended.

High-assurance controls such as fire safes used for data backups and archives are engineered to meet, and often certified against, the relevant standards or requirements. Provided they are installed and used in an approved manner (*e.g.* appropriately located, fixed in place, locked shut when not being used, with the keys and codes being properly secured), there is high confidence and little doubt that they are performing correctly (in accordance with their specifications). They can be trusted to do what they are meant to do.

Low-assurance controls such as most trust-based controls rely largely on faith, the *belief* that they are effective: it can be difficult to gain confidence let alone prove beyond doubt that they are

performing correctly, especially in situations where those involved stand to gain by undermining or breaking the control, and have the motivation and opportunity to conceal inappropriate activities.

There is a substantial middle ground for controls offering partial assurance - controls that are somewhat but not entirely trustworthy and dependable. This attribute may therefore be used to prioritise or weight certain control options over others under consideration. A definite business requirement for high assurance (*e.g*. for 'key controls' whose failure is untenable due to severe consequences) may drive the search for, or design of, appropriate controls, discounting others that offer insufficient assurance. Alternatively, it may be appropriate to select additional controls to compensate for a given control's weaknesses (*e.g*. more frequent or thorough audits and reviews of key controls lacking in assurance, or improving the control monitoring arrangements).

3.2 Complexity

Relatively simple or basic controls are typically favoured over more advanced and complex ones because they are easier to understand, use, manage *etc*. They tend to be cheaper and more reliable. Where basic controls are inadequate, the principle extends to more advanced controls in terms of preferring simplicity of design and operation, also known as elegance - a security engineering concept. Given the choice, controls that have clearly been well engineered (properly thought-through, explicitly specified and designed and proven to meet the requirements) have greater credibility and value, provided they satisfy the organization's control objectives (which may not fully align with the design requirements).

There are two slightly different aspects to this. Some controls are:

- Inherently complex *e.g*. advanced cryptography; and/or
- Complicated or difficult to adopt, use and manage within the organization's existing control framework or structure.

Having been designed and developed by specialists, the complexities inherent in some controls are often handled deep within the associated technologies (*e.g*. smart cards and other cryptographic modules, functions or subsystems). Nevertheless, they may not be easy to implement, use and manage in practice, with the additional risk that mistakes or shortcuts might weaken the controls.

Even relatively simple controls can be problematic to adopt if (for some reason) the organization is not prepared (*i.e.* willing and able) to change accordingly. Any change can be quite tough for mature, stable, large organizations, particularly if there are strong compliance obligations or other external constraints imposed upon them. The pace and scale of change are factors that often constrain an organisation's ability to cope with additional demands, such as implementing an **I**nformation **S**ecurity **M**anagement **S**ystem.

Familiar information security controls are mostly 'compound' or 'complex' controls comprised of multiple 'atomic' or 'elemental' controls. For example, ISO/IEC 27002:2022 clause 8.30 "Outsourced development" recommends that activities related to outsourced system development should be directed, monitored and reviewed: some would consider those to be three distinct controls at a finer level of analysis. Further controls are identified in the supporting details (*e.g.* acceptance testing), and still more are implied (*e.g*. the information risks relating to those 'activities relating to outsourced system development', plus the security requirements or criteria for software testing, would typically be analysed and specified in the form of policy statements, contractual clauses, audit/review checklists *etc*.).

Information security controls are intended to be used in an organizational context comprising general controls. The information security management system specified by ISO/IEC 27001, for instance, comprises a number of governance, management, process and compliance/assurance controls designed to act in harmony, supporting and enabling the information security controls.

Dependencies are a further consideration. A few controls work independently or separately from others ('standalone'), whereas most interoperate with and to some extent rely upon other controls, forming integral parts of systems or frameworks of control. Whereas interoperable, integrated control systems are generally preferred (*e.g.* for ease of use and management), they tend to be more complex, hence simpler, single-purpose, standalone controls may be preferred in critical situations, perhaps as fallback controls providing a relatively basic but reliable backstop in case more advanced, complex controls fail.

3.3 Contextual-fit

Information security controls don't exist and operate in a vacuum, other than in academic studies or theoretical situations. The culture, context or environment into which they fit is important.

For example, industries and organizations that have evolved strong security or compliance cultures have a greater capacity to absorb and fulfil the requirements of new information security-related laws, regulations or policies than those that value innovation, creativity and self-direction. In some cases, workers may be permitted or encouraged to bend (interpret) or even break the rules if that is in the best interests of the organization and its stakeholders, whereas conversely strict compliance may be more appropriate in other situations, leaving little latitude for workers.

There are cultural factors associated with the information technologies an organisation uses.

Considering the industry, business, technology, culture, competitiveness, regulation and cultural factors may help select more appropriate controls – those deemed to be a strong contextual fit, rather than controls that are dissonant or disconcerting in some way. Conventional controls that are widely used within the organisation's industry segment, for instance, are more likely to be suggested and adopted, partly because of the implication that they have proven valuable for peer organisations. However, by the same token, unconventional/innovative and potentially valuable controls may be disregarded without due consideration.

3.4 Control systems

At face value, a mechanical door lock is a simple example of a 'standalone' control that is physically distinct from its environment. The internal design and construction of the lock affects its ability to withstand lock-picks, drills and blunt force attacks.

However, in reality, there are several other relevant factors:

- The mechanical strength of the door and frame into which it is fitted, plus the surrounding walls;
- How many keys there are;
- Who issues, holds and controls the keys;
- What prevents the keys being copied;
- Wear and tear on the lock mechanism over time;

- Whether the lock is actually used properly, which brings in yet more factors such as training, management oversight and governance arrangements (*e.g*. who, exactly, is accountable for correct use of the lock, and what does that even mean?).

Clearly, then, even such a basic physical access control as a mechanical door lock is, in reality, just a component part of a bigger, more complex control environment or system. The situation gets more complicated still in the case of keyed-alike locks or locks with master and sub-master keys, electro-mechanical locks (particularly those coupled with card or biometric authentication, perhaps inter-locked with intruder or fire alarms), networked smart locks and so forth.

'Systems thinking' broadens the perspective when selecting all forms of control, taking into consideration not just the control itself but related factors such as the control lifecycle (*e.g*. novel controls are inherently less familiar to those using, monitoring, administering, maintaining and reviewing/auditing them, than more conventional controls).

3.5 Failure modes

Controls that fail silently, giving little if any indication that they have failed, can be problematic if the organisation assumes and relies upon the controls being effective, perhaps only discovering the failure when incidents occur. This is precisely what happens in a substantial proportion of cases: information security and privacy incidents are often the first indication of serious trouble, by which time it is clearly too late to prevent them occurring. Compounding this issue are control failure indications (warning signs) that are neglected or ignored for various reasons, such as those who should be monitoring controls failing to do so, perhaps believing some other party is responsible. In some attacks, distractions may be generated deliberately to conceal the true nature of the attack and delay, mislead or evade the response.

Some controls fail suddenly and dramatically, without warning, whereas others fail gracefully or in a manner that indicates problems, potentially in time for the organisation to react and perhaps avoid or at least prepare for complete failure.

Generally speaking, controls that fail in a safe, locked or secure condition (strengthening the control) are preferred over those that fail unsafe, open or insecure (weakening the control), but sometimes conflicting requirements (such as maintaining the availability of vital information) take precedence. This consideration is clearly context-dependent. An overloaded firewall at the perimeter of a bank's network, for instance, should probably not drop the filtering rules and pass all traffic, whereas it may be appropriate for an overloaded internal network firewall to suspend security filtering until the peak load subsides, especially if various other controls can be relied upon to mitigate the most important risks meanwhile.

The probability, nature and severity or significance of control failures is clearly important in the case of 'key controls' mitigating serious or substantial information risks. Less obviously, failures of non-key/supporting controls may be problematic due to the complex, dynamic relationships between controls, risks and information.

3.6 Formality

Some of the information controls being managed through an ISO/IEC 27001 ISMS are **mandatory** in the sense of being demanded or formally imposed by third parties *e.g*. in applicable laws, regulations, contracts and agreements. The organisation has no choice over their use.

In the military/governmental sphere, mandatory controls are strictly enforced by the technologies and processes to the extent that they may be theoretically impossible to circumvent. Formal methods for designing security controls aim to address all conceivable operating conditions or situations, hence the controls can be *proven* effective, at least within a tightly-defined operating environment or situation.

Conversely, informal designs typically demand less, presuming more latitude in their implementation, although that laxity can result in additional control failure modes, some of which are unanticipated.

Most controls are **discretionary,** meaning someone has the discretion, choice or possibility to use or not use them in a certain way. Even controls that are 'patently essential' to support and enable achievement of the organisation's business objectives (*e.g*. business continuity arrangements to maintain critical information systems, services and processes for any organisation that is critically dependent on information) are discretionary since management can conceivably opt not to adopt them for business reasons.

Strictly speaking, few if any information security controls are formally demanded by ISO/IEC 27001, although some are strongly associated with the management controls in the main body of the standard (*e.g.* the formalities of maintaining ISMS documentation such as policies, procedures, records *etc*. *implies* the need for suitable access and change controls, plus awareness and compliance controls which are not specified in detail and mandated).

Formality also varies in the implementation, use and management of controls – information risk and security-related processes for example may be entirely undocumented (understood or passed-on by word of mouth), informally documented (guidance notes), formally documented (procedures), mandated (procedures supporting policies), and perhaps monitored and actively enforced.

In some circumstances, informal guidance, explanation, encouragement and support may even achieve better outcomes than the imposition of formal information security policies with strict or harsh compliance enforcement – although sometimes *both* approaches are appropriate.

3.7 Incident scenarios

ISO/IEC 27002 section A.2 notes *"if an organization has constructed its risk treatment plans [see ISO/IEC 27001:2013, 6.1.3 e)] based on events, it can wish to associate a risk scenario attribute to each control in this document. The benefit of such an attribute is to speed up the process of fulfilment of ISO/IEC 27001 requirement related to risk treatment, which is to compare the controls determined through the process of risk treatment (referred to as "necessary" controls), with those in ISO/IEC 27001:2013, Annex A (which are issued from in this document) to ensure that no necessary control has been overlooked"*.

Considering information security events, incidents or scenarios ranging from probable/highly likely to improbable/highly unlikely is one technique to identify information risks and information security controls that are (to some extent) relevant to the organisation. It is fairly common in business continuity planning as a way of 'grounding' those involved in the process, helping them envisage, discuss and plan realistically for credible situations, particularly if steps are taken to counter the natural bias towards the types of situation that are known to have previously occurred rather than unrecognised or novel future situations, regardless of the true risks.

Each scenario typically involves several information risks requiring several treatments, including a number of information security controls to mitigate unacceptable risks. Each control can therefore be associated with one or more scenarios.

3.8 Integration

Integral controls are designed and built into information systems and processes from the earliest stages, forming an integral part of the whole, whereas supplementary controls are separate, being added-on later.

Although built-in controls tend to be stronger and more reliable, supplementary controls can provide valuable additional functionality and can produce a more rounded and effective control environment. They may, for instance, reduce customers' reliance on the designers and suppliers of security technologies.

Notice that integration is not an inherent attribute of certain controls, rather it depends on the particular systems or processes under consideration, and their implementations. Whereas two IT systems might both offer, say, backup capabilities, these may be built-in to one but added-on to another, with differing implications such as costs. It may be appropriate to combine both types of control if they are complementary in some way, with implications for the control system as a whole.

3.9 Maturity

The accumulation of practical experience with mature, conventional controls that have been used for a long time – hundreds or even thousands of years in some cases – and/or widely implies that inherent weaknesses have either been reduced/eliminated or at least are reasonably well understood.

Security controls that are documented in published mainstream standards tend to be more mature, better understood and arguably more valuable than those that are undocumented, or whose documentation is relatively limited and obscure. As a class, 'privacy controls' are fairly well covered by extant standards that are widely accepted and used, whereas 'IoT security controls' (for example) are still developing.

Novel controls that the organisation and/or its people have never used before implies a lack of experience and expertise, increasing the possibility of implementation problems and costs. However, implementation problems can also occur even with controls that are already in place (*e.g.* if they are changed in some way), despite the experience and expertise on hand.

Novel controls can present novel risks, such as unanticipated failure modes and limitations. However, innovation in the design and application of security controls is often necessary when dealing with new situations and information risks, such as those arising from new technologies.

Popularity is a loosely-related attribute: workers are more likely to be familiar with controls that are popular and commonplace, than with those that are unpopular and hence rare.

If newly-designed or applied controls are particularly important in order to mitigate significant information risks, more effort should be invested to ensure they are soundly engineered (specified, designed, developed and tested), as well as being properly implemented, used, monitored and managed, to minimise the risk of control failures and maximise their value.

3.10 Measurability

The 'measurability' attribute concerns the ability to monitor and measure controls in operation, particularly concerning their effectiveness. Controls that are buried deep in technology (such as the kernel security functions of computer operating systems, hardware security modules, cryptographic processors *etc*.) are often *designed* to be opaque, reducing the attack surface and keeping their inner workings confidential to protect cryptographic keys and other secrets. In contrast, procedural controls are comparatively obvious and easy to oversee/monitor, although measuring them objectively is a different matter. The strength of physical security controls may typically be measured using physical test rigs and equipment such as strain gauges. However, such testing is generally destructive and requires specialist skills, hence is generally left to the manufacturer or some form of quality control or compliance assessment, confirming that product samples satisfy specified requirements. That in turn raises costs and further concerns *e.g*. are those specified requirements both applicable and adequate for the organisation's intended application of the control? To what extent can the product testing or certification processes be trusted to guarantee the capabilities of a control once implemented?

3.11 Multi-functionality

Some information security controls have a relatively narrow purpose or specific function, addressing particular information risks exclusively. Others, however, are more broadly applicable, perhaps addressing several of the incident scenarios of concern to the organisation.

Contrast, for instance, 'access controls' against 'oversight'.

At a deeper level of analysis, 'access control' indicates not a specific mono-functional control but a class of similar controls, all of which are intended to permit or deny access by various subjects to various objects under various circumstances. Access controls of different kinds are valuable in different situations, and sometimes are useful in support of each other (*e.g*. physical access controls to a computer system can reduce the possibility of technological access controls being disabled).

'Oversight' comprises an even larger class of controls that have application in many situations, ranging from monitoring and supervision, to direction, compliance and assurance. It is hard to think of *any* business circumstance that would be better off without oversight of some form. It is an integral and valuable part of sound governance and management.

There is a paradoxical relationship between this attribute and control strength. Generally speaking, multi-functional controls are less effective at mitigating specific risks than mono-functional controls. Nevertheless, multi-functional controls are extremely common because they have such broad application, plus a long history, meaning substantial appreciation of both their utility and their limitations. They *tend* to be cheaper given minimal marginal costs to apply them to additional situations.

Security awareness and training neatly illustrate the distinction. An effective security awareness program leads to a widespread, general understanding and appreciation of information security throughout the organisation. Training, in contrast, focuses on specific aspects for those who need additional guidance and skills. These are *complementary*, not *alternative* controls.

3.12 Origin

The origin or source of security controls can be a valuable characteristic to consider. Possible values for this attribute include:

- Standards or advisories, whether international, national, industry-specific or professional. The 'cybersecurity concepts' attribute in ISO/IEC 27002 is one such example.
- Laws and regulations which may vary between jurisdictions.
- Generally-accepted methods (such as COBIT) and good practices, begging questions if controls do not in fact appear to have been documented anywhere.
- Unique/custom controls designed in-house.
- Commercial security products, including system security controls built in to firmware and operating systems, plus most security software.
- Free/open-source security products.
- Customised controls - controls that have been modified by or on behalf of the organisation for some business purpose.
- Unknown – a category hinting at the need to search for the original source, perhaps learning important caveats or suggesting creative adaptations.

This attribute relates to the organisation's ability to determine or at least influence changes in a security control. A control that is conceived, developed and managed entirely within the organisation is potentially more readily updated to suit changing requirements or situations, without reference to or dependence on external parties such as vendors – provided the organisation has the skills and resources necessary to do so, or can readily obtain them. The complex nature of many information security controls suggests the need for competent specialists, who may or may not be employees.

The design of security controls within commercial and open-source products are driven by the requirements of a wider community of customers/users, hence an individual organisation (especially a small one) has limited influence. On the other hand, broad community engagement *can* bring many eyes to the problem area - although that in itself is no guarantee that all issues in fact are duly addressed, particularly in relation to highly technical controls.

3.13 Regulation

The extent to which a control is, or could be, regulated is an attribute.

Some controls are effectively self-regulated: flexibility in their design and operation allows them to adapt naturally to the situation at hand, often without overt intervention. For example, corporate policies and procedures are flexible in that they are interpreted and applied in real-world situations that may not have been entirely anticipated when they were drafted and approved. With appropriate governance arrangements in place, the people involved in the operation and monitoring of procedural control activities have some latitude to adapt their behaviours according to circumstances, while awareness and training, plus compliance reinforcement and enforcement, help distinguish acceptable from unacceptable activities. Such control systems typically incorporate feedback loops that dynamically respond to changing circumstances, performance, threats *etc*.

In contrast, some controls are strictly mandated and enforced, typically by an authority with strong powers. Laws and regulations are relatively unambiguous, with specific boundaries carefully

defined, leaving little if any room for interpretation in practice. This, in turn, implies a level of formality plus a raft of associated controls that increases the cost of such controls, relative to policies and procedures. Compliance monitoring in the legal and regulatory sphere can be onerous and costly to operate, aside from the obvious possibility of penalties being imposed for noncompliance. Furthermore, compliance alone may be insufficient to achieve the intended control action.

Finally, some controls are more-or-less unregulated. These may be advised, recommended or optional, and compliance may be purely a matter of choice and ethics. Only a fraction of legally-binding intellectual property rights, for instance, are actively monitored and enforced, and in some cultures they are openly flaunted, disrespected or ignored, despite the law.

3.14 Strength

'Strength' in this context may refer to a control's robustness, durability, resilience, dependability and/or other factors relating to its ability to mitigate substantial information risks.

Accurately and objectively measuring the strength of an information security control is no easy task, even for physical controls such as locks. However, based on general experience, controls can be classified roughly into categories ranging from very weak (probably unsuitable for critical applications involving significant information risks) to very strong (better suited to critical, high-risk applications but perhaps excessive for medium to low-risk ones).

A subjective strength measure might use a continuous scoring scale such as this, with typical criteria indicating several positions on the scale:

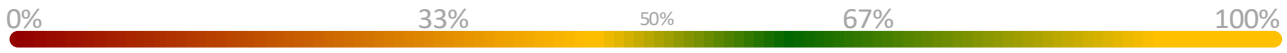

3.15 Targets

Information security controls are intended to mitigate particular aspects or elements of information risk. ISO/IEC 27005, for instance, discusses vulnerabilities, threats and impacts.

Most information security controls reduce **vulnerabilities** - security patching, for instance, prioritises the implementation of software changes that address known vulnerabilities. Concealing, reducing or closing off vulnerabilities can reduce the possibility of their being exploited.

A few controls reduce **threats** - deterrent controls mostly, plus security awareness and compliance enforcement or reinforcement.

Whereas controls that reduce threats or vulnerabilities make incidents less likely to occur, some controls target **impacts**, making incidents less damaging if they do occur. Backups, for instance, don't prevent data loss from the primary storage media, but facilitate its recovery from backup media, enabling systems to be rebuilt, data to be recovered and information services to be restored to the business.

General purpose controls often address multiple aspects. The knowledge imparted through security awareness and training, for instance, helps workers identify and react to potential or actual incidents more efficiently and effectively than if they had remained ignorant. Workers' recognition of the threat and sensible reactions (such as not clicking dubious hyperlinks) can prevent some incidents (such as phishing attacks) in the first place, otherwise appropriate responses (such as

reporting security warnings or IT system anomalies promptly to the Help Desk) can minimise the damage caused.

3.16 Transparency

Being able to view the internal design of a control (figuratively or literally) can be useful for potential users evaluating its suitability and strength, or monitoring its operation. However, in the same manner, full transparency may also enable adversaries to identify exploitable vulnerabilities, if they are exposed to those adversaries. In both cases, understanding a complex control's inner workings sufficiently well to spot vulnerabilities takes uncommon skills and expertise, assuming the control was competently designed and constructed anyway.

3.17 Value

Security controls vary widely in the associated costs throughout their lifecycle:

- Clearly the development or purchases costs differ markedly between controls;
- Some security controls are relatively difficult and costly to implement (install and put into service), typically because they are inherently complex (*e.g*. advanced technologies) or require significant changes (e.g. new processes involving multiple departments and people);
- Whereas some controls are either maintenance-free or low-maintenance (*e.g*. set-and-forget antivirus software that automatically updates itself), others require constant care and attention to keep them working effectively (*e.g*. Security Incident and Event Monitoring software). Automation can help here, provided the automation itself is effective and easily maintained;
- There are similar considerations regarding the organisation's ability to monitor and manage security controls. Complex, dynamic, multifaceted, multifunctional controls can be challenging (costly) to oversee and direct effectively, whereas simple, single-purpose, stable controls require comparatively little effort and expense.

Likewise, the business benefits vary between controls and over time, hence the net value of controls differ. When selecting/designing and deciding whether to introduce new or changed controls, one should ideally consider all aspects of the entire control lifecycle – not just the initial investment needed to procure/develop and install them.

In practice, however, both the costs and benefits can be hard to quantify in advance of implementation, and may be tough to quantify and account-for once in operation (*e.g*. what is the value in reducing information risks? If a given control is not implemented, how many and how severe might the corresponding incidents have been?).

3.18 Other controls, other attributes

ISO/IEC 27002:2022 categorises a good range of commonplace information security controls, but does not aim to be totally comprehensive. Other controls (from other sources) can be added to an organisation's control catalogue, and may be characterised or categorised using the same attributes as those already listed.

Likewise it may be appropriate to categorise controls using other attributes, features or characteristics that suit the organisation's purposes at that time. For instance, when finances are tight and security budgets threatened, more creative definitions of 'value' may help identify

worthwhile controls that might otherwise have been neglected, or poor value controls that are candidates to be modified or dropped (retired or replaced).

Neither ISO/IEC 27002 nor this paper claims to be comprehensive. Controls and attributes are many and varied, hence an organisation may well be considering or using others.

4 Using control attributes in practice

This section offers pragmatic advice on how to use various attributes, themes or views when considering, shortlisting, discussing, evaluating and selecting controls, ideally in the context of a structured and systematic approach to the management of information risks and security controls such as an ISO/IEC 27001 **I**nformation **S**ecurity **M**anagement **S**ystem.

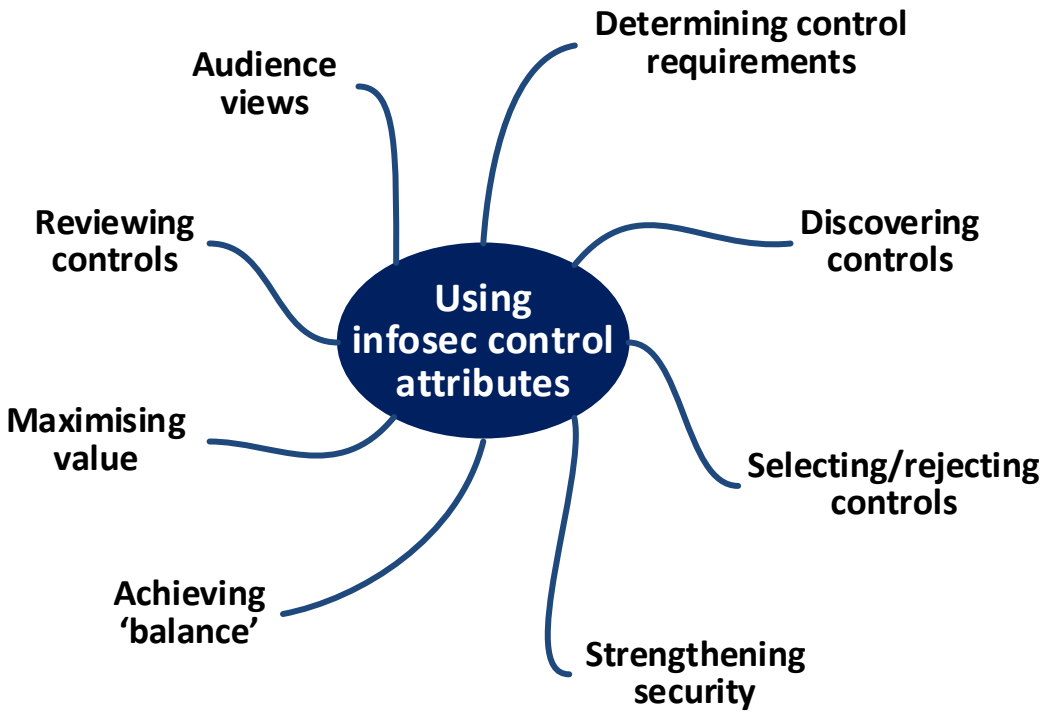

4.1 Determining information security control requirements

Determining which attributes might be relevant to a situation implies the need to understand the organisation's information risks and hence the control objectives, along with the risk appetite and the resources available to treat the risks. The associated information gathering and analysis may be as simple as someone preparing a brief outline of whichever attributes seem most relevant and useful, as a prompt to consider various aspects – a viable initial approach. Workshops, discussion papers and brainstorming through collaborative working/social media tools may be appropriate for some organisations, or at different stages in the process. An in-depth study or review would generate additional, deeper insight but at greater cost: it is for management to decide whether the benefits justify the investment, choosing the most valuable approach/es.

This should all be seen in the context of ongoing activities to monitor and respond to ever-changing information risks and maintain the organisation's information security controls, ideally through an ISO/IEC 27001 **I**nformation **S**ecurity **M**anagement **S**ystem or a similar structured framework.

'Specify control requirements' is itself a multi-functional control with value in *many* different situations. In reality, however, control requirements are rarely specified in practice, at least not explicitly. Therefore, management may wish to review the organisation's suite of information security controls to determine whether the corresponding control requirements are adequately understood and specified/captured in some form, as well as whether they are fully satisfied by the selected controls. A review of applicable control attributes may be the mechanism and stimulus to reconsider the control arrangements from different perspectives.

4.2 Discovering relevant controls potentially worth adopting

ISO/IEC 27002 is, in effect if not in name, a controls catalogue - a structured collection of information security controls. Studying the standard may indicate factors relevant to an organisation's information risks that had not been entirely appreciated, including those relating to the attributes identified in - or those omitted from - the standard.

For example, threat intelligence (control 5.7 in ISO/IEC 27002:2022) involves feeding information about threats through to the organisation's risk analyses in order to improve understanding and so lead to better-informed decision-making. The control's attributes and attribute values are identified as follows:

Control type	Information security properties	Cybersecurity concepts	Operational capabilities	Security domains
#Preventive #Detective #Corrective	#Confidentiality #Integrity #Availability	#Identify #Detect #Respond	#Threat and vulnerability management	#Defence #Resilience

The three columns on the left state that this control is relevant to all possible values of those three attributes, indicating a multi-purpose control. However, several values of the 'operational capabilities' and 'security domains' attributes are *not* listed [see the introduction section or the standard for details], begging the question "Why not?". Answering questions of that nature involves exploring and considering the control in some depth.

Considering any control's attributes or characteristics may prompt creative thought leading to innovative approaches *e.g.* searching for other controls with similar, complementary or contrasting attributes, or additional attributes that are relevant to the control requirement.

4.3 Selecting or rejecting controls

Obviously, control attributes can be used to select tagged controls addressing particular requirements. For example, if maintaining the confidentiality of sensitive information is a primary control objective arising from the risk analysis, controls tagged with the 'confidentiality' attribute (such as encryption) are most likely to be suitable.

Reviewing controls categorised by other attributes may suggest alternatives or additional (supplementary or complementary) controls, such as tamper-resistant enclosures (physical control) housing cryptographic modules (technological control), plus key management and other associated processes (people, procedural or administrative controls).

The converse also applies, in that controls may be discounted or rejected outright on the basis of having undesirable attributes or characteristics. If the budget is tight, for instance, costly controls may be rejected outright simply on the basis that they are presently unaffordable.

In a given situation, controls with similar attributes *may* be redundant and hence might be candidates for rationalisation, unless there are particular redeeming features, hinting at a previously-unrecognised or under-appreciated control requirement.

4.4 Strengthening information security (improving control)

Understanding the requirements (control objectives derived from the need to mitigate various aspects of unacceptable information risks) and the controls available leads to better, more informed decision making when initially selecting controls, and when using and managing them. For example, a control's fragility and propensity to fail-insecure *may* suggest the need to choose a more robust control, or add further controls to detect promptly and react appropriately to control failures, thereby reducing the organisation's reliance upon it.

Whereas some control attributes are inherent, fixed characteristics, it may be feasible to modify or exploit certain attributes or characteristics of a control in order to enhance its utility and value in a specific application. For example, 'formalising' controls typically involves analysing and documenting them in more detail, reviewing/testing and authorising them, thereby increasing their quality, strength, robustness *etc*. Conversely, procedural controls may be *deliberately* de-formalised or relaxed in order to improve overall effectiveness by giving workers more latitude in practice and making compliance less onerous, less costly.

You may wish to group related controls together, perhaps distinguishing 'core' from 'peripheral' or 'supporting' controls - yet another attribute, although this one relates to the manner in which the control is to be implemented and used. A security engineering approach may involve designing a control system with the appropriate combination of controls from scratch, or reviewing and redesigning an existing arrangement (*e.g*. following one or more incidents caused by control inadequacies).

It is reasonable to expect that information risk levels are reflected in the strength, robustness, assurance *etc*. of the mitigating controls. *Significant* information risks typically require mitigation using *substantial* or *key* controls. Conversely, relatively minor information risks may not require mitigation using controls (if other forms of risk treatment are more appropriate). Therefore, the strength-related control attributes should broadly align with the risk levels: obvious anomalies (*e.g*. highly robust controls addressing negligible risks) may be worth investigating, potentially leading to updates in the controls framework. This can be a powerful approach, both from the information risk and security management perspective (ensuring that significant risks are properly treated) and from the business perspective (reducing costs and freeing-up resources by not over-controlling negligible risks).

4.5 Achieving a 'balanced' suite of controls

Although achieving 'balance' is not necessarily an appropriate or rational goal *per se*, categorising and then reviewing the number of controls with various attributes may reveal *apparent* shortages in some areas and excesses elsewhere.

ISO/IEC 27002:2022 notes that the 'control types' attribute (#Preventive, #Detective, #Corrective) *"can be used as a means for the organization to check the balance of determined controls; for*

example, there can be adequate controls to detect information security events but insufficient controls to prevent information security incidents."

Similarly, a preponderance of technological controls with rather few physical and procedural controls *may* indicate a bias in the way controls have been selected, leading to weaknesses in the control framework and hence opportunities for improvement.

Analysis may even reveal distinct gaps in the organisation's information security control suite - for instance, over-reliance on mandatory controls required for compliance with imposed obligations and a lack of discretionary controls satisfying the organisation's own internal/business objectives.

Regardless of the details, seeking 'balance' again illustrates the value of considering the organisation's information security controls as a whole. Whether the analysis is prompted by new/changed information risks or compliance or business requirements, by incidents and near-misses that have occurred, or by reviewing the distribution of controls with various attributes, doesn't particularly matter: the analysis is what counts.

4.6 Maximising the value of existing and proposed controls

Whereas a given information security control may have been adopted in order to satisfy a particular requirement, its attributes may suggest other potential applications. An organisation may be using a **P**ublic **K**ey **I**nfrastructure, for instance, principally for authentication purposes (hence it might be categorised as a technological integrity control), but the same PKI may have value for encryption (a technological confidentiality control).

Some information security controls have very specific purposes perhaps mitigating a particular aspect of a specific information risk. Multi-purpose or multi-functional controls, in contrast, can potentially be applied to mitigate various aspects of numerous information risks. Re-using multi-functional controls in a wider range of applicable scenarios may squeeze more value from them where it makes business sense to do so. The organisation gains expertise and experience in the process, becoming more familiar with the control's strengths and weaknesses in practice.

4.7 Reviewing, assessing and auditing controls

Several of the suggested uses for control attributes in this paper involve comparing the organisation's existing and/or potential/proposed controls against lists of controls tagged with certain attributes. Such comparisons may be conducted or guided by specialists with knowledge in this area, potentially including external advisors and auditors with the advantage of deeper knowledge, broader experience and independence complementing insiders' knowledge of the organisation's business, information risks, resources, policies, priorities, culture *etc*.

ISO/IEC TS 27008 notes that "Review procedures can be tailored by … Selecting the review method depth and coverage attribute values necessary to meet the review expectations based upon the characteristics of the controls being reviewed and the specific determinations to be made …".

4.8 Audience views

Annex A of ISO/IEC 27002 mentions that the 'operational capabilities' attribute "c*an be used when the organization wants to classify controls from the practitioner's perspective; for example, when the organization wants to assign responsible departments within the organization based on these*

attribute values." Clause 4.2 notes that "*Attributes can be used to filter, sort or present controls in different views for different audiences*".

Various individuals, groups or teams have differing perspectives and concerns relating to information risk and security *e.g.*:

- Technologists (principally IT specialists);
- Senior management, C-suite, directors;
- Middle/junior management, supervisors and team leaders;
- Specialists in risk, security, compliance *etc*. such as security architects and testers, risk managers, business continuity managers;
- Procurement, HR, Finance and other specialists.

So, for example, procedural controls relating to the use of cloud computing services may be of most interest to IT, HR and procurement specialists. Turning that on its head, a given audience (*e.g*. senior management) is likely to be most concerned about certain types or categories of control (*e.g*. governance and compliance) and information risks (only the most significant ones). Attributes are therefore one way of identifying potentially relevant aspects or issues for various audiences and purposes, with application in security metrics and reporting.

5 Conclusion

Control attributes are a surprisingly powerful and flexible tool for information security management purposes, a novel way to select and improve an organisation's approach to mitigating unacceptable information risks, supplementing more traditional methods. It will be fascinating to see how useful they turn out to be in practice.

6 References

ISACA COBIT (originally 'Control Objectives in Information Technology').

ISO/IEC 27001:2013 "Information technology — Security techniques — Information security management systems — Requirements".

ISO/IEC 27002:2022 "Information security, cybersecurity and privacy protection — Information security controls".

ISO/IEC 27005:2018 "Information technology — Security techniques — Information security risk management".

ISO/IEC TS 27008:2019 "Information technology — Security techniques — Guidelines for the assessment of information security controls".

ISO/IEC 27028 "Information security, cybersecurity and privacy protection — Guidelines for ISO/IEC 27002 attributes" (currently at the early stages of drafting).

ISO/IEC TS 27110:2021 "Information security, cybersecurity and privacy protection — Cybersecurity framework development guidelines".

National **I**nstitute of **S**tandards and **T**echnology (2014) "Framework for Improving Critical Infrastructure Cybersecurity" version 1.0.

Appendix: examples

Attributes & values	Information security controls & applications
Assurance	
High assurance	Competent, independent, accredited certification
Low assurance	Marketing claims and collateral, self-certification or assertion
Complexity	
Simple/atomic	Validation
Complex/compound	Authentication
Dependent	Authentication depends on identification; identification depends on validation
Contextual-fit	
Business/industry	Classification in government and military organisations
Technology	Linux based systems in a UNIX-centric organisation
Corporate culture	Worker discretion and self-direction in start-ups
National culture	Compliance and regulation in US organisations
Control systems	
Core	Technology
Peripheral/supporting	Processes, documentation, training …
Failure modes	
Obvious/known	Fire alarm systems with self-checking and regular inspections
Cryptic/unknown	Fire alarm systems without self-checking and regular inspections
Safe/secure	Fireproofing, fire resistant/low smoke materials, intrinsically safe equipment, heat-activated fire sprinklers
Unsafe/insecure	Fire prevention and avoidance, generally
Formality	
Formal	Contracts, laws, regulations; policies; obligations, requirements; mandatory access control; audits, third party inspections
Informal	Agreements; guidelines, advisories, recommendations; standards; discretionary access control; reviews
Incident scenarios	
Power failure	Uninterruptible power supply; power failure alarm; regular UPS maintenance and testing
Hardware failure	Redundant/spare equipment; routine equipment inspection and maintenance; contracted hardware support

Attributes & values	Information security controls & applications
Maturity	
Conception	Post-quantum cryptography
Immature	Cloud-based crypto-systems
Mature	Device-based crypto-systems
Decrepit	Encoding, Caesar cypher
Measurability	
Difficult/costly to measure	Prevention and detection of malware infections
Easy/cheap to measure	Responding to and resolving identified malware threats
Multi-functionality	
Multi-functional	Oversight, supervision, management and peer review
Mono-functional	Completion of checklists/tick-lists, review of logs and audit trails
Origin	
Laws and regulations	GDPR and other privacy laws; company law
Policies, contracts/agreements and guidelines	Privacy policies, both as published and as actually used in practice
Proprietary/bespoke (in-house)	Supplier security questionnaire and relationship management
Commercial	Periodic accredited PCI-DSS security assessments
Public/open-source	International standards compliance
Customised/adapted	Fraud controls to detect and deter fraud
Regulation	
Independently-regulated	Standalone IT system performance and capacity
Self-regulated	Cloud system performance and capacity
Unregulated	The Internet
Strength	
Strong	Biometric authentication with all the appropriate controls around enrolment, management, validation, proof-of-life *etc*.
Intermediate	Multifactor authentication using security tokens, out-of-band communications *etc*.
Weak	Passwords, pass-phrases, PIN codes

Attributes & values	Information security controls & applications
Targets	
Vulnerabilities	Security patching
Threats	Warnings about monitoring and prosecution
Impacts	Backups
Mixed	Security awareness
Transparency	
Transparent	Published cryptographic algorithms
Translucent	Proprietary cryptographic algorithms
Opaque	Algorithms embedded in tamper-resistant crypto-modules or smart cards
Value	
Invaluable	The most relevant controls in ISO/IEC 27002 (depends on the organisation's information risks and hence security requirements)
Mediocre	Other controls in ISO/IEC 27002
Valueless	Deprecated controls such as enforced password lifetimes and weak cryptographic algorithms