2. *Information Security Governance and the Law*

Learning objectives of this chapter: Principles and practices of information security governance. Relevant policies and programmes. Laws and regulations, and policies and programmes for/affecting information security. Governance, risk management and compliance.

Information security governance is a core responsibility of the upper management of an organization (board, executive management) to ensure that the organization's information systems are well protected, by proper risk assessment and determination of strategic goals while also ensuring proper alignment of these security goals with the organization's overall corporate governance and IT needs, and accordingly committing resources to meet the goals. This includes investment in enabling tools, personnel and (business) processes to meet the identified security needs, and well defined organizational structure, roles and responsibilities with well defined tasks, as well as mechanisms to review and measure performance, and continuously carry out necessary reassessment of the security goals and means to realize any necessary course correction.

Guiding principles

Robust information security is arguably achieved as much through art as science — in that there is no perfect recipe for it, but a few basic do's and don'ts are nevertheless well understood. These help in charting out some guiding principles we elaborate next, before delving into a more comprehensive discussion on information security governance within organizations, as well as national and international laws are regulations to both facilitate organizations and hold them accountable to carry out prudent information security governance.

CONTROLS: Controls are countermeasures to prevent the exploitation of a system's vulnerabilities.

Controls can be broadly put into three categories, depending on their purpose and at what stage of a security incident they are applicable, namely — preventive, detective and corrective controls, their meanings being rather intuitive. Mechanisms to prevent security events fall under the category of preventive controls. A simple and ubiquitous example is a login/password to prevent access to a system, or an encryption technique to protect the confidentiality of data, and so on. However, if despite preventive measures, some event nevertheless does occur, then recognizing that a vulnerability has somehow been exploited, and preventive controls have been contravened is essential to contain the effects. Detective controls are mechanisms to recognize security events - during or after such events occur. Real time malware detection, intrusion detection, or identification that a network router is down, or a data disk has encountered a failure, and triggering necessary fall back actions, would all fall under the category of detective controls. Corrective controls are the fall back actions required to limit the extent of damage caused by an event. For instance, when credit card information are stolen, revoking those cards and issuing new cards to the affected customers would fall under the category of corrective control.

Controls can also be classified based on the nature of the controls, namely — physical, procedural, technical or legal and & regulatory compliance controls. An example of physical control is to determine who can or not access a specific section of an organization or a data center. Another could be to prevent people from bringing in their personal devices, or carrying out usb devices in/out of a premise. Procedural control is to precisely codify the do's and don't to realize different preventive, detective and corrective controls. This would include incident response processes, management oversight, security awareness and training, and so on. Technical controls refer to the technological solutions being applied to realize security. Use of multi-factor authentication (say, password and biometrics), malware detection software, data backup and restoration system, etc. are all examples of technological controls. Privacy laws (for instance, Personal Data Protection Act of Singapore, the still being deliberated the General Data Protection Regulation (GDPR) in European Union, etc.), as well as industry specific standards and regulations, such as Payment Card Industry Data Security Standard (PCI DSS), Banking Act or Casino Control Act in Singapore, are some representative examples of legislative and regulatory controls.

DESIGNING CONTROLS: Controls need to be designed with the

understanding that there is no perfect security. An immediate and obvious ramification is that, preventive controls by themselves is not going to suffice. The other corollary of this inconvenient truth is that ***Defense in Depth*** is required — whereby multiple layers of security controls (defense) are deployed throughout a system — introducing redundancy, so that, should one control against a specific vulnerability, or protecting a specific resource fail, there are other controls which still continue to protect the resource. Ultimately, even if the attacker succeeds in countermanding multiple controls, a defense in depth approach also helps the defender buy some time, to put in place further controls, or plan contingencies. The idea of defense in depth for information security and assurance is a doctrine borrowed from military strategy, and was conceived by the National Security Agency (NSA) of United States of America.

It is worth emphasizing at this juncture, that even as multiple controls are put in place, it is critical to ensure that the controls are easy to use, and do not significantly deteriorate the performance of the core functionalities of a system. Otherwise, there is a risk that the controls won't be properly used or configured, causing more harm than good. Likewise, even though there are multiple controls, best effort must be made that each control is as robust as possible, rather than be complacent when deploying individual controls. Furthermore, care must be taken to ensure that the multiple controls do not share same vulnerability, which can then be exploited to circumvent all the controls.

Defense in Depth: When designing and deploying a defense in depth approach to achieve information security assurance, it is vital to first identify the potential adversaries, their motivations and the kind of resources the adversaries would have, and the class of attacks they are likely to launch. Likewise, the security goals - such as confidentiality, integrity, availability, authentication, non-repudiation, need to be compiled. This diligence is required because, typically, there is a budget constraint on the amount of resource that can be allocated for the purpose of defense, and the risks associated with the failure to withhold a specific security objective helps prioritize and allocate resources. This means, an executive decision to forego some controls, and accept certain risks may have to be taken. It is also essential to plan controls across the gamut of protection **(prevention), detection and reaction (correction)**. And the depth needs to be realized through a robust and integrated set of information assurance measures and actions across the three primary elements, viz. **People, Technology and Operations** of an organization, and its information system infrastructure.

People: An organization needs to demonstrate a clear management level commitment to information security and assurance, and invest in its human resources accordingly. This includes a Chief Information Officer (CIO) or an even more specialized Chief Information Security Officer (CISO) carrying out a proper assessment of perceived threats, laying out proper processes and procedures accordingly, assigning accountable roles and responsibilities to individuals, and enabling them through proper allocation of resources and necessary training.

Technology: The organization needs to identify, acquire and keep up to date necessary technology to meet its security needs, in coherence with its policies and procedures. This would include use of third party solutions with proven credibility (for instance, market recognition or industry certification), apply recommended guidance, set up proper communication channels to communicate newly detected vulnerabilities, and apply necessary patches in a timely manner.

The defense also needs to be (i) at multiple places, and (ii) layered. Multiple places refers to defending different targets. Layered defense refers to applying multiple defense to protect any specific target. The need of the later (layered security) stems from the realization that there is no perfect security. No specific individual piece of technology or security product can provide absolute security over the life-cycle of a target, and eventually it is bound to be circumvented. Consequently, multiple, independent defenses mitigating attacks against a specific attack are desired. In the best case scenario, even if some of the defenses are defeated, the target still remains protected. In also provides more time for the defenders to respond.

Let's elaborate the ideas of defense at multiple places and layered defense with some examples. Say that an organization implements an authentication based data access through a web based interface. The ultimate target of an attacker being the data. An attacker may gain access to it either by compromising the authentication system itself, or gaining access to sniff network messages, or gain physical access to the storage device on which the data is being stored. An example of layered defence would be to apply two/multi-factor authentication, say using password with biometric and/or a hardware token based one-time password. This will ensure that even if the attacker gains access to a legitimate user's password, say, using a key logger, then it will not be adequate to defeat the authentication system. However, having layered defence for authentication will not be adequate if the attacker can just walk in to the server room, and make a full copy of the backend data base. The later situation elaborates why the controls need to be applied at multiple places, and also, in this

case, demonstrate the physical nature of the particular control, to complement the technology controls.

Operations: Having the technology in place, the procedures and policies laid out, and the necessary people to effectuate the same, an organization needs to ensure that on a day to day basis, the adopted security posture is adhered to, and adapted as and when the threat landscape evolves, and complacency does not set in. This includes responding to an actual attack in a timely and effective manner, as well as other aspects of assurance such as applying timely updates and patches, conducting audits and checking readiness with security drills (for instance, by carrying out a mock phishing attack, conducting penetration testing, etc.), regular certification and accreditation, periodic reassessment of threats, and so on and so forth.

Principle of easiest penetration & weakest link

Say data is stored in a server in an encrypted manner, in a server room with restricted physical access. Further, lets assume that it is accessed using a web based application running on another server (which also stores the decryption key) that requires multi-factor authentication, then decrypts the data, ultimately providing access to only specific subsets of the data based on the access control policies that have been set up in the system for the particular user based on the credential s/he has logged in with. However, if the actual implementation of the web application allows a SQL injection attack, or say, as in the case of the Target's 2013 data breach incident we studied in the previous chapter, the application does not prevent an user from uploading an executable file to run some malicious codes on the server, then all the other defenses in multiple places and multiple layers of defense will fall short of fulfilling the ultimate security objective. This (contrived) example is to elaborate the fact that despite defense in multiple places, and applying layered defense, security is ultimately no stronger than the weakest link.

Though the idea of weakest link was elaborated above using an example considering only technical weaknesses in the system, frequently, people end up being the weakest link. Be it being errors made in configuring a system, or failing to follow a process or respond too slowly (as was the case again with Target's 2013 customer and credit card data breach incident), or losing a laptop with confidential data, or falling victim to a phishing scam, human errors oftentimes provide the first point of entry to an attacker. Best of the technology notwithstanding, the fallibility of the human in the loop creates a sense of foreboding, which any security manager ought not to take lightly.

Need to Know

Since the weakest link in a system may expose the easiest path for penetration, but it is never clear a priori who or which system component or processes may be most vulnerable to an attack, it is desirable to limit the potential damage. One mechanism to do so is to allow an entity (be in people or processes) access to only information on the basis of a 'need to know'. If and when an entity does not need access to a specific information to carry out tasks that have been assigned to the entity, then the entity should be prevented access to said information (also known as 'least privilege access'). Need to know based information sharing is often augmented with compartmentalization when assigning tasks - where different entities have access to different subsets of information, so that as an ensemble, a given task can be accomplished, but most individual entities do not have access to the whole set of information. This paradigm is again inspired from military doctrine. A specific example of facilitating need to know information sharing in the context operating systems is to apply discretionary access control, where the owner of a file can determine whether specific people can access the said file. Mandatory access controls augments such an approach, where an explicit authorisation for access of a particular resource is determined. A layman example of discretionary access control will be where a project manager decides to share a particular contract (say with a client) document with a programmer working on the client's project, because he thinks it will be useful for the programmer to understand all the client needs. However, if the official policy of the organization is that personnel below managerial positions should in general not have access to legal documents such as contracts, then, the programmer will still need an explicit authorisation for access enforced through mandatory access control.

Implementing a need to know access paradigm may be wrought with problems of practicality — it is not always obvious what might adequate information to carry out a task, and if and when more information is sought, it may be difficult to assess and grant access in a timely manner not to disrupt functionality. In the best case situation, it may render the system inefficient, while at the worst, it may compromise functionality altogether. Thus, an organization needs to carry out proper risk and benefit analysis and determine when and up to what extent a paradigm like need to know is to be applied.

Guidelines & Standards

Security guidelines and standards are aimed to articulate best practices for achieving cybersecurity. They typically enumerate security concepts, policies, tools and technologies, safeguards, action plan, risk management strategies, training and auditing mechanisms, guidance on technology and configuration management, and so on, which, if adhered to, ensure a level of consistency.

In contrast to guidelines or principles, standards are typically actionable. Thus, a guideline may have a recommendation of the nature 'Use a good random number generator', without necessarily saying what constitutes a good random number generator or laying out a mechanism to achieve it. In contrast, a standard would typically lay out specific ways to carry out a function, or evaluate a system and so on. For instance, the standard ISO/IEC 18031:2011 specifies the characteristics of the main elements requires for a deterministic or a non-deterministic random bit generator, and also establishes the security requirements for both.

Often, standards do not have any legal mandate or enforcement, but nevertheless, market forces (e.g., public relations, or competitive disadvantage if not compliant of or certified for a specific standard, etc) may motivate compliance. There are many agencies and (inter-)national organizations which overlook the creation and maintenance of such guidelines and standards. National Institute of Standards and Technology (NIST) in the United States, International Organization for Standardization (ISO), The Internet Engineering Task Force (IETF), International Telecommunication Union (ITU), Information Technology International Library (ITIL) and Common Criteria (CC) are some organizations which have Information Security guidelines and standards within their purview. There are also other groupings, typically industry specific, aimed to fill in a void not catered by the standardization bodies, or to unify somewhat ad-hoc practices among the major industry players, who are custodian of one or just a few proprietary standards particularly relevant for a specific purpose. Payment Card Industry Data Security Standard (PCI DSS) to increase controls around card and holder data to reduce credit card fraud via exposure of the said data is one such example proprietary and industry specific standard.

Major industry players often contribute and negotiate the drafting of a standard. There are multiple motivations at play. They can push for something they may already be working on, and would have a advantage once the standard is adopted. In all cases, it allows an organization to see how the standard is evolving and better prepare for it. It also provides leadership recognition and influence.

Though there is often no legal enforcement of a standard, there are a multitude of market forces that incentives following a standard. Particularly in the context of security, it can be a public relationship disaster not to be compliant to popular recommendations and standards, particularly given that standards are often seen as a minimum common denominator (the least you ought to be doing) for security. Standardisation also helps access diverse markets.

When an organization or a product adheres to all the requirements of a standard, it is said to be compliant. If this compliance is actually vetted by a neutral and competent third party (typically by a widely accepted accreditation agency) then one is said to be certified. Ultimately, having a certified product, or being a certified organization can help enhance customer confidence, and facilitate in its marketing.

ISO/IEC standards

ISO/IEC 27000-series: International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) jointly publish a series of information security standards - also known as the **Information Security Management Systems** (ISMS) family of standards. There are multiple standards within the family, some finalised, some still under deliberation — spanning a wide gamut of aspects, from specifying security requirements, to code of practice,articulating how an implementation ought to be done, or how to carry out an audit on an ISMS, or provide certification, and so on. The 27000 series encompasses also a wide range of technical topics: some standards delve into specific aspects such as network security or storage security, while there are also several sector specific standards - for instances, related to health organization and health information, financial services, and so on. Refer to 'Figure 1 - ISMS Family of Standards Relationships' in the ISO/IEC 27000 document provided separately on NTULearn as supplementary reading material.

The individual standards are also updated regularly - for instance, the latest version of the ISMS overview standard ISE/IEC 27000 makes the first version published in 2009 version obsolete, and introduces new notions and details, for example, pertaining risk treatment which was absent in the original version. Likewise, the 2013 version of ISO/IEC 27001 standard renders the 2005 version obsolete, and adds numerous new controls, including on system security testing, assessment of and response to information security events, etc.

Recommended supplementary reading: ISO/IEC 27000 and 27001

ISO/IEC 15408 COMPUTER SECURITY CERTIFICATION STANDARD

The **Common Criteria** for Information Technology Security Evaluation (abbreviated as Common Criteria or CC) is an international standard (ISO/IEC 15408), along with the Common Methodology for Information Technology Security Evaluation (CEM) facilitates the international agreement Common Criteria Recognition Arrangement (CCRA) to ensure that security products can be evaluated by competent and independent (licensed) bodies to determine and certify whether a specific product fulfils particular security properties. Because of a well defined and accepted set of procedures to carry out the certification (by a properly licensed agency), the certificate is then recognised globally irrespective of which particular agency actually conducted the scrutiny — providing a degree of assurance regarding the product and its certification.

In particular, with respect to different kinds of functionalities, there is a well enumerated list of security properties that have been laid out in the standards documents. For instance, the class FCO: Communication (elaborated in ISO/IEC 15408-2:2008(E)) is aimed at assuring the identity of a party participating in data exchange, be it as the originator of transmitted information (proof of origin assuring non-repudiation of origin) or the recipient of information (non-repudiation of receipt). Likewise, Privacy is dealt with in another class FPR: Privacy, which specify requirements for protecting a user against discovery and misuse of identity by other users. And so on. The ISO/IEC Information Technology Task Force (ITTF) web site `http://standards.iso.org/ittf/PubliclyAvailableStandards/index.html` provides the complete documentation, and readers are referred to the same for an exhaustive list of the classes of security properties enumerated in ISO/IEC 15408.

Specifically, an end user can use the Common Criteria framework to specify their security functional (SFR) and assurance (SAR) requirements using protection profiles (PPs). Vendors can then implement or propose an existing product (or module) meeting the said protection profile. This product is then to be evaluated — Target of Evaluation (TOE). The TOE typically includes the actual product as well as associated documentation and administrator guides. The evaluation is then carried out in accordance. Depending on the rigour of evaluation, different Evaluation Assurance Levels (EAL: EAL1 through EAL7)[1] can be assigned. EAL1 (Functionality Tested) is the minimal requirement, where only the correctness of operations (say, claimed in the associated documentation) are evaluated, without further scrutiny of threats to security. EAL4 (Methodically designed, tested and reviewed) is typically expected in many products such as popular operating systems since it sits at a

[1] `https://en.wikipedia.org/wiki/Evaluation_Assurance_Level`

sweet spot of feasibilty and cost-effectiveness on one hand, while guaranteeing that a principled design using security best practices has been applied, and so on. A portal containing a list of all certified products, along with the assurance level is maintained at https://www.commoncriteriaportal.org/products/.

In principle, Common Criteria is a good endeavour, particularly in facilitating a certification process with a degree of consistency that can be compared and thus recognised globally, and thus acting as a guide for developers, evaluators as well as for organizations procuring IT/security products.

However the Common Criteria mechanism has also been criticised for several drawbacks: The evaluation primarily is centered around the documentation, rather than evaluating the actual technical correctness or security of a product, and furthermore, the process is time, effort and cost intensive – both in terms of preparing (for the vendor) and then vetting (for the certifying laboratory) documentary evidences of evaluation, requiring an upfront investment with no tangible guarantee of creating a more secure product (though the testing process may actually help identify some issues overlooked during the normal implementation and debugging process cycle). The certification is also meaningful only in the context of the security properties that were evaluated. Thus, the certification and corresponding assurance level has no meaning in isolation from the particular protection profile (PP) for which the evaluation was carried out, and end-users need to be careful in interpreting the same.

Alternatives to the Common Criteria include free and open source software following agile software development paradigm, where the openness of the source code allows thorough scrutiny, though of-course, such alternative paradigms are also full of their own shortcomings.

Laws & Regulations

Recommendations and standards typically do not have any explicit legal mandate or mechanisms for enforcement. In contrast, Laws and regulations come with mandate from the legislative arm of a government, or a regulatory board or agency, creating legal obligations that may be enforced by relevant enforcement agencies. In the context of information & cyber security, there are many national level (federal and state level) laws as well as some international laws and agreements, that are directly related, for instance, Computer Misuse and Cybersecurity Act (2007) of Singapore, while there are other laws and regulations where the core focus is on some other issues, but nevertheless they may have strong information security implications,

for example, Personal Data Protection Act (2012), Banking Act (2008) or Casino Control Act of Singapore, which has direct implications to issues such as data confidentiality (e.g., financial details of customers), integrity (e.g., for audit purposes), availability (e.g., for audit purposes or for fraud detection, and so on). Some relevant US laws include Federal Information Security Management Act (2002) outlining the necessary security for IT infrastructure used by any federal government body, while acts like Sarbanes-Oxley Act of 2002 is for corporate accountability to check financial and accounting irregularity, or Health Insurance Portability and Accountability Act (HIPAA) 1996 to protest patient health information acts serve specific (other) purposes that however rely heavily on proper information security management. Most countries have or are in the process of legislating similar, as well as other laws and regulations which often have strong information & cyber security implications, simply because of the ubiquity of information systems in all walks of modern life.

Many cybercrimes are however borderless in nature, and hard to address because of a multitude of issues, including heterogeneity of national laws, some of which are occasionally conflicting in nature, lack of a global law enforcement agency with jurisdiction over all sovereign entities, as well also because many criminal (organizations) are stateless, operating from dispersed geographic locations, and also because many criminal acts are in fact orchestrated by state actors, or state sponsored actors, albeit in disguise, to be conveniently denied.

Some examples where states or state sponsored entities have been alleged (varied degree of evidence, but unsurprisingly, without any convictions at any court of law) to have carried out cybercrime (Cyber attacks/wars, Cyber espionage) include the massive denial of service attack on Estonia's IT infrastructure in 2007, the 2010 Stuxnet worm incident, NSA's PRISM surveillance programme (revealed by Edward Snowden in 2013), to name a few prominent ones.

The conflicting interests notwithstanding, there are efforts to streamline international efforts to stem cybercrime. The Budapest convention on cybercrime (2001) is an international treaty aimed to address internet and computer crimes by harmonizing national laws, improving and coordinating investigations, and facilitate cooperation among nations. Even where there are intent and in principle agreements on part of countries, it is not always feasible to harmonize all relevant laws, particularly when they are in contradiction to other laws of a country. A well documented instance is the conflict posed by US's first amendment for free speech, based on which a ban on virtual child pornography had already been struck down from the Child Pornography Prevention Act of 1996 in US by the U.S. Supreme Court's ruling (2002) in Ashcroft v. Free Speech Coalition [2],

[2] https://en.wikipedia.org/wiki/Ashcroft_v._Free_Speech_Coalition

even though the international convention seeks to ban virtual child pornography as well. In addition to global level efforts, there are also numerous multilateral and regional efforts, most significantly in the European Union, but also across other entities - which have direct or indirect information security implications.

When it comes to inter/national laws, we also need to understand that even legitimate businesses operating across multiple countries may find themselves in legal quandary, even as they try to satisfy conflicting requirements. For instance, in July 2014, Microsoft was ordered in an US court to hand over emails of EU residents, and stored in a data center in Ireland, operated by a subsidiary, which, if Microsoft abides by, will put them squarely in violation with EU privacy laws.

We shall next use some national level laws, typically from Singapore and the US, mainly for the purpose of elaboration of key concepts and ideas. Many other countries may have similar laws. It is however worth putting in a disclaimer at this juncture, that the treatment of the topics here is at a very high level, and may carry some inaccuracies, and for any legal purposes the readers are referred to the original legislative documentation and their precise reading in respective jurisdictions.

Some representative laws and regulations from Singapore

The Computer Misuse and Cybersecurity Act, originally enacted in 1993 and revised subsequently in 2003, as the name would suggest, has direct bearings to Information and Cyber Security issues. For instance, it has provisions regarding unauthorized access - which foremost has (among the CIA triad security goals) confidentiality implications. Another section (specifically, Section 5) on unauthorized modification of computer material likewise has integrity ramifications, and so on and so forth. Laws generally not only enumerate precisely what are various illegal activities, but also stipulate the penalties (which may be viewed as corrective measures) if an illegal activity is carried out, along with other practicalities as jurisdiction, etc.

Some other laws, such as Personal Data Protection Act (PDPA) of 2012 and Evidence Act (1997) on the other hand are aimed for somewhat different purposes, but with very immediate and obvious concomitant of Information security. PDPA stipulates whether and what information about a person can be collected and retained by any entity, and how such personal information may be used, and restricts transfer of said data outside of Singapore. Many other countries also have, or are discussing similar — though some stricter, and some laxer, privacy protection laws. Evidence Act requires

It is worth mentioning that EU's privacy laws are particularly robust. An interesting and somewhat controversial provision is the 'right to be forgotten' which requires that data collectors remove data that is 'inadequate, irrelevant, or no longer relevant'. While journalistic and news websites would thus be exempt from the requirement, a search Engine such as Google is obliged to systematically eliminate such information from their search results. A major precedent was set by a 2014 decision by the Court of Justice of the European Union (CJEU) in the Google Spain v AEPD and Mario Costeja González case, wherefrom Google (and search engines in general) are obliged to consider requests from individuals to remove links to even freely accessible web pages resulting from a search on their name. There are several controversies around such a provision, starting with its conflicts with freedom of speech, but also its practicality. Ironically, after having removed the original documents from its query results, Google was brought to task again, when in August 2015, UK's data watchdog ruled that Google must now delist a new set of links referencing articles about right to be forgotten link removals and carrying the original name and information about Mario Costeja González which needs to be forgotten!

preservation and integrity of information, prohibiting information tampering and requires adequate measures to not lose data relevant as evidence for any legal proceeding.

Other laws and regulations, often specific to certain industry, may at a first appearance not seem to have anything directly to do with Information security, but ultimately the implications are equally strong. For instance, the Banking Act (revised in 2008) provides for the licensing and regulation of the business of banking and similar financial institutions. Within, there are requirements which necessarily rely on robust Information security management. For example, sections 25 and 43 on publication and exhibition of audited balance-sheet, and inspection of banks respectively can be facilitated only when integrity of information is guaranteed. Banking secrecy is aimed at providing customers privacy through confidentiality of their financial records.

There are also many industry specific regulatory agencies which both facilitate as well as hold accountable organizations' information security related activities (among other aspects). For instance, the Casino Regulatory Authority (CRA) of Singapore derives its mandate through the Casino Control Act, and monitors various activities in Casinos. Among these is included the obligation on a Casino's part not to carry out any activity in absence of proper video surveillance, where the surveillance and data retention plans need to be vetted and approved by CRA, and any deviations from the approved plan are likewise monitored and accounted for by the regulatory agency.

Likewise, the Monetary Authority of Singapore (MAS) overlooks activities of financial organizations. In order to ensure that information security is managed properly, MAS has prepared a set of Technology Risk Management Guidelines[3], which again are centered majorly around Information security, and is rather comprehensive and holistic. It encompasses many diverse aspects, including acquisition of information systems and their source code review or certification requirements, issues pertaining system reliability and data backup, power backup, access control (physical control, segregation of tasks, etc), IT audit requirements, and so forth. To elaborate further on the nature of these MAS technology risk management guidelines, let's use a few specific examples. For data backup reliability, resiliency and recoverability, the MAS guidelines suggest 'processes should be in place to review the architecture and connectivity of sub disk storage systems for single points of failure and fragility in functional design and specifications' and 'carry out periodic testing and validation of the recovery capability of backup media and assess if the backup media is adequate and sufficiently effective'. Likewise, encryption is recommended for offsite storage. In a similar vein,

[3] See downloadable resources under http://www.mas.gov.sg/regulations-and-financial-stability/regulatory-and-supervisory-framework/risk-management/technology-risk.aspx

the guidelines discuss IT outsourcing risks and things to consider if cloud computing service is engaged. Similarly, for internal security, numerous defense in depth best practices are catalogued explicitly. These include — 'never alone principle', stipulating that more than one person be required when conducting certain sensitive or critical activities, segregation of duties and job rotations to mitigate weakest links or perpetuity of single points of failures or fraudulent activities, and application of principle of least privilege (need to know) based on job responsibility and necessity.

We conclude by reiterating that though the above laws and regulations have been derived from Singapore, they reflect upon what may be deemed necessary or desirable in terms of their consequent Information security implications in general, and have been discussed for the purpose of instruction and elaboration. Nevertheless, it must also be remembered that often the actual laws of a country also depend on the values that the corresponding societies put to various aspects of life. For instance, there could be conflicts between two ideals — privacy rights as right to be forgotten versus freedom of speech and right to know and transparency.

Some example American laws

Given its economic dominance, and its leadership in the IT industry, laws in United States with information security implications often have impact at a global scale, and we will thus discuss some key such laws.

Following several financial and accounting scandals around 2001 (for example, the ones involving Enron, Tyco International, Adelphia, Peregrine Systems, WorldCom) highlighting audit failures, in 2002, the federal law Sarbanes-Oxley Act (SOX) was enacted, creating a set of new or expanded requirements for all U.S. public company boards, management and public accounting firms. Though the essence of the law is to realize better transparency, accuracy and accountability of financial activities to protect investors, it has several information system and security management repercussions. The Section 302 of SOX states that the Chief Executive Officer (CEO) and the Chief Financial Officer (CFO) must personally certify that financial reports are accurate and complete. They must also assess and report on the effectiveness of internal controls around financial reporting. This section clearly places responsibility for accurate financial reporting on the highest level of corporate management. CEOs and CFOs now face the potential for criminal fraud liability. For them to carry out their tasks effectively, CEOs and CFOs will require enabling information systems which maintain proper log of all activities, guaranteeing integrity and provenance of the information. The Section 404 of SOX

states that a corporation must assess the effectiveness of its internal controls and report this assessment annually to the Securities and Exchange Commission (SEC). The assessment must also be reviewed and judged by an outside auditing firm. This again requires proper monitoring and logging of all information access and manipulation within the system, and ensuring that the recorded information itself is stored reliably withholding integrity and availability. In order to comply with the regulatory requirements, organizations need to deploy an enabling framework, for instance, Control Objectives for Information and Related Technology (COBIT) created by Information Systems Audit and Control Association (ISACA). The latest version of this framework is COMIT 5 from 2012. The framework has many aspects to it. These include mechanisms to link business goals to IT goals, associating responsibilities of business and IT process owners, as well as operationalizing the processes by dividing IT into four aspects Plan and Organize, Acquire and Implement, Deliver and Support, and Monitor and Evaluate.

Another far reaching US act, also from 2002, is the Federal Information System Management Act (FISMA) aimed at strengthening information security at US Federal Government agencies and organizations. In itself, the law essentially states a high level wish list, mandating NIST (National Institute of Standards and Technology) in turn to translate the high level wish list to practice through development of standards for mandatory information security risk management. Specific follow-up and ongoing activities on FISMA are made available by NIST http://csrc.nist.gov/groups/SMA/fisma/index.html.

Another, sector specific example is the US federal Health Insurance Portability and Accountability Act (HIPAA) of 1996. The provisions for patient privacy and confidentiality within require proper processes and health information system infrastructure, as well as adequate training of personnel to ensure that there are no inadvertent or unintended violations of patient privacy. As an example[4], lets say Lisa and Carol work in the same medical insurance company but in different departments. While Lisa is a part of the Accounts department, Carol works in the Administrative wing. One day, when Lisa was away from her desk to get some water, Carol came looking for her to go out for lunch. Carol happened to notice medical billing records of her neighbor, Anne on the computer screen. She was shocked to see that Anne was HIV positive and had underwent an abortion.

Did something wrong happen here? Though Lisa is in the Accounts section of the insurance company and thus may, as per her job requirements, is entitled to know that Anne is a customer of the company for accounting purposes, she may not have the need and

[4] Acknowledgement: The following example situation is adapted from http://blog.rsystems.com/aspects-of-security-in-healthcare-usa-2/

thus right to know the actual medical condition(s) Anne has. This is thus an example of breach of Anne's privacy as a patient.

How can such breach be mitigated? The breach in this event did not happen due to deliberate malice on either Lisa or Carol's part. Nevertheless, as per HIPAA requirements, PHI (Patient Health information) ought to be protected from even accidental disclosures. The lapse here has been at the company's process and/or operations aspects. For instance, Lisa should have used a screen saver, locked her computer or closed the application before leaving her desk. Likewise, Carol entered the accounts department. Likewise, if a policy of restricted physical access was in place, then the PHI disclosure would not have occurred.

This example also goes on to demonstrate that even a small slippage, whether unknowingly or by mistake could risk the confidentiality and integrity of the organization and thus lead to major implications.

The intent of the discussion here is not so much as getting into specifics of what is required and how it is achieved in the IT space to comply with SOX, FISMA or HIPAA (or other laws which have implications on an organization's information security governance) but to emphasize that all acquired third party as well as built in-house IT tools and systems used in relevant sectors are immediately affected by the corresponding acts, and be it the executives of the organizations that need to assess their IT infrastructure and acquire the necessary tools, or be the IT product vendors catering to these entities, a lot of information system design, implementation, operations and monitoring has/had to undergo consequent changes and adaptations, and the effect sometimes transcend the specific industry for which the original law was intended. For instance, a private sector organization, or an entity even outside US may still also benefit by using FISMA compliant products.

Governance, Risk Management & Compliance (GRC)

Governance: The management approach used by senior executives to direct and control an organization.

Risk management: Process to identify, analyze and when necessary act, to mitigate risks that otherwise would affect the business objectives.

Compliance: Conforming with stated requirements - set out internally, or through extrinsic requirements like laws, regulations, contracts, etc.

So far, we have laid out guiding principles to approach information security, and how different instruments such as guidances, recommendations and standards and legal checks and balances like laws and regulations provide a roadmap for proper governance. But ultimately, it is the upper management of an organization which has to effectuate information security governance as an integral part of the overall governance of the organization. This is to be done by aligning the information security objectives among others, with the organization's business goals and obligations - legal, e.g., laws, regulations, contracts, etc. as well as self-regulation (internal policies) e.g.,

decision to be compliant to certain internal processes, best practices or industry standards and recommendations, etc. In the process, the organization will need to identify risks, and identify which if these needs to be mitigated, and determine the controls accordingly, as well as choose to accept certain risks. This will ultimately determine the organization's overall security plan (security posture) on how to carry out risk management. Finally, having adopted a security posture, and laid out the mechanisms and resources to achieve the planned objectives, the organization will need to continuously assess whether they are compliant with the aforementioned goals and obligations, determine the efficacy (measure performance) of the current practices and tools, and use this assessment to continuously refine the governance. Such continuous monitoring and refinement may also be necessitated by a changing realities - for instance, emergence of new technology (say, proliferation of cloud computing, which was not pervasive in 2005) or legislation of a new law (say, PDPA, which was not in effect in 2005). GRC is then an umbrella term encompassing these three aspects of Governance, Risk Management and Compliance.

Over the life-cycle of an organization, GRC may be seen as a continuous sets of cycles of Identification (what are the threats?), Assessment (how serious are these threats and which of them should be addressed), Enforcement (putting in place and applying the controls to actualize the decisions based on prior steps of Identification and Assessment) and finally derive Feedbacks to revisit the Identification and follow-up activities, repeating the cycle. This cycle of four steps can further be seen as a superimposition of the three logically different activities of Governance, Risk Management and Compliance occurring simultaneously. Thus, in the Governance phase, the Identification step is based on business needs and extrinsic factors such as environmental threats, laws and regulations. The compliance phase applies audit tools, adhering to and aided by standards and laws/regulations, determining both whether the due diligence in the governance phase was adequate (in framing its policies and processes, applying necessary tools, and operationalising everything), and also ascertain if there is any shortcomings in the actual realization - be it intentional, accidental, out of complacence or any other reason. Risk management phase accounts for the dynamic nature of the IT landscape in general, and also that of consequent security implications. As business needs change (e.g., a new product being launched, new kind of customer data being stored in the system, and so on), or extrinsic factors change (e.g., a new law has been legislated), a security incident has already occurred, a cheaper IT solution has become available, etc., there is need to continuously manage risks. Naturally, there is feedback based interactions among these three

logical phases.

Compliance

Depending on its domain of business, an organization is likely subject to industry specific regulations. For instance, health sector businesses, for instance hospitals, medical insurance companies, etc. are legally obliged to adhere to Health Insurance Portability and Accountability Act (HIPAA), Publicly traded companies are required to follow the provisions of Sarbanes Oxley Act (SOX), and so on. There are various (audit) frameworks laying out the specifics of how to ascertain if an organization is compliant with the required legal requirements. A popular audit framework is COBIT (Control Objectives for Information and Related Technology), which is suitable for a multitude of industries including publicly traded companies (regulated under SOX), Banks and other organizations active in the financial sector (regulated under Gramm-Leach-Biley Act, etc), credit card merchants (governed under PCI) and so on.

The focus of the res of our discussions will be centred not so much on the specific details of an audit framework, but an overview of the major activities typically carried out in multiple stages under an IT security auditing exercise.

Audit planning and preparation: The foremost task is to carry out planning and preparatory tasks. This involves charting out known and typical possible areas of concerns, review organisational chart and job descriptions, review the companies IT policies, procedures and planning documentation, attack response and disaster recovery plans, inventory all equipments, software applications and operating systems, and so on.

One of the cardinal outcome of the audit would eventually be whether the actual operations of the organization are streamlined with respect to the documented intent of its security posture.

Establish audit objectives: Next step is to establish well defined objectives for the audit. It may be limited in scope and accordingly selective (say, for due diligence), or be comprehensive (for instance, to certify compliance with respect to a standard, or to determine adherence to government laws). Limited scope objectives include determining the need for (re)training personnel, determining if the processes are aligned with industry specific as well as general regulations, if the IT infrastructure and assets are properly maintained and patched with latest security updates, data is backed up regularly, and so on. A comprehensive audit will naturally encompass all these example aspects, and many more - typically dictated by the regulatory requirements.

Perform the review: The actual review processes for the audit may rely on various means. It could be manual assessments, including interviewing staff, performing security vulnerability scans and penetration testing, analyzing physical access to systems, determining if the latest updates have been installed in all equipments, etc. Use of computer-aided audit tools (CAAT) for automated monitoring, applying statistical and data mining tools for anomaly detection, etc. are also on the rise.

The final outcome of a review needs to detail the methodology and steps taken to carry out the review, describe the findings on the performance of the organization, as well as make recommendations advocating corrective follow-up actions.

We conclude with a commentary on measuring security and its performance. Though security is intrinsically measureless, there are numerous attempts to quantify security performance, typically by measuring different manifestations of security (lapses). For instance, though it can still be be tedious, one can determine the number of infected (by known malware) machines, or number of machines without patches for known vulnerabilities. One can also record known occurrences of security incidents. Likewise, an organization can qualitatively measure capability, particularly in a relative manner with respect to some known benchmarks (say, relative to the industry best practices). One can also align security measurement with risk management by determining the costs/benefits and return of investment to determine whether to establish specific security capabilities, or to assign a value to the said security.