Information Asset Value for Security

The information-centric approach to security architecture specification is the way forward

Mike Barwise

Although the term 'information security' is habitually used in discussion, businesses have, for the most part, traditionally implemented IT security, a techno-centric subset of the whole gamut of information security. Security has traditionally been deployed within budgets based on a proportion of overall IT spend, and has primarily focused on the protection of physical infrastructure and data as technical assets, without significant recognition of their individual intrinsic values to the business.

While physical infrastructure and data are concrete objects which are easy to identify and to surround with conventional perimeter defences, it is difficult to directly establish their true business value. Where information value has been included in security decision-making, it has mostly been categorised loosely ('low', 'medium', 'high') using subjective assessments based on broad classifications derived from the way it is managed from a technical perspective (e.g. 'the customer database'). Only infrequently has the interaction between specific information assets and business processes been considered. As a result, both budgets and solutions in traditional techno-centric security have tended to be driven primarily by the security product marketplace. Often both have corresponded quite loosely with the value to the business of the protected business information. It has therefore been difficult to defend the distribution and prioritisation of the security spend, and to ensure the effectiveness of the resulting implementations.

Such techno-centric solutions have also tended to emphasise reactive countermeasures to external threats (anti-virus, spam filters, intrusion detection), rather than security, in terms of proactive measures for robustness. While such countermeasures will always be necessary minima, they may well prove insufficient on their own in the face of the growing penetration of information processing into every facet of business, coupled with the ever increasing sophistication of security threats we face today and in the future.

Robustness against the unexpected is an increasing requirement as technical threats evolve ever faster, so proactive security implementations that emphasise the protection of assets with the highest business value are becoming a must. Security strategy has to become information-centric. It must evolve from being grounded in an IT perspective to being driven at top level by business information value.

We need much more finely-tuned security architectures that mirror the connections and segregations between business processes in order to provide demonstrable levels of protection and also to minimise collateral damage when incidents do occur (as they always will, despite our best efforts).

The key questions for those budgeting for and implementing future information security must be, how much to spend on what kind of security, and where? This requires two changes to traditional thinking: first, acceptance of the need for a multi-stage, rigorous valuation process to replace familiar single-step intuitive (and necessarily subjective) judgements; and second, formalisation of decision-making methods to ensure reliability. The overriding concern, once we start developing and implementing these necessarily more complex and granular security architectures, is the precision and consistency with which the business requirement is translated into the technical solution. If assumptions are invalid at any stage, potential exposures may fail to be addressed, and if communication is ambiguous, excellent solutions may be provided to the wrong problems. A framework of basic principles is needed to maintain the correctness and integrity of intent throughout the security provisioning process from specification of business requirements through to delivery of technical solutions.

First, the information content of data has value depending on the business context in which it is used. This underlies the concept of an 'information asset'. An information asset is a piece of data with a specific function in a specific business context, and its information value to that process is the cost to the business should the process be affected by a detriment to the information. The data from which the information is derived will have an aggregate value to the business, which is the sum of the information values to all the business processes that use it as information.

This value can be surprisingly high (frequently considerably in excess of the intuitive values that emerge from traditional security management brainstorming sessions). Second, business information threat definitions must clearly correspond with business, rather than purely technical, hazards. They must be stable and have a very long lifetime, making use of generic terminology that does not depend on the details of current technical threats. Third, all threat definitions must have unique unambiguous meanings to both the business and technical echelons to ensure that business security requirements are reliably translated into technical security provisions. They must also, as

far as possible, be mutually orthogonal: their meanings should not overlap. Fourth, reliable, consistent, and auditable valuation methods must be implemented to ensure that the valuation process and its results are repeatable and robust in the face of limited or varying evidence.

The first principle most often causes concern to businesses considering information-centric security for the first time. It implies that there must be a complete and accurate record of all business processes, the information they make use of and the way in which it is processed. A surprising number of even quite large enterprises still do not maintain information asset inventories, although their value to the business in numerous spheres of operation seems incontestable. Whilst security standards advocate concentrating on the 'most significant' information assets, it is unclear how one can safely determine relative significance without reviewing the whole portfolio. So this is indeed critical metadata, which must include descriptions of the business processes, the information used, the nature of its processing, storage, and transmission, plus the derived data: threat exposures and asset values.

Threat definitions conversant with the second and third principles are not difficult to develop. We all accept the three traditional information attributes (confidentiality, integrity, availability) as starting points for security discussion. Each has a stable high-level meaning, and each can be explained consistently to the business echelon in terms of an expectation and to the technical echelon in terms of an implementation to fulfil that expectation. In contrast, 'malware' is a poor threat definition for this purpose. It has no consistent durable meaning either to the business echelon in terms of its impact or to the technical echelon in terms of protection, as the malware threat changes in nature over time. Furthermore, the variability of its business impact makes the term imprecise as a communication tool, dependent on the currency of the knowledge of each party to the consultation. So the traditional three information attributes are a good starting point, but in the modern business context they are probably not entirely sufficient for the complex ways we use information. A typical lexicon might therefore be:

Confidentiality: only the right people accessing only the right information;
Integrity: freedom from corruption and unauthorised alteration of the information;
Availability: assurance of authorised access to the information when required;
Authenticity: assurance of the authorised origin and validity of the information;
Retention: fulfilment of both business and statutory archival requirements;
Ownership: control over the continued authorised possession of the information;
Regulation: fulfilment of purely regulatory obligations for information management such as UK Data Protection or statutory audit trail maintenance.

An information-centric threat is an event that jeopardises one of these information attributes. It is crucial to recognise that a given information asset may have quite different values in respect of different attributes: the confidentiality of information may be critical regardless of its completeness and accuracy, or its immediate availability may be paramount for continued operation of some key business process. In the context of medical primary care for example, these two scenarios often seem to conflict, clearly demonstrating the need for fine granularity. To achieve this, the valuation process for information-centric security needs to be multi-stage. Although information asset value depends both on the impact on the business process of threats to the information, and on the value of the business process to the enterprise as a whole, these are necessarily separate enquiries drawing on separate evidential sources.

Once information asset values have been established, the identified assets can be mapped back to the business and technical infrastructure to determine concentrations of value for the different information attributes. On the basis of this mapping, informed decisions can be made concerning the appropriate expenditure on, and nature of, security measures at different points on the infrastructure. Of course, not all security measures will prove to be technological. Many exposures may be best managed by policies, particularly where manual information processing remains necessary. Interestingly, this point tends to be missed in traditional techno-centric security architecture definition, where policies frequently emerge as afterthoughts to mop up the problems that remain once technological solutions have been taken as far as they can go.

The outcome of the information-centric approach is a finely-tuned, layered security architecture that maps closely to the needs of business processes and applies the most appropriate proactive controls to the protection of information assets on the basis of their most important attributes to the business. Such controls are relatively long-lived and stable, and by virtue of their fine granularity, tend to contain any incident to its point of origin. On the other hand, the traditional techno-centric security model, aiming as it does at reactive protection against a current but changeable pattern of attack, needs constant intervention to stay ahead.

Furthermore, the general and non-granular nature of the controls tends to militate against containment, permitting the effects of breaches to spread within the enterprise. On the basis of such considerations, as information processing becomes ever more diverse and integral to business, there is a clear indication that the information-centric approach to security architecture specification is the way forward.

Mike Barwise
An independent consultant specialising in information security strategy
m.barwise@bcs.org.uk