Improving Assurance of Information Security RoI

Michael D. Barwise[1]

[1]Integrated InfoSec
6 Maple Green, Hemel Hempstead, Hertfordshire HP1 3PY, UK
mbarwise@bcs.org.uk

Abstract

Changing business expectations of information infrastructures have imposed new demands on security architectures. Established technocentric perimeter-oriented security architectures are yielding ground to business-driven deperimeterised architectures that assume extensive information and resource sharing and global virtualisation. These changes provide the opportunity to take a new approach to security architecture specification, based at its highest level not on costing of reactive countermeasures to current technical threats, but on prioritising the allocation of resources to robustly and proactively protect business information assets against business-oriented exposures. This permits tighter specification of both requirements and budgeting with a concomitant improvement in RoI, but depends on a new approach to management described here.

1 Changing Security Architectures

Until quite recently, information security budgeting was a relatively simple affair. Perimeter-oriented security technologies and personnel awareness campaigns were the essence of the solution for the majority of businesses. The components, and thus the component costs, of such a security implementation were generally well-defined and their selection was driven almost exclusively by techno-centric architectural decisions. However, the resulting architectures (essentially an assemblage of secured cells each consisting of a hard shell surrounding a soft centre, coupled by hardened data corridors to similar cells) have never been optimum. They generally bear no close relationship to business structure, essentially echoing instead the geographical distribution of business premises, data centres and facilities, they suffer from single points of failure, and their greatest weakness has always been at the user interface endpoints.

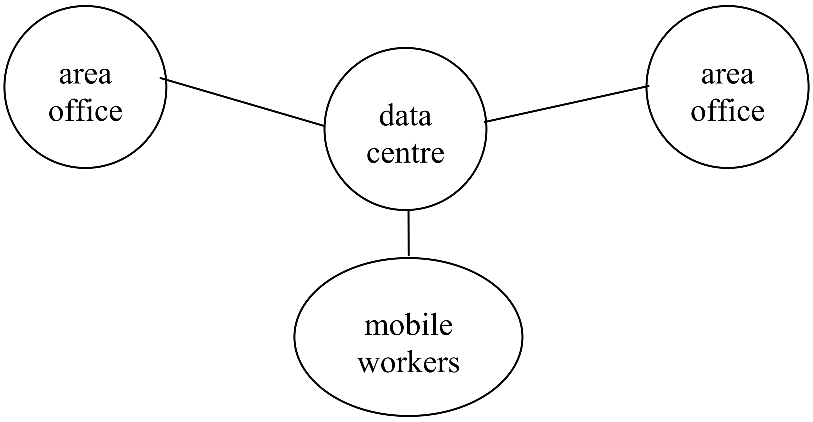

Figure 1: Traditional Perimeter-oriented Architecture

Such architectures also tend to reinforce a reactive approach to security, with the emphasis on countering technical threats from outside the perimeter as they emerge. The reactive stance makes the relationship between implemented security technologies and levels of protection achieved hard to demonstrate, and there is seldom any strong correlation between the value of the protected assets and the cost of protecting them.. Uncertainty both in predicting the nature of emerging threats and in ensuring that reactions are appropriate and proportional both contribute to this. Typical of the problems faced are "zero-day exploits" (software vulnerabilities that are exploited by an attacker immediately they are publicised, prior to any remedy being developed). Such threats are extremely difficult to predict, and in the absence of proactive management for robustness can cause losses that are not only potentially large but also very difficult to quantify precisely. Clearly, although the budgeting process is relatively simple to execute for traditional perimeter-oriented architectures, assurance of security and cost-effectiveness are both low, so poor RoI is to be expected.

1.1 New Demands

The expansion of technological facilities driving twenty-first century business has created new departures in infrastructure, not only in terms of scale but also in nature. The increasing ubiquity of web services, virtual networks and grids, distributed and shared data, mobile working and a host of other business demands that tend to destructure corporate network boundaries result in the requirement for tight integration of security solutions with business processes, rather than primarily with network physical infrastructures as in the past.

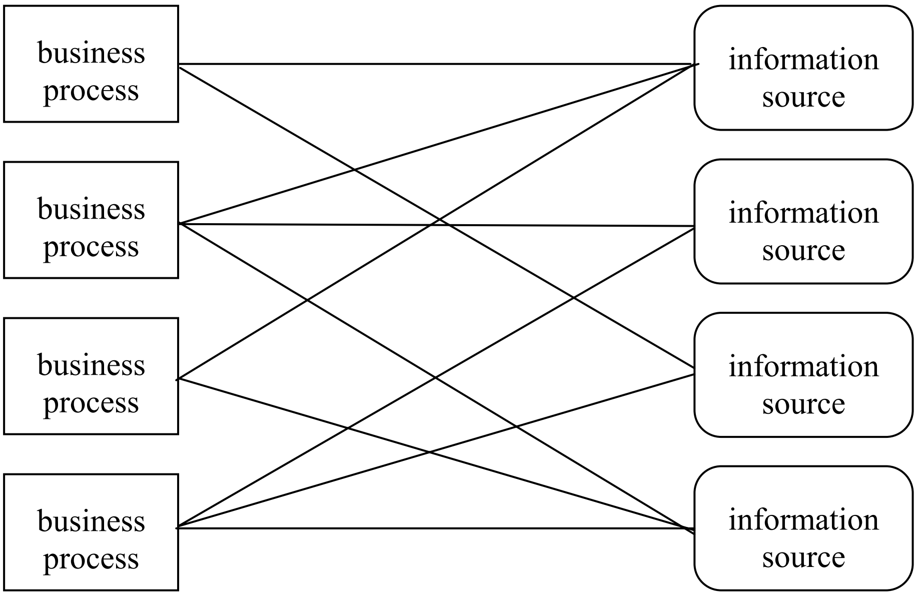

Figure 2: Business-oriented Architecture

Information security must become an intrinsic component of enterprise architecture. This need is driven by ever more complex mappings between data sources and users. Users may reside on partner networks over which data providers have little or no security authority, and business information is likely to be assembled from numerous data sources not necessarily under common control. These developments force us to take a more proactive approach to security architecture decision-making based on business structure and processes. Nevertheless, the ultimate security implementation remains a technical issue: the creation of a network of secure channels between segregated and secured assets and segregated and secured users that has no soft centre and has a much less well-defined corporate perimeter or, ultimately, even no perimeter. The change of emphasis provides an associated opportunity to improve the quality of the budgeting process and consequently both the cost-effectiveness and the assurance of security, but it requires a new management approach that encompasses both business expectations and technical solutions, and a rigorous methodology allowing business and technical personnel to collaborate effectively towards the common goal of security without loss of focus.

2 A New Approach

For the technical security architect, this blurring of boundaries has made the primary requirement one of understanding the assemblage of business processes and their associated information flows. For security management and budgeting personnel, the aim has becomes to prioritise the allocation of a security budget with reference to the business value of the various information assets to be protected, the first task being to place values on those assets in the context of the business processes they serve and the hazards to which they are potentially exposed. For both parties this must be in the first instance a business-centric matter, not a techno-centric one, so we must speak of information, not of data, must consider detriments in business rather than in solely technical terms, and must define hazards within a business, not a purely technical, frame of reference.

Whereas in the past techno-centric security architectures were generally derived from a single-stage technical requirements analysis, the need is now for a multi-stage analysis that bridges the business-technical divide on the lines of a Zachman framework [Zach99], but, and critically, with an explicit mechanism for maintaining conceptual integrity between layers.

The components of this analysis process are

1. recognition of business structure
2. discovery of business processes and their boundary communications
3. discovery of the information assets used by each business process
4. for the identified information assets, discovery of the business detriment caused by specific hazards
5. costing of identified business detriments
6. assigning values to the information assets for each attribute by mapping back from the discovered detriments
7. mapping the information assets to the physical infrastructure
8. aggregating the information asset values by attribute against the infrastructure
9. specifying proportional budgets for securing infrastructure components against breach of the attributes
10. delivery of technical solutions with reference to the required attributes and technical threats, and within the proportional budgets

Clearly this is a more complex specification process than has generally been deemed necessary in the past, but it must be recognised that it has to fulfil a much more sophisticated brief, not least in bridging the business/technical cultural divide to ensure that the necessarily complex business security vision is accurately reflected in the final technical implementation. It is also imperative to recognise that the way the process is performed will have an overriding influence on the quality of the results. Most extant security management and architecture development methodologies gloss over this distinction, but experience shows that it is absolutely critical to recognise it if consistency is desired. Uncertainty, and consequent variability in judgement quality, must be minimised as far as possible by some formalised process of debiasing [KaST99]. The key contributors to uncertainty in the current context are subjectivity and linguistic imprecision, which have been widely discussed in other spheres of risk judgement [MoHe90]. They will primarily affect components 3, and 4, 5 in cases where finite data are not available (for example when assessing regulatory liabilities), and possibly 2, depending on the complexity of the business process set. But these two causes of uncertainty can be minimised by making use of well-defined parameter sets, standard semantics and rigorously defined investigative methods.

2.1 Defining the Parameters

The precise definitions of business structure, processes and information assets will obviously be dependent on the nature of the enterprise. It is highly probable that existing business process management systems and tools can provide support here, remembering that for a complete security implementation any manual processes and information sources must also be included. However, the distinction between data and information is important. An information asset is defined jointly by the information used and the specific business process that uses it. The same actual information used by a different process constitutes a different information asset. This allows for the possibility of differing impact of a given breach depending on the nature of the business process that is affected.

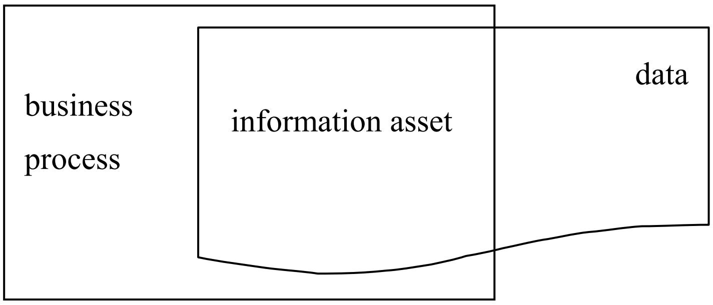

Figure 3: Venn Diagram of an Information Asset

2.2 Criteria for the Semantic Structure

To consistently assess information asset values, it is necessary to work within a well-defined semantic structure that expresses exposures and detriments unequivocally. Linguistic imprecision is one of the major contributors to inconsistency in risk decision-making, and much care must be exercised to minimise it. Any such structure must be transportable across the business-technical divide without ambiguity. Consequently it must be small in scale, and must exclusively make use of terminology that is simply expressed in both spheres. Considering first the exposure vocabulary, concepts such as “loss of reputation” from the business perspective, or “viruses” from the technical, are not useful as they are vague and their implications are subjective.

2.2.1 Information Attributes

The established information security trilogy of information attributes: “confidentiality”, “integrity”, “availability”; are a good starting point for an exposure structure. Each can be explained simply in terms of both business expectations and technical requirements. But this minimal vocabulary has to be extended to cover the range of current needs. A core vocabulary should include

- Confidentiality - only the right people accessing only the right information
- Integrity - freedom from corruption and unauthorised alteration of information
- Authenticity - assurance of authorised origin and validity of information
- Availability - assurance of authorised access to information when required
- Retention - fulfilment of both business and statutory archival requirements
- Ownership - control over the possession of supplied information
- Regulation - fulfilment of purely regulatory obligations for information management

2.2.2 Defining a Breach

The other crucial component of the exposure semantics is the definition of a breach. As we are not at this level of the framework addressing the incidence or likelihood of a breach (which is a technical consideration based on the driving events), we can usefully assume a Boolean operator: it either occurs or it does not occur. Thus any breach is deemed to be a total breach. This makes the process more manageable without losing any important detail. It also essentially ensures we work to worst-case scenarios in the first instance. The breach of an attribute in respect of a specific information asset is termed *exposure*.

2.2.3 Defining Business Detriments

The business detriment vocabulary will depend to a great extent on the nature of the enterprise, but core terms will include such concepts as “unable to process order”, “delayed delivery”, “customer lost”, “fraudulent transaction enabled”, “legal liability”, and so on. There may of course be sub-categories, as for example in the case of fraudulent transactions and legal liabilities, the potential scale of which will depend on the context of specific business processes.

The key attributes of a good detriment term will be its freedom from ambiguity and independence, i.e. its lack of intersection with other terms in the set. This is why terms such as “loss of reputation” are inadequate, as they are purely qualitative.

2.2.4 Costing Business Detriments

Having established a set of business detriments appropriate to the enterprise, costs can be placed globally on the detriments, based on recent business records for the enterprise in consultation with accounting, legal, audit and actuarial. It is strongly recommended that a minimum, median and maximum expected cost for each detriment, as the distribution may vary widely. In the case of simple detriments such as lost orders, accounts data for a suitable period can be analysed to directly establish values. Where the relationship between the detriment and the cost is less well-defined (for example contingent losses, legal costs) a formal enquiry method should be used as when identifying the detriments themselves. Time constraints such as process turn-round and throughput per time period must be determined and factored into the costings. In all cases adequate documentation must be preserved to show how the result was obtained.

2.3 The Process

Having established the semantic framework, how do we proceed? In the first stage, the information assets are unambiguously allocated identities that refer back to their data type and the business process they are derived from. This metadata is then carried forward to the business enquiry phase.

2.3.1 Business Enquiry Phase

Staff who operate the business process are asked to consider which of the defined detriments the process would suffer if each attribute of each information asset they make use of were breached in turn. The output of this stage is a list of information assets by business process against which a detriment term is associated with each of the hazard terms. The list can be then merged with the established detriment cost data to yield values attached directly to the attributes of information assets. Where more than one possible detriment may derive from an information asset attribute breach, the costs of all should be aggregated with strict regard for the information attribute, although it will prove beneficial to keep a record of the breakdown as well. This aggregation is a useful simplification for the next (mapping) stage without loss of functionality. The metadata now consists of a list of information assets, each associated with a value for each attribute. At this stage, the value metadata can usefully be quantised in to a limited number of bands. Often, such quantisation is limited to three bands (high, medium and low priority) but experience has shown that the optimum is five bands (noted but not for action, low priority action, medium priority action, high priority action, immediate imperative).

2.3.2 Technical Mapping Phase

The attribute-costed information assets are next mapped back to the technical infrastructure as data in terms of storage and user nodes and transport pathways, not forgetting that manual processes are part of that infrastructure and manually handled information also constitutes information assets. This allows concentrations of exposure in the infrastructure to be identified in respect of each information attribute. It is here that the primary benefit of this approach becomes clearest. Because the information attributes (until now considered only with reference to business expectations) have direct technical meaning in terms of generic security solutions, the transition across the business/technical divide is unambiguous. A technical architecture is automatically derivable, and budgets can be allocated to specific business/technical requirements with good assurance that the spend will be proportional to the value of the assets being protected and will be allocated to covering the real exposures of those assets.

3 Methods

It is important to use verifiably reliable methods when gathering the necessarily large and diverse body of metadata required for this process. Necessarily a multi-stage process is open to considerable opportunities for error if inadequately robust methods are used at any stage, and the greater the number of stages the greater the cumulative potential for error.

The most significant characteristic of any method chosen to deliver this process is its ability to ensure consistency. Absolute (numerical) accuracy is much less important, as the primary contribution to the budgeting task made by this process is the optimisation of security spending priorities, i.e. the optimum distribution of the available budget. Therefore high confidence in relative measures is the essential goal.

3.1 Interviewing

The way enquiries and metadata gathering are conducted has been widely recognised in other fields of decision-making as having a strong influence on the reliability of results, although this is not yet common knowledge in the information security arena. Put simply, a major contribution to quality can be made by asking the right people the right questions in the right way. Although this seems obvious, it often overlooked. Security management committees still tend to look inwards for answers from their membership, and consultancy analyses still tend to concentrate their enquiries at senior management level. Experience suggests however that it is most effective to start asset and process analysis by interviewing the lowest echelon that handles any process or structure being investigated, the next tier above being resorted to if answers conflict or respondents are uncertain. The lowest echelon is most likely to be aware what is actually done, whereas higher echelons will generally be more guided in their responses by what is documented as the official procedure. In general, for best assurance, at least three individuals should be interviewed at any tier and the answers of each validated anonymously with the others (a simple variant of the Delphi method [Gord94]). In accord with the precepts of Delphi, standard question and answer sets must be carefully prepared both for consistency and to eliminate ambiguity. To this end, generic answers based on the detriment vocabulary are used to validate free-form answers from the interviewees in a two-stage process. The interviewee is first asked “what would happen if ...” and their answers are noted. When all questions have been answered, the interviewer returns to the start of the question set and agrees with the interviewee which of the standard answers best accords with their free-form answer. In general there should be good conformity, but significant discrepancies will indicate cases requiring further analysis. A Delphi approach should be applied in a similar manner to the detriment costing exercise wherever finite data are not available from which to draw an objective numerical response.

3.2 Metadata Management

Management of the necessarily large volume of metadata requires careful consideration. It is an extremely valuable resource for the business that can be brought to bear on other areas of decision-making than security, but it needs to be carefully structured. Experience suggests that a good approach is a virtual tree of business processes with the root at the enterprise and the information assets forming the leaves, as this permits information assets to be uniquely identified by parsing the tree. Metadata structural considerations will depend on local preferences, as often existing business management tools can be drawn on to support the process. It is indeed in many cases very possible than some of the business structure metadata is already available and will not need to be recreated. Common standards will eventually emerge, but as a minimum a consistent presentational standard should be provided throughout the enterprise. Collaboration from interested specialists would be welcomed.

The final important management consideration is the security requirement for the metadata set itself. It will be an extremely sensitive descriptor of the business as a whole, and therefore will quite possibly become the most highly costed information asset, particularly in respect of confidentiality.

4 Return on Investment

It must be asked whether implementing such a system compares favourably with extant security budgeting and specification methodologies from the perspective of RoI. RoI is recognised as notoriously difficult to determine in information security, so this is an extremely important question. There are several aspects to the response. First, the described approach improves quality of fit between business expectations and technical implementations. This must by definition lead to better allocation of resources to security implementations than blanket security applied across the board. Second, assurance of that fit is improved by the evidential support for decisions that is inherent in the methodology. Third, many extant methodologies are designed to be delivered as project-oriented consultancy, whereas the described method is intended to be implemented as an embedded component of ongoing business management on the lines of CMM. As when implementing any business management system, initial investment may be considerable, depending entirely on the maturity of business management currently in place and the extent of business process intelligence already available. In most cases, bulk metadata capture will be the most costly component of initiating the method. once in place, it should operate as a CMM level 5 self-sustaining business function with little further overhead except periodic updates when business changes occur.

5 Conclusion

A process has been described for proactively prioritising the allocation of information security budgets based on potential business losses resulting from detriments to defined business information attributes. It differs from previous approaches in that

- it makes used of a well-defined semantic structure to ensure conceptual integrity across the business/technical divide, thereby ensuring the correct problems are solved
- it is business-centric rather than techno-centric at the budgeting and architecture decision points

Its key advantages are that it promises to improve long-term RoI in information security management by optimising the fit between security requirements and solutions, that it offers improved assurance by relating security budget distribution directly to the value and distribution of protected assets, and that the process metadata provide an objective and auditable underpin for decision-making. The same metadata will also support additional purposes such as business continuity planning and process streamlining, further improving overall RoI.

References

[Zach99] Zachman, JA: A Framework for Information Systems Architecture. In: IBM Systems Journal, IBM, 1999, p. 454-470.

[KaST99] Kahneman, D Slovic, P Tversky, A: Judgement under Uncertainty: Heuristics and Biases. Cambridge University Press, 1999, p. 422-444.

[MoHe90] Morgan, MG and Henrion M: Uncertainty. Cambridge University Press, 1990, p. 56-60.

[Gord94] Gordon, TJ: The Delphi Method. Futures Group AC/UNU Millennium Project, 1994.

Keywords

security, budgeting, proactive, business, information, asset, deperimeterisation, architecture, uncertainty