

The Information Security Policy Hierarchy

Developing A Governing Policy & Subsidiary Policies

A Maturing Field: As the discipline of information security becomes more sophisticated, codified, standardized, and mature, it is not surprising that the old-fashioned approach to information security policy writing is no longer appropriate. We are talking here about the “one-size-fits-all” information security policy that is supposed to apply to all workers in a specific organization. Different people within an organization have different things that they need to know from an information security policy. This diverse set of readers should not be required to wade through a lot of irrelevant material in order to find the sought-after information.

More and more organizations are breaking their single information security policy document into various information security policies. What we often see is an umbrella information security policy relevant to all readers, accompanied by policies intended for specific readers only. In the latter category we see policies for systems developers, quality control engineers, and other functional groups. Most readers do not need to read the latter type of narrowly scoped policies, so it’s best if this information is separated from the main “everyone has to read this” material.

Getting User Friendly: After this separation between an umbrella policy and subsidiary policies, on a level of sophistication scale, the next stage is breaking down information security policies by job title. A very large American bank did this with great success via an intranet, and the workers really appreciated knowing what exactly they were responsible for, and also what they didn’t have to worry about. Using a more progressive perspective, it would be better to structure this document breakdown by specific cross-departmental business process. As information security policies continue to expand in size, and as they become ever more detailed, this type of audience targeting is increasingly necessary.

On one more sophisticated level still, a level to which very few organizations have presently gone, is a breakdown of policies into very brief statements relevant to a specific task. For example, if someone wanted to gain access to a new computerized business application, a privilege that they currently didn’t have, the organization could have built a series of web forms that such a person could fill out to submit a request. Pop-ups would appear instructing them as they fill out these forms. On selected pop-ups, and also available as links (to be followed as desired), they would see a paragraph or two of information security policy material, but only material relevant to this specific task. Unfortunately this approach takes a lot of effort, is rather time consuming, and in some instances can’t be done at all (if off-the-shelf packages are used for example). Nonetheless, this approach integrates information security with automated processes, and in that respect is desirable because it communicates the message that “information security is a normal part of how things are done around here.”

This last approach eliminates the whole question of people ignoring information security policies, for example because they failed to consult the policies that were found in a separate place. Instead the policies are merged with business processes, and compliance is achieved via one or more action-forcing mechanisms. An example might be a digital signature from a departmental manager being required before access to a specific system privilege is granted. A

systems administrator would be blocked from changing the privileges for a normal end user unless that digital signature has first been obtained and confirmed as legitimate.

Long Term View: If one uses an umbrella information security policy, sometimes called a governing policy, and then develops specific policies that fall under that umbrella policy, you will find that maintenance and updates will be considerably easier. Approval of a short and narrowly-scoped document will be conceptually easier for many people, especially non-technical managers. Breaking things down in this way also supports making a clear and sharp distinction between different types of information security documentation, for instance distinguishing between policies, guidelines, procedures, technical standards, and contingency plans. Clearly differentiating between these document types allows information security policies to be kept on a high-level of abstraction, and thus at least potentially be in force for five years without modification (although policies should be reviewed annually for relevance and needed changes).

In an umbrella policy, we will typically see a statement of objectives for information security, which explains how these objectives support organizational goals. We would also typically see a statement from the CEO stating his or her expectation that everyone working at the organization comply with all information security requirements (policies being just one of these). We would additionally expect to see human resource related matters applicable to all readers. For example, a discussion of the disciplinary actions that will be taken in response to a violation would be found in an umbrella policy. Training and awareness matters, such as a required annual refresher course, those too would be addressed in an umbrella policy. Structures used in other policies, ways of looking at information security that everybody needs to understand, such as the user-custodian-owner model, these would also typically be explained in an umbrella policy.

Links to Specific Policies: In an umbrella policy, we would furthermore expect to see links to more specific policies, sometimes called technical policies. These more specific policies could address generic areas like access control, user authentication, system logging, physical security for computers, and encryption. Although some organizations have chosen to organize their more specific policies along the lines of vendor technologies, like the Windows operating system, this author recommends against such an approach. Information is not confined to only one operating environment, and a consistent approach to security is needed across all vendor technologies, across all operating systems, and for that matter across all organizations that have access to the information in question. It is far better to take an information sensitivity oriented approach, for example breaking statements about required controls down by a data classification system. This approach reflects the perspective that technology should follow business needs, and of course information security is a business need.

So when structuring the subsidiary policies, you can use the traditional role-based approach, you can use a business process based approach, you can use an information sensitivity based approach, or you can use an issue based approach. With the latter approach, in one subsidiary policy document, we would for example talk about what to do after there has been an intrusion. That document would address who makes executive decisions such as shutting the affected system down, what information about the intrusion must be recorded for legal and insurance purposes, how to gather and properly store evidence, who acts as a public spokesperson, etc.

Subsidiary Components: More specific subsidiary policies should for example address matters such as: (a) who is responsible for buying, renting or leasing new systems, (b) who must approve of the security measures on new production systems, (c) who is responsible for managing the security on these systems, and (d) how these systems must comply with standardized configurations. The operating system configurations can and should be defined in separate documents, for example a set-up procedure for systems administrators. In general, if a separate person is going to handle the more detailed matters, such as configuring a new computer, then this is a good point at which to have a separate document.

Using this — or a similar — rigorous top-down hierarchical approach will bring a discipline to information security policy writing that will be much appreciated down the road. This author has seen far too many cases of spaghetti-style policy writing, where everything is hopelessly interconnecting and overlapping, and it's very hard to figure out what the policies actually require. Of course, in the latter case, update and maintenance is a nightmare. Often, the lowest-cost and most-expedient approach is to replace the whole spaghetti-style document with an entirely new set of clearly-structured documents.

A Bonus: Another significant benefit of breaking things down as suggested here is that access control, based on the principle of “least privilege” can easily be maintained. For example, if a contractor is going to help with the systems design of the accounts receivable system, then only the security policies applicable to the accounts receivable and collections areas need to be disclosed to him. And for many other people too, both inside and outside an organization, both separation of duties and dual control can be better supported if we have a combination of multiple narrowly-scoped policy documents, and access control restrictions at the document level.

Whatever your reasons for structuring information security policies the way you do, make sure they reflect the business needs of the organization in question. More specifically, before you make a decision about the appropriate policy document structure, make sure your organization has a recent risk assessment that talks about the most pressing information security issues confronting the organization. It's important to use that information to then create a structure for policy documents that fits the prevailing organizational structure, vulnerable business processes, and important information security tasks.