



No Illusions: Rethinking Information Security Policies and Standards

Andrew Stewart

Abstract

As a profession we over-estimate the value of security policies and standards. In this paper I describe the myths of security policies and discuss the common problems that we face when attempting to employ policies and standards in their current form. A number of techniques are proposed that compensate for some of the weaknesses in the traditional approach, and a philosophical approach to the use of security policies and standards within a modern information security programme is described.

Terminology

When I talk about security policies and standards in this paper I am referring to the types of policies and standards that we would expect to see as part of a comprehensive information security programme within an organization. I am not referring to capabilities based standards such as the Common Criteria or the Rainbow Series.

What are we Trying to Achieve?

"Do you have information security policies within your organization and who is responsible for them?" That is the classic opening line on the subject of security policies that we've all heard countless times from auditors and security consultants alike.

But when thinking about security policies, rather than starting with "where are they?" and "who is responsible for them?" we should take a step back and ask "what is it that we are actually trying to achieve?". An organization will pursue the development of security policies and standards for a mix of the following five reasons:

Regulatory compliance – An organization may have to comply with a regulatory act such as the Health Insurance Portability and Accountability act (HIPAA) or Graham-Leach-Bliley (the Privacy of Consumer Financial Information act). Having policies and standards is usually part of the compliance criteria.

Corporate governance – A high-level policy statement that is ratified by senior management can provide the security function within an organization with the authority to carry out the various tasks associated with a security programme.

Certification – Insurance companies will sometimes offer a lower premium for "cyber-insurance" if an organization has been certified to comply with a security standard such as ISO 17799. Certification also has a role in creating the perception of having taken a diligent approach to security (more on this later).

Due diligence – An organization will often desire compliance with best practices and industry standards which are defined by the security industry as a whole. These are usually presented in the form of standards documents.

SECURITY POLICIES

To provide direction – The majority of the motivation for creating policies and standards falls into this category. The thinking is that by providing the organization with say, a Windows Server 2003 security standard, then the subsequent rollout of that technology will have had input from the security team.

The first three items in this list are valid reasons for developing security policies and standards. For the last two items in the list, we really need to re-assess the value that policies and standards deliver.

Information security efforts have to be business-driven, and so we need to look rationally at the benefit that each security measure can return. This paper performs that analysis for policies and standards, beginning with a description of the myths surrounding policies and standards.

Policy Myths

Our perception of the value that security policies and standards deliver is remarkably over-inflated. The reason for this misconception is that there are a number of myths associated with security policies and standards, and the security profession has institutionalized these myths.

“Start With Security Policies”

Starting “from the top down” by writing policies is a tenet of information security doctrine, but it can often lead to the mistaken idea that the mere existence of security policies proves that a structured approach to security is being taken. I think that we’ve all seen organizations that have very thorough policies but a technical security program that is disproportionately weak. Quite often, these organizations will pass security audits because the security auditors take a paper-based approach that consists of examining policies!

We are generally told that “without security policies, you have no general security framework”. But any security programme that takes that approach is taking a security-centric view of the world. That path leads to countless difficulties, as I will explain later.

Bruce Schneier has joked that for any given security standard he could create a product that 1) met the standard and 2) was still insecure. Reading over typical security standards you have to realize that this is usually easily done, and so we have to ask ourselves: to what extent does the creation of security policies and standards translate into real gain?

“Deliver Policies Into Your Organization”

In many organizations, the notion of a policy document sitting on a shelf gathering dust is literally true (I have seen it more than once). There is a tremendous difference in difficulty between writing a policy and then attempting to deliver that policy into an organization.

Most security problems arise out of mistakes, not malice. Education and security awareness would seem to be a logical approach, but to quote from popular culture: there’s a difference between knowing the path and walking the path. No amount of security awareness activities can equip a typical user with the knowledge to make security decisions at the level that’s usually necessary. And that is really the root cause of the problem: users are placed in a position where they *can* make security decisions.

When you really want to solve a problem you tackle its roots. Consider for example which of the following is the more effective strategy in avoiding buffer overflows: a) the use of a high-level language such as Java or b) a security standard that tries to describe how to avoid buffer overflows in C? Security policies and standards do not attack the roots of the problems that we collectively face.

Problems with Policies

Alongside those policy myths are several difficulties with the way we attempt to deliver policies into an organization and try to make them effective. These difficulties are described below.

Policies and Set Theory

A friend of mine once told me a story about a company that had network points in meeting rooms to enable employees to access the company intranet while they were in meetings. The problem was that the meeting rooms were on the ground floor of the building and had very little physical security. Almost anyone could walk in from the street, go into an empty meeting room and plug directly into the company network.

That problem could only have been foreseen by taking into consideration both the virtual and physical security domains. A typical taxonomy of security policies runs something like “Network, Host, Application, Physical, Remote Access, etc.” That distribution is very vertically stratified and doesn’t lend itself well to considering those kinds of “blended” threats. We need a solution that can dynamically define security goals when presented with a list of requirements (or technologies). Picture a series of Venn diagrams with each circle representing a traditional policy area; it is the areas of intersection that we are interested in.

Handling Exceptions to Policy

With the traditional approach, once a security policy has been defined a process has to be created for managing exceptions to that policy. This is necessary because inevitably there will come a time where a valid business case for non-compliance is presented. The exemption process usually involves the project owner signing a form that states that they acknowledge and accept the risk, and for those of us in security who have

worked through that process it is always a very unsatisfying result.

An exception is really an example of a situation where the participant in a process has to step outside that process. The goal of a security programme should be to become part of business processes so that, in effect, it is impossible for an 'exception' to occur.

The Irony of Best Practices

How much security is enough? We may not know the answer but realistically we can say that the same level of security is not suitable for every organization. And yet with traditional policies and standards we persist in trying to apply a "one size fits all" approach.

The majority of policies and standards have their roots in environments that are nothing like our own. Policies that have originated in large, monolithic, mainframe environments are of very little use to most modern organizations. Likewise, security standards that come out of very sensitive, restrictive environments are of almost no practical use. The NSA guides to locking down Windows 2000 come to over a thousand pages; who's got the time?

The phrase "best practices" can itself be misleading where "best" security is taken to mean "most" security. Setting the bar too high creates the problem of how to deal with policy exceptions and it forces the security organization into a never-ending reactionary spiral that unfortunately most of us are familiar with.

New Approaches

Many of the issues I have described have their roots in how security is perceived within an organization and indeed how security is delivered within an organization. I have found the following approaches to be successful in tackling these problems.

Keep in Mind the Law of Diminishing Returns

Don't waste time and money on elements of security policies and standards that have no real return. A policy statement that is necessary for regulatory compliance, certification or corporate governance reasons has attributable value. But requirements such as these need to be mapped to your policy effort, otherwise the tendency is to write voluminous policy documents that will never translate into real gains in security.

Presentation is Key

The problem of having to handle exceptions to policy is actually an example of a general problem that occurs over and over again in different aspects of information security. For example, if your IDS policy is too paranoid you'll get lots of false positives but if it's too lax you'll miss attacks. If your password policy is too simplistic then passwords will be easy to crack but if the

policy is too complex then users will write down their passwords.

We approach these types of problem by retaining an element of flexibility and by applying multiple overlapping techniques. Our network IDS might miss some attacks but we've got integrity checking software on the servers and we review logs, perform backups, and so on. Part of the password file can probably be cracked but we use two-factor authentication, monitor user activity patterns, etc.

Unfortunately that flexibility doesn't exist with traditional policies where we take a prescriptive approach of stating security goals in black and white terms. Both the prescriptive approach itself and the format (typically a document) contribute to this problem.

As a field we understand role-based concepts, but for some reason we have not applied that idea to policies and standards. The nature of many of the new technologies that we are engaged with means that a developer who is building say, a .NET application, would have to read the network security standard, the application security standard, the relevant operating system security standard, the authentication policy, the database security standard, possibly also the gateway security standard... the list goes on.

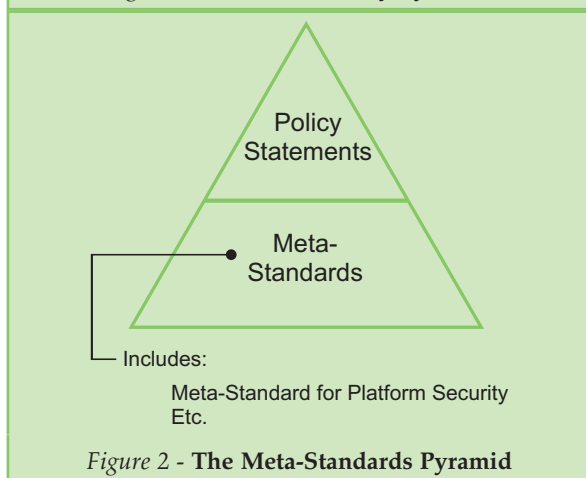
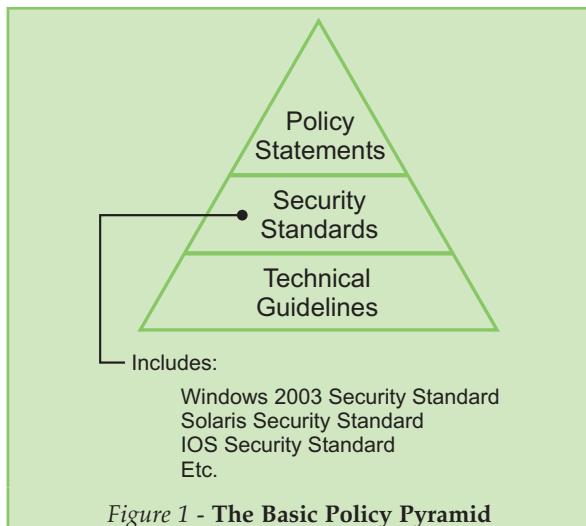
The security advice and direction we provide has to be categorized in a role-based way, and it should be delivered in a form that can be navigated easily and cross-referenced interactively. The web provides a way to achieve these goals, and a web-based query-engine would allow the developer in my example to quickly navigate the breadth of security advice that is available in different categories and see what parts are relevant to his or her particular project. Fall-through logic can be used to predict when additional areas of advice might be required (branch prediction).

Traditional policies are all "push" and no "pull". The technique I've described above shifts that ratio into more favorable territory through its ease of use and by creating the perception that security is delivering a service rather than mandating requirements.

Where's the SLA for Security?

Let's be honest, in many organizations security and audit are viewed the same way: as an annoyance. But almost certainly the network team isn't viewed in that way, nor are other teams such as the application development team. What causes this disparity? By focussing on policies and standards we perpetuate the idea that security functions is an audit-like role. The tension that often exists between security teams and other operational functions such as engineering teams has its roots in that perception: that security is an inhibitor required mostly for the sake of compliance.

SECURITY POLICIES



Your business considers the network team to be providing a service, the parameters of which are codified in a *Service Level Agreement* (SLA). Security should be no different. By defining “what security does” in terms of quantifiable, measurable services it becomes easier to agree expectations and to demonstrate ongoing value. A hybrid approach is also possible in which only certain aspects of a security programme (usually technical services) are described in an SLA.

From Standards to Meta-Standards

At the security standard level we’re usually concerned with ensuring that there is an appropriate level of security in an operating system build or in a class of application such as web servers or databases. The problem is that with each new product or with each new release of a product we need to either rewrite the security standard or write a new one. Your security standard for Windows Server 2003 will only share about 25% of its content with your Windows NT standard because of the large difference in functionality between those products.

Also consider that the people in your organization who best understand the Windows platform are in the Windows team, not your security team. Traditional standards do little to empower

the technologists who are actually performing the work because they don’t usually explain the concepts *behind* the security requirements.

The use of meta-standards is an approach that takes these factors in consideration. Meta-standards transcend individual technologies, freeing your security team from having to perpetually rewrite security standards as underlying technologies evolve. Meta-standards also encourage the delivery of security knowledge into other teams.

Consider the classic policies and standards pyramid shown in Figure 1.

The meta-standard approach replaces multiple security standards and technical guidelines in each functional area with a single document. Technical guidelines are no longer necessary. See Figure 2.

This can be achieved because it is possible to describe security requirements in a platform-agnostic way. Consider requirements that you would typically expect to see in platform-specific security standards such as requirements for log-on controls, backups, the permissible use of remote control software, file system security, auditing, system database security, and so on. Rather than explaining which check-box to select in Windows NT to enable strong passwords, a meta-standard will specify each requirement, explain the reasoning behind each requirement, and leave the implementation to the reader.

This is a real difference in approach. Help people to do the right thing and work with them to do it, not demand that they do it and deliver those demands in a rule book. Of course, this new approach requires more investment in face-time with the other teams but it results in those teams being able to understand and apply security principles themselves. A key benefit here is in the perception of ownership. People want to do the right thing with respect to security but they also want to own their investment in pursuing the right thing to do. A rule-book doesn’t allow them to get that feeling of ownership.

As a general observation, the security field is obsessed with technicalities (if you enjoy reading about countless variations on a small number of underlying concepts then you’ll love BugTraq). As security professionals we would do more good if we were to step back from the minutiae of technical security and focus on communicating the guiding principles of the field. This applies to many aspects of technical information security such as vulnerability assessment and writing secure software. Although we might consider them trite, teaching concepts such as *practice defense in depth, be reluctant to trust and follow the principle of least privilege* adds far more value in the long term than being able to say that a particular buffer overflow exists in a particular version of a piece of software (who cares?).

Summary of Actions

- Map your compliance requirements and certification goals to your high-level policy effort. Take a pragmatic approach to considering the value that each policy returns, and spend time only where it is necessary. A minimalist set of policies is almost always as effective as a verbose set.
- Save time and effort by transitioning any policy or standard that you expect to have to update within the next year into a meta-standard. Focus on the principles and ideas that underpin security, not the specifics of technologies.
- Work with the relevant teams to integrate the lessons in your meta-standards into their active work processes. For example, work with the NT team and the UNIX team to integrate your meta-standard for platform security into their standard build processes. Once those teams are thinking in terms of the direction described in the meta-standard, your job in that area is essentially done and you can move into an oversight role.
- We traditionally deliver security policies and standards in the form of point-by-point prescriptive documents in distinct subject areas. Your goal should be to create a delivery mechanism that is the exact opposite (role-based, web-based and query-able).
- Your relationship with project managers within your organization (and not just IT project managers) is as important as your relationship with system and network engineers. Only by integrating within your business to the extent that security is no longer considered a separate activity can security become truly effective.

Conclusion

By all means, employ security policies as part of your security programme. Creating the perception of security is often a valuable and necessary exercise. But be under no illusion that following the way we approach policies and standards today will usually provide the appearance of security alone. In the form of prescriptive, inflexible documents, traditional policies and standards may actually be detrimental to the effective delivery of a security programme within an organization.

An effective security programme is business-driven, and we should not exempt security policies and standards from that business-driven approach. The use of meta-standards and a focus on the delivery of security concepts into business activities are mechanisms to compensate for the weaknesses in the traditional approach.

I believe it is both healthy and necessary to explore how we can evolve all aspects of our field. I would welcome the opportunity to engage in a discussion of these ideas in an open forum. Please feel free to contact me on my personal email address to discuss this paper: andrew_j_stewart@mac.com.

About the Author

Andrew Stewart is senior consultant with a professional services company based in Atlanta Georgia. He has previously held information security positions at Deutsche Bank, Internet Security Systems, Reuters and Barclays Capital. His professional interests are in the areas of security programme development and enterprise security.

He is certified as a CISSP (Certified Information Systems Security Professional) by the Information Systems Security Certifications Consortium and is a member of the Internet Society, ACM, IEEE, USENIX and ISSA. Andrew received a Bachelors of Science in Computing from Oxford Brookes University.

Visit our exciting new Web Site at

www.isb-online.net/

to read the **daily news** about products, companies, persons, developments, vulnerabilities and other issues of interest to the information security professional – and make you own contributions!

**Information Security is where it happens -
isb-online.net is where you read about it**