Instant Compliance for a Grand

Michael Barwise, Integrated InfoSec
mbarwise@intinfosec.com

I've just received yet another email offering me a "complete set of pre-written information security policies for your organisation." For around a thousand pounds (plus VAT) I can solve my security problems in one stroke. Compliance is assured and data leakage will be a thing of the past. I should coco.

I freely admit I wasn't prepared to spend a grand to find out the exact content of this particular wonder pack just so I could write this paper. But I have instead investigated a set of open standard specimen policies from a respected source on the web, on the fairly safe assumption they'd probably be comparable. The one that impressed me most was a "password policy".

The Document

This document starts reasonably enough with a short introduction stating that strong passwords are a good idea, and specifying its scope as *"all systems users."* It then identifies some general requirements, about half of which are narrow technical items relevant and comprehensible to only a very small proportion of the supposedly universal user base. Then there's a list of nine prohibited things that make a password weak, including using *"a word found in any dictionary."* The ninth is any of the eight already prohibited things *"preceded or followed by a digit."* Six recommended rules for creating strong passwords follow. The last of these is actually two quite separate rules rolled into one, accompanied by a wordy explanation of how to achieve the second. But not one of these prohibitions or requirements includes any explanation of how it contributes to security.

In case you're not sufficiently confused yet, that was just the *"Password Construction"* section. The next section - *"Password Protection"* - starts with a preamble that assumes ordinary users can distinguish between Windows and Unix systems they access and requires them to use different passwords for each. So goodbye single sign-on. Then comes a "list of don'ts" - half a dozen quite sensible prohibitions and three more that aren't actually "don'ts" but things you're instead expected to do. One of the "do's" in this list of supposed "don'ts" rather threw me: *"change passwords at least once every six months, except system-level passwords which must be changed quarterly. The recommended change interval is every four months."*

The final joy is the "enforcement" section, consisting of the single sentence *"Any member of staff found to have violated this policy may be subject to disciplinary action up to and including dismissal."*

Putting myself for a moment in the place of a typical office worker faced with changing their password, I think I'll go nuts. Not only do I have to be aware of all dictionaries ever published in every language (please Sir, does that include dead languages and obscure tribal dialects?) but also to remember their entire contents as well, and if I get it wrong I could be sacked. So excuse me while I take an extended sabbatical in the British Library. I might be away a bit longer than the recommended password change interval though.

But even if I mutter "stuff that stupid rule" - or just ignore it - I still have to negotiate eight more (or was it seven?) prohibitions and six (or was it seven?) recommendations, plus two apparently conflicting suggestions of how to resolve the lot into an acceptable password. That's quite a challenge for starters. Then I have to remember the "compliant" result as I'm not allowed to write it down.

The Principle

Cutting through all this detail, what's the fundamental flaw here? This document puts almost the entire onus for password authentication security squarely in the lap of the general computer user, which is the worst place it could possibly land. There's no mention of locking out those who make too many mistakes in short order, there's no mention of audit trails or alerts, and most importantly there's no mention of protecting servers against illicit access to the master password files they store. Not even in the section *"for developers"*.

Now by far the most common ways passwords are breached are individually by people sharing them voluntarily or negligently, and remotely in bulk via illicit access to password files on authentication servers. In all fairness this document does mandate against sharing passwords, but none of the complexity rules it includes - indeed no complexity rule on Earth - will protect against breaches resulting from sharing. These rules merely offer an indeterminate level of protection against password cracking. And making the end user responsible for the resistance of their password to offline cracking attacks on password files stolen from the server is not only unreasonable. It's also unlikely to succeed for very long in thwarting an attacker already in possession of those files. In any case, if your systems can be breached from outside to the extent that your password files can be stolen, the passwords themselves are the least of your worries.

The only other realistic source of password breaches is live attacks on password interfaces - a movie-makers' favourite, but in reality quite uncommon because you can easily get caught carrying it out. But even were a business to be attacked in this way, there's no evidence at all that the complicated rules specified in this document - even if fully understood and complied with - would render password authentication substantially more secure than other simpler alternative rule sets do, provided some basic monitoring and a lockout rule were in place. In fact there's growing evidence from rigorously conducted research - for example by Weir and others[1] - that they don't.

So this pseudo-security burden on users doesn't protect against sharing - the only source of breaches over which the user can legitimately be expected to exercise primary control - despite making it much harder for conscientious users to comply. Nor does it protect against the two main attack types which your IT folks could forestall with some basic technological countermeasures that this document completely fails to mention. Furthermore, it's probably legally indefensible in most civilised jurisdictions to sack an employee for breaking such a confusing and self-contradictory set of rules - even if it could by some miracle be shown that a specific non-compliance was the proximate cause of an actual breach.

But I think I understand how these complexity rules were arrived at. They're unchallenged received wisdom, based on guidance found, among other places, in the supposedly authoritative NIST handbook on electronic authentication[2] and swallowed whole without due consideration. But wherever its actual origin in this case, it's poor guidance - the application of a naive interpretation of Shannon Entropy to an inappropriate problem. Shannon Entropy is fine for modelling the transmission capacity of a comms channel - what it was designed for - but it's entirely unsuitable for modelling modern password cracking, which is highly heuristic. So although they look impressive to the uninitiated, these rules won't contribute much if anything to real security. They're attempting to solve the wrong problem.

The Failure

So what's the bottom line? First, this document is not a policy at all - it's a set of procedures. A policy should define objectives and their rationale and set standards for compliance and efficacy - not get bogged down in the details of technical implementations.

But to create even a robust and effective set of procedures, you have to understand, and build on, the first principles underlying the actual problems you're attempting to solve. Which means you've got to be clear what those problems really are. The authors of this document obviously didn't and weren't, and the result is a top notch five star recipe for failure from both the technical and legal standpoints. But it has nevertheless been widely accepted as an example of "best practice".

The truth is, this document is waste paper - indeed it's worse than waste paper. It offers a semblance of control while exercising effectively none over the real hazards - providing a false sense of security that's vastly more dangerous than a known exposure. I suggest that if it did no more than mandate the missing controls I've outlined above, user passwords - and consequently the rules it imposes - could be substantially simpler without any loss of security. Security might even be improved. But for every control there would have to be a documented justification in terms of the specific threat it's supposed to protect against and how it achieves that protection, so those tasked with maintaining the document can ensure the set of procedures remains functional and relevant as a whole as the threat landscape evolves. It might even be a good idea to offer some non-technical explanation of why each of the controls is present - compliance might be better if everyone understood the reasons for it. But this document still wouldn't be a policy.

At the end of the day, the sole purpose of a policy - or even a set of procedures - is to minimise the chances of an adverse event occurring, so it had better cover all the real bases effectively, succinctly and intelligibly. It must say clearly what it means and clearly mean what it says. And it must demonstrably actually work.

References:

[1] Weir M *et al*. Testing Metrics for Password Creation Policies by Attacking Large Sets of Revealed Passwords http://sites.google.com/site/reusablesec/Home/presentations-and-papers/Defcon09v2.pdf

[2] NIST Special Publication 800-63. Version 1.0.2. Electronic Authentication csrc.nist.gov/publications/nistpubs/800-63/SP800-63V1_0_2.pdf

---

Author:

Michael Barwise is a UK-based policy and risk management consultant with a background in information assurance and systems engineering. He has contributed to international policy and national legislation on Internet and information risk, has lectured on policy development at Masters level, and is widely published online. His special interest is assisting business to improve the quality of risk decision-making.