Too Many, Too Often, Too Serious

Michael Barwise

At the end of August, a server on the Linux Kernel Archive infrastructure was compromised, and its custodians have been busy verifying whether the repositories were contaminated. A few weeks earlier, Vasco reported a breach of the Certificate Authority infrastructure of its Dutch subsidiary DigiNotar, which resulted in the malicious creation of a so far unconfirmed number of spoof SSL and EVSSL certificates. And it's only a few short months since a breach at RSA resulted in a huge loss of confidence in SecurID, leading the company to offer to replace all seed keys.

Ok, I know breaches happen all the time - on average several times a day world-wide, and that's bad enough in itself. But it's the potential repercussions on the web at large that really matter in cases such as these. The number and scale of breaches in such services is really disturbing.

The worst of it is the apparent simplicity of some of the attacks and the looseness of control exhibited by the victims in these cases. The scale of the DigiNotar breach was apparently not fully identified for a week and the fraudulent certificates included one for the Mozilla add-on site and another for any Google domain; the Linux archive was supposedly penetrated via a "compromised credential"; RSA was breached via user-triggered malware in an email attachment.

So the inevitable conclusion must be that security-critical services are not themselves being protected by their providers much better than rank and file web sites. The quality of the defence is much worse than the quality of the adversary - but the adversary hardly has to lift a finger in many of these cases. We're still leaving the keys under the door mat if not actually in the lock any more. Of course that's your option if you run an online shop - you only have your own customers to answer to - but it's just not good enough when you're providing core services to others who might be running anything from an online shop to a government. For then a breach can affect all of us - potentially world-wide.

So I'm making a proposal right now that all providers of security-critical offerings that support the web services of others should be required to publish the outcome of regular security tests. Not the detailed results (that would be stupid) - but notification of an audited pass or failure. The credit card security standard already imposes some such testing on those who process payments, and for very good reason. They have a moral duty to ensure they merit our trust. I'm suggesting this duty should be both extended to all security-related services and supported by mandatory disclosure of evidence of security. It's not a panacea, as testing can't guarantee to cover all the bases. But at least it would be a start.

Originally appeared on the Infosecurity Network, September 2011