Outsourcing penetration testing? Go local!

It does not matter if you have an internal security team in place or how good they are, penetration testing should always be outsourced. The rational for this is exactly the same for using external financial auditors: they offer an unbiased perspective as well as being able to bring in particular expertise. One option is to use offshore outsourcing of highly skilled IT work, network security testing included. This is the current issue amongst many IT professionals and whilst there are certainly some cost-saving advantages, the major arguments against it include the lack of company transparency i.e. who is the company and what will they do with the information and what can you really do if dissatisfied with the service. Some also quote a negative impact of offshore information security outsourcing on the national security and various political issues.

In addition to these concerns, we will show that there are also technical arguments against cross-border IT security outsourcing and argue that wherever possible network security audits should be a local procedure, ideally from within the same town or city.

Lost on-route ?

First, let us examine various negative effects of the extended packet path to the targeted hosts. Additional hops on the path to the assessed network do introduce additional problems – sometimes so many that the audit itself becomes worthless. With the increased security awareness of ISP's, many providers filter out ports involved in common attacks and frequently abused by worms. These often include NetBIOS and common trojan, worm or botnet server ports. The longer the path to a target, the higher the probability of encountering such filters in place. Ports filtered anywhere else than the corporate firewall (even if it is their providers gateway) that show up as closed in the penetration testing report create a false sense of security and can strongly reduce the quality of the testing procedure. Tools such as hping, lft or even a simple traceroute can help to determine whether filtering occurs at the tested company firewall, their ISP gateway or further up the router chain. Another potential problem is a service transparently redirected through a proxy, or port forwarding enabled on one of the routers on the path. Such a port could be shown as open on the scan report and trigger a negative alert or create a false positive. It can even get attacked by a less skilled auditor or, more likely, by an automated audit tool or script, thus unknowingly “exploiting” the system outside the tested domain and potentially breaching the law. Again, running hping to check TCP timestamps combined with using scanrand from the Paketto Kerietsu suit helps. Nevertheless, none of the listed tools can guarantee a 100 per cent correct result, and employing them introduces additional auditing time, workload and expense for the audited company.

Another issue introduced by a longer packet path is, of course, delay and packet loss. This is particularly evident when traffic is sent from an offshore site connected via a slow and/or unreliable link using lower end routers. The packet loss can easily make the result of a wide reconnaissance sweep unreliable, in particular if an automated scanning tool is left to run unsupervised for a long time. As for the delay, imagine scanning several dozens of IP addresses to examine all 65535 TCP and, particularly, UDP ports via a single PVC link of an oversubscribed and noisy line. Unacceptable delay may even force the penetration tester to use the standard /etc/services or nmap-services ports instead of the full port range (again, this would specifically apply to the already not-so-reliable UDP scanning). This, in turn, can lead to overlooked problems and an incomplete audit. Also, in some cases where the availability of a service is a highly critical issue, prearranged DoS resilience testing can be included as a part of the external assessment. High delay and bottleneck routers will greatly reduce the efficiency of, for example, syn flooding the audited services and won't show how successful an attacker more close to the company's network can be. And when a shellcode is sent to a target service in more than a single packet using a stateless protocol, such as UDP, a single dropped packet will prevent the exploit from working. Ten minutes later the bottleneck may go away, and the exploit would work successfully, but “service not exploitable in practice” entry is already in the audit report! Finally, bottlenecks and other connectivity issues can make remote fuzzing completely unfeasible, thus depriving penetration testers from using this highly popular and efficient assessment methodology or twisting it's results. Some of the test cases fired at the target service over a stateless protocol can get lost on the way, and a delay from the service response (or lack of thereof) can introduce confusion making it harder to determine which particular test case has caused the crash.

Putting close range threats into perspective

The problems outlined above can be summarised as the inevitable reduction in precision or introduction of more randomness with every additional hop between the auditors and the tested network. You can obviously argue that these cases can be solved with applying higher auditor skills and increasing the assessment time, while the disadvantages brought by traffic delay and possible packet loss could be still outweighed by all the money saved by not hiring a more expensive local IT security team. However, there are some attacks which demand a limited amount of hops between the attacker and targets. Their success can be fatal and there isn't a way to evaluate risks they present from a remote site many hops away.

One example of such a “close range” attack is RIP route insertion aimed at redirecting the traffic to a sniffing/hijacking machine. This can be done employing Nemesis, any other custom RIP packet generation tool or even running a rogue routed/quagga on the attacker’s box. Of course, 15 hops is still quite a distance, but don't be so sure. The distance from our gateway in the UK securityfocus.com is 21 hops and it is obviously higher for Asian or Australian networks. A much better example is exploiting loose and strict source routing IP options. It can only be done within the 8 hop distance, limiting the attacker and auditor alike to the same or neighbouring ISP with their target. Considering the great potential for abuse (check out the old good Todd's lsrtunnel if you have any doubts!), we view testing the audited networks for enabled loose/strict source routing as essential for a complete remote network security audit.

What about attacks launched from a single hop distance? If the determined cracker cannot bypass the firewall, he is likely to go after the ISP gateway or hosts which belong to the same network together with the “dirty side” of the company firewall/router. The later is particularly problematic for cable networks. In approximately 70% of cases the providers gateway would be a Cisco router or switch. Even though many people think that having control over such a device is only good for DoS or unencrypted traffic snooping (debug ip packet “insert the ACL defining the traffic of interest here*), traffic mirroring from the “owned” Cisco to a Linux machine loaded with all the tools necessary for launching man-in-the-middle, connection hijacking and other attacks sorts out this problem. The methodologies for such mirroring are well-documented, see our “Hacking Exposed Cisco Networks” Chapter 10 for a variety of examples. But even when attackers cannot obtain control over the gateway box, a variety of efficient traffic redirection/man-in-the-middle attacks can be launched from controlled hosts on the same network with the target company or even the ISP gateway. These can range from the “traditional” ARP-based man-in-the-middle attacks to abusing ICMP redirects and router advertisements (irdpresponder from the Phenoelit's IRPAS would do a good job), exploiting weaknesses in link state routing protocols or even BGP. While little can be done to prevent such attacks, since both the ISP gateway router and the network hosts are outside the tested administrative/legal domain, measures can be taken to evaluate the resilience of the company gateway/traffic to these threats. These measures may include:

- evaluating traffic leaking out through the gateway and giving away unnecessary information about the internal network. Such traffic may include DHCP, STP, routing protocols updates, NTP, HSRP and probably many other protocols one can think about.
- testing if the gateway will accept and follow illicit ICMP redirects and faked routing updates
- checking the strength of ciphers used to encrypt bypassing traffic when security protocols are in use
- auditing the resilience of security protocols (PPTP, SSL/TLS, SSH, IPSec) implementations employed against man-in-the-middle attacks (dsniff, ettercap, omen).

To test the audited gateway/network for these and similar security flaws we have proposed the so-called “semi-remote assessment”. It may involve at least two different approaches. In a more of a testing lab-style setup an auditor can plug a laptop loaded with all necessary tools into a spare WAN port or an equivalent on the corporate gateway. Alternatively, a redundant stand-by gateway can be assessed instead. A black box (and potentially – black hat) variety of the semi-remote testing can target servers in data centres by trying to buy and (ab)use rack space on the same network segment with the targets. That will inevitably require some social engineering effort which is well worth it. Apart from the fact that checking out resilience to “close range” or “single hop” man-in-the-middle and other attacks becomes possible, portscanning etc. can be performed at a much greater speed, saving time for both the auditor and the client. Besides, auditing the firewall ACL rules including RFC 1917 IP's filtering testing with client-server tools such as Firewall Tester becomes far more precise and efficient. Of course, all kinds of semi-remote testing can be only performed on-site.

Walk the walk but talk the talk

Coming to an end of the discussion about different benefits of outsourcing IT security audits locally, we have to accent the so-called “social element” of penetration testing (and no, we aren't talking about social engineering this time). This element is needed to fill the gap between the dryness of the official report and flexibility of human contact, allowing greater interaction between penetration testers and client system administrators/other personnel in a due course of the audit, as well as afterwards. It does not necessarily have to apogee in one of the local pubs, but it may well do. We do not question the importance of the report that describes methodology, procedures, actions, vulnerabilities, conclusions and recommendations given as the ultimate outcome of the penetration test. However, it is not possible to fit everything into a limited number of pages. The information in the report is usually condensed and prepared on the assumption that the client company's IT personnel has sufficient knowledge of information security to understand, analyse and act upon the findings. It should never take the form of “apply that patch” or “update to a next version of the software”, rather if the system administrator doesn't fully understand a paragraph, s/he should always be able to ask for additional explanations. And again, the personal contact between penetration testers and client company IT team is essential for such communication efficiency.

To summarise, before thinking of outsourcing your external security audits, analyse the arguments provided in this discussion and, perhaps, you will turn your attention to the folks next door who may charge a bit more but offer indispensable services that no remote security professional can provide doesn't matter how skilful he or she is. Also, think of scalability and future requirements and plans. You don't want to have different companies performing external, internal and wireless audits, do you?