WESTPOINT
internet reconnaissance services

External Security Testing

Introduction

This document discusses the requirements for independent and objective security testing of Internet facing infrastructure; describes some of the common approaches to external security testing; contrasts their relative merits; and provides some guidance on selecting an appropriate strategy.

Scope

The scope of this document is constrained to the security testing of Internet facing IT infrastructures typically associated with modern day penetration testing and active vulnerability assessment. It does not include discussion of the assessment of the physical security controls protecting IT infrastructure or passive techniques used to determine system vulnerabilities.

Audience

The target audience for this document is considered to be technical management or those interested in the differences between security testing services and the benefits each are likely to provide an organisation.

Drivers for Security Testing

The drivers behind organisations engaging third parties to perform security testing of their Internet facing IT assets are specific to that organisation's business and the risks it perceives. Some common drivers include:

- Compliance with
  - Legislation and enforced regulation, such as the Payment Card Industry's Data Security Standard; Financial Services Authority regulation; Sarbanes-Oxley; Basel II Capital Accord; Turnbull; Gramm-Leach-Bliley; Health Insurance Portability and Accountability Act; other country specific privacy and data protection legislation.
  - Adopted organisational operating practices and standards, such as ISO13569, ISO17799/ISO27001, ITIL, ISF's Standard of Good Practice for Information Security; public sector guidelines such as the UK government's GSI code of connection; and internally created standards and policies.
- As part of effective (corporate/IT/security) governance processes and the need for expert independent and objective assessment.
- Brand protection and the avoidance of adverse media coverage.
- Revenue protection and cost control
  - Minimising the risks related to availability and responsiveness of on-line services
  - Avoiding costs incurred as part of recovery from a breach
  - Reducing the likelihood of direct and downstream legal liability
- Protection of intellectual property and competitive edge.

There are costs in complying with legislation, regulation and protecting an organisation's IT assets: in the case of legislation and enforced regulation such costs, although not optional, should provide value for money. For organisations where there is less legislation or in less heavily regulated industries the investment they make in protecting their assets must be cost effective, in which case the combined costs of its countermeasures must be less than the cost of the impacts of any breaches it may suffer.

For many organisations this can be a difficult figure to estimate. For instance, what would be the monetary impact of poor publicity? It could be very little or knock tens of millions off a company's

market capitalisation. A more detailed view of the losses an organisation may suffer can be determined by performing business impact assessment. One of the best sources for determining impact costs would be the organisation's own history: have incidents previously occurred and, if so, what were the impacts. Other factors can assist in this process too, such as knowing what similar organisations are doing, or the findings of independent surveys. According to the 2005 CSI/FBI Computer Crime/Security Survey (up to 699 respondents interviewed) 56% of respondents had suffered a security breach in the previous twelve months; 95% of respondents had experienced more than 10 website incidents; and the average financial loss per organisation was $204,000. According to the 2004 NHTCU Survey (up to 201 respondents interviewed) 84% had experienced 'high tech' crime; the average financial impact was £1,200,000; and the average financial impact of a website breach was £86,000.

Reasons Behind Exposure

Threat Sources

Typical threat sources faced by an organisation's Internet infrastructure include automated worms and viruses; malicious hackers; organised criminals; (industrial) spies; and terrorists. A problem for all organisations is that the motivation behind these threat sources range from the highly targeted, such as the industrial spy trying to gain competitive advantage, to the casual opportunism of a hacker, to the highly indiscriminate of the computer worm. In the cases of the hacker and computer worm it becomes very difficult for an organisation to assert: "I'm too small/insignificant to be a target", or question: "Why would anybody bother with me?".

Once an organisation chooses to offer services via the Internet it will generally have very little, if any, control over the threat sources present within that environment. However, it is often able to exert some control over the vulnerabilities these threat sources exploit and hence protect itself.

Vulnerabilities

Broadly speaking a vulnerability is a weakness in a system that if exploited by some threat will realise a risk that can have adverse business consequences.

In 2003 Gartner produced a paper entitled Taxonomy of Software Vulnerabilities which, to paraphrase, categorised vulnerabilities as:

- Design flaws – the system operates as intended but the design was remiss. For example, an application designed to service internal users is subsequently deployed into a different environment servicing external users (a web front end bolted on to a legacy mainframe application and rolled out to a different population of users).
- Implementation flaws – the system does not operate as designed, exhibiting behaviours not specified. For example, a coder puts a back door into an application; poor coding leading to the existence of buffer overrun exploits; software prototypes 'becoming' release code due to the pressure of delivery deadlines.
- Operational flaws – incurred as a result of system installation, configuration, operation or administration. For example, systems offering unnecessary network services or loose access controls; the degeneration of firewall rule-sets over time; simple configuration errors; the difficulties of managing an effective patch management program in a heterogeneous environment.

Gartner forecast that 65% of successful attacks would be attributable to avoidable operational flaws – a projection qualified by Carnegie Mellon University who report that greater than 95% of intrusions result from exploitation of known vulnerabilities or configuration errors where countermeasures where available.

In order to gauge the volume of vulnerabilities discovered in computer software CERT records the number of unique vulnerabilities reported to it each year. In 2006 the number of vulnerability alerts issued by CERT rose to 8064, up by 2074 from 5990 in 2005, a 34% increase. These figures are for all software platforms and a typical organisation will not be running all software platforms, certainly not within its Internet presence. How many vulnerabilities will an organisation have to manage on an ongoing basis? There are many factors that affect such a calculation but consider the following: suppose an organisation's Internet facing presence is only affected by 0.5% of CERTs vulnerability alerts - that still equates to three relevant vulnerabilities per month. In addition to this, vulnerabilities arising from operational flaws and in-house developed applications etc. still need to be considered. The important point here is not so much how many vulnerabilities affect an organisation as the realisation that managing vulnerabilities is an ongoing process and not a one-off activity.

The first step in managing vulnerabilities is discovering/confirming their presence. Once they are known a remediation strategy can then be developed, implemented and its effectiveness monitored. Ongoing security testing assists with the discovery of vulnerabilities and measuring the success of an organisation's remediation strategy.

Approaches to Security Testing

Security testing the networks, systems and applications that compose an organisation's Internet presence can range from running an automated vulnerability testing tool to a penetration test and full source code audit[1]. Security testing approaches can be conveniently split into three categories:

- Automated vulnerability assessment.
- Blended vulnerability assessment.
- Penetration testing.

Regardless of the security testing approaches adopted by an organisation they should all be able to provide a similar set of benefits, albeit dramatically ranging in the quality and depth of their findings:

- Independent and objective assessment.
- Discovery of the latest security vulnerabilities.
- Verification of effectiveness of remediation activities.
- Verification of effectiveness of other security controls and processes (e.g. incident response procedure, technical security policies, etc).
- Satisfies regulators, auditors and stakeholders.
- Assists with demonstrating effective risk management.

Unfortunately there are no real standards available for vulnerability assessment although having it is a requirement for organisations that need to adhere to the Payment Card Industry's Data Security Standard. The PCI Security Standards Council[2] maintain a list of "Approved Scanning Vendors" whose services meet their scanning requirements. The situation with penetration testing is a little more mature. For instance, the UK Government's CESG issue a certificate to vendor's staff capable of performing network penetration tests to their CHECK standard. The recent launch of CREST (Council for Registered Ethical Security Testers) has resulted in the production of security testing standards and guidance with a view to regulating the competency of its members. Other security testing guidelines (not certifiable standards) are available from the Open Source community:

- OWASP, The Open Web Application Security Project[3] provides information and tools regarding the development and testing of secure web applications.

[1] Source code auditing is considered beyond the scope of this paper.
[2] https://www.pcisecuritystandards.org/
[3] http://www.owasp.org/

- OSSTMM, The Open Source Security Testing Methodology Manual[4] provides broad guidance for security testing.

Security testing does have some implications that must be managed by the organisation. Common to all the testing approaches is that they will run port scanning tools to detect what services are being offered by a system (e.g. web, mail). Scanning a system for the presence of all such services can generate a fair amount of traffic and connection requests, so testing will consume some network bandwidth and system resources. Before taking a security test consider the following guidelines that may prove helpful:

- Ensure you are clear about the scope of the security test and, especially with a penetration test, if any destructive testing is to occur.
- Ensure you are satisfied with the credentials, experience, standing and processes of the vendor.
- Ensure the testers have a methodology to which they adhere. This ensures consistency of results and is particularly important where manual testing occurs (blended vulnerability assessments and penetrations tests).
- Ensure testing occurs over a bandwidth-limited link and your infrastructure can accommodate the increase in traffic.
- Be careful with older systems and legacy applications that potentially may not be able to handle a security test. If this is the case have an initial test scheduled at a non-critical time or during a maintenance window.
- If all connections are proxied or address-translated through a security gateway ensure it has sufficient capacity to cope with the traffic, volume of connections, and any state it may have to maintain.
- If you have any intrusion prevention or port scan blocking in place[5] ensure the security testers' source IP addresses have been added to any whitelists.

Automated Vulnerability Assessment

What is it?

Automated vulnerability assessment consists of running software tools to:

- Discover what services are being offered by a system, e.g. web, mail, etc.
- Discover what vulnerabilities are present in 'off the shelf' applications bound to those services and recommend appropriate remediation.

This type of assessment is, as the name suggests, totally automated and other than perhaps starting and stopping the process has no manual involvement. The tool used can be an open source one, or a commercial offering, or it can be offered by an online security service provider driven via a web interface, in which case the tool often becomes an 'engine'[6]. The onus of running the tool or engine falls to the end user. However in the case of the latter the service provider will ensure the engine has the latest bug fixes, security tests, maintenance updates, etc.

Reporting varies depending on the vendor but is virtually all computer generated.

Merits and Drawbacks

- Commoditised, and hence low cost. Run frequently.
- Quick turnaround of results.
- Generally a 'black box' test.
- Non-destructive/non-intrusive.

[4] http://www.isecom.org/osstmm/
[5] and the purpose of your test is not just to assess its effectiveness
[6] marketing-speak for a tool.

- Can exercise other security controls.
- Accommodates vast numbers of systems.
- User drives tool/service.
- Prone to false positives.
- Generally reliant on a single vulnerability testing tool/engine.
- Unable to detect 'contextual' vulnerabilities (e.g. if information is 'sensitive').
- Poor at detecting certain classes of vulnerability, such as SQL injection, Cross-Site Scripting, or vulnerabilities in bespoke applications, etc.
- Geared towards standalone reports of systems and hence little, if any, historical context to a system's prior vulnerability/remediation status.

Blended Vulnerability Assessment

What is it?

Blended vulnerability assessment is a service, not a stand-alone application. It consists of running one or more security testing tools/engines and additional phases of manual testing and vulnerability verification to discover:

- What services are being offered by a system, e.g. web, mail, etc.
- What vulnerabilities are present in the applications bound to those services and recommend appropriate remediation.

The main difference between blended and automated vulnerability assessment is the ability of the blended assessment to further manually eliminate false positives from reports, detect vulnerabilities beyond the scope of automated testing (e.g. those in bespoke applications or ones where context is relevant). The amount of manual testing effort is much less than found in penetration tests yet is sufficient to discriminate itself in the market as a cost effective and worthwhile service.

Quality of reporting can vary between vendors: some go for a penetration test-like written report whilst some go for an html archive offering the ability to order results in various ways, track system vulnerability histories, group results according to organisational (or other) structure, metrics and trending.

Blended vulnerability assessment is both cost effective and well suited to organisations that have a very large Internet presence and many bespoke applications. For such organisations this technology can summarise/prioritise vulnerabilities and highlight risk 'hotspots' across the whole organisation. On this scale penetration testing can be prohibitively expensive.

Merits and Drawbacks

- Non-destructive/non-intrusive.
- Service provider drives service.
- Crystal- or black-box test.
- Improved result quality and fewer false positives – due to phases of manual testing/analysis and lack of reliance on single test tool/engine.
- Good at detecting virtually all classes of vulnerability, such as SQL injection, Cross-Site Scripting, or vulnerabilities in bespoke applications, contextual vulnerabilities, etc.
- Ability to have tests, and in some cases reporting, customised.
- Remediation advice can contain customised information, e.g. URLs or comments, to better illustrate and explain issues.
- Accommodates vast numbers of systems.
- Can exercise other security controls.
- Low to medium cost. Run frequently.

- Testing takes longer than an automated vulnerability assessment.

Penetration Testing

What is it?

Penetration testing is different from vulnerability assessment in many ways. Rather than just supplying information about target IP addresses, ranges and hosts it is more likely that a scoping exercise will be conducted between the organisation and testing vendor to determine, for example, which portion of the infrastructure or its applications will be examined; the amount of time to be spent testing; how invasive/destructive the testing may be; the business value of the components under test; and a host of other requirements/limitations.

Penetration testing also has similarities with vulnerability testing in that the penetration tester will run a number of automated tools to discover the configuration of networks/systems/applications and their potential attack vectors. However once this phase is complete the rest of the test will predominantly consist of manual effort with the penetration testers using their knowledge and experience to uncover security holes. With blended vulnerability assessment the amount of time spent on manual investigation may be constrained to a couple of hours per system. With penetration testing this time can be of the order of a couple of days per system.

Once security holes are discovered penetration testing again diverges from vulnerability testing and further examines how far into the network/system/application leverage/exploitation of the initial security hole will allow the tester to progress. If subsequent, 'deeper', security holes are found these will similarly be leveraged with the penetration tester often capturing 'evidence' of the depth to which the network/system/application could be penetrated. Production of such evidence often provides shock value resulting in the organisation's senior management releasing the resources necessary to rectify problems and improve the target's security.

As the bulk of a penetration test consists of manual effort the extent of its findings can be very dependent on the experience and quality of the individual tester and the amount of time they are able to spend on the investigation. This is why it is important for penetration testing vendors to have a testing methodology; experienced and, ideally, qualified testers; and recognised security testing credentials and support procedures (e.g. secure handing, management and destruction of the data obtained during their test).

One issue to consider before taking a penetration test is that if an organisation's infrastructure contains fundamental security holes most of the subsequent findings of the penetration test may well be of little value, or even invalid.

Penetration test reports are usually delivered as written documents complete with executive summaries, impact scenarios, technical detail and appendices of low level tool output and findings.

Merits and Drawbacks

- Generally intrusive.
- Destructive or non-destructive.
- Crystal- or black-box testing.
- Best result quality in terms of lowest incidence of false positives/negatives and greatest number of security issues uncovered.
- Ability to uncover security issues in networks/systems/application behind the perimeter.
- Ability to uncover all classes of security vulnerability through crafting of bespoke exploits.
- Potential to detail business impacts and costs of intrusions.
- Shock value of findings.

- High cost. Run occasionally and hence more of a security/risk 'snapshot'.
- Geared towards smaller numbers of systems/installations.
- Findings dependent on quality of tester.
- Less useful if the organisation has fundamental security issues.
- Generally standalone and does not focus on target history.
- Longest test time.

Conclusions

Organisations considering security testing have a number of options available to them to assist their risk assessment efforts. Undoubtedly these will be constrained by their tolerance towards risk and their budget:

- Very little budget:
  - consider automated or blended vulnerability assessment.
- Limited budget, only enough for one testing approach:
  - consider blended vulnerability assessment.
- Some budget:
  - consider blended vulnerability assessment with annual/bi-annual/quarterly penetration tests.
- Unrestricted budget/security is of paramount importance:
  - consider monthly penetration tests and near-continuous blended vulnerability assessment.

The third option, "Some budget", is an effective strategy combining the detailed 'snapshot' assessment of the penetration test with regular blended vulnerability assessment, thereby providing ongoing and detailed assessment in a reasonably cost effective manner. If this approach is considered a favoured candidate then the vulnerability assessment should be conducted first. This will result in discovery of fundamental vulnerabilities and, if they are fixed, this will, in turn, ensure (i) penetration test effort is not wasted in the discovery of basic problems; and (ii) penetration testers will be focused on the more complex security issues not discoverable by vulnerability assessment.