Security Vulnerability Management

Reliable prioritisation is needed to reduce the business exposure

Mike Barwise

Protecting business IT against the exploitation of operating system and application vulnerabilities is a sad fact of life, as security bugs continue to surface in software on a regular basis. However, in a large-scale diverse IT environment the workload can become very high, and it is essential that those vulnerabilities presenting the greatest business exposure are dealt with first. A reliable prioritisation method is therefore necessary, but at first sight it seems difficult to achieve. From within the IT department it is easy to ignore the bigger picture; that there are two components of business exposure. The first and most familiar (technical exposure) is a combination of the absolute exploitability of a vulnerability and its technical impact on systems. The second, and ultimately more significant, (operational exposure) is the resultant collateral damage to the business, which may be manifested at more than one level. Traditionally, as patching has been managed entirely within the IT department, technical exposure has dominated thinking on vulnerability management, but as the criticality of IT to business processes increases it becomes ever more necessary to consider operational exposure when prioritising protection.

Numerous vendors have developed proprietary vulnerability ranking systems, but these have some key failings in common. Firstly, they all address only technical exposure. Secondly, individual vendors' systems are based on differing criteria, and the resultant rankings may therefore not be comparable across vendors. Thirdly, none of them are very granular, usually consisting of three or four categories on a verbal scale such as "low" through "critical". Priorities are, therefore, difficult to establish in the field as there will often be numerous supposedly "critical" vulnerabilities reported by different vendors simultaneously vying for attention. We may call this situation a "collision".

The Common Vulnerability Scoring System (CVSS) was introduced in 2005 by the US-based Forum of Incident Response and Security Teams, representing government, corporate, and vendor interests. Since 2007 it has been completely revised. CVSS was intended to address the specific failings of proprietary vulnerability rankings by providing a single consistent, vendor-neutral ranking system. For any vulnerability, it offers three rankings: a "base score" that describes the vulnerability in abstract (for example, does it require authentication to exploit, can it be exploited remotely); a "temporal score" derived from the base score that contributes additional factors (such as whether exploits already exist and whether patches have been issued); and a method for end users to calculate an "environmental score" that attempts to introduce factors relating to local impact.

Values of the variables that feed into the base and temporal scores are specified in relatively unambiguous natural language, referenced to lookup tables of discrete numerical values that are applied to the equations. The base and temporal scores will normally be calculated at source by whoever announces the vulnerability. The temporal score is probably the most immediately familiar, as it is functionally equivalent to the best of extant vendor rankings. However, it has the additional advantages of being numeric and quite granular.

The CVSS approach to environmental scoring is less satisfactory. In the first (2005) version, the environmental score calculation was simple to perform but rather primitive. It merely multiplied the temporal score by two locally-derived factors: "collateral damage potential" referred to a four point scale of "none" through "high", and "target distribution" on a quantised scale referred to percentage of systems affected. Each of these factors has major weaknesses. The guidance provided on estimating collateral damage referred only to "property damage or loss", which failed to allow for impact on business processes. Equally, "target distribution" as defined is a poor criterion for operational exposure, as it completely fails to take into account the existence of small numbers of specialised systems providing business-critical services within a large physical infrastructure.

The CVSS was revised in the summer of 2007. All three equations were rendered more complicated, and practically all the numerical input values were modified. Some improvement in the base and temporal scores has resulted in an observed reduction in residual collisions, but the revised environmental equation retains the failings of its predecessor and is also now much more difficult to apply. The scale is now on five levels, but the calculation process is much more involved. Where previously the published temporal score could be fed by an end user into a simple formula to derive the environmental score, it is now necessary to recalculate the base and temporal scores locally, and then feed the revised results into the environmental calculation. To render this possible, published CVSS rankings now need to be accompanied by a vector containing the values of all the

© Photos.com

parameters contributing to the original base and environmental scores. The guidance has also been updated, now stating that the environmental metric "... measures the potential for loss of life or physical assets through damage or theft of property or equipment. The metric may also measure economic loss of productivity or revenue", and that "... each organisation must determine for themselves the precise meaning of "slight, moderate, significant, and catastrophic"

Of course, the determination of these precise meanings is the most critical and potentially complex component of vulnerability management. So given a consistent technical vulnerability ranking such as is provided by the CVSS temporal score, how is operational exposure to be factored in? Clearly, an operational exposure index has purely local significance. It must be based on the actual exposure resulting from compromise of the systems that serve the business. That requires accurate recognition of the significance to business processes of IT assets and the criticality to the enterprise of those business processes themselves. The method for determining operational exposure must be well-defined and its results must be consistent. It must include consideration of both the direct and indirect business losses that would result from the asset being compromised. Responsibility for managing the development of the criteria, exposure scale, and calculations for specific assets should be borne by the CIO rather than by the IT department, although input from IT Management will be required, not least to correlate business functions and services with physical assets. Business Management must contribute to the determination of the actual operational exposure values. Input will also be needed from accounting, actuarial, and business process representatives. This valuation process is a not inconsiderable operation to undertake *ab initio*, but once it has yielded its results, the process of maintaining them up-to-date becomes an easily manageable background process. The goal is a numeric value assigned to each physical asset, defining its overall criticality to the business. Where a single physical asset serves more than one business process, the operational exposure values should be summed.

Because it is a continuously recurring task, the process of factoring the resulting operational exposure values into vulnerability rankings must be extremely simple, quick, and transparent. The final operational exposure scale should therefore be designed to accommodate this: for example, by referring calculated operational exposure values to a relative linear scale of zero through one, by which the CVSS temporal score may simply be multiplied.

Appropriate process management is also critical. To make an environmental metric work in the long-term, there must be well defined and functional communication between business management, IT management, and the board via the CIO. Information on changes to infrastructure, business processes, services, and liabilities must all be pooled to keep the operational exposure metrics up-to-date. There may also be a need for culture change at the coalface. The long-established "low", "medium", "high" technical exposure rankings tend to encourage processes whereby a time limit is assigned to each category, patching being considered successful if a given vulnerability is patched within the time specified for its ranking. This approach, although it benefits workload management by introducing an element of predictability into the task, is becoming increasingly fragile as a real protective measure as the threat space hots up. Vulnerabilities ranked "low" may well be exploited before their scheduled remediation comes round. Surprisingly, given the high granularity offered by the CVSS, its own documentation proposes such a reduction to four time-determinate priority levels. A much more satisfactory approach that takes advantage of the granularity intrinsic to the CVSS is to order outstanding vulnerabilities by rank on a continuous basis, and address them from the top of the list as swiftly as possible even if their absolute rankings are all low.

Of course, the remedy will not always be a patch. With increasing frequency, the availability of patches lags significantly behind the announcement and even the exploitation of vulnerabilities. It is, therefore, necessary to consider alternative remediation, whether this be platform- or application-specific workarounds or controls elsewhere in the infrastructure, such as modifications to access rights or protocol blocking, or even, in a worst case, service or system quarantine. This requirement implies close integration of vulnerability management into operational IT management at least at the tactical level. Indeed, a significant contribution to security can be made at the architecture level by ensuring that business services are segregated as far as possible from each other and from the outside world without impeding their business function.

At a glance

- Mike Barwise is an independent consultant specialising in information security strategy

Mike Barwise
Independent Consultant
m.barwise@bcs.org.uk