

Domain Based Security – From Soft and Goopy to Hard and Crunchy

Abstract:

Either consciously or unconsciously, most companies have employed a domain based security approach to the design of their information systems architecture. Corporate information systems generally comprise of one or more trusted domains - usually the office networks - with a firewall providing a barrier to the untrusted domain - the internet. This approach was promoted in the mid to late nineties as 'hard and crunchy on the outside and soft and goopy on the inside', implying that if your firewall provided sufficient strength at the gateway then there was little need for protection on the internal corporate networks. Recent reappraisal of the risks to corporate information systems suggests that this approach is partially flawed and that the concept should really be 'hard and crunchy on the outside *and* hard and crunchy on the inside too.' This paper discusses why this may be desirable and identifies what makes a network 'hard and crunchy'.

1 Introduction

Prior to the uptake of the internet, corporate information systems were built as standalone systems incorporating Commercial Off-The-Shelf servers and workstations. As the internet revolution prompted many companies to connect to the internet, a range of unexpected security concerns arose, the most widely publicised being that of hackers on the internet and the devastating effects of viruses. In an attempt to combat these threats, many companies introduced barrier devices (routers and firewalls) and deployed anti-virus systems.

In the mid to late nineties the concept of Domain Based Security became popular; whereas traditional information systems architectures revolved around the requirements for general interconnectivity, Domain Based Architectures focused on the information exchange requirements between identified 'domains'. Typically, a company would primarily recognise two domains, the trusted internal network, and the untrusted internet.

The information exchange requirements between these two domains were initially email and web browsing, supported by domain name services. Very quickly a third domain was recognised, that of the De-Militized Zone, a space that was not entirely trusted, nor entirely distrusted. Web, email, and domain name servers were situated within the DMZ to act as an intermediate step between the internal network and the internet. Once an organisation had introduced a firewall to control the connections between the three domains, the concept of the 'hard and crunchy' exterior had been implemented¹.

The protection offered by the firewall generally provided a sense of security against the hackers on the internet. As it was felt that the systems were secure, little additional security appeared to be required internally. Most organisations simply employed the controls offered by the operating systems, such as passwords and access control lists. This concept within domain based security has become known as the 'soft and goopy' interior.

¹ The firewall would typically allow tcp ports 25 and 80 (email and web), and udp port 53 (dns) between the DMZ and each of the internal and internet domains; no traffic would typically be allowed to pass directly between the internal and internet domains.

² To complete the 'hard and crunchy' exterior, the servers in the DMZ would be clamped down and have all security patches applied.

Recent media reports of internal hackers and the increasing complexity of malicious software (viruses, trojan horses, and worms) imply that the 'soft and gooey' interior may not be sufficiently protected by the 'hard and crunchy' exterior and that things need to change. As domain based security matures, systems may need to become 'hard and crunchy' throughout.

2 Threats and Risks to Information Systems

There are many threats to information systems; this paper will only discuss network based technical threats. There are four distinct threat groups that pose risks to the internal network:

1. **External hackers** – these are people who attempt to break into corporate systems via the internet, via modems, or via wireless networks. Whilst firewalls provide good protection for the internal network, they do not provide absolute protection and it is always possible that a hacker may breach them either directly, or by leap-frogging via the DMZ. Modems provide remote dial-in access for employees which often provide password-controlled access to the internal network; again, these may be breached by a skilled or lucky hacker. Wireless networks often suffer from inadequate encryption making them additional entry points for hackers³.
2. **Malicious Software** – these include viruses, trojan horses, and worms. Due to the mechanics of email replication, most of these need to reach a recipients' desktop on the internal network prior to their execution. If the anti-virus software is out of date or if the malicious software is not listed in the signature files then there is a likelihood that these can be executed on the internal network. As malicious software becomes more complex and devastating, it is likely that new strains will begin to actively attack the internal networks. One virus tomorrow could take an entire system off-line through a Denial Of Service attack.
3. **Disgruntled (or bored) employees** – these are legitimate users of your internal network. If they wish to attack your network, either out of interest or because of a complaint, then it is likely that they will have access to all the tools they will need to gain superuser status. For most companies it takes at least one internal breach before the systems and rules are tightened up.
4. **Visitors, contractors, partners, cleaners** – these are people with legitimate access to your offices but not necessarily with legitimate access to your systems. If any of these people turn against your company it is likely that they will have the physical access to your network required to mount an attack. As with the rise of the script kiddie, your potential attackers no longer need advanced skills to cause havoc; they simply need a small laptop, some free tools, and network access.

The 'soft and gooey' interior of many corporate systems will not counter the threats posed by the above attack groups. Internal systems that are not patched regularly, that contain default installations of applications and services, that include 'trusts' to other systems, that make use of 'simple' passwords, that use clear-text protocols, or that rely on inadequate security products are all at risk from these potential attack groups. Whilst there is always the option of managing or accepting these risks, a little thought and a little work can help to mitigate them.

3 Mitigating the Risks to the Internal Systems

If you consider that your legitimate users – your own employees – are a threat to your systems, either directly through intentional hacking or indirectly through unintentional execution of malicious software, then the only way to protect your systems is by removing the

³ Wired Equivalent Privacy (WEP) does not currently provide adequate encryption as the encryption key is usually computable in less than a day, often in a few hours.

vulnerabilities that they may exploit. Whilst this may seem obvious, many companies are loathe to consider their employees as a potential threat citing either that if they want to damage the company there are far easier methods, or that employees have to be considered as trustworthy as it is too difficult to secure the systems. Both of these assumptions are flawed as research has shown that inside attacks are responsible for a large number of security breaches, and that securing systems is possible with the right skills and desire.

So how does one secure the internal systems? There are four basic elements required to secure a system: the first is a solid, secure design; the second is the maintenance of security patches; the third is good system security management; and the fourth is regular, trusted, third-party testing.

Patching systems is seen as being relatively straight forward: most weaknesses in common applications and operating systems are reported to central internet bodies⁴, with patches following shortly after. One should not underestimate the amount of regression testing required to apply patches, however, as it is always possible that a patch may stop a system working, effectively resulting in a denial of service. Strong system management practices can mitigate the risk of a patch damaging a system through good configuration management and the maintenance of a reference rig (a system that mimics the live, operational system).

The creation and maintenance of a solid, secure design is a little more complicated. Many internal systems are built to support business functions, often with little knowledge of the business requirements beyond the provision of office tools, email and web browsing. Because of this, many systems administrators do not clamp-down installations and remove software and services that are not needed in case a business function relies on them or later claims to require them. During the lifetime of a system, many more requirements will be placed upon it, often resulting in a vast quantity of software, using protocols of questionable security, and system-to-system trusts in support of the panacea of Single Sign-On (SSO).

The first step to securing an operational system is to remove the services that do not appear to be required, patch those that are required to the latest security patch, and then hire a trusted, third-party to test the system for vulnerabilities⁵. The testers should be able to identify all the vulnerabilities remaining and provide guidance on securing them. It is important to use an external company for these tests for an unbiased report, free from internal politics. Once the remedial action has been taken the system should be re-tested to confirm that the system is now secure.

Maintaining the technical security of the system is then a matter of maintaining the latest security patch level, and regularly testing the configuration. These repeat tests can be performed in-house with the right training, with an external company only required to perform an audit. The audits are essential to ensure new and existing measures are tested correctly.

Good systems security management is required to maintain operational security; this includes accounting and auditing, user account management, the concept of least privilege, and regular backups. To perform these activities well, the administrators will require specialist skills and the benefits of a dedicated security management course should not be underestimated.

⁴ Good CERTs (Computer Emergency Response Teams) and lists include UNIRAS (the UK CERT, www.uniras.gov.uk), securityfocus.com, and bugtraq.

⁵ There are numerous Penetration Test Companies, offering services that vary in quality. In the UK, the only recognised approval scheme is CESG CHECK (www.cesg.gov.uk). CHECK is often mandated for UK Government departments and ensures that the testers meet a baseline standard.

4 Addressing The Business Requirements

The most important element of any system is its ability to meet the business requirements in a secure fashion. The balance between capability and security is often tricky to negotiate as security is often seen as a reduction in capability. This does not necessarily need to be the case if a risk management approach is taken. When one considers that the only truly secure systems are those that are switched off, then one accepts that every operational system carries with it some level of risk. If the levels of risk associated with different capabilities can be quantified then it is possible for the system owners to make informed decisions about the level of risk they wish to accept for any given capability.

In order to develop systems capable of meeting the business requirements one needs to fully understand the business processes that are to be supported by the systems. These processes can be characterised as 'Information Access Models' and 'Information Flow Models'.

Information Access Models describe in precise terms the people and departments that need to access the different types of information to be stored by any given system. In order to assess these requirements, one needs to first identify the different groups of people who will be using the system, then identify the different types of information to be stored, and finally identify the access profiles of groups of people to information. The resulting information can be depicted as a lattice or graph and should be agreed with the system owner. This lattice will form the basis of the Access Control Lists that will be implemented on the final system.

Information Flow Models describe the movement of information within a system, and between a system and the systems connected to it. Information Flows should include all email, web browsing, video conferencing, collaborative working, digital input/output procedures, and printing requirements. The task to create an Information Flow Model should not be too onerous if the business processes are well understood. If the business processes are not well understood then perhaps a preliminary task of fully documenting the business processes is required.

Once the Information Access Model and Information Flow Model have been agreed for a given system, it should be a relatively simple procedure to design a high-level Information System to meet the requirements. Each of the components required to support the elements of the business process can then be designed with options that provide differing levels of security for differing levels of cost and usability. The system designer can assess these risks using information gathered from the internet or by consulting a security specialist. If the risks can be described in business terms (i.e. the likelihood of a breach resulting from a particular component, weighed against the cost to the business should a breach occur, and considering the user implications of the particular technology) then a cost-benefit analysis can be successfully conducted by the system owner.

By using this process to design an Information System, the responsibility for deciding the level of risk that a business may wish to take will lie with the system owner. The benefits of this are two-fold: the system owner should be in a much better position to consider the risks that will be acceptable to the business; and the design decisions that put the business at risk are no longer the sole responsibility of the system designer, who may not fully appreciate the business implications of any risks in the design.

5 Designing Secure Domain Based Security Architectures

The definition of Secure Domain Based Security Architectures is 'hard and crunchy throughout'. This section addresses the design of these systems.

All Information Systems should be designed to meet the business requirements, as proposed in section 4. For new Information Systems this is simply a case of following the process, agreeing the risk, and implementing the final design.

It is recognised, however, that many companies will already have information systems in place and will not wish to completely replace them. It is possible to redesign the current systems with security in mind and perform a migration to new, secure designs. If one assumes that any information system that is currently operational meets the specified business requirements, and that any new, proposed system will be designed to meet the same requirements, then the task of migration may be performed one component, or business process, at a time.

The system should be redesigned as if it was a new system being designed for the first time. This should remove any legacy insecure aspects of the current design whilst justifying the presence of all the existing secure components. The new system design should be subject to the same cost-benefit risk analysis as a new system except the costs incurred will be that of changing or migrating as opposed to the creation or purchase of new components.

Once a system design has been agreed with the system owner, the design should be documented in terms of the services required, including all of the supporting services. This should lead to a system diagram that matches the business requirements against the Information Flow Model, and the underlying software services and applications. If this diagram is correct (some design iteration may be required) then only the services specified will need to be implemented, and each should only accept connections from the applications or services that require them. This will ensure that the design is driven by the business requirements and will not include any services that are not needed, fulfilling the requirement for service minimisation.

Migration from the current system to the new secure system can be performed one component at a time by testing the replacement component on a reference rig and then agreeing with the system owner the time-scales, potential downtime, and any new training requirements. The latest security patches for the new components should be incorporated on day one, such that upon completion of the migration, the system should not be prone to vulnerabilities in the implemented services.

Once implemented or migrated, the operational security will rely heavily on maintaining the latest security patch level, tested by regular penetration tests. The 'hard and crunchy' interior should then be complete, meeting the desired business requirements and carrying only the risks deemed acceptable.

The most important aspects of modern systems security are recognising the risks to a system and fully understanding them in the context of the business requirements. Whilst it is perfectly acceptable to manage these risks and continue to maintain a 'soft and gooey' interior, this is only possible if the risks are well understood. Once the risks are understood, however, I believe most companies will opt to migrate towards a 'hard and crunchy' interior to mitigate the risks rather than accept them.