Secure Network Zones

Peter Kai Wimmer

atsec information security GmbH
Steinstraße 70, 81667 Munich, Germany
peter.wimmer@atsec.com

.....

Abstract

Large networks, which are often distributed over physically separate locations, require a coherent security approach. This paper introduces the concept of secure network zones, arranged in "onion-like" layers, providing increasing security levels towards the inner, more secure zones.
Increased security is provided by both protective layers around sensitive networks and additional (cumulative) security measures, ranging from basic measures such as hardening and firewalls to more sophisticated techniques such as intrusion detection and encryption of transmitted and stored data.
The implementation of secure network zones is described, including classification of data, assignment of applications to zones, and data flow. A path for the migration of existing environments is discussed and recommendations for special use cases are provided.

1 Introduction

Internal networks used to be flat entities, separated only from the Internet by a single firewall. With the increasing use of electronic services such as e-mail and web, internal networks grew rapidly, and were – at least logically – separated along department boundaries. A growing awareness of the value of confidential information (e.g., design specifications, financial data) and of increased dependency on the electronic infrastructure led to implementation of further security measures, such as access control and strong authentication mechanisms.

Malware with the ability to spread without user interaction by replicating itself to other systems on a network (a.k.a. worms) imposes a tremendous threat to flat network structures. Humans, either acting as insiders (e.g., employees) or outsiders (“hackers”, spies, script kiddies, etc.), pose a similar threat. Thus, additional protection for networks against each other, as well as for systems within these networks, is required.

In addition, legacy applications that cannot be patched with security fixes due to restrictions from the vendor (e.g., support only for a specific configuration) or simply due to the sheer number of missing updates are considered a major risk for the whole environment. By separating such applications into a protected subnet, both the risk of being compromised and the impact of compromise on the surrounding systems are efficiently reduced.

The secure network zones model provides a sophisticated and granular approach to protecting assets, focusing on *information* as the most valuable (electronic) asset. Sensitive data is surrounded by additional layers of protection, providing both network and logical security measures such as access control, confidentiality protection, and intrusion prevention. The segmentation of networks efficiently restricts vulnerabilities and the associated threats and risks to a

limited environment. The secure network zones model implements *defense in depth* through its layered model and *diversity in defense* through an adequate protection profile for each zone.

The author originally developed a secure network zones model in 2003, and some similar approaches exist. However, few related papers focus on data security and restrictions for the traversal of zones. Indeed, most papers proposing an extended zone model for secure networks discuss “classic” network separation, e.g., isolating the server LAN from client LANs. This separation model also references “zones”; however, such zones usually represent a flat separation of subnets with differing security requirements.

The work of Bell and La Padula, as well as the “hierarchical protection domains” model (see chapter 5, Related work) for protecting data and resources, serve as foundation for the secure zones model introduced in this paper.

2 Architecture

A “secure network zone” is a dedicated network segment, with a well-defined communication flow to other zones and implementing specific security measures.

Three internal zones are defined, and these zones are nested, i.e., the innermost zone is protected by the surrounding zones (see **Fig. 1**). An internal zone may consist of various subnets and even span several locations.

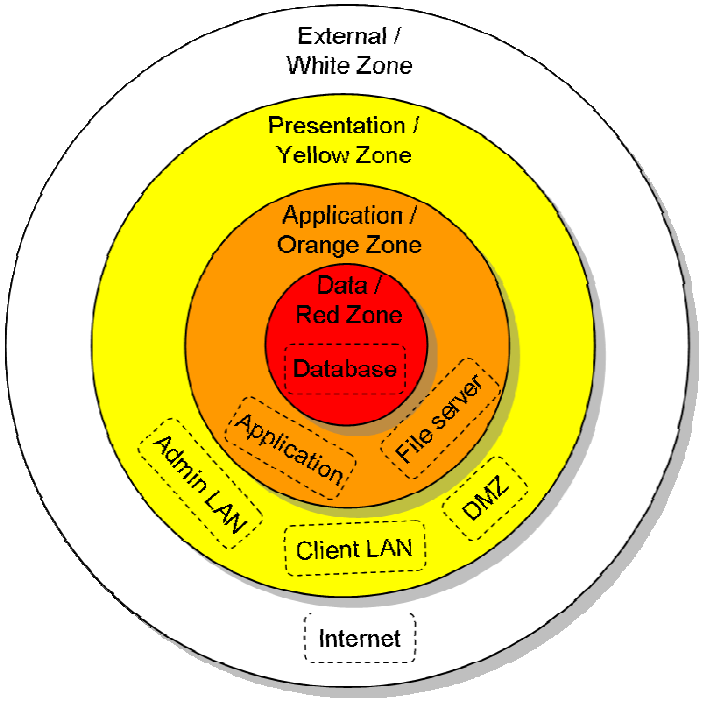

Fig. 1: Secure network zones

An external zone represents all networks that are not controlled by the organization; this includes not only the Internet but also "attached" networks of outsourcing partners, suppliers, and service providers.

Zones may be divided into segments (labeled “zone instances”, implemented as subnets), since further separation is often desired, e.g., for a DMZ and client LANs. This structure also

supports the protection of critical applications by isolating them from other parts of the network.

Availability and accountability are individually defined for each segment, independent from the data security level. For example, a web server providing information classified as "public" may require high availability, while an internal system with confidential data may only need to be available during business hours.

2.1 Classification

The focus of this approach is data security, i.e., the confidentiality and integrity requirements of information. Therefore, it is necessary to classify data according to the information classification guidelines of the organization. The three essential levels are *public*, *internal* and *confidential*. Further granularity usually may be reduced to one of these three levels. For information not (yet) classified, a reasonable approach is to consider such data as internal.

Although confidential data must be stored in the inner, most secure zone, a subset of this data is typically processed in an application in the zone "below", which in turn forwards part of this (processed) data to the presentation layer in the adjacent lower zone.

Classification also determines whether the confidentiality of information must be protected while being transferred or stored; such protection is therefore independent of the zone in which data is currently processed.

2.2 Users

Users are usually members of the organization, but may be customers, partners or anonymous users. Users may remain anonymous for certain services, while being required to authenticate for others, usually when access to restricted information or restricted resources is requested. Access to non-public data or resources (i.e., the orange or red zone) always mandates authentication and authorization.

Therefore, users are either authenticated or anonymous (see **Table 1** in chapter 2.3).

2.3 Assignment

The secure zone model postulates a three-tier application architecture, typically assigning databases to the inner (red) zone, applications to the middle (orange) zone and web servers, (reverse) proxies, etc. to the outer (yellow) zone (see **Table 1**).

Table 1: Zone and user level assignments

Zone / Layer	Typical contents	User level
Internet (white)	Public Customers Partners	Anonymous
Presentation (yellow)	Client LANs (Reverse) proxy Web server VPN endpoint DMZ	Authenticated

Zone / Layer	Typical contents	User level
Application (orange)	Applications Admin jump station	Authenticated
Data (red)	Databases	Authenticated

The location where information is actually *stored* determines the assignment to zones (see also 2.1).

Examples of common assignments:

- Web server providing public information only - yellow zone
- File server providing internal documents, such as guidelines - orange zone
- Database with confidential customer data, such as a CRM - red zone, and the application server - orange zone

Clients may be in either the yellow or white zones; for access to internal and confidential information, clients must be authenticated and authorized (see also chapter 2.4).

2.4 Data flow and zone traversal

Data flow is restricted to and from *adjacent* zones only, thus restricting direct access to sensitive data from networks considered insecure (see **Fig. 2**). This restriction applies both from outer to inner layers, as well as vice versa. Although the inner zones are considered "more secure" (actually, better protected), direct connection from an inner (e.g., database) zone to the Internet would allow a Trojan to transfer sensitive data without any further obstacles.

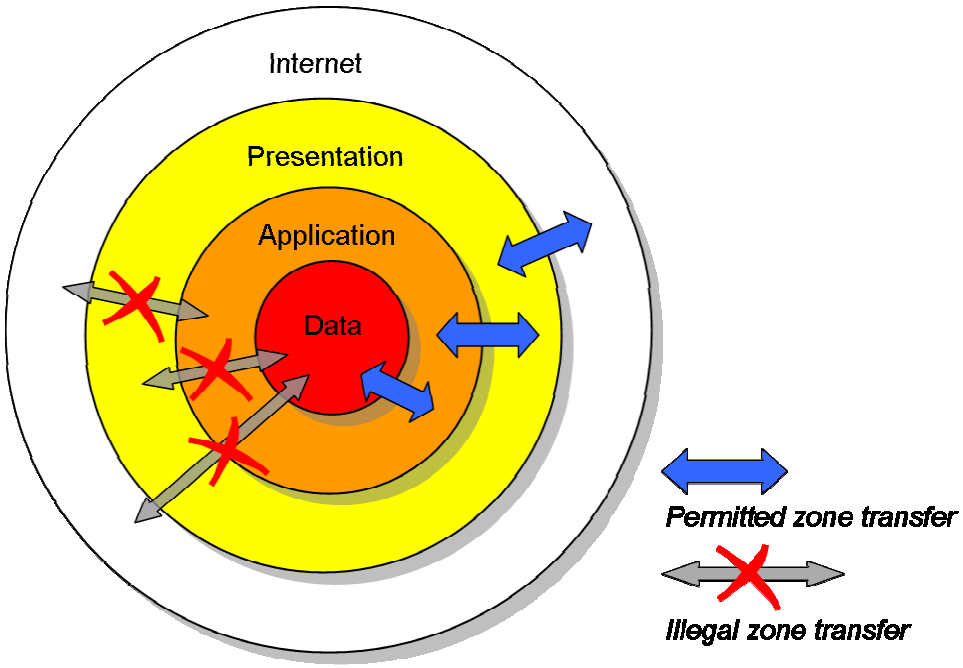

Fig. 2: Zone traversal

Any user accessing resources or data in the orange or red zone must be authenticated and authorized.

Applications in the orange zone may need to be accessed from the Internet, e.g., by members of the organization working in remote locations. A number of such applications do not pro-

vide a separate presentation layer due to the limitations of a two-tier architecture. As no direct access from the Internet to the orange zone is allowed, a gateway is required in the yellow zone (internal users are already located in the yellow zone). Depending on the frontend provided, such gateways include reverse proxies (for web services) or terminal servers for Windows applications.

For additional security, these gateways may implement content filtering and virus scanning, to prevent malicious code from entering higher zones.

2.5 Security measures

In addition to the firewalls between secure zones, the security of the systems and applications themselves must be maintained. Security measures to be applied are authentication, logging (for accountability), and virus scanning. A reverse proxy may provide authentication for external users; a web application gateway — a.k.a. web application firewall, usually acting as a (transparent) reverse proxy — restricts traffic that is considered malicious, e.g., containing SQL injection attacks. Organizational security measures include user account management and user authorization, as well as physical protection of the data center, such as fenced-in premises, a guarded entrance, and no unaccompanied physical access to systems.

Table 2 provides an overview of which zones must implement which technical security measures, and references the subchapter where the respective measure is described. The measures are cumulative, i.e., each zone implements all measures of the zone "below" and possibly additional measures.

Table 2: Security measures

Zone	Security Measure	Ref.	Remarks
Presentation (yellow)	Firewalls	2.5.1	
	Hardening	2.5.2	
	Data transfer encryption	2.5.3	
	Virus scanners	2.5.4	
	Audit trail	2.5.5	
Application (orange)	Access control	2.5.6	
	IDS	2.5.7	
Data (red)	Data storage encryption	2.5.8	optional

2.5.1 Firewalls

Each zone implements its own filtering, typically using firewalls with several network interfaces to accommodate a number of zone instances (subnets) within that zone. The firewalls protect zones by filtering incoming and outgoing connections as well as traffic between zone instances according to the documented communication flow for the applications within the zone.

As an example, the firewall for the orange zone controls access both from the yellow and red zones, and also between zone instances (subnets) within the orange zone. It denies all other access, e.g., from the white zone.

2.5.2 Hardening

System and application hardening are *the* most important step towards protecting data and assets.

Security updates for the operating system and all software packages must be installed as soon as they become available. Furthermore, services not required on a system must be uninstalled or at least disabled.

The configuration of the applications and the operating system must follow best practices for security. Tools to determine insecure configurations are often provided by the manufacturer, as well as information on secure configuration.

Therefore, hardening is not a one-time measure at the time of deployment of a system, but must be implemented as an ongoing process.

2.5.3 Data transfer encryption

Confidential data that is transferred over a network must always be encrypted, regardless of the zone(s) it traverses.

The amount of (sensitive) information typically becomes less from the database layer (red) towards the Internet, since the application layer (orange) usually queries more data than it actually needs from the database and then passes on a subset of this data to the presentation layer (yellow); see also 2.6. Nevertheless, such sensitive information is also less protected towards the outer zones, and therefore must be encrypted all the way.

Table 3 lists some replacements for unencrypted or unauthenticated protocols.

Table 3: Data transfer encryption

Insecure protocol	Secure replacement	Remarks
HTTP	HTTPS	via SSL / TLS
FTP	SFTP	Secure FTP
	scp	secure copy
telnet	ssh	secure shell
SQLNET	SSL	e.g., Oracle and DB/2 provide several authentication and encryption mechanisms

The use of insecure protocols such as ODBC for the transfer of sensitive information must be prohibited. In case there is no authentication and encryption mechanism intrinsic to a specific type of transfer, end-to-end communication may also be embedded, e.g., in an SSH or SSL tunnel between the source and target components.

2.5.4 Virus scanners

Virus scanners are mandatory on internal clients and must also be implemented on gateway or proxy servers, such as web and ftp proxies, as well as mail servers.

Depending on the applications used, it may also be feasible to install a virus scanner on specific servers. For example, a system used to convert files (pictures, MS Word docs, etc.) for

anonymous users may want to verify that no buffer overflow occurs due to a deliberately manipulated file.

2.5.5 Audit trail

For accountability and non-repudiation purposes, logging of critical events is mandatory:

- Important application events like start / stop, critical errors, etc.
- Security relevant events, i.e., login / logout of users, including failed logins, as well as configuration changes

Write access to log data must be restricted to the application that is the source of the audit trail. A remote log host is highly recommended to protect the audit trail from manipulation and to provide centralized log analysis.

Log data must not include confidential information, like passwords. Log entries must always include date and time, as well as the source of the event. If available, user name and IP address should also be part of a log entry.

2.5.6 Access control

Access control is the ability to permit or deny the use of a resource by an entity. Access control includes *identification* and *authentication*, *authorization* and an *audit trail* (see 2.5.5).

Access can be granted or denied based on a wide variety of arbitrary criteria, such as the network address of the communication partner, the time of day, type of request, etc. These criteria may bear no reference to the attributes of a particular request.

2.5.6.1 Identification and authentication

Applications that process internal or confidential data must only be accessed by an authenticated entity for which specific access rights have been defined. Humans and also applications may communicate with a service; therefore, the following two types of authentication apply:

- **User authentication**
  Individual accounts (instead of shared logins) must be set up for each end user to ensure accountability. Depending on the sensitivity of information, strong authentication mechanisms are recommended.
- **Service authentication**
  Services usually authenticate using a “technical user”, which must only be used in this context, for example, access to a database or to a bus in an SOA environment.
  Especially for web services, certificate-based authentication is recommended.

It is assumed that *applications* implement access control mechanisms controlling end users’ access to resources. Therefore, other services like databases have to rely on access control mechanisms implemented in the application from which a query originates. This ensures that end user access control only has to be implemented once, in the application the end user is communicating with.

For large environments with a substantial number of users, a single sign-on service is recommended, which greatly reduces the various passwords to be remembered while ensuring that no stale accounts remain after a user leaves the organization or is transferred to another department.

2.5.6.2 Authorization

Authorization is the process of providing and restricting access to resources. As such, it is very much credential-focused and dependent on specific rules and access control lists preset by the application administrator(s) or data owners. Typical authorization checks involve querying for membership in a particular (user) group, possession of a particular clearance, or an entry in the approved access control list of a resource.

Any access control mechanism is clearly dependent on effective and forge-resistant authentication controls used for authorization.

2.5.7 IDS

An intrusion detection system (IDS) alerts administrators if an intrusion attempt or a successful compromise occurs within a network or on a host. Since an IDS requires considerable resources to implement and maintain, an IDS is only required for the red zone, which typically consists of only a few systems. A host-based IDS provides more reliable information about whether an intrusion attempt was successful, while a network-based IDS also sees undirected traffic (e.g., reconnaissance), which is expected to be rather limited within a red zone.

2.5.8 Data storage encryption

Data storage encryption is strongly recommended at the application level (see 2.5.8.1).

2.5.8.1 Application data encryption

In order to protect the confidentiality of sensitive data against intruders or malevolent administrators on the database server, application-level encryption is recommended, where available.

2.5.8.2 File system encryption

If application-level encryption is not an option, file system encryption at least protects confidential data against physical theft or loss, e.g., when old or non-functional hard drives are improperly disposed of.

2.6 Example – online banking

In **Fig. 3**, an online banking application is illustrated as an example of a typical implementation in a secure zones environment.

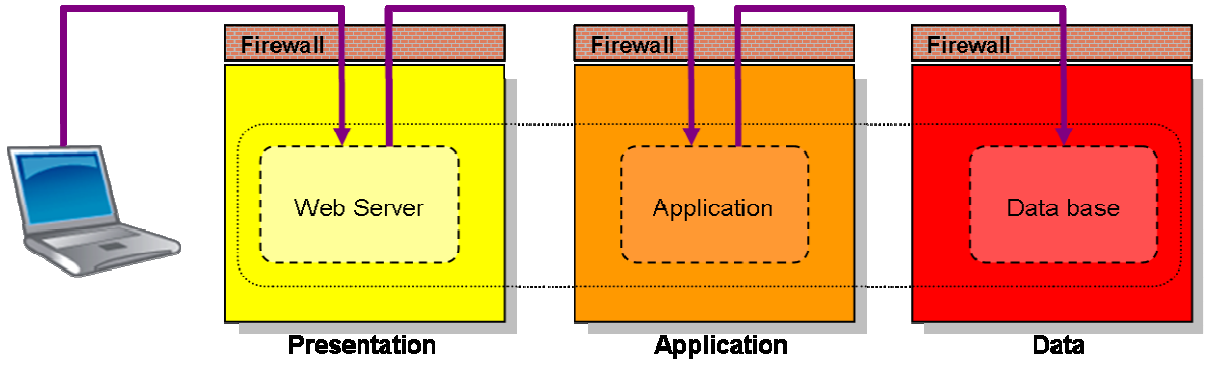

Fig. 3: Online banking example

The client is somewhere in the Internet, the white zone, and accesses the online banking front-end, which executes on a web server in the yellow zone of the bank. The web frontend authenticates the customer, and only after the user's credentials are verified successfully, forwards the requests to the actual online banking application in the orange zone. If the online banking application needs further data to fulfill the request, it queries the database in the red zone, processes the data, and returns a dynamic web page via the web server in the yellow zone to the customer's client.

Since all banking data is considered confidential, each communication link is encrypted.

3 Special cases

This paper also discusses the compliance of several use cases, such as system administration, MAN / WAN interconnectivity, backup, and small enterprises with this security model.

3.1 System administration

For system administration, a "jump station" is placed in the orange zone, from where it is allowed to access this zone, as well as the red and yellow zones (see **Fig. 4**). Administrative access to other systems is only permitted from the jump host, requiring administrators to log on to this jump host first.

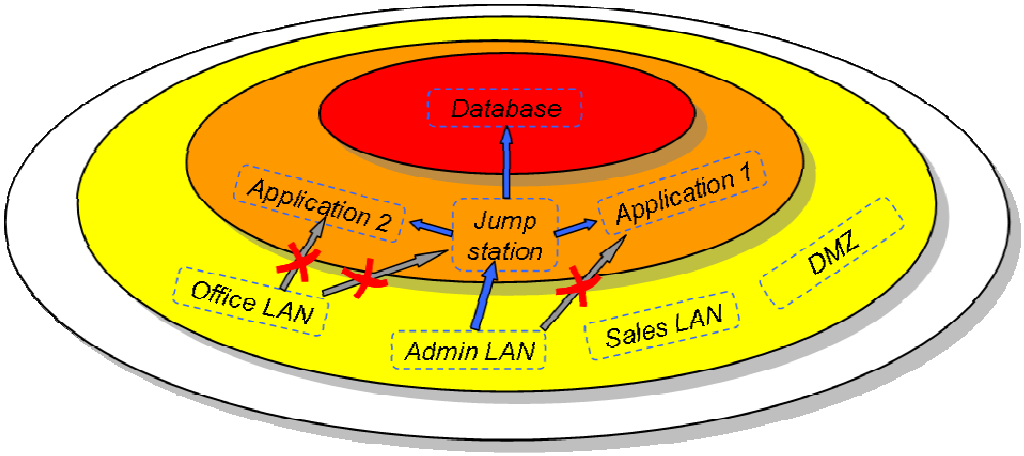

Fig. 4: System administration with a jump station

The jump station holds ssh keys and certificates required to authenticate to other systems, providing central control over credentials, as well as accountability via logging of user activity.

According to best practice, unencrypted or insecure protocols, such as telnet, rsh, or http, must not be used for system administration. Also, trust relationships such as `.rhost` never were a valid security concept and must not be employed.

3.2 Interconnectivity

The interconnectivity between remote locations is not logically different from local networks organized in network zones. Routers and transfer networks are not considered part of zone instances, but part of the network infrastructure (see **Fig. 5**).

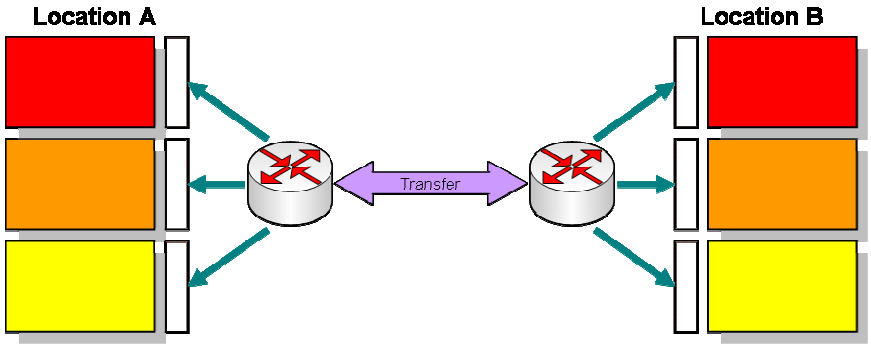

Fig. 5: Location interconnectivity

The zone traversal paradigm (see 2.4) applies to interconnected zones. Zones are considered adjacent (as defined in **Fig. 2**) across transfer networks. Zones of the same level may connect directly, e.g., red zone in location A with red zone in location B.

3.3 Backup

Existing (additional) backup networks often interconnect all hosts to be backed up, thus circumventing intermediate firewalls. Furthermore, backup media often is not clearly classified, and backup procedures intermix non-sensitive and confidential data, usually resulting in insufficient safeguarding of sensitive data.

As a result, backup systems must be placed in a secure zone with at least the same confidentiality level as the data being backed up.

A separate backup network may be used to avoid congestion of production networks, but must not compromise the secure zones. Adequate protection of the backup network and separation from the productive network may allow transferring data unencrypted to the backup servers for performance.

3.4 Small enterprises

For small and medium enterprises, a simplified approach with only two internal zones (yellow and orange) may be feasible. Systems with confidential data, which are normally assigned to a red zone, are placed into the orange zone.

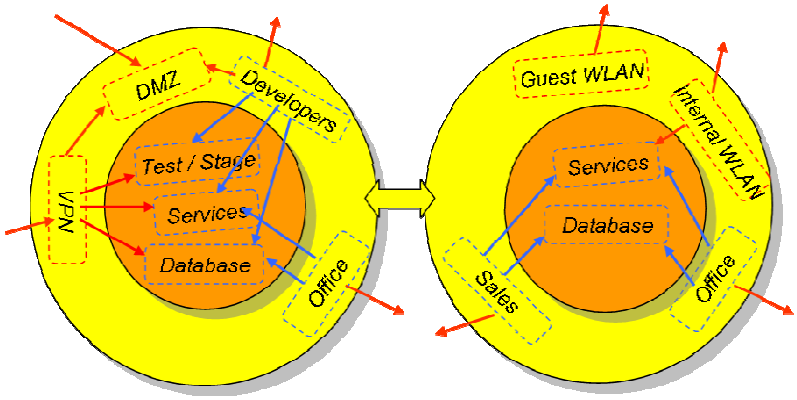

Fig. 6: Secure zones for small and medium environments

Fig. 6 shows an example of an SME environment, where the data subnet is part of the orange zone.

4 Migration

A secure zones model is typically deployed in existing data centers with a large number of legacy applications in a flat network structure. Therefore, a migration plan must be designed, beginning with a pilot migration as proof of concept, to get acquainted with the new architecture and identify pitfalls caused by the existing network and application structure.

Applications need to be separated into their three tiers (presentation, logic and data); the logic tier usually is an application server. The communication flow for the application must be documented, i.e., source, target, protocol, and content must be identified. This information is required to define access controls (see 2.5.1), as well as for classification of data (see 2.1).

For the data tier, it is feasible to implement a centralized database service which provides table space to the applications. A centralized approach also reduces maintenance costs, while concentrating on a secure and redundant implementation of that service.

The presentation layer either is a client in a local LAN or a (reverse) proxy providing access for external clients. Such a proxy authenticates external clients, forwards client requests and delivers server responses, and optionally filters incoming data for malicious content. In addition, data transfer mechanisms may have to be adapted to provide encryption.

Some protocols are quite easily switched to their encrypted equivalent, e.g., HTTP to HTTPS (see 2.5.3). However, some (proprietary) protocols may not provide any option for encryption at all, or the application can not be changed (e.g., no source code available). Therefore, there are always some legacy applications that cannot be migrated and that will remain in the insecure, flat network until they are obsolete.

5 Related work

Bell and La Padula ([BeLP73], [BeLP76]) developed a security policy for access control for military and government applications. This model is based on security labels on information objects and clearances on subjects. The security labels are classifications, such as “secret”, whereas the clearances represent roles and the corresponding (access) rights. The Bell and La Padula focus on data confidentiality and data flow was adapted for the secure zones model.

A mechanism to protect data and resources called “hierarchical protection domains” or “protection rings” has been developed to provide layers of privilege. Typically this is implemented in hardware, e.g., CPU architectures that use “supervisor” mode (a.k.a. “kernel” mode) and “user” mode for different levels of access. This hierarchical approach was adopted for the secure zones.

6 Conclusion

Secure network zones provide in-depth protection for sensitive data and vital systems by implementing several layers of increased security. The impact of malicious code or attacks is effectively contained, and the standardized approach for securely deploying applications also improves data center operation efficiency.

References

[BeLP73] Bell, David Elliott and La Padula, Leonard J.: Secure Computer Systems: Mathematical Foundations. MITRE Corporation, 1973.

[BeLP76] Bell, David Elliott and La Padula, Leonard J.: Secure Computer System: Unified Exposition and Multics Interpretation. MITRE Corporation, 1976.

[LaHM84] Landwehr, C.E., C.L. Heitmeyer, and J. McLean, "A Security Model for Military Message Systems," ACM Trans. on Computer Systems Vol. 9, No. 3 (Aug. 1984), pp. 198-222.

[Zelt00] Zeltser, Lenny: Firewalls, Perimeter Protection, and VPNs. GCFW Practical Assignment, SANS, December 2000, p. 13-35.

[Zelt02] Zeltser, Lenny: Firewall Deployment for Multitier Applications. http://www.informit.com/articles/article.aspx?p=26254, informIT, April 5, 2002.

[BCF+07] Buecker, Axel; Carreno, Ana Veronica; Field, Norman; Hockings, Christopher; Kawer, Daniel; Mohanty, Sujit; Monteiro, Guilherme: Enterprise Security Architecture. IBM Redbook, IBM International Technical Support Organization, August 2007, p. 29-39.

Index