

Special Focus on Computer Security Model

Erick Franco



**SPECIAL FOCUS ON
COMPUTER SECURITY
MODEL**

SPECIAL FOCUS ON COMPUTER SECURITY MODEL

Erick Franco



Special Focus on Computer Security Model
by Erick Franco

Copyright© 2022 BIBLIOTEX

www.bibliotex.com

All rights reserved. No part of this book may be reproduced or used in any manner without the prior written permission of the copyright owner, except for the use brief quotations in a book review.

To request permissions, contact the publisher at info@bibliotex.com

Ebook ISBN: 9781984664297



Published by:

Bibliotex

Canada

Website: www.bibliotex.com

Contents

Chapter 1	Computer Security and Design Model	1
Chapter 2	Modern Network Devices for Security	30
Chapter 3	Computer Security Model	73
Chapter 4	Capability-Based Security	126
Chapter 5	Computer Information Security	139
Chapter 6	Data Security Model	171

1

Computer Security and Design Model

Computer security is a branch of computer technology known as information security as applied to computers and networks. The objective of computer security includes protection of information and property from theft, corruption, or natural disaster, while allowing the information and property to remain accessible and productive to its intended users. The term computer system security means the collective processes and mechanisms by which sensitive and valuable information and services are protected from publication, tampering or collapse by unauthorized activities or untrustworthy individuals and unplanned events respectively. The strategies and methodologies of computer security often differ from most other computer technologies because of its somewhat elusive objective of preventing unwanted computer behaviour instead of enabling wanted computer behaviour.

SECURITY BY DESIGN

The technologies of computer security are based on logic. As security is not necessarily the primary goal of most computer applications, designing a programme with security in mind often imposes restrictions on that program's behaviour. There are 4 approaches to security in computing, sometimes a combination of approaches is valid:

1. Trust all the software to abide by a security policy but the software is not trustworthy (this is computer insecurity).
2. Trust all the software to abide by a security policy and the software is validated as trustworthy (by tedious branch and path analysis for example).
3. Trust no software but enforce a security policy with mechanisms that are not trustworthy (again this is computer insecurity).
4. Trust no software but enforce a security policy with trustworthy hardware mechanisms.

Many systems have unintentionally resulted in the first possibility. Since approach two is expensive and non-deterministic, its use is very limited. Approaches one and three lead to failure. Because approach number four is often based on hardware mechanisms and avoids abstractions and a multiplicity of degrees of freedom, it is more practical. Combinations of approaches two and four are often used in a layered architecture with thin layers of two and thick layers of four. There are various strategies and techniques used to design security systems. However there are few, if any, effective strategies to enhance security

after design. One technique enforces the principle of least privilege to great extent, where an entity has only the privileges that are needed for its function. That way even if an attacker gains access to one part of the system, fine-grained security ensures that it is just as difficult for them to access the rest. Furthermore, by breaking the system up into smaller components, the complexity of individual components is reduced, opening up the possibility of using techniques such as automated theorem proving to prove the correctness of crucial software subsystems. This enables a closed form solution to security that works well when only a single well-characterized property can be isolated as critical, and that property is also assessible to math. Not surprisingly, it is impractical for generalized correctness, which probably cannot even be defined, much less proven. Where formal correctness proofs are not possible, rigorous use of code review and unit testing represent a best-effort approach to make modules secure.

The design should use “defense in depth”, where more than one subsystem needs to be violated to compromise the integrity of the system and the information it holds. Defense in depth works when the breaching of one security measure does not provide a platform to facilitate subverting another. Also, the cascading principle acknowledges that several low hurdles does not make a high hurdle. So cascading several weak mechanisms does not provide the safety of a single stronger mechanism. Subsystems should default to secure settings, and wherever possible should be designed to “fail secure” rather than “fail insecure”. Ideally, a secure system should require a deliberate, conscious, knowledgeable and

free decision on the part of legitimate authorities in order to make it insecure. In addition, security should not be an all or nothing issue. The designers and operators of systems should assume that security breaches are inevitable. Full audit trails should be kept of system activity, so that when a security breach occurs, the mechanism and extent of the breach can be determined. Storing audit trails remotely, where they can only be appended to, can keep intruders from covering their tracks. Finally, full disclosure helps to ensure that when bugs are found the “window of vulnerability” is kept as short as possible.

SECURITY ARCHITECTURE

Security Architecture can be defined as the design artifacts that describe how the security controls (security countermeasures) are positioned, and how they relate to the overall information technology architecture. These controls serve the purpose to maintain the system’s quality attributes, among them confidentiality, integrity, availability, accountability and assurance.” Hardware mechanisms that protect computers and data Hardware based or assisted computer security offers an alternative to software-only computer security. Devices such as dongles may be considered more secure due to the physical access required in order to be compromised.

SECURE OPERATING SYSTEMS

One use of the term computer security refers to technology to implement a secure operating system. Much of this technology is based on science developed in the 1980s and

used to produce what may be some of the most impenetrable operating systems ever. Though still valid, the technology is in limited use today, primarily because it imposes some changes to system management and also because it is not widely understood. Such ultra-strong secure operating systems are based on operating system kernel technology that can guarantee that certain security policies are absolutely enforced in an operating environment. An example of such a Computer security policy is the Bell-LaPadula model. The strategy is based on a coupling of special microprocessor hardware features, often involving the memory management unit, to a special correctly implemented operating system kernel. This forms the foundation for a secure operating system which, if certain critical parts are designed and implemented correctly, can ensure the absolute impossibility of penetration by hostile elements. This capability is enabled because the configuration not only imposes a security policy, but in theory completely protects itself from corruption. Ordinary operating systems, on the other hand, lack the features that assure this maximal level of security. The design methodology to produce such secure systems is precise, deterministic and logical.

Systems designed with such methodology represent the state of the art of computer security although products using such security are not widely known. In sharp contrast to most kinds of software, they meet specifications with verifiable certainty comparable to specifications for size, weight and power. Secure operating systems designed this way are used primarily to protect national security information, military secrets, and the data of international

Special Focus on Computer Security Model

financial institutions. These are very powerful security tools and very few secure operating systems have been certified at the highest level (Orange Book A-1) to operate over the range of “Top Secret” to “unclassified” (including Honeywell SCOMP, USAF SACDIN, NSA Blacker and Boeing MLS LAN.) The assurance of security depends not only on the soundness of the design strategy, but also on the assurance of correctness of the implementation, and therefore there are degrees of security strength defined for COMPUSEC. The Common Criteria quantifies security strength of products in terms of two components, security functionality and assurance level (such as EAL levels), and these are specified in a Protection Profile for requirements and a Security Target for product descriptions. None of these ultra-high assurance secure general purpose operating systems have been produced for decades or certified under Common Criteria.

In USA parlance, the term High Assurance usually suggests the system has the right security functions that are implemented robustly enough to protect DoD and DoE classified information. Medium assurance suggests it can protect less valuable information, such as income tax information. Secure operating systems designed to meet medium robustness levels of security functionality and assurance have seen wider use within both government and commercial markets. Medium robust systems may provide the same security functions as high assurance secure operating systems but do so at a lower assurance level (such as Common Criteria levels EAL4 or EAL5). Lower levels mean we can be less certain that the security functions

are implemented flawlessly, and therefore less dependable. These systems are found in use on web servers, guards, database servers, and management hosts and are used not only to protect the data stored on these systems but also to provide a high level of protection for network connections and routing services.

SECURE CODING

If the operating environment is not based on a secure operating system capable of maintaining a domain for its own execution, and capable of protecting application code from malicious subversion, and capable of protecting the system from subverted code, then high degrees of security are understandably not possible. While such secure operating systems are possible and have been implemented, most commercial systems fall in a 'low security' category because they rely on features not supported by secure operating systems (like portability, et al.). In low security operating environments, applications must be relied on to participate in their own protection. There are 'best effort' secure coding practices that can be followed to make an application more resistant to malicious subversion.

In commercial environments, the majority of software subversion vulnerabilities result from a few known kinds of coding defects. Common software defects include buffer overflows, format string vulnerabilities, integer overflow, and code/command injection. It is to be immediately noted that all of the foregoing are specific instances of a general class of attacks, where situations in which putative "data" actually contains implicit or explicit, executable instructions

are cleverly exploited. Some common languages such as C and C++ are vulnerable to all of these defects (see Seacord, “*Secure Coding in C and C++*”). Other languages, such as Java, are more resistant to some of these defects, but are still prone to code/command injection and other software defects which facilitate subversion.

Recently another bad coding practice has come under scrutiny; dangling pointers. The first known exploit for this particular problem was presented in July 2007. Before this publication the problem was known but considered to be academic and not practically exploitable. Unfortunately, there is no theoretical model of “secure coding” practices, nor is one practically achievable, insofar as the variety of mechanisms are too wide and the manners in which they can be exploited are too variegated. It is interesting to note, however, that such vulnerabilities often arise from archaic philosophies in which computers were assumed to be narrowly disseminated entities used by a chosen few, all of whom were likely highly educated, solidly trained academics with naught but the goodness of mankind in mind. Thus, it was considered quite harmless if, for (fictitious) example, a FORMAT string in a FORTRAN programme could contain the J format specifier to mean “shut down system after printing.” After all, who would use such a feature but a well-intentioned system programmer? It was simply beyond conception that software could be deployed in a destructive fashion. It is worth noting that, in some languages, the distinction between code (ideally, read-only) and data (generally read/write) is blurred. In LISP, particularly, there is no distinction whatsoever between code and data, both

taking the same form: an S-expression can be code, or data, or both, and the “user” of a LISP programme who manages to insert an executable LAMBDA segment into putative “data” can achieve arbitrarily general and dangerous functionality. Even something as “modern” as Perl offers the eval() function, which enables one to generate Perl code and submit it to the interpreter, disguised as string data.

CAPABILITIES AND ACCESS CONTROL LISTS

Within computer systems, two security models capable of enforcing privilege separation are access control lists (ACLs) and capability-based security. The semantics of ACLs have been proven to be insecure in many situations, e.g., the confused deputy problem. It has also been shown that the promise of ACLs of giving access to an object to only one person can never be guaranteed in practice. Both of these problems are resolved by capabilities. This does not mean practical flaws exist in all ACL-based systems, but only that the designers of certain utilities must take responsibility to ensure that they do not introduce flaws. Capabilities have been mostly restricted to research operating systems and commercial OSs still use ACLs. Capabilities can, however, also be implemented at the language level, leading to a style of programming that is essentially a refinement of standard object-oriented design. An open source project in the area is the E language. First the Plessey System 250 and then Cambridge CAP computer demonstrated the use of capabilities, both in hardware and software, in the 1970s. A reason for the lack of adoption of capabilities may be that ACLs appeared to offer a ‘quick

fix' for security without pervasive redesign of the operating system and hardware. The most secure computers are those not connected to the Internet and shielded from any interference. In the real world, the most security comes from operating systems where security is not an add-on.

APPLICATIONS

Computer security is critical in almost any technology-driven industry which operates on computer systems. Computer security can also be referred to as computer safety. The issues of computer based systems and addressing their countless vulnerabilities are an integral part of maintaining an operational industry.

CLOUD COMPUTING SECURITY

Security in the cloud is challenging, due to varied degree of security features and management schemes within the cloud entities. In this connection one logical protocol base need to evolve so that the entire gamut of components operates synchronously and securely.

IN AVIATION

The aviation industry is especially important when analyzing computer security because the involved risks include human life, expensive equipment, cargo, and transportation infrastructure. Security can be compromised by hardware and software malpractice, human error, and faulty operating environments. Threats that exploit computer vulnerabilities can stem from sabotage, espionage, industrial competition, terrorist attack, mechanical malfunction, and

human error. The consequences of a successful deliberate or inadvertent misuse of a computer system in the aviation industry range from loss of confidentiality to loss of system integrity, which may lead to more serious concerns such as data theft or loss, network and air traffic control outages, which in turn can lead to airport closures, loss of aircraft, loss of passenger life. Military systems that control munitions can pose an even greater risk. A proper attack does not need to be very high tech or well funded; for a power outage at an airport alone can cause repercussions worldwide. One of the easiest and, arguably, the most difficult to trace security vulnerabilities is achievable by transmitting unauthorized communications over specific radio frequencies.

These transmissions may spoof air traffic controllers or simply disrupt communications altogether. These incidents are very common, having altered flight courses of commercial aircraft and caused panic and confusion in the past. Controlling aircraft over oceans is especially dangerous because radar surveillance only extends 175 to 225 miles offshore. Beyond the radar's sight controllers must rely on periodic radio communications with a third party. Lightning, power fluctuations, surges, brown-outs, blown fuses, and various other power outages instantly disable all computer systems, since they are dependent on an electrical source. Other accidental and intentional faults have caused significant disruption of safety critical systems throughout the last few decades and dependence on reliable communication and electrical power only jeopardizes computer safety.

NOTABLE SYSTEM ACCIDENTS

In 1994, over a hundred intrusions were made by unidentified crackers into the Rome Laboratory, the US Air Force's main command and research facility. Using trojan horse viruses, hackers were able to obtain unrestricted access to Rome's networking systems and remove traces of their activities. The intruders were able to obtain classified files, such as air tasking order systems data and furthermore able to penetrate connected networks of National Aeronautics and Space Administration's Goddard Space Flight Center, Wright-Patterson Air Force Base, some Defense contractors, and other private sector organizations, by posing as a trusted Rome center user.

COMPUTER SECURITY POLICY

UNITED STATES

CYBERSECURITY ACT OF 2010

On April 1, 2009, Senator Jay Rockefeller (D-WV) introduced the "Cybersecurity Act of 2009 - S. 773" (full text) in the Senate; the bill, co-written with Senators Evan Bayh (D-IN), Barbara Mikulski (D-MD), Bill Nelson (D-FL), and Olympia Snowe (R-ME), was referred to the Committee on Commerce, Science, and Transportation, which approved a revised version of the same bill (the "Cybersecurity Act of 2010") on March 24, 2010. The bill seeks to increase collaboration between the public and the private sector on cybersecurity issues, especially those private entities that own infrastructures that are critical to national security

interests (the bill quotes John Brennan, the Assistant to the President for Homeland Security and Counterterrorism: “our nation’s security and economic prosperity depend on the security, stability, and integrity of communications and information infrastructure that are largely privately-owned and globally-operated” and talks about the country’s response to a “cyber-Katrina”), increase public awareness on cybersecurity issues, and foster and fund cybersecurity research. Some of the most controversial parts of the bill include Paragraph 315, which grants the President the right to “order the limitation or shutdown of Internet traffic to and from any compromised Federal Government or United States critical infrastructure information system or network.” The Electronic Frontier Foundation, an international non-profit digital rights advocacy and legal organization based in the United States, characterized the bill as promoting a “potentially dangerous approach that favours the dramatic over the sober response”.

INTERNATIONAL CYBERCRIME REPORTING AND COOPERATION ACT

On March 25, 2010, Representative Yvette Clarke (D-NY) introduced the “International Cybercrime Reporting and Cooperation Act - H.R.4962” (full text) in the House of Representatives; the bill, co-sponsored by seven other representatives (among whom only one Republican), was referred to three House committees. The bill seeks to make sure that the administration keeps Congress informed on information infrastructure, cybercrime, and end-user protection worldwide. It also “directs the President to give

priority for assistance to improve legal, judicial, and enforcement capabilities with respect to cybercrime to countries with low information and communications technology levels of development or utilization in their critical infrastructure, telecommunications systems, and financial industries” as well as to develop an action plan and an annual compliance assessment for countries of “cyber concern”.

PROTECTING CYBERSPACE AS A NATIONAL ASSET ACT OF 2010 (“KILL SWITCH BILL”)

On June 19, 2010, United States Senator Joe Lieberman (I-CT) introduced a bill called “Protecting Cyberspace as a National Asset Act of 2010 - S.3480” (full text in pdf), which he co-wrote with Senator Susan Collins (R-ME) and Senator Thomas Carper (D-DE). If signed into law, this controversial bill, which the American media dubbed the “*Kill switch bill*”, would grant the President emergency powers over the Internet. However, all three co-authors of the bill issued a statement claiming that instead, the bill “[narrowed] existing broad Presidential authority to take over telecommunications networks”.

TERMINOLOGY

The following terms used in engineering secure systems are explained below.

- Authentication techniques can be used to ensure that communication end-points are who they say they are.
- Automated theorem proving and other verification tools can enable critical algorithms and code used in

secure systems to be mathematically proven to meet their specifications.

- Capability and access control list techniques can be used to ensure privilege separation and mandatory access control. This section discusses their use.
- Chain of trust techniques can be used to attempt to ensure that all software loaded has been certified as authentic by the system's designers.
- Cryptographic techniques can be used to defend data in transit between systems, reducing the probability that data exchanged between systems can be intercepted or modified.
- Firewalls can provide some protection from online intrusion.
- A microkernel is a carefully crafted, deliberately small corpus of software that underlies the operating system *per se* and is used solely to provide very low-level, very precisely defined primitives upon which an operating system can be developed. A simple example with considerable didactic value is the early '90s GEMSOS (Gemini Computers), which provided extremely low-level primitives, such as "segment" management, atop which an operating system could be built. The theory (in the case of "segments") was that—rather than have the operating system itself worry about mandatory access separation by means of military-style labeling—it is safer if a low-level, independently scrutinized module can be charged solely with the management of individually labeled segments, be they memory "segments" or file system

“segments” or executable text “segments.” If software below the visibility of the operating system is (as in this case) charged with labeling, there is no theoretically viable means for a clever hacker to subvert the labeling scheme, since the operating system *per se* does not provide mechanisms for interfering with labeling: the operating system is, essentially, a client (an “application,” arguably) atop the microkernel and, as such, subject to its restrictions.

- Endpoint Security software helps networks to prevent data theft and virus infection through portable storage devices, such as USB drives.

Some of the following items may belong to the computer insecurity article:

- Access authorization restricts access to a computer to group of users through the use of authentication systems. These systems can protect either the whole computer – such as through an interactive logon screen – or individual services, such as an FTP server. There are many methods for identifying and authenticating users, such as passwords, identification cards, and, more recently, smart cards and biometric systems.
- Anti-virus software consists of computer programmes that attempt to identify, thwart and eliminate computer viruses and other malicious software (malware).
- Applications with known security flaws should not be run. Either leave it turned off until it can be

patched or otherwise fixed, or delete it and replace it with some other application. Publicly known flaws are the main entry used by worms to automatically break into a system and then spread to other systems connected to it. The security website Secunia provides a search tool for unpatched known flaws in popular products.

- Backups are a way of securing information; they are another copy of all the important computer files kept in another location. These files are kept on hard disks, CD-Rs, CD-RWs, and tapes. Suggested locations for backups are a fireproof, waterproof, and heat proof safe, or in a separate, offsite location than that in which the original files are contained. Some individuals and companies also keep their backups in safe deposit boxes inside bank vaults. There is also a fourth option, which involves using one of the file hosting services that backs up files over the Internet for both business and individuals.
 - o Backups are also important for reasons other than security. Natural disasters, such as earthquakes, hurricanes, or tornadoes, may strike the building where the computer is located. The building can be on fire, or an explosion may occur. There needs to be a recent backup at an alternate secure location, in case of such kind of disaster. Further, it is recommended that the alternate location be placed where the same disaster would not affect both locations. Examples of alternate disaster recovery sites being compromised by the same

disaster that affected the primary site include having had a primary site in World Trade Center I and the recovery site in 7 World Trade Center, both of which were destroyed in the 9/11 attack, and having one's primary site and recovery site in the same coastal region, which leads to both being vulnerable to hurricane damage (e.g. primary site in New Orleans and recovery site in Jefferson Parish, both of which were hit by Hurricane Katrina in 2005). The backup media should be moved between the geographic sites in a secure manner, in order to prevent them from being stolen.

- Encryption is used to protect the message from the eyes of others. Cryptographically secure ciphers are designed to make any practical attempt of breaking infeasible. Symmetric-key ciphers are suitable for bulk encryption using shared keys, and public-key encryption using digital certificates can provide a practical solution for the problem of securely communicating when no key is shared in advance.
- Firewalls are systems which help protect computers and computer networks from attack and subsequent intrusion by restricting the network traffic which can pass through them, based on a set of system administrator defined rules.
- Honey pots are computers that are either intentionally or unintentionally left vulnerable to attack by crackers. They can be used to catch crackers or fix vulnerabilities.

Special Focus on Computer Security Model

- Intrusion-detection systems can scan a network for people that are on the network but who should not be there or are doing things that they should not be doing, for example trying a lot of passwords to gain access to the network.
- Pinging The ping application can be used by potential crackers to find if an IP address is reachable. If a cracker finds a computer, they can try a port scan to detect and attack services on that computer.
- Social engineering awareness keeps employees aware of the dangers of social engineering and/or having a policy in place to prevent social engineering can reduce successful breaches of the network and servers.
- File Integrity Monitors are tools used to detect changes in the integrity of systems and files.

SECURITY-FOCUSED OPERATING SYSTEM

This is an alphabetical list of operating systems with a sharp security focus. Their order does not imply rank. In our context, “Security-focused” means that the project is devoted to increasing the security as a major goal. As such, something can be secure without being “security-focused.” For example, almost all of the operating systems mentioned here are faced with security bug fixes in their lifetime; however, they do all strive to consistently approach all generic security flaws inherent in their design with new ideas in an attempt to create a secure computing environment.

BSD

BSD is a family of Unix variants derived from a code base originating at the University of California, Berkeley. All derived BSD operating systems are released under the terms of a BSD-style license. There are several BSD variants, with only one being heavily focused on security.

OPENBSD

OpenBSD is an open source BSD operating system that is known to be concerned heavily with security. The project has completed rigorous manual reviews of the code and addressed issues most systems have not. OpenBSD also supplies an executable space protection scheme known as WX (memory is writable xor executable), as well as a ProPolice compiled executable base.

TRUSTEDBSD

TrustedBSD is a sub-project of FreeBSD designed to add trusted operating system extensions, targeting the Common Criteria for Information Technology Security Evaluation. Its main focuses are working on access control lists, event auditing, extended attributes, mandatory access controls, and fine-grained capabilities. Since access control lists are known to be confronted with the confused deputy problem, capabilities are a different way to avoid this issue. As part of the TrustedBSD project, there is also a port of the NSA's FLASK/TE implementation to run on FreeBSD. Many of these trusted extensions have been integrated into the main FreeBSD branch starting at 5.x.

LINUX

Linux itself is inherently security-focused; however, many distributions and projects attempt to make Linux more secure.

ANNVIX

Annvix was originally forked from Mandriva to provide a security-focused server distribution that employs ProPolice protection, hardened configuration, and a small footprint. There have been plans to include full support for the RSBAC Mandatory access control system. However, Annvix seems to be a dormant operating system with the last version being released December 30, 2007.

ENGARDE SECURE LINUX

EnGarde Secure Linux is a secure platform designed for servers. It has boasted a browser-based tool for MAC using SELinux since 2003. Additionally, it can be accompanied with Web, DNS, and Email enterprise applications, specifically focusing on security without any unnecessary software. The community platform of EnGarde Secure Linux is the bleeding-edge version freely available for download.

FEDORA

Fedora is a free, Red Hat sponsored community developed Linux distribution. It is one of those mainstream Linux distribution, with a concentrated effort to improve system security, as a consequence it boasts a fully integrated SELinux MAC and fine-grained executable memory permission system (Exec Shield) and all binaries compiled

with GCC's standard stack-smashing protection, as well as focusing on getting security updates into the system in a timely manner.

HARDENED GENTOO

Hardened Gentoo is a subproject of the Gentoo Linux project. Hardened Gentoo offers a ProPolice protected and Position Independent Executable base using exactly the same package tree as Gentoo. Executable space protection in Hardened Gentoo is handled by PaX. The Hardened Gentoo project is an extremely modular project, and also provides subprojects to integrate other intrusion-detection and Mandatory access control systems into Gentoo. All of these can be optionally installed in any combination, with or without PaX and a ProPolice base.

HARDENED LINUX

Hardened Linux is a small distribution for firewalls, intrusion detection systems, VPN-gateways and authentication jobs that is still under heavy development. It includes GRSecurity, PaX and GCC stack smashing protection.

IMMUNIX

Immunix is a commercial distribution of Linux focused heavily on security. They supply many systems of their own making, including StackGuard; cryptographic signing of executables; race condition patches; and format string exploit guarding code. Immunix traditionally releases older versions of their distribution free for non-commercial use. Note that

the Immunix distribution itself is licensed under two licenses: The Immunix commercial and non-commercial licenses. Many tools within are GPL, however; as is the kernel.

OPENWALL PROJECT

Owl by a developer known as Solar Designer was the first distribution to have a non-executable userspace stack, /tmp race condition protection and access control restrictions to /procdata, by way of a kernel patch. It also features a per-user tmp directory via the pam_mktemp PAM module, and supports Blowfish password encryption.

RED HAT ENTERPRISE LINUX

Red Hat Enterprise Linux - offers the same security benefits as Fedora with the additional support of back-porting security fixes to the released versions of the packages (particularly the kernel) so the sys-admin does not have to perform a significant (and risky) upgrade to get a security fix.

UBUNTU

Like Fedora and Red Hat Enterprise Linux, Ubuntu provides swift security fixes for stable releases. It also has AppArmor installed by default and supports SELinux. Ubuntu locks the root account by default.

SOLARIS

Solaris is a Unix variant created by Sun Microsystems. Solaris itself is not inherently security-focused. Majority of Solaris source code has been released via the OpenSolaris project, mostly under the Common Development and

Distribution License. Enhancements to OpenSolaris, both security related and others, are backported to the official Solaris when Sun certifies their quality.

TRUSTED SOLARIS

Trusted Solaris is a security-focused version of the Solaris Unix operating system. Aimed primarily at the government computing sector, Trusted Solaris adds detailed auditing of all tasks, pluggable authentication, mandatory access control, additional physical authentication devices, and fine-grained access control. Trusted Solaris is Common Criteria certified. The most recent version, Trusted Solaris 8, received the EAL4 certification level augmented by a number of protection profiles.

SOLARIS 10 AND TRUSTED FUNCTIONALITY

Trusted Solaris functionality has now been added to the mainstream version of Solaris. In the 11/06 update to Solaris 10, the *Solaris Trusted Extensions* feature adds mandatory access control and labelled security. Introduced in the same update, the *Secure by Default Networking* feature implements less services on by default compared to most previous releases which had most services enabled. RBAC, found in both mainstream Solaris and Trusted Solaris, dramatically lessens the need for using root directly by providing a way for fine grained control over various administrative tasks.

SECURITY ARCHITECTURE

Security provided by IT Systems can be defined as the IT system's ability to be able to protect confidentiality and

integrity of processed data, as well as to be able to provide availability of the system and data. “IT Architecture” may be defined as a set of design artifacts, that are relevant for describing an object such that it can be produced to requirements (quality) as well as maintained over the period of its useful life (change). The design artifact describe the structure of components, their inter-relationships, and the principles and guidelines governing their design and evolution over time. Consequently the definition of “IT Security Architecture” may be considered as:

The design artifacts that describe how the security controls (= security countermeasures) are positioned and how they relate to the overall IT Architecture. These controls serve the purpose to maintain the system’s quality attributes, among them confidentiality, integrity and availability.

Security qualities are often considered as Non-functional requirements when systems are designed. In other words they are not required for the system to meet its functional goals such as processing financial transactions, but are needed for a given level of assurance that the system will perform to meet the functional requirements that have been defined. In recent years there has been a trend towards a hierarchy of control objectives, controls and specific technical implementations of controls, which are implemented within a given security architecture in order to meet the security requirements.

SECURE BY DESIGN

Secure by design, in software engineering, means that the software has been designed from the ground up to be

secure. Malicious practices are taken for granted and care is taken to minimize impact when a security vulnerability is discovered or on invalid user input. Generally, designs that work well do not rely on being secret. It is not mandatory, but proper security usually means that everyone is allowed to know and understand the design *because it is secure*. This has the advantage that many people are looking at the code, and this improves the odds that any flaws will be found sooner (Linus's law). Of course, attackers can also obtain the code, which makes it easier for them to find vulnerabilities as well. Also, it is very important that everything works with the least amount of privileges possible (principle of least privilege). For example a Web server that runs as the administrative user (root or admin) can have the privilege to remove files and users that do not belong to itself. Thus, a flaw in that programme could put the entire system at risk. On the other hand, a Web server that runs inside an isolated environment and only has the privileges for required network and filesystem functions, cannot compromise the system it runs on unless the security around it is in itself also flawed. A perfect authentication system for logins does not allow anyone to log in at all, because the user could be a threat to the system. However, some designs can never be perfect. Passwords, biometrics, and such are never perfect.

SECURITY BY DESIGN IN PRACTICE

Many things, especially input, should be distrusted by a secure design. A fault-tolerant programme could even distrust its own internals. Two examples of insecure design

are allowing buffer overflows and format string vulnerabilities. The following C programme demonstrates these flaws:

```
int main()
{
  char buffer[100];
  printf("What is your name?\n");
  gets(buffer);
  printf("Hello, ");
  printf(buffer);
  printf("!\\n");
  return 0;
}
```

Because the `gets` function in the C standard library does not stop writing bytes into `buffer` until it reads a newline character or EOF, typing more than 99 characters at the prompt constitutes a buffer overflow. Allocating 100 characters for `buffer` with the assumption that almost any given name from a user is no longer than 99 characters doesn't prevent the user from actually *typing* more than 99 characters. This can lead to arbitrary machine code execution. The second flaw is that the programme tries to print its input by passing it directly to the `printf` function. This function prints out its first argument, replacing conversion specifications (such as `"%s"`, `"%d"`, et cetera) sequentially with other arguments from its call stack as needed. Thus, if a malicious user entered `"%d"` instead of his name, the programme would attempt to print out a non-existent integer value, and undefined behaviour would occur. A related mistake in Web programming is for an online

script not to validate its parameters. For example, consider a script that fetches an article by taking a filename, which is then read by the script and parsed. Such a script might use the following hypothetical URL to retrieve an article about dog food:

```
http://www.example.net/cgi-bin/  
article.sh?name=dogfood.html
```

If the script has no input checking, instead trusting that the filename is always valid, a malicious user could forge a URL to retrieve configuration files from the webserver:

```
http://www.example.net/cgi-bin/  
article.sh?name=../../../../etc/passwd
```

Depending on the script, this may expose the `/etc/passwd` file, which on Unix-like systems contains (among others) user IDs, their login names, home directory paths and shells.

SERVER/CLIENT ARCHITECTURES

In server/client architectures, the programme at the other side may not be an authorised client and the client's server may not be an authorised server. Even when they are, a man-in-the-middle attack could compromise communications. Often the easiest way to break the security of a client/server system is not to go head on to the security mechanisms but instead to go around them. A man in the middle attack is a simple example of this, because you can use it to collect details to impersonate a user. Which is why it is important to consider encryption, hashing, and other security mechanisms in your design to ensure that

Special Focus on Computer Security Model

information collected from a potential attacker won't allow access. Another key feature to client-server security design is general good-coding practices. For example, following a known software design structure such as client and broker can help in designing a well built structure with a solid foundation. Further more that if the software is modified in the future it is even more important that it follows a logical foundation of separation between the client and server. This is because if a programmer comes in and can not clearly understand the dynamics of the programme they may end up adding or changing something that can add a security flaw. Even with the best design this is always a possibility, but the better standardized the design the less chance there is of this occurring.

2

Modern Network Devices for Security

Modern network devices are complex entities composed of both silicon and software. Thus, designing an efficient hardware platform is not, by itself, sufficient to achieve an effective, cost-efficient and operationally tenable product.

The control plane plays a critical role in the development of features and in ensuring device usability.

Although progress from the development of faster CPU boards and forwarding planes is visible, structural changes made in software are usually hidden, and while vendor collateral often offers a list of features in a carrier-class package, operational experiences may vary considerably. Products that have been through several generations of software releases provide the best examples of the difference made by the choice of OS. It is still not uncommon to find routers or switches that started life under older, monolithic

software and later migrated to more contemporary designs. The positive effect on stability and operational efficiency is easy to notice and appreciate.

However, migration from one network operating system to another can pose challenges from non-overlapping feature sets, noncontiguous operational experiences and inconsistent software quality. These potential challenges make it is very desirable to build a control plane that can power the hardware products and features supported in both current and future markets. Developing a flexible, long-lasting and high-quality network OS provides a foundation that can gracefully evolve to support new needs in its height for up and down scaling, width for adoption across many platforms, and depth for rich integration of new features and functions. It takes time, significant investment and in-depth expertise.

Most of the engineers writing the early releases of Junos OS came from other companies where they had previously built network software. They had firsthand knowledge of what worked well, and what could be improved. These engineers found new ways to solve the limitations that they'd experienced in building the older operating systems.

Resulting innovations in Junos OS are significant and rooted in its earliest design stages. Still, to ensure that our products anticipate and fulfil the next generation of market requirements, Junos OS is periodically reevaluated to determine whether any changes are needed to ensure that it continues to provide the reliability, performance and resilience for which it is known. Contemporary network operating systems are mostly advanced and specialized

branches of POSIX-compliant software platforms and are rarely developed from scratch. The main reason for this situation is the high cost of developing a world-class operating system all the way from concept to finished product.

By adopting a general purpose OS architecture, network vendors can focus on routing-specific code, decrease time to market, and benefit from years of technology and research that went into the design of the original (donor) products.

FIRST-GENERATION OS: MONOLITHIC ARCHITECTURE

Typically, first-generation network operating systems for routers and switches were proprietary images running in a flat memory space, often directly from flash memory or ROM. While supporting multiple processes for protocols, packet handling and management, they operated using a cooperative, multitasking model in which each process would run to completion or until it voluntarily relinquished the CPU.

All first-generation network operating systems shared one trait: They eliminated the risks of running full-size commercial operating systems on embedded hardware. Memory management, protection and context switching were either rudimentary or nonexistent, with the primary goals being a small footprint and speed of operation.

Nevertheless, first-generation network operating systems made networking commercially viable and were deployed on a wide range of products. The downside was that these systems were plagued with a host of problems associated with resource management and fault isolation; a single

runaway process could easily consume the processor or cause the entire system to fail. Such failures were not uncommon in the data networks controlled by older software and could be triggered by software errors, rogue traffic and operator errors.

Legacy platforms of the first generation are still seen in networks worldwide, although they are gradually being pushed into the lowest end of the telecom product lines.

SECOND-GENERATION OS: CONTROL PLANE MODULARITY

The mid-1990s were marked by a significant increase in the use of data networks worldwide, which quickly challenged the capacity of existing networks and routers. By this time, it had become evident that embedded platforms could run full-size commercial operating systems, at least on high-end hardware, but with one catch: They could not sustain packet forwarding with satisfactory data rates. A breakthrough solution was needed. It came in the concept of a hard separation between the control and forwarding plane—an approach that became widely accepted after the success of the industry's first application-specific integrated circuit (ASIC)-driven routing platform, the Juniper Networks M40. Forwarding packets entirely in silicon was proven to be viable, clearing the path for next generation network operating systems, led by Juniper with its Junos OS.

Today, the original M40 routers are mostly retired, but their legacy lives in many similar designs, and their blueprints are widely recognized in the industry as the second-generation reference architecture.

Second-generation network operating systems are free from packet switching and thus are focused on control plane functions. Unlike its first-generation counterparts, a second-generation OS can fully use the potential of multitasking, multithreading, memory management and context manipulation, all making systemwide failures less common. Most core and edge routers installed in the past few years are running second-generation operating systems, and it is these systems that are currently responsible for moving the bulk of traffic on the Internet and in corporate networks. However, the lack of a software data plane in second-generation operating systems prevents them from powering low-end devices without a separate (hardware) forwarding plane. Also, some customers cannot migrate from their older software easily because of compatibility issues and legacy features still in use.

These restrictions led to the rise of transitional (generation 1.5) OS designs, in which a first-generation monolithic image would run as a process on top of the second-generation scheduler and kernel, thus bridging legacy features with newer software concepts. The idea behind “generation 1.5” was to introduce some headroom and gradually move the functionality into the new code, while retaining feature parity with the original code base. Although interesting engineering exercises, such designs were not as feature-rich as their predecessors, nor as effective as their successors, making them of questionable value in the long term.

THIRD-GENERATION OS: FLEXIBILITY, SCALABILITY AND CONTINUOUS OPERATION

Although second-generation designs were very successful, the past 10 years have brought new challenges.

Increased competition led to the need to lower operating expenses and a coherent case for network software flexible enough to be redeployed in network devices across the larger part of the end-to-end packet path. From multiple terabit routers to Layer 2 switches and security appliances, the “best-in-class” catchphrase can no longer justify a splintered operational experience—true “network” operating systems are clearly needed. Such systems must also achieve continuous operation, so that software failures in the routing code, as well as system upgrades, do not affect the state of the network. Meeting this challenge requires availability and convergence characteristics that go far beyond the hardware redundancy available in second-generation routers.

Another key goal of third-generation operating systems is the capability to run with zero downtime (planned and unplanned). Drawing on the lesson learned from previous designs regarding the difficulty of moving from one OS to another, third-generation operating systems also should make the migration path completely transparent to customers. They must offer an evolutionary, rather than revolutionary upgrade experience typical to the retirement process of legacy software designs.

BASIC OS DESIGN CONSIDERATIONS

Choosing the right foundation (prototype) for an operating system is very important, as it has significant implications for the overall software design process and final product quality and serviceability. This importance is why OEM vendors sometimes migrate from one prototype platform to another midway through the development process, seeking

a better fit. Generally, the most common transitions are from a proprietary to a commercial code base and from a commercial code base to an open-source software foundation.

Regardless of the initial choice, as networking vendors develop their own code, they get further and further away from the original port, not only in protocol-specific applications but also in the system area. Extensions such as control plane redundancy, in-service software upgrades and multi chassis operation require significant changes on all levels of the original design. However, it is highly desirable to continue borrowing content from the donor OS in areas that are not normally the primary focus of networking vendors, such as improvements in memory management, scheduling, multi core and symmetric multiprocessing (SMP) support, and host hardware drivers. With proper engineering discipline in place, the more active and peer-reviewed the donor OS is, the more quickly related network products can benefit from new code and technology.

This relationship generally explains another market trend—only two out of five network operating systems that emerged in the routing markets over the past 10 years used a commercial OS as a foundation.

Juniper's main operating system, Junos OS, is an excellent illustration of this industry trend. The basis of the Junos OS kernel comes from the FreeBSD UNIX OS, an open-source software system. The Junos OS kernel and infrastructure have since been heavily modified to accommodate advanced and unique features such as state

replication, nonstop active routing and in-service software upgrades, all of which do not exist in the donor operating system. Nevertheless, the Junos OS tree can still be synchronized with the FreeBSD repository to pick the latest in system code, device drivers and development tool chains, which allows Juniper Networks engineers to concentrate on network-specific development.

COMMERCIAL VERSUS OPEN-SOURCE DONOR OS

The advantage of a more active and popular donor OS is not limited to just minor improvements—the cutting edge of technology creates new dimensions of product flexibility and usability. Not being locked into a single-vendor framework and roadmap enables greater control of product evolution as well as the potential to gain from progress made by independent developers. This benefit is evident in Junos OS, which became a first commercial product to offer hard resource separation of the control plane and a real-time software data plane. Juniper-specific extension of the original BSD system architecture relies on multicore CPUs and makes Junos OS the only operating system that powers both low-end software-only systems and high-end multiple-terabit hardware platforms with images built from the same code tree. This technology and experience could not be created without support from the entire Internet-driven community. The powerful collaboration between leading individuals, universities and commercial organizations helps Junos OS stay on the very edge of operating system development. Further, this collaboration works both ways:

Juniper donates to the free software movement, one example being the Juniper Networks FreeBSD/MIPS port.

FUNCTIONAL SEPARATION AND PROCESS SCHEDULING

Multiprocessing, functional separation and scheduling are fundamental for almost any software design, including network software. Because CPU and memory are shared resources, all running threads and processes have to access them in a serial and controlled fashion. Many design choices are available to achieve this goal, but the two most important are the memory model and the scheduling discipline.

MEMORY MODEL

The memory model defines whether processes (threads) run in a common memory space. If they do, the overhead for switching the threads is minimal, and the code in different threads can share data via direct memory pointers.

The downside is that a runaway process can cause damage in memory that does not belong to it.

In a more complex memory model, threads can run in their own virtual machines, and the operating system switches the context every time the next thread needs to run. Because of this context switching, direct communication between threads is no longer possible and requires special Inter Process Communication (IPC) structures such as pipes, files and shared memory pools.

SCHEDULING DISCIPLINE

Scheduling choices are primarily between cooperative and preemptive models, which define whether thread switching

happens voluntarily. A cooperative multitasking model allows the thread to run to completion, and a preemptive design ensures that every thread gets access to the CPU regardless of the state of other threads.

VIRTUAL MEMORY/PREEMPTIVE SCHEDULING PROGRAMMING MODEL

Virtual memory with preemptive scheduling is a great design choice for properly constructed functional blocks, where interaction between different modules is limited and well defined. This technique is one of the main benefits of the second-generation OS designs and underpins the stability and robustness of contemporary network operating systems. However, it has its own drawbacks.

Notwithstanding the overhead associated with context switching, consider the interaction between two threads, A and B, both relying on the common resource R. Because threads do not detect their relative scheduling in the preemptive model, they can actually access R in a different order and with varying intensity. For example, R can be accessed by A, then B, then A, then A and then B again. If thread B modifies resource R, thread A may get different results at different times—and without any predictability. For instance, if R is an interior gateway protocol (IGP) next hop, B is an IGP process, and A is a BGP process, then BGP route installation may fail because the underlying next hop was modified midway through routing table modification. This scenario would never happen in the cooperative multitasking model, because the IGP process would release the CPU only after it finishes the next-hop maintenance. This

problem is well researched and understood within software design theory, and special solutions such as resource locks and synchronization primitives are easily available in nearly every operating system. However, the effectiveness of IPC depends greatly on the number of interactions between different processes. As the number of interacting processes increases, so does the number of IPC operations. In a carefully designed system, the number of IPC operations is proportional to the number of processes (N). In a system with extensive IPC activity, this number can be proportional to N^2 .

Exponential growth of an IPC map is a negative trend not only because of the associated overhead, but because of the increasing number of unexpected process interactions that may escape the attention of software engineers.

In practice, overgrown IPC maps result in systemwide “IPC meltdowns” when major events trigger intensive interactions. For instance, pulling a line card would normally affect interface management, IGP, exterior gateway protocol and traffic engineering processes, among others. When interprocess interactions are not well contained, this event may result in locks and tight loops, with multiple threads waiting on each other and vital system operations such as routing table maintenance and IGP computations temporarily suspended. Such defects are signatures of improper modularization, where similar or heavily interacting functional parts do not run as one process or one thread.

The right question to ask is, “Can a system be too modular?” The conventional wisdom says, “Yes.” Excessive modularity can bring long-term problems, with code

complexity, mutual locks and unnecessary process interdependencies. Although none of these may be severe enough to halt development, feature velocity and scaling parameters can be affected. Complex process interactions make programming for such a network OS an increasingly difficult task.

On the other hand, the cooperative multitasking, shared memory paradigm becomes clearly suboptimal if unrelated processes are influencing each other via the shared memory pool and collective restartability. A classic problem of first-generation operating systems was systemwide failure due to a minor bug in a nonvital process such as SNMP or network statistics. Should such an error occur in a protected and independently restartable section of system code, the defect could easily be contained within its respective code section.

This brings us to an important conclusion. No fixed principle in software design fits all possible situations. Ideally, code design should follow the most efficient paradigm and apply different strategies in different parts of the network OS to achieve the best marriage of architecture and function. This approach is evident in Junos OS, where functional separation is maintained so that cooperative multitasking and preemptive scheduling can both be used effectively, depending on the degree of IPC containment between functional modules.

GENERIC KERNEL DESIGN

Kernels normally do not provide any immediately perceived or revenue-generating functionality. Instead, they

perform housekeeping activities such as memory allocation and hardware management and other system-level tasks. Kernel threads are likely the most often run tasks in the entire system. Consequently, they have to be robust and run with minimal impact on other processes.

In the past, kernel architecture largely defined the operating structure of the entire system with respect to memory management and process scheduling. Hence, kernels were considered important differentiators among competing designs.

Historically, the disputes between the proponents and opponents of lightweight versus complex kernel architectures came to a practical end when most operating systems became functionally decoupled from their respective kernels.

Once software distributions became available with alternate kernel configurations, researchers and commercial developers were free to experiment with different designs.

For example, the original Carnegie-Mellon Mach microkernel was originally intended to be a drop-in replacement for the kernel in BSD UNIX and was later used in various operating systems, including mkLinux and GNU FSF projects. Similarly, some software projects that started life as purely microkernel-based systems later adopted portions of monolithic designs. Over time, the radical approach of having a small kernel and moving system functions into the user-space processes did not prevail. A key reason for this was the overhead associated with extra context switches between frequently executed system tasks running in separate memory spaces. Furthermore, the benefits associated with restart ability of essentially all

system processes proved to be of limited value, especially in embedded systems. With the system code being very well tested and limited to scheduling, memory management and a handful of device drivers, the potential errors in kernel subsystems are more likely to be related to hardware failures than to software bugs. This means, for example, that simply restarting a faulty disk driver is unlikely to help the routing engine stay up and running, as the problem with storage is likely related to a hardware failure (for example, uncorrectable fault in a mass storage device or system memory bank).

Another interesting point is that although both monolithic and lightweight kernels were widely studied by almost all operating system vendors, few have settled on purist implementations.

For example, Apple's Mac OS X was originally based on microkernel architecture, but now runs system processes, drivers and the operating environment in BSD-like subsystems. Microsoft NT and derivative operating systems also went through multiple changes, moving critical performance components such as graphical and I/O subsystems in and out of the system kernel to find the right balance of stability, performance and predictability. These changes make NT a hybrid operating system. On the other hand, freeware development communities such as FSF, FreeBSD and NetBSD have mostly adopted monolithic designs (for example, Linux kernel) and have gradually introduced modularity into selected kernel sections (for example, device drivers). So what difference does kernel architecture make to routing and control?

RECEIVING MULTICAST MESSAGES

Receiving a multicast message is a bit trickier. Since multicast messages are generally ignored unless someone has explicitly registered an interest in a particular address, there is a bit of setup that needs to be done.

Here's the code for receiving multicast messages:

```
require 'socket'
require 'ipaddr'
MULTICAST_ADDR = "225.4.5.6"
PORT = 5000
ip = IPAddr.new(MULTICAST_ADDR).hton +
IPAddr.new("0.0.0.0").hton
sock = UDPSocket.new
sock.setsockopt(Socket::IPPROTO_IP,
Socket::IP_ADD_MEMBERSHIP, ip)
sock.bind(Socket::INADDR_ANY, PORT)
loop do
  msg, info = sock.recvfrom(1024)
  puts "MSG: #{msg} from #{info[2]} (#{info[3]})/
#{info[1]} len #{msg.size}"
end
```

The tricky part was figuring out the right `setsockopt` options and values needed to register interest in our multicast address. I had to do a little reading in the Unix man pages on the C level `setsockopt()` function call. The third option to the C function is a structure that contains two 4-byte IP addresses. The first IP address is the multicast address, and the second IP address is the address of the local host adapter that we wish to use to listen for the multicast. The 0.0.0.0 address means use any of the local network adapters. `IPAddr` handles parsing the human readable form of the IP address and returns a string of 4 bytes in the order needed by the C level `setsockopt()` function.

USAGE

Save the above code in files named `send.rb` and `rcv.rb`. In one console window, type:

```
ruby rcv.rb
```

In another console window on the same or different machine (on the same local network), type:

```
ruby send.rb This is a test.
```

UPDATE

I added the Time To Live option on send.

POINT-TO-POINT COMMUNICATIONS

A point-to-point connection is a dedicated communication link between two systems or processes. Think of a wire that directly connects two systems. The systems use that wire exclusively to communicate. The opposite of point-to-point communications is broadcasting, where one system transmits to many. A telephone call is a circuit-oriented, point-to-point link between two phones. However, calls are usually multiplexed across telephone company trunks; so, while the circuit itself may be virtual, the users are engaging in a point-to-point communication session.

An end-to-end connection refers to a connection between two systems across a switched network. For example, the Internet is made up of a mesh of routers. Packets follow a hop-by-hop path from one router to the next to reach their destinations.

Each hop consists of a physical point-to-point link between routers. Therefore, a routed path consists of multiple point-to-point links. In the ATM and frame relay

environment, the end-to-end path is called a virtual circuit that crosses a predefined set of point-to-point links.

A shared LAN such as Ethernet provides a form of point-to-point communications. Keep in mind that on shared LANs, all nodes listen to signals on the cable, so broadcasting is supported. However, when one node addresses frames to another node and only that node receives the frames, one could say that the two nodes are engaged in point-to-point communications across a shared medium.

Point-to-multipoint connections are possible over multidrop links. A mainframe and its terminals is an example. The device that provides the multipoint connection is usually an intelligent controller that manages the flow of information from the multiple devices attached to it.

Point-to-point communications is defined in the physical and data link layers of the OSI protocol stack.

Process one (`myrank = 1`) receives this message with the receive operation `MPI_RECV`. The message to be received is selected according to the value of its envelope, and the message data is stored into the receive buffer. In the example above, the receive buffer consists of the storage containing the string message in the memory of process one. The first three parameters of the receive operation specify the location, size and type of the receive buffer. The next three parameters are used for selecting the incoming message. The last parameter is used to return information on the message just received.

The next sections describe the blocking send and receive operations. We discuss send, receive, blocking communication semantics, type matching requirements, type

conversion in heterogeneous environments, and more general communication modes. Non-blocking communication is addressed next, followed by channel-like constructs and send-receive operations. We then consider general datatypes that allow one to transfer efficiently heterogeneous and non-contiguous data. We conclude with the description of calls for explicit packing and unpacking of messages. The basic communication mechanism of MPI is the transmittal of data between a pair of processes, one side sending, the other, receiving. We call this "point to point communication." Almost all the constructs of MPI are built around the point to point operations and so this chapter is fundamental. It is also quite a long chapter since: there are many variants to the point to point operations; there is much to say in terms of the semantics of the operations; and related topics, such as probing for messages, are explained here because they are used in conjunction with the point to point operations.

MPI provides a set of send and receive functions that allow the communication of typed data with an associated tag. Typing of the message contents is necessary for heterogeneous support - the type information is needed so that correct data representation conversions can be performed as data is sent from one architecture to another. The tag allows selectivity of messages at the receiving end: one can receive on a particular tag, or one can wild-card this quantity, allowing reception of messages with any tag. Message selectivity on the source process of the message is also provided. A fragment of C code appears in Example for the example of process 0

sending a message to process 1. The code executes on both process 0 and process 1. Process 0 sends a character string using `MPI_Send()`. The first three parameters of the send call specify the data to be sent: the outgoing data is to be taken from `msg`; it consists of `strlen(msg)+1` entries, each of type `MPI_CHAR` (The string "Hello there" contains `strlen(msg)=11` significant characters. In addition, we are also sending the `tex2html_html_special_mark_quot` string terminator character). The fourth parameter specifies the message destination, which is process 1. The fifth parameter specifies the message tag. Finally, the last parameter is a communicator that specifies a communication domain for this communication. Among other things, a communicator serves to define a set of processes that can be contacted. Each such process is labeled by a process rank. Process ranks are integers and are discovered by enquiry to a communicator (see the call `MPI_Comm_rank()` to `MPI_COMM_WORLD`). `MPI_COMM_WORLD` is a default communicator provided upon start-up that defines an initial communication domain for all the processes that participate in the computation. Much more will be said about communicators in Chapter .

We have already said rather a lot about a simple transmittal of data from one process to another, but there is even more. To understand why, we examine two aspects of the communication: the semantics of the communication primitives, and the underlying protocols that implement them. Consider the previous example, on process 0, after the blocking send has completed. The

question arises: if the send has completed, does this tell us anything about the receiving process? Can we know that the receive has finished, or even, that it has begun?

Such questions of semantics are related to the nature of the underlying protocol implementing the operations. If one wishes to implement a protocol minimizing the copying and buffering of data, the most natural semantics might be the "rendezvous" rendezvous version, where completion of the send implies the receive has been initiated (at least). On the other hand, a protocol that attempts to block processes for the minimal amount of time will necessarily end up doing more buffering and copying of data and will have "buffering" semantics.

The trouble is, one choice of semantics is not best for all applications, nor is it best for all architectures. Because the primary goal of MPI is to standardize the operations, yet not sacrifice performance, the decision was made to include all the major choices for point to point semantics in the standard.

The above complexities are manifested in MPI by the existence of modes for point to point communication. Both blocking and non-blocking communications have modes. The mode allows one to choose the semantics of the send operation and, in effect, to influence the underlying protocol of the transfer of data.

In standard mode the completion of the send does not necessarily mean that the matching receive has started, and no assumption should be made in the application programme about whether the out-going data is buffered by MPI. In buffered mode the user

can guarantee that a certain amount of buffering space is available. The catch is that the space must be explicitly provided by the application programme. In synchronous mode synchronous mode a rendezvous semantics between sender and receiver is used. Finally, there is ready mode. ready mode This allows the user to exploit extra knowledge to simplify the protocol and potentially achieve higher performance. In a ready-mode send, the user asserts that the matching receive already has been posted. Modes are covered in Section .

INTERNET PROTOCOL ADDRESSING

An IP (Internet Protocol) address is a unique identifier for a node or host connection on an IP network. An IP address is a 32 bit binary number usually represented as 4 decimal values, each representing 8 bits, in the range 0 to 255 (known as octets) separated by decimal points. This is known as "dotted decimal" notation. Every IP address consists of two parts, one identifying the network and one identifying the node. The Class of the address and the subnet mask determine which part belongs to the network address and which part belongs to the node address.

There are 5 different address classes. You can determine which class any IP address is in by examining the first 4 bits of the IP address.

- Class A addresses begin with 0xxx, or 1 to 126 decimal.
- Class B addresses begin with 10xx, or 128 to 191 decimal.

- Class C addresses begin with 110x, or 192 to 223 decimal.
- Class D addresses begin with 1110, or 224 to 239 decimal.
- Class E addresses begin with 1111, or 240 to 254 decimal.
- Addresses beginning with 01111111, or 127 decimal, are reserved for loopback and for internal testing on a local machine; Class D addresses are reserved for multicasting; Class E addresses are reserved for future use. They should not be used for host addresses.

PRIVATE SUBNETS

There are three IP network addresses reserved for private networks. The addresses are 10.0.0.0, Subnet Mask 255.0.0.0, 172.16.0.0, Subnet Mask 255.240.0.0, and 192.168.0.0, Subnet Mask 255.255.0.0. These addresses are also notated 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16; this notation will be explained later in this tutorial. They can be used by anyone setting up internal IP networks, such as a lab or home LAN behind a NAT or proxy server or a router. It is always safe to use these because routers on the Internet by default will never forward packets coming from these addresses. These addresses are defined in RFC 1918.

IP ADDRESSING AND SUBNETTING

This document gives you basic information needed to configure your router for routing IP, such as how addresses are broken down and how subnetting works. You learn how

to assign each interface on the router an IP address with a unique subnet. There are many examples to help tie everything together.

Prerequisites

Requirements

There are no specific prerequisites for this document.

ADDITIONAL INFORMATION

If definitions are helpful to you, use these vocabulary terms to get you started:

- Address-The unique number ID assigned to one host or interface in a network.
- Subnet-A portion of a network sharing a particular subnet address.
- Subnet mask-A 32-bit combination used to describe which portion of an address refers to the subnet and which part refers to the host.
- Interface-A network connection.

If you have already received your legitimate address(es) from the Internet Network Information Center (InterNIC), you are ready to begin. If you do not plan to connect to the Internet, Cisco strongly suggests that you use reserved addresses from RFC 1918 .

NETWORK MASKS

A network mask helps you know which portion of the address identifies the network and which portion of the address identifies the node. Class A, B, and C networks have default masks, also known as natural masks, as shown here:

Class A: 255.0.0.0

Class B: 255.255.0.0

Class C: 255.255.255.0

An IP address on a Class A network that has not been subnetted would have an address/mask pair similar to: 8.20.15.1 255.0.0.0. To see how the mask helps you identify the network and node parts of the address, convert the address and mask to binary numbers.

8.20.15.1 = 00001000.00010100.00001111.00000001
255.0.0.0 = 11111111.00000000.00000000.00000000

Once you have the address and the mask represented in binary, then identifying the network and host ID is easier. Any address bits which have corresponding mask bits set to 1 represent the network ID. Any address bits that have corresponding mask bits set to 0 represent the node ID.

8.20.15.1 = 00001000.00010100.00001111.00000001
255.0.0.0 = 11111111.00000000.00000000.00000000
net id | host id
netid = 00001000 = 8
hostid = 00010100.00001111.00000001 = 20.15.1

APPLICATION OF LAYERS

APPLICATION LAYER 7

It is employed in software packages which implement client-server software. When an application on one computer starts communicating with another computer, then the Application layer is used. The header contains parameters that are agreed between applications. This header is often only sent at the beginning of an application operation. Examples of services within the application layer include:

- FTP
- DNS
- SNMP

- SMTP gateways
- Web browser
- Network File System (NFS)
- Telnet and Remote Login (rlogin)
- X.400
- FTAM
- Database software
- Print Server Software

PRESENTATION LAYER 6

This provides function call exchange between host operating systems and software layers. It defines the format of data being sent and any encryption that may be used, and makes it presentable to the Application layer. Examples of services used are listed below:

- MIDI
- HTML
- GIF
- TIFF
- JPEG
- ASCII
- EBCDIC

SESSION LAYER 5

The Session layer defines how data conversations are started, controlled and finished. The Session layer manages the transaction sequencing and in some cases authorisation. The messages may be bidirectional and there may be many of them, the session layer manages these conversations and

creates notifications if some messages fail. Indications show whether a packet is in the middle of a conversation flow or at the end. Only after a completed conversation will the data be passed up to layer 6. Examples of Session layer protocols are listed below:

- RPC
- SQL
- NetBIOS names
- Appletalk ASP
- DECnet SCP

TRANSPORT LAYER 4

This layer is responsible for the ordering and reassembly of packets that may have been broken up to travel across certain media. Some protocols in this layer also perform error recovery. After error recovery and reordering the data part is passed up to layer 5. Examples are:

- TCP
- UDP
- SPX

NETWORK LAYER 3

This layer is responsible for the delivery of packets end to end and implements a logical addressing scheme to help accomplish this. This can be connectionless or connection-oriented and is independent of the topology or path that the data packets travel. Routing packets through a network is also defined at this layer plus a method to fragment large packets into smaller ones depending on MTUs for different media (Packet Switching). Once the data from layer 2 has been

received, layer 3 examines the destination address and if it is the address of its own end station, it passes the data after the layer 3 header to layer 4. Examples of Layer 3 protocols include:

- Appletalk DDP
- IP
- IPX
- DECnet

DATA LINK LAYER 2

This layer deals with getting data across a specific medium and individual links by providing one or more data link connections between two network entities. End points are specifically identified, if required by the Network layer Sequencing. The frames are maintained in the correct sequence and there are facilities for Flow control and Quality of Service parameters such as Throughput, Service Availability and Transit Delay.

Examples include:

- IEEE 802.2
- IEEE 802.3
- 802.5 - Token Ring
- HDLC
- Frame Relay
- FDDI
- ATM
- PPP

The Data link layer performs the error check using the Frame Check Sequence (FCS) in the trailer and discards the

frame if an error is detected. It then looks at the addresses to see if it needs to process the rest of the frame itself or whether to pass it on to another host. The data between the header and the trailer is passed to layer 3.

The MAC layer concerns itself with the access control method and determines how use of the physical transmission is controlled and provides the token ring protocols that define how a token ring operates. The LLC shields the higher level layers from concerns with the specific LAN implementation.

PHYSICAL LAYER 1

This layer deals with the physical aspects of the media being used to transmit the data. The electrical, mechanical, procedural and functional means This defines things like pinouts, electrical characteristics, modulation and encoding of data bits on carrier signals.

It ensures bit synchronisation and places the binary pattern that it receives into a receive buffer. Once it decodes the bit stream, the physical layer notifies the data link layer that a frame has been received and passes it up.

Examples of specifications include:

- V.24
- V.35
- EIA/TIA-232
- EIA/TIA-449
- FDDI
- 802.3
- 802.5
- Ethernet

- RJ45
- NRZ
- NRZI

You will notice that some protocols span a number of layers (*e.g.* NFS, 802.3, etc.). A benefit of the seven layer model is that software can be written in a modular way to deal specifically with one or two layers only, this is often called Modular Engineering. Each layer has its own header containing information relevant to its role.

This header is passed down to the layer below which in turn adds its own header (encapsulates) until eventually the Physical layer adds the layer 2 information for passage to the next device which understands the layer 2 information and can then strip each of the layers' headers in turn to get at the data in the right location.

ASYNCHRONOUS TRANSFER MODE

The next generation of computer and telecommunication networks will use the asynchronous transfer mode (ATM). ATM networks use small fixed size cells to transmit information. This allows them to share the same network for voice, video, and data at a wide range of distances. Most computer and telecommunication companies are working on ATM products and services.

In this tutorial seminar, we will give a brief history and overview of the ATM technology. The key features that allow this technology to work better than the current packet networks will be presented including why this technology allows wider ranges of distance and number of nodes. Various

networking interfaces that are being standardized will be explained.

ATM, or Asynchronous Transfer Mode is a solid option for high-bandwidth network deployments that require guaranteed service levels and robust reliability. Backed by years of technology maturity, ATM now offers a widely supported platform for building global networks as well as high bandwidth, inter-regional trunks.

Pacnet ATM Service is a high speed, scalable networking service that enables service providers or enterprise customers to transport multiple network applications (voice, data and video). ATM offers quality of service (QoS) over a single connection between geographically distant locations via a public network of shared circuits, that is scalable from 2Mbps to 155Mbps.

With Pacnet's ATM service, you get access to a state-of-the-art ATM network that spans all key business centers in more than 25 countries.

FEATURES

- Wide range of ATM port speed, including nxT1/E1, DS3 and STM1.
- Any-to-any end-point connectivity through PVC implementation.
- Full internetworking support between ATM and Frame Relay to support smaller branch connections.
- Supports different application mixes through multiple Classes-of-Service, including CBR for real-time applications, VBR-rt for time-sensitive applications, VBR-nrt for non-real time applications with bursty

traffic characteristics and UBR, best effort service for non-critical applications.

- Upgrade and downgrade options to match market requirements.
- Fully managed services portfolio and industry-leading SLAs with the following components.
 - Network Availability
 - Network Transit Delay
 - Cell Delivery
 - Mean Time-To-Repair
 - Installation
- Expert support by network design and provisioning professionals focused on both PoP-to-PoP and last mile access.

TRANSMISSION OF DATA IN TUNNELING

DEFINITION

Tunneling, also known as "port forwarding," is the transmission of data intended for use only within a private, usually corporate network through a public network in such a way that the routing nodes in the public network are unaware that the transmission is part of a private network. Tunneling is generally done by encapsulating the private network data and protocol information within the public network transmission units so that the private network protocol information appears to the public network as data. Tunneling allows the use of the Internet, which is a public network, to convey data on behalf of a private network.

One approach to tunneling is the Point-to-Point Tunneling Protocol (PPTP) developed by Microsoft and several other companies. The PPTP keeps proprietary data reasonably secure, even though part of the path(s) between or among end users exists in public communication channels. The PPTP makes it possible for authorized users to gain access to a private network - called a virtual private network (VPN) - through an Internet service provider (ISP) or online service.

Another commonly used tunneling protocol is generic routing encapsulation (GRE), developed by Cisco Systems. There are numerous, less common tunneling protocols. Tunneling, and the use of a VPN, is not intended as a substitute for encryption/decryption. In cases where a high level of security is necessary, the strongest possible encryption should be used within the VPN itself, and tunneling should serve only as a convenience.

EXPLANATION

Tunneling, or port forwarding, is a way to forward otherwise insecure TCP traffic through SSH Secure Shell for Workstations. You can secure for example POP3, SMTP and HTTP connections that would otherwise be insecure.




Fig. Encrypted SSH2 tunnel

The tunneling capability of SSH Secure Shell for Workstations is a feature that allows, for example, company employees to access their e-mail, company intraweb pages

and shared files securely by even when working from home or on the road.

Tunneling makes it possible to access e-mail from any type of Internet service (whether accessed via modem, a DSL line or a cable connection, or a hotel Internet service). As long as the user has an IP connection to the Internet she can get her mail and access other resources from anywhere in the world securely.

This often is not the case with more traditional IPsec based VPN technologies because of issues with traversing networks that are implementing Network Address Translation (NAT) - this is especially the case in hotels. NAT breaks an IPsec connection unless special protocols such as NAT-Traversal are implemented on the client and gateway.

The client-server applications using the tunnel will carry out their own authentication procedures, if any, the same way they would without the encrypted tunnel. The protocol/application might only be able to connect to a fixed port number (*e.g.* IMAP 143). Otherwise any available port can be chosen for port forwarding.

HTTP-TUNNELING

If you answer yes to any of these, HTTP-Tunnel is the answer:

- Need to bypass any firewall
- Need secure internet browsing
- Need to use favourite programmes with out being monitored by work, school, ISP or gov.
- Extra security for online transactions
- Encrypt all your Internet traffic.

- Need play online games
- Visit sites that you are previously blocked
- Prevent 3rd party monitoring or regulation of your Internet browsing and downloads
- Use your favourite applications previously blocked
- Hide your IP address
- Make it next to impossible for you to identify online.
- Free unlimited data transfer
- Compatible with most major Internet applications
- Secure and virus-free servers
- 99 per cent uptime
- No spam, pop-ups, or banners

APPLICATIONS

A list of some of the more popular supported applications can be found here. Also, user guides to help you configure your HTTP-Tunnel Client can be found.

LOW BANDWIDTH SERVICE FEATURES

- Free option for moderate users
- Sufficient for Instant Messaging programmes (Trillian, ICQ, etc.)
- Good for users who only need to access blocked web sites
- No spam, pop-ups, or banners

HIGH BANDWIDTH SERVICE FEATURES

- High speed option for power users
- No bandwidth limitation

- Ideal for users behind a high speed internet connection (Cable/DSL modem, T1, etc.)
- Good for applications which are transfer-intensive (peer-to-peer programmes, FTP, etc.)
- No spam, pop-ups, or banners

If any technical problems are encountered, it is recommended to visit the support page for assistance prior to asking for a refund.

OS REQUIREMENTS

HTTP-Tunnel runs on a variety of Windows operating systems (Windows 98, ME, NT 4.X, 2000, XP and Vista). There are currently no releases available for Unix or MacIntosh platforms.

Tunneling is a way in which data is transferred between two networks securely. All the data that is being transferred are fragmented into smaller packets or frames and then passed through the tunnel. This process is different from a normal data transfer between nodes. Every frame passing through the tunnel will be encrypted with an additional layer of tunneling encryption and encapsulation which is also used for routing the packets to the right direction. This encapsulation would then be reverted at the destination with decryption of data which is later sent to the desired destined node.

A tunnel is a logical path between the source and the destination endpoints between two networks. Every packet is encapsulated at the source will be de-capsulated at the destination. This process will keep happening as long as the logical tunnel is persistent between the two endpoints.

MONOLITHIC VERSUS MICROKERNEL NETWORK OPERATING SYSTEM DESIGNS

In the network world, both monolithic and microkernel designs can be used with success.

However, the ever-growing requirements for a system kernel quickly turn any classic implementation into a compromise. Most notably, the capability to support a real-time forwarding plane along with stateful and stateless forwarding models and extensive state replication requires a mix of features not available from any existing monolithic or microkernel OS implementation.

This lack can be overcome in two ways.

First, a network OS can be constrained to a limited class of products by design. For instance, if the OS is not intended for mid- to low-level routing platforms, some requirements can be lifted. The same can be done for flow-based forwarding devices, such as security appliances. This artificial restriction allows the network operating systems to stay closer to their general-purpose siblings—at the cost of fracturing the product lineup. Different network element classes will now have to maintain their own operating systems, along with unique code bases and protocol stacks, which may negatively affect code maturity and customer experience.

Second, the network OS can evolve into a specialized design that combines the architecture and advantages of multiple classic implementations.

This custom kernel architecture is a more ambitious development goal because the network OS gets further away from the donor OS, but the end result can offer the benefits

of feature consistency, code maturity, and operating experience. This is the design path that Juniper selected for Junos OS.

JUNOS OS KERNEL

According to the formal criteria, the Junos OS kernel is fully customizable. At the very top is a portion of code that can be considered a microkernel. It is responsible for real-time packet operations and memory management, as well as interrupts and CPU resources. One level below it is a more conventional kernel that contains a scheduler, memory manager and device drivers in a package that looks more like a monolithic design.

Finally, there are user-level (POSIX) processes that actually serve the kernel and implement functions normally residing inside the kernels of classic monolithic router operating systems. Some of these processes can be compound or run on external CPUs (or packet forwarding engines). In Junos OS, examples include periodic hello management, kernel state replication, and protected system domains (PSDs). The entire structure is strictly hierarchical, with no underlying layers dependent on the operations of the top layers.

This high degree of virtualization allows the Junos OS kernel to be both fast and flexible.

However, even the most advanced kernel structure is not a revenue-generating asset of the network element.

Uptime is the only measurable metric of system stability and quality. This is why the fundamental difference between the Junos OS kernel and competing designs lies in the focus

on reliability. Coupled with Juniper's industry-leading nonstop active routing and system upgrade implementation, kernel state replication acts as the cornerstone for continuous operation. In fact, the Junos OS redundancy scheme is designed to protect data plane stability and routing protocol adjacencies at the same time. With in-service software upgrade, networks powered by Junos OS are becoming immune to the downtime related to the introduction of new features or bug fixes, enabling them to approach true continuous operation.

Continuous operation demands that the integrity of the control and forwarding planes remains intact in the event of failover or system upgrades, including minor and major release changes. Devices running Junos OS will not miss or delay any routing updates when either a failure or a planned upgrade event occurs.

This goal of continuous operation under all circumstances and during maintenance tasks is ambitious, and it reflects Juniper's innovation and network expertise, which is unique among network vendors.

PROCESS SCHEDULING IN JUNOS OS

Innovation in Junos OS does not stop at the kernel level; rather, it extends to all aspects of system operation.

As mentioned before, there are two tiers of schedulers in Junos OS, the topmost becoming active in systems with a software data plane to ensure the real-time handling of incoming packets. It operates in real time and ensures that quality of service (QoS) requirements are met in the forwarding path.

The second-tier (non-real-time) scheduler resides in the base Junos OS kernel and is similar to its FreeBSD counterpart. It is responsible for scheduling system and user processes in a system to enable preemptive multitasking.

In addition, a third-tier scheduler exists within some multithreaded user-level processes, where threads operate in a cooperative, multitasking model. When a compound process gets the CPU share, it may treat it like a virtual CPU, with threads taking and leaving the processor according to their execution flow and the sequence of atomic operations.

This approach allows closely coupled threads to run in a cooperatively multitasking environment and avoid being entangled in extensive IPC and resource-locking activities.

ARCHITECTURE AND INFRASTRUCTURE

PARALLELISM

Advances in multicore CPU development and the capability to run several routing processors in a system constitute the basis for increased efficiency in a router control plane. However, finding the right balance of price and performance can also be very difficult.

Unlike the data mining and computational tasks of supercomputers, processing of network updates is not a static job. A block of topology changes cannot be prescheduled and then sliced across multiple CPUs. In routers and switches, network state changes asynchronously (as events happen), thus rendering time-based load sharing irrelevant. Sometimes vendors try to solve this dilemma by statically sharing the load in functional, rather than

temporal, domains. In other words, they claim that if the router OS can use separate routing processors for different tasks (for example, OSPF or BGP), it can also distribute the bulk of data processing across multiple CPUs.

To understand whether this is a valid assumption, let's consider a typical CPU utilization capture. What is interesting here is that the different processes are not computationally active at the same time—OSPF and BGP do not compete for CPU cycles. Unless the router runs multiple same-level protocols simultaneously, the well-designed network protocol stack stays fairly orthogonal. Different protocols serve different needs and seldom converge at the same time.

Show Processes CPU Seconds Unicast

<i>Date</i>	<i>Average</i>	<i>RIP</i>	<i>OSPF</i>	<i>BGP</i>	<i>RIPng</i>	<i>OSPF6</i>	<i>BGP4+</i>
01/22 15:48:19	3	0 0	0 0	0 0	0	0 0	0
01:22 15:48:20	3	0 1	0 0	0	0	0 0	0
01/22 15:49:18	3	0 0	1 0	0	0	0 0	0

Typical CPU times capture (from NEC 8800 product documentation).

For instance, an IGP topology change may trigger a Dijkstra algorithm computation; until it is complete, BGP nexthop updates do not make much sense. At the same time, all protected MPLS LSPs should fall on precomputed alternate paths and not cause major RSVP activities.

Thus, the gain from placing different processes of a single control plane onto physically separate CPUs may be limited, while the loss from the overhead functions such as synchronization and distributed memory unification may be significant.

Does this mean that the concept of parallelism is not applicable to the routing processors? Not at all. Good coding practice and modern compilers can make excellent use of

multicore and SMP hardware, while clustered routing engines are indispensable when building multichassis (single control and data plane spanning multiple chassis) or segmented (multiple control and data planes within a single physical chassis) network devices. Furthermore, high-end designs may allow for independent scaling of control and forwarding planes, as implemented in the highly acclaimed Juniper Networks JCS1200 Control System.

With immediate access to state-of-the-art processor technology, Juniper Networks engineers heavily employ parallelism in the Junos OS control plane design, targeting both elegance and functionality.

A functional solution is the one that speeds up the control plane without unwanted side effects such as limitations in forwarding capacity. When deployed in a JCS1200, Junos OS can power multiple control plane instances (system domains) at the same time without consuming revenue-generating slots in the router chassis. Moreover, the Junos OS architecture can run multiple routing systems (including third-party code) from a single rack of routing engines, allowing an arbitrary mix-and-match of control plane and data plane resources within a point of presence (POP).

These unique capabilities translate into immediate CAPEX savings, because a massively parallel control plane can be built independent of the forwarding plane and will never confront a limited common resource (such as the number of physical routers or a number of slots in each chassis).

Elegance means the design should also bring other technical advantages: for instance, bypassing space and

power requirements associated with the embedded chassis and thus enabling use of faster silicon and speeding up the control plane. Higher CPU speed and memory limits can substantially improve the convergence and scaling characteristics of the entire routing domain.

The goal of Juniper design philosophy is tangible benefits to our customers—without cutting corners.

FLEXIBILITY AND PORTABILITY

A sign of a good OS design is the capability to adapt the common software platform to various needs. In the network world, this equates to the adoption of new hardware and markets under the same operating system.

The capability to extend the common operating system over several products brings the following important benefits to customers:

- Reduced OPEX from consistent UI experience and common management interface
- Same code for all protocols; no unique defects and interoperability issues
- Common schedule for software releases; a unified feature set in the control plane
- Accelerated technology introduction; once developed, the feature ships on many platforms.

Technology companies are in constant search of innovation both internally and externally. New network products can be developed in-house or within partnerships or acquired. Ideally, a modern network OS should be able to absorb domestic (internal) hardware platforms as well as foreign (acquired) products, with the latter being gradually

folded into the mainstream software line. The capability to absorb in-house and foreign innovations in this way is a function of both software engineering discipline and a flexible, well-designed OS that can be adapted to a wide range of applications.

On the contrary, the continuous emergence of internally developed platforms from the same vendor featuring different software trains and versions can signify the lack of a flexible and stable software foundation.

For example, when the same company develops a core router with one OS, an Ethernet switch with another and a data centre switch with a third, this likely means that in-house R&D groups considered and rejected readily available OS designs as impractical or unfit.

Although partial integration may still exist through a unified command-line interface (CLI) and shared code and features, the main message is that the existing software designs were not flexible enough to be easily adapted to new markets and possibilities. As a result, customers end up with a fractured software lineup, having to learn and maintain loosely related or completely unrelated software trains and develop expertise in all of them—an operationally suboptimal approach.

3

Computer Security Model

A computer security model is a scheme for specifying and enforcing security policies. A security model may be founded upon a formal model of access rights, a model of computation, a model of distributed computing, or no particular theoretical grounding at all. For a more complete list of available articles on specific security models, see [Category:Computer security models](#).

NETWORK SECURITY

In the field of networking, the area of network security consists of the provisions and policies adopted by the network administrator to prevent and monitor unauthorized access, misuse, modification, or denial of the computer network and network-accessible resources.

THE FIRST STEP TO INFORMATION SECURITY

The term network security and information security are often used interchangeably. Network security is generally taken as providing protection at the boundaries of an organization by keeping out intruders (hackers). Information security, however, explicitly focuses on protecting data resources from malware attack or simple mistakes by people within an organization by use of data loss prevention (DLP) techniques. One of these techniques is to compartmentalize large networks with internal boundaries.

NETWORK SECURITY CONCEPTS

Network security starts from authenticating the user, commonly with a username and a password. Since this requires just one thing besides the user name, i.e. the password which is something you 'know', this is sometimes termed one factor authentication. With two factor authentication something you 'have' is also used (e.g. a security token or 'dongle', an ATM card, or your mobile phone), or with three factor authentication something you 'are' is also used (e.g. a fingerprint or retinal scan). Once authenticated, a firewall enforces access policies such as what services are allowed to be accessed by the network users. Though effective to prevent unauthorized access, this component may fail to check potentially harmful content such as computer worms or Trojans being transmitted over the network. Anti-virus software or an intrusion prevention system (IPS) help detect and inhibit the action of such malware. An anomaly-based intrusion detection system may also monitor the network and traffic for unexpected (i.e.

suspicious) content or behaviour and other anomalies to protect resources, e.g. from denial of service attacks or an employee accessing files at strange times. Individual events occurring on the network may be logged for audit purposes and for later high level analysis. Communication between two hosts using a network could be encrypted to maintain privacy. *Honeypots*, essentially decoy network-accessible resources, could be deployed in a network as surveillance and early-warning tools as the honeypot will not normally be accessed. Techniques used by the attackers that attempt to compromise these decoy resources are studied during and after an attack to keep an eye on new exploitation techniques. Such analysis could be used to further tighten security of the actual network being protected by the honeypot.

SECURITY MANAGEMENT

Security Management for networks is different for all kinds of situations. A small home or an office would only require basic security while large businesses will require high maintenance and advanced software and hardware to prevent malicious attacks from hacking and spamming.

SMALL HOMES

- A basic firewall or a unified threat management system.
- For Windows users, basic Antivirus software. An anti-spyware programme would also be a good idea. There are many other types of antivirus or anti-spyware programmes out there to be considered.

Special Focus on Computer Security Model

- When using a wireless connection, use a robust password. Also try to use the strongest security supported by your wireless devices, such as WPA2 with AES encryption.
- If using Wireless: Change the default SSID network name, also disable SSID Broadcast; as this function is unnecessary for home use. (However, many security experts consider this to be relatively useless. <http://blogs.zdnet.com/Ou/index.php?p=43>)
- Enable MAC Address filtering to keep track of all home network MAC devices connecting to your router.
- Assign STATIC IP addresses to network devices.
- Disable ICMP ping on router.
- Review router or firewall logs to help identify abnormal network connections or traffic to the Internet.
- Use passwords for all accounts.
- Have multiple accounts per family member, using non-administrative accounts for day-to-day activities. Disable the guest account (Control Panel> Administrative Tools> Computer Management> Users).
- Raise awareness about information security to children.

MEDIUM BUSINESSES

- A fairly strong firewall or Unified Threat Management System
- Strong Antivirus software and Internet Security Software.

Special Focus on Computer Security Model

- For authentication, use strong passwords and change it on a bi-weekly/monthly basis.
- When using a wireless connection, use a robust password.
- Raise awareness about physical security to employees.
- Use an optional network analyzer or network monitor.
- An enlightened administrator or manager.

LARGE BUSINESSES

- A strong firewall and proxy to keep unwanted people out.
- A strong Antivirus software package and Internet Security Software package.
- For authentication, use strong passwords and change it on a weekly/bi-weekly basis.
- When using a wireless connection, use a robust password.
- Exercise physical security precautions to employees.
- Prepare a network analyzer or network monitor and use it when needed.
- Implement physical security management like closed circuit television for entry areas and restricted zones.
- Security fencing to mark the company's perimeter.
- Fire extinguishers for fire-sensitive areas like server rooms and security rooms.
- Security guards can help to maximize security.

SCHOOL

- An adjustable firewall and proxy to allow authorized users access from the outside and inside.

Special Focus on Computer Security Model

- Strong Antivirus software and Internet Security Software packages.
- Wireless connections that lead to firewalls.
- Children's Internet Protection Act compliance.
- Supervision of network to guarantee updates and changes based on popular site usage.
- Constant supervision by teachers, librarians, and administrators to guarantee protection against attacks by both internet and sneakernet sources.

LARGE GOVERNMENT

- A strong firewall and proxy to keep unwanted people out.
- Strong antivirus software and Internet Security Software suites.
- Strong encryption.
- Whitelist authorized wireless connection, block all else.
- All network hardware is in secure zones.
- All host should be on a private network that is invisible from the outside.
- Put web servers in a DMZ, or a firewall from the outside and from the inside.
- Security fencing to mark perimeter and set wireless range to this.

PROACTIVE CYBER DEFENCE

Proactive Cyber Defence means acting in anticipation to oppose an attack against computers and networks. It

represents the dynamic between purely offensive and defensive action; interdicting and disrupting an attack or a threat's preparation to attack, either pre-emptively or in self-defence. Proactive cyber defence will most often require operationalizing upstream security (security from the Cloud) mechanisms of the telecommunications/Internet providers. Some of the compelling reasons for a proactive defence strategy are about cost and choice. Decisionmakers have few choices after an impact and that all of them are costly. Proactive defence is key to mitigating operational risk.

BACKGROUND

In the Fifth century, B.C., Sun Tzu advocated “foreknowledge” or predictive analysis as part of a winning strategy. He warned that planners must have a precise understanding of the active threat and not “remain ignorant of the enemy’s condition.” The thread of proactive defence is spun throughout his teachings. Psychiatrist Viktor Frankl was likely the first to use of the term proactive in his 1946 book *Man’s Search for Meaning* to distinguish the act of taking responsibility for one’s own circumstances rather than attributing one’s condition to external factors. Later in 1982, the United States Department of Defense (DoD) used “proactive” as a contrary concept to “reactive” in assessing risk. In the framework of risk management ‘proactive’ meant taking initiative by acting rather than reacting to threat events. Conversely “reactive” measures respond to a stimulus or past events rather than predicting the event. In military science, then and now considers defence is the science-art of thwarting an attack.

Furthermore doctrine poses that if a party attacks an enemy who is about to attack this could be called active-defence. Defence is also a euphemism for war but does not carry the negative connotation of an offensive war. Usage in this way has broadened the term to include most military issues including offensive, which is implicitly referred to as active-defence. Politically the concept of national self-defence to counter a war of aggression refers to a defensive war involving pre-emptive offensive strikes and is one possible criterion in the 'Just War Theory'. Proactive defence has moved beyond theory. It has been put into practice in theatres of operation. In 1989, Stephen Covey's *The Seven Habits of Highly Effective People*, published by Free Press, transformed the meaning "to act before a situation becomes a source of confrontation or crisis." From that day "proactive" has been placed in opposition to the words "reactive" or "passive." Cyber is derived from "Cybernetics", a word originally coined by a group of scientists led by Norbert Wiener and made popular by Wiener's book of 1948, *Cybernetics or Control and Communication in the Animal and the Machine*. Cyberspace typically refers to the vast and growing logical domain composed of public and private networks; independently managed networks linked together through the lingua franca of the Internet, the Internet Protocol (IP).

The definition of Cyberspace has been extended to include all network-space which at some point, through some path, may have eventual access to the public internet. Under this definition, cyberspace becomes virtually every networked device in the world, which is not devoid of a network interface

entirely. There is no air-gap anymore between networks. The origins of cyber defence undoubtedly evolved from the original purpose of the Internet which was to harden military networks against the threat of a nuclear strike. Later cyber defence was coveted by the tenets of information warfare and information operations. The rapid evolution of information warfare operations doctrine in the 1990's embraced a proactive pre-emptive cyber defence strategy. "Information Warfare is an emergent reality that comes from a self-organization process that has never seen before. The problem is that we talk about it using terms that have well known connotations. And it is difficult to talk about something completely new using words that bring with them specific understanding and expectancies. The early period of the automobile faced a similar situation. At one time it was called a "horseless carriage" as this was the only way to define its essential quality. The car is more than a carriage without a horse. This is the dilemma we face when we discuss Information Warfare.

The danger is that the uses of familiar words misrepresent and mask the true extend of the revolution that will have to take place if we are to be able to retain a military capacity in a new physical, social and cognitive space." - Dr. Garigue, 1994. The National Strategy to Secure Cyberspace was published in February 2003 to outline an initial framework for both organizing and prioritizing efforts to secure the cyberspace. It highlighted the necessity for public private partnerships. Proactive threads include the call to deter malicious activity and prevent cyber attacks against America's critical infrastructures. The hype-cycle of

discussion reached its peak in 1994. Present-day proactive cyber defence strategy was conceived within the context of the rich discussion that preceded it, existing doctrine and real proactive cyber defence programmes that have evolved globally over the past decade. Dr. Robert John Garigue, a computational epistemologist and father of information warfare in Canada, published *Information Warfare, Developing a Conceptual Framework*. This was a landmark document in 1994 and genesis for proactive cyber defensive theory in Canada.

“Effective cyber defenses ideally prevent an incident from taking place. Any other approach is simply reactive. FedCIRC, the NIPC, the NSIRC, the Department of Defense and industry components realize that the best [action] is a pre-emptive and proactive approach.” - Sallie McDonald, the Assistant Commissioner for the Office Of Information Assurance and Critical Infrastructure Protection, Federal Technology Service and General Services Administration; in offering testimony with regard to the National Infrastructure Protection Center (NIPC) and the Federal Computer Incident Response Center or FedCIRC; before The Subcommittee on Terrorism Technology and Government Information Committee on Judiciary and the United States Senate July 25, 2001. The notion of a Proactive Pre-emptive Operations Group (P2OG) emerged from a report of the Defense Science Board (DSB), 2002 briefing. The briefing was reported by Dan Dupont in *Inside the Pentagon* on September 26, 2002 and was also discussed by William M. Arkin in the *Los Angeles Times* on October 27, 2002. The *Los Angeles Times* has subsequently quoted US Secretary of Defence Donald Rumsfeld revealing

the creation of the 'Proactive, Pre-emptive Operations Group.' The mission of the P2OG is reportedly to conduct Aggressive, Proactive, Pre-emptive Operations to interdiction and disruption the threat using: Psychological operations, Managed Information Dissemination, Precision Targeting, Information Warfare Operations, and SIGINT...

The proactive defence strategy is meant to improve information collection by stimulating reactions of the threat agents, provide strike options and to enhance operational preparation of the real or virtual battle space. The P2OG has been recommended to be constituted of "one hundred 'highly specialized people with unique technical and intelligence skills such as information operations, PSYOPS, network attack, covert activities, SIGINT, HUMINT, SOF, influence warfare/deception operations and to report to the National Security Council with an annual budget of \$100 million." The group would be overseen by the White House's deputy national security adviser and would carry out missions coordinated by the secretary of defense or the CIA director. "The proposal is the latest sign of a new assertiveness by the Defense Department in intelligence matters, and an indication that the cutting edge of intelligence reform is not to be found in Congress but behind closed doors in the Pentagon." - Steven Aftergood of the Federation of American Scientists. DoD doctrinally would initiate a 'pre-emptive' attack on the basis of evidence that an enemy attack is imminent. Proactive measures, according to DoD are those actions taken directly against the preventive stage of an attack by the enemy. Strike back doctrine aligns with pre-emptive and counter-attack tactics of a proactive cyber

defence strategy. The notion of 'proactive defence' has a rich history. The hype of 'Proactive cyber defence' reached its zenith around 1994. This period was marked by intense 'hype' discussions under the auspices of Information Warfare. Much of the current doctrine related to proactive cyber defence was fully developed by 1995.

A number of programmes were initiated then, and advanced to full operation by 2005 including those of hostile states. Meanwhile the public discussions diminished until the most recent resurgence in proactive cyber defence 2004-2008. Now most of the discussions around proactive defence in the literature are much less 'proactive' than the earlier discussions in 1994 or existing operational programmes. 'Proactive' is often used to hype marketing of security products or programmes, in much the same way that 'extreme' or 'quality' adjectives have been misused.

INTERNET PRIVACY

Internet privacy involves the desire or mandate of personal privacy concerning transactions or transmission of data via the Internet. It also involves the exercise of control over the type and amount of information revealed about a person on the Internet and who may access said information. Internet privacy forms a subset of computer privacy. A number of experts within the field of Internet security and privacy believe that privacy doesn't exist; "Privacy is dead – get over it" This should be more encouraged according to Steve Rambam, private investigator specializing in Internet privacy cases. In fact, it has been suggested that the "appeal of online services is to broadcast personal information on

purpose.” On the other hand, in his essay *The Value of Privacy*, security expert Bruce Schneier says, “Privacy protects us from abuses by those in power, even if we’re doing nothing wrong at the time of surveillance.”

LEVELS OF PRIVACY

People with only a casual concern for Internet privacy need not achieve total anonymity. Internet users may achieve an adequate level of privacy through controlled disclosure of personal information. The revelation of IP addresses, non-personally-identifiable profiling, and similar information might become acceptable trade-offs for the convenience that users could otherwise lose using the workarounds needed to suppress such details rigorously. On the other hand, some people desire much stronger privacy. In that case, they may try to achieve *Internet anonymity* to ensure privacy — use of the Internet without giving any third parties the ability to link the Internet activities to personally-identifiable information (P.I.I.) of the Internet user. In order to keep your information private, people need to be careful on what they submit and look at online. When filling out forms and buying merchandise, that becomes tracked and because your information was not private, companies are now sending you spam and advertising on similar products. The Sanders decision relied heavily on another California decision from a year ago. In *Shulman v. Group W Productions*, the court concluded that two people injured in a car accident could sue for invasion of privacy because a cameraman recorded emergency care given in a rescue helicopter. According to the court, while the accident victims

could not claim a reasonable expectation of privacy at the accident scene (where they were recorded by the same cameraman), they could claim a reasonable expectation of privacy in the rescue helicopter, even if they expected that their conversations in the helicopter would be overheard.

Because these cases make it more difficult to determine under what circumstances undercover reporting would violate a reasonable expectation of privacy-thus exposing journalists to liability-news organizations may think twice about their approach to investigative reporting. In California at least, as a result of these recent decisions, trial judges will be reluctant to throw out cases before trial, allowing them to go before a jury. And because media lawyers are uncertain about whether jurors would think that a privacy invasion was justified by a legitimate need to gather news, they are likely to offer conservative advice and deter stations from engaging in certain investigations. Resulting, once again, in a chilling effect on the media. Related State Laws Privacy of Personal Information: Nevada and Minnesota require Internet Service Providers to keep information private regarding their customers. This is only unless a customer approves their information being given out. According to the National Conference of State Legislators, the following states have certain laws on the personal privacy of its citizens.

Minnesota Statutes §§ 325M.01 to .09 -Prohibits Internet service providers from disclosing personally identifiable information, including a consumer's physical or electronic address or telephone number; Internet or online sites visited; or any of the contents of a consumer's data storage devices. Provides for certain circumstances under which information

Special Focus on Computer Security Model

must be disclosed, such as to a grand jury; to a state or federal law enforcement officer acting as authorized by law; pursuant to a court order or court action. Provides for civil damages of \$500 or actual damages and attorney fees for violation of the law. Nevada Revised Statutes § 205.498 - In addition, California and Utah laws, although not specifically targeted to on-line businesses, require all nonfinancial businesses to disclose to customers, in writing or by electronic mail, the types of personal information the business shares with or sells to a third party for direct marketing purposes or for compensation. Under the California law, businesses may post a privacy statement that gives customers the opportunity to choose not to share information at no cost. There are also certain laws for employees and businesses and privacy policies for websites. California, Connecticut, Nebraska and Pennsylvania all have specific privacy policies regarding websites, these include:

“California (Calif. Bus. & Prof. Code §§ 22575-22578) California’s Online Privacy Protection Act requires an operator, defined as a person or entity that collects personally identifiable information from California residents through an Internet Web site or online service for commercial purposes, to post conspicuously its privacy policy on its Web site or online service and to comply with that policy. The bill, among other things, would require that the privacy policy identify the categories of personally identifiable information that the operator collects about individual consumers who use or visit its Web site or online service and third parties with whom the operator may share the information.

Special Focus on Computer Security Model

Connecticut (Conn. Gen Stat. § 42-471) Requires any person who collects Social Security numbers in the course of business to create a privacy protection policy. The policy must be “publicly displayed” by posting on a web page and the policy must (1) protect the confidentiality of Social Security numbers, (2) prohibit unlawful disclosure of Social Security numbers, and (3) limit access to Social Security numbers. Nebraska (Nebraska Stat. § 87-302(14)) Nebraska prohibits knowingly making a false or misleading statement in a privacy policy, published on the Internet or otherwise distributed or published, regarding the use of personal information submitted by members of the public. Pennsylvania (18 Pa. C.S.A. § 4107(a)(10)) Pennsylvania includes false and misleading statements in privacy policies published on Web sites or otherwise distributed in its deceptive or fraudulent business practices statute.” There are also at least 16 states that require government websites to create privacy policies and procedures or to include machine-readable privacy policies into their websites. These states include Arizona, Arkansas, California, Colorado, Delaware, Iowa, Illinois, Maine, Maryland, Michigan, Minnesota, Montana, New York, South Carolina, Texas, Utah, and Virginia.

RISKS TO INTERNET PRIVACY

In today’s technological world, millions of individuals are subject to privacy threats. Companies are hired not only to watch what you visit online, but to infiltrate the information and send advertising based on your browsing history. People set up accounts for Facebook; enter bank and credit card

information to various websites. Those concerned about Internet privacy often cite a number of *privacy risks* — events that can compromise privacy — which may be encountered through Internet use. These methods of compromise can range from the gathering of statistics on users, to more malicious acts such as the spreading of spyware and various forms of bugs (software errors) exploitation. Privacy measures are provided on several social networking sites to try to provide their users with protection for their personal information. On Facebook for example privacy settings are available for all registered users. The settings available on Facebook include the ability to block certain individuals from seeing your profile, the ability to choose your “friends,” and the ability to limit who has access to your pictures and videos. Privacy settings are also available on other social networking sites such as E-harmony and MySpace. It is the user’s prerogative to apply such settings when providing personal information on the internet. In late 2007 Facebook launched the Beacon programme where user rental records were released on the public for friends to see. Many people were enraged by this breach in privacy, and the *Lane v. Facebook, Inc.* case ensued.

HTTP COOKIES

An HTTP cookie is data stored on a user’s computer that assists in automated access to websites or web features, or other state information required in complex web sites. It may also be used for user-tracking by storing special usage history data in a cookie. Cookies are a common concern in the field of privacy. As a result, some types of cookies are

classified as a *tracking cookie*. Although website developers most commonly use cookies for legitimate technical purposes, cases of abuse occur. In 2009, two researchers noted that social networking profiles could be connected to cookies, allowing the social networking profile to be connected to browsing habits. Systems do not generally make the user explicitly aware of the storing of a cookie. (Although some users object to that, it does not properly relate to Internet privacy. It does however have implications for computer privacy, and specifically for computer forensics. The original developers of cookies intended that only the website that originally distributed cookies to users so they could retrieve them, therefore returning only data already possessed by the website. However, in practice programmers can circumvent this restriction. Possible consequences include:

- the placing of a personally-identifiable tag in a browser to facilitate web profiling, or,
- use of cross-site scripting or other techniques to steal information from a user's cookies.

Some users choose to disable cookies in their web browsers – as of 2000 a Pew survey estimated the proportion of users at 4%. Such an action eliminates the potential privacy risks, but may severely limit or prevent the functionality of many websites. All significant web browsers have this disabling ability built-in, with no external programme required. As an alternative, users may frequently delete any stored cookies. Some browsers (such as Mozilla Firefox and Opera) offer the option to clear cookies automatically whenever the user closes the browser. A third

option involves allowing cookies in general, but preventing their abuse. There are also a host of wrapper applications that will redirect cookies and cache data to some other location. The process of *profiling* (also known as “tracking”) assembles and analyzes several events, each attributable to a single originating entity, in order to gain information (especially patterns of activity) relating to the originating entity. Some organizations engage in the profiling of people’s web browsing, collecting the URLs of sites visited. The resulting profiles can potentially link with information that personally identifies the individual who did the browsing.

Some web-oriented marketing-research organizations may use this practice legitimately, for example: in order to construct profiles of ‘typical Internet users’. Such profiles, which describe average trends of large groups of Internet users rather than of actual individuals, can then prove useful for market analysis. Although the aggregate data does not constitute a privacy violation, some people believe that the initial profiling does. Profiling becomes a more contentious privacy issue when data-matching associates the profile of an individual with personally-identifiable information of the individual. Governments and organizations may set up honeypot websites – featuring controversial topics – with the purpose of attracting and tracking unwary people. This constitutes a potential danger for individuals.

FLASH COOKIES

Flash cookies, also known as Local Shared Objects, work the same ways as normal cookies and are used by the Adobe Flash Player to store information at the user’s computer.

They exhibit a similar privacy risk as normal cookies, but are not as easily blocked, meaning that the option in most browsers to not accept cookies does not affect flash cookies. One way to view and control them is with browser extensions or add-ons.

EVERCOOKIES

An Evercookie is a JavaScript-based application which produces cookies in a web browser that actively “resist” deletion by redundantly copying themselves in different forms on the user’s machine (e.g.: Flash Local Shared Objects, various HTML5 storage mechanisms, window.name caching, etc.), and resurrecting copies are missing or expired.

PHOTOGRAPHS ON THE INTERNET

Today many people have digital cameras and post their photos online. The people depicted in these photos might not want to have them appear on the Internet. Some organizations attempt to respond to this privacy-related concern. For example, the 2005 Wikimania conference required that photographers have the prior permission of the people in their pictures. Some people wore a ‘no photos’ tag to indicate they would prefer not to have their photo taken. The Harvard Law Review published a short piece called “In The Face of Danger: Facial Recognition and Privacy Law,” much of it explaining how “privacy law, in its current form, is of no help to those unwillingly tagged.” Any individual can be unwillingly tagged in a photo and displayed in a manner that might violate them personally in some way, and by the time Facebook gets to taking down the photo,

many people will have already had the chance to view, share, or distribute it. Furthermore, traditional tort law does not protect people who are captured by a photograph in public because this is not counted as an invasion of privacy.

The extensive Facebook privacy policy covers these concerns and much more. For example, the policy states that they reserve the right to disclose member information or share photos with companies, lawyers, courts, government entities, etc. if they feel it absolutely necessary. The policy also informs users that profile pictures are mainly to help friends connect to each other.

However, these, as well as other pictures, can allow other people to invade a person's privacy by finding out information that can be used to track and locate a certain individual. In an article featured in ABC news, it was stated that two teams of scientists found out that Hollywood stars could be giving up information about their private whereabouts very easily through pictures uploaded to the Internet. Moreover, it was found that pictures taken by iPhones automatically attach the latitude and longitude of the picture taken through metadata unless this function is manually disabled.

SEARCH ENGINES

Search engines have the ability to track a user's searches. Personal information can be revealed through searches including search items used, the time of the search, and more. Search engines have claimed a necessity to retain such information in order to provide better services, protect against security pressure, and protect against fraud.

DATA LOGGING

Many programmes and operating systems are set up to perform data logging of usage. This may include recording times when the computer is in use, or which web sites are visited. If a third party has sufficient access to the computer, legitimately or not, the user's privacy may be compromised. This could be avoided by disabling logging, or by clearing logs regularly.

PRIVACY WITHIN SOCIAL NETWORKING SITES

Prior to the social networking site explosion over the past decade, there were early forms of social network technologies that included online multiplayer games, blog sites, news groups, mailings lists and dating services. These all created a backbone for the new modern sites, and even from the start of these older versions privacy was an issue. In 1996, a young woman in New York City was on a first date with an online acquaintance and later sued for sexual harassment as they went back to her apartment after when everything became too real. This is just an early example of many more issues to come regarding internet privacy. Social networking sites have become very popular within the last five years. With the creation of Facebook and the continued popularity of MySpace many people are giving their personal information out on the internet. These social networks keep track of all interactions used on their sites and save them for later use. Most users are not aware that they can modify the privacy settings and unless they modify them, their information is open to the public. On Facebook privacy settings can be accessed via the drop down menu under account in the top

right corner. There users can change who can view their profile and what information can be displayed on their profile. In most cases profiles are open to either “all my network and friends” or “all of my friends.” Also, information that shows on a user’s profile such as birthday, religious views, and relationship status can be removed via the privacy settings. If a user is under 13 years old they are not able to make a Facebook or a MySpace account, however, this is not regulated. Social networking has redefined the role of Internet privacy. Since users are willingly disclosing personal information online, the role of privacy and security is somewhat blurry. Sites such as Facebook, Myspace, and Twitter have grown popular by broadcasting status updates featuring personal information such as location. Facebook “Places,” in particular, is a Facebook service, which publicizes user location information to the networking community. Users are allowed to “check-in” at various locations including retail stores, convenience stores, and restaurants. Also, users are able to create their own “place,” disclosing personal information onto the Internet. This form of location tracking is automated and must be turned off manually. Various settings must be turned off and manipulated in order for the user to ensure privacy.

According to epic.org, Facebook users are recommended to: (1) disable “Friends can check me in to Places,” (2) customize “Places I Check In,” (3) disable “People Here Now,” and (4) uncheck “Places I’ve Visited.”. Moreover, the Federal Trade Commission has received two complaints in regards to Facebook’s “unfair and deceptive” trade practices, which are used to target advertising sectors of the online

community. “Places” tracks user location information and is used primarily for advertising purposes. Each location tracked allows third party advertisers to customize advertisements that suit one’s interests. Currently, the Federal Trade Commissioner along with the Electronic Privacy Information Center are shedding light on the issues of location data tracking on social networking sites. Recently, Facebook has been scrutinized for having a variety of applications that are considered to be invasive to user privacy. “The Breakup Notifier” is an example of a Facebook “cyberstalking” app that has recently been taken down. Essentially, the application notifies users when a person breaks up with their partner through Facebook, allowing users to instantly become aware of their friend’s romantic activities. The concept became very popular, with the site attracting 700,000 visits in the first 36 hours; people downloaded the app 40,000 times. Just days later, the app had more than 3.6 million downloads and 9,000 Facebook likes.

There are other applications that border on “cyberstalking.” An application named “Creepy” can track a person’s location on a map using photos uploaded to Twitter or Flickr. When a person uploads photos to a social networking site, others are able to track their most recent location. Some smart phones are able to embed the longitude and latitude coordinates into the photo and automatically send this information to the application. Anybody using the application can search for a specific person and then find their immediate location. This poses many potential threats to users who share their information with a large group of

followers. Facebook recently updated its profile format allowing for people who are not “friends” of others to view personal information about other users, even when the profile is set to private. However, As of January 18, 2011 Facebook changed its decision to make home addresses and telephone numbers accessible to third party members, but it is still possible for third party members to have access to less exact personal information, like one’s hometown and employment, if the user has entered the information into Facebook. EPIC Executive Director Marc Rotenberg said “Facebook is trying to blur the line between public and private information.

And the request for permission does not make clear to the user why the information is needed or how it will be used.” Similar to Rotenberg’s claim that Facebook users are unclear of how or why their information has gone public, recently the Federal Trade Commission and Commerce Department have become involved. The Federal Trade Commission has recently released a report claiming that Internet companies and other industries will soon need to increase their protection for online users. Because online users often unknowingly opt in on making their information public, the FTC is urging Internet companies to make privacy notes simpler and easier for the public to understand, therefore increasing their option to opt out. Perhaps this new policy should also be implemented in the Facebook world. The Commerce Department claims that Americans, “have been ill-served by a patchwork of privacy laws that contain broad gaps,”. Because of these broad gaps, Americans are more susceptible to identity theft and having

their online activity tracked by others. Spokeo Spokeo is a “people-related” search engine with results compiled through data aggregation. The site contains information such as age, relationship status, estimated personal wealth, immediate family members and home address of individual people. This information is compiled through what is already on the internet or in other public records, but the website does not guarantee accuracy.

Spokeo has been faced with potential class action law suits from people who claim that the organization breaches the Fair Credit Reporting Act. In September, 2010, Jennifer Purcell claimed that the FCRA was violated by Spokeo marketing her personal information. Her case is pending in court. Also in 2010, Thomas Robins claimed that his personal information on the website was inaccurate and he was unable to edit it for accuracy. The case was dismissed because Robins did not claim that the site directly caused him actual harm. On February 15, 2011, Robins filed another suit, this time stating Spokeo has caused him “imminent and ongoing” harm. Twitter Case - In January 2011, the government recently obtained a court order to force the social networking site, Twitter, to reveal information applicable surrounding certain subscribers involved in the WikiLeaks cases. This outcome of this case is questionable because it deals with the user’s First Amendment rights. Twitter moved to reverse the court order, and supported the idea that internet users should be notified and given an opportunity to defend their constitutional rights in court before their rights are compromised.

Facebook Friends Study - A study was conducted at Northeastern University by Alan Mislove and his colleagues at the Max Planck Institute for Software Systems, where an algorithm was created to try and discover personal attributes of a Facebook user by looking at their friend's list. They looked for information such as high school and college attended, major, hometown, graduation year and even what dorm a student may have lived in. The study revealed that only 5% of people thought to change their friend's list to private. For other users, 58% displayed university attended, 42% revealed employers, 35% revealed interests and 19% gave viewers public access to where they were located. Due to the correlation of Facebook friends and universities they attend, it was easy to discover where a Facebook user was based on their list of friends. This fact is one that has become very useful to advertisers targeting their audiences but is also a big risk for the privacy of all those with Facebook accounts.

Law enforcement prowling the networks - The FBI has dedicated undercover agents on Facebook, Twitter, MySpace, LinkedIn. The rules and guidelines to the privacy issue is internal to the Justice Department and details aren't released to the public. Agents can impersonate a friend, a long lost relative, even a spouse and child. This raises real issues regarding privacy. Although people who use Facebook, Twitter, and other social networking sites are aware of some level of privacy will always be compromised, but, no one would ever suspect that the friend invitation might be from a federal agent whose sole purpose of the friend request was to snoop around. Furthermore, Facebook, Twitter, and

MySpace have personal information and past posts logged for up to one year; even deleted profiles, and with a warrant, can hand over very personal information. One example of investigators using Facebook to nab a criminal is the case of Maxi Sopo. Charged with bank fraud, and having escaped to Mexico, he was nowhere to be found until he started posting on Facebook. Although his profile was private, his list of friends were not, and through this vector, they eventually caught him.

In recent years, some state and local law enforcement agencies have also begun to rely on social media websites as resources. Although obtaining records of information not shared publicly by or about site users often requires a subpoena, public pages on sites such as Facebook and MySpace offer access to personal information that can be valuable to law enforcement. Police departments have reported using social media websites to assist in investigations, locate and track suspects, and monitor gang activity. Teachers and MySpace - Teachers' privacy on MySpace has created controversy across the world. They are forewarned by The Ohio News Association that if they have a MySpace account, it should be deleted. Eschool News warns, "Teachers, watch what you post online." The ONA also posted a memo advising teachers not to join these sites. Teachers can face consequences of license revocations, suspensions, and written reprimands. The *Chronicle of Higher Education* wrote an article on April 27, 2007, entitled "A MySpace Photo Costs a Student a Teaching Certificate" about Stacy Snyder. She was a student of Millersville University of Pennsylvania who was denied her teaching

degree because of an unprofessional photo posted on MySpace, which involved her drinking with a pirate's hat on and a caption of "Drunken Pirate". As a substitute, she was given an English degree. Internet privacy and Blizzard Entertainment - On July 6, 2010, Blizzard Entertainment announced that it would display the real names tied to user accounts in its game forums.

On July 9, 2010, CEO and cofounder of Blizzard Mike Morhaime announced a reversal of the decision to force posters' real names to appear on Blizzard's forums. The reversal was made in response to subscriber feedback. Internet privacy and Google Maps - In Spring 2007, Google improved their Google Maps to include what is known as "Street View". This feature gives the user a 3-D, street level view with real photos of streets, buildings, and landmarks. In order to offer such a service, Google had to send trucks with cameras mounted on them and drive through every single street snapping photos. These photos were eventually stitched together to achieve a near seamless photorealistic map. However, the photos that were snapped included people caught in various acts, some of which includes a man urinating on the street, nude people seen through their windows, and apparently, a man trying to break into someone's apartment, etc; although some images are up to interpretation. This prompted a public outburst and sometime after, Google offered a "report inappropriate image" feature to their website. Internet privacy and Facebook advertisements The illegal activities on Facebook are very wild, especially "phishing attack" which is the most popular way of stealing other people's passwords.

The Facebook users are led to land on a page where they are asked for their login information, and their personal information is stolen in that way. According to the news from *PC World Business Center* which was published on April 22, 2010, we can know that a hacker named Kirillos illegally stole and sold 1.5 million Facebook IDs to some business companies who want to attract potential customers by using advertisements on the Facebook. Their illegal approach is that they used accounts which were bought from hackers to send advertisements to friends of users. When friends see the advertisements, they will have opinion about them, because “People will follow it because they believe it was a friend that told them to go to this link,” said Randy Abrams, director of technical education with security vendor Eset. There were 2.2232% of the population on Facebook that believed or followed the advertisements of their friends. Even though the percentage is small, the amount of overall users on Facebook is more than 400 million worldwide. The influence of advertisements on Facebook is so huge and obvious. According to the blog of Alan who just posted advertisement son the Facebook, he earned \$300 over the 4 days. That means he can earn \$3 for every \$1 put into it. The huge profit attracts hackers to steal users’ login information on Facebook, and business people who want to buy accounts from hackers send advertisements to users’ friends on Facebook.

INTERNET SERVICE PROVIDERS

Internet users obtain Internet access through an Internet service provider (ISP). All data transmitted to and from

users must pass through the ISP. Thus, an ISP has the potential to observe users' activities on the Internet. However, ISPs are usually prevented from participating in such activities due to legal, ethical, business, or technical reasons. Despite these legal and ethical restrictions, some ISPs, such as British Telecom (BT), are planning to use deep packet inspection technology provided by companies such as Phorm in order to examine the contents of the pages that people visit. By doing so, they can build up a profile of a person's web surfing habits, which can then be sold on to advertisers in order to provide targeted advertising. BT's attempt at doing this will be marketed under the name 'Webwise'. Normally ISPs do collect at least *some* information about the consumers using their services. From a privacy standpoint, ISPs would ideally collect only as much information as they require in order to provide Internet connectivity (IP address, billing information if applicable, etc). Which information an ISP collects, what it does with that information, and whether it informs its consumers, pose significant privacy issues. Beyond the usage of collected information typical of third parties, ISPs sometimes state that they will make their information available to government authorities upon request. In the US and other countries, such a request does not necessarily require a warrant.

An ISP cannot know the contents of properly-encrypted data passing between its consumers and the Internet. For encrypting web traffic, https has become the most popular and best-supported standard. Even if users encrypt the data, the ISP still knows the IP addresses of the sender and of the recipient. (However, see the IP addresses section for

workarounds.) An Anonymizer such as I2P – The Anonymous Network or Tor can be used for accessing web services without them knowing your IP address and without your ISP knowing what the services are that you access. General concerns regarding Internet user privacy have become enough of a concern for a UN agency to issue a report on the dangers of identity fraud. While signing up for internet services, each computer contains a unique IP, Internet Protocol address. This particular address will not give away private or personal information, however, a weak link could potentially reveal information from your ISP. Social networking has redefined the role of Internet privacy. Since users are willingly disclosing personal information online, the role of privacy and security is somewhat blurry.

Sites such as Facebook, Myspace, and Twitter have grown popular by broadcasting status updates featuring personal information such as location. Facebook “Places,” in particular, is a Facebook service, which publicizes user location information to the networking community. Users are allowed to “check-in” at various locations including retail stores, convenience stores, and restaurants. Also, users are able to create their own “place,” disclosing personal information onto the Internet. This form of location tracking is automated and must be turned off manually. Various settings must be turned off and manipulated in order for the user to ensure privacy. According to epic.org, Facebook users are recommended to: (1) disable “Friends can check me in to Places,” (2) customize “Places I Check In,” (3) disable “People Here Now,” and (4) uncheck “Places I’ve Visited.”. Moreover, the Federal Trade Commission has

received two complaints in regards to Facebook's "unfair and deceptive" trade practices, which are used to target advertising sectors of the online community. "Places" tracks user location information and is used primarily for advertising purposes. Each location tracked allows third party advertisers to customize advertisements that suit one's interests. Currently, the Federal Trade Commissioner along with the Electronic Privacy Information Center are shedding light on the issues of location data tracking on social networking sites.

LEGAL THREATS

Use by government agencies of an array of technologies designed to track and gather Internet users' information are the topic of much debate between privacy advocates, civil libertarians and those who believe such measures are necessary for law enforcement to keep pace with rapidly changing communications technology.

SPECIFIC EXAMPLES

- Following a decision by the European Union's council of ministers in Brussels, in January, 2009, the UK's Home Office adopted a plan to allow police to access the contents of individuals' computers without a warrant. The process, called "remote searching", allows one party, at a remote location, to examine another's hard drive and Internet traffic, including email, browsing history and websites visited. Police across the EU are now permitted to request that the British police conduct a remote search on their behalf.

The search can be granted, and the material gleaned turned over and used as evidence, on the basis of a senior officer believing it necessary to prevent a serious crime. Opposition MPs and civil libertarians are concerned about this move toward widening surveillance and its possible impact on personal privacy. Says Shami Chakrabarti, director of the human rights group Liberty, “The public will want this to be controlled by new legislation and judicial authorisation. Without those safeguards it’s a devastating blow to any notion of personal privacy.”

- The FBI’s Magic Lantern software programme was the topic of much debate when it was publicized in November, 2001. Magic Lantern is a Trojan Horse programme that logs users’ keystrokes, rendering encryption useless.

LAWS FOR INTERNET PRIVACY PROTECTION

USA PATRIOT ACT

The purpose of this act, enacted on October 26, 2001 by former President Bush, was to enhance law enforcement investigatory tools, investigate online activity, as well as to discourage terrorist acts both within the United States and around the world. This act reduced restrictions for law enforcement to search various methods and tools of communication such as telephone, e-mail, personal records including medical and financial, as well as reducing restrictions with obtaining of foreign intelligence.

ELECTRONIC COMMUNICATIONS PRIVACY ACT (ECPA)

This act makes it unlawful under certain conditions for an individual to reveal the information of electronic communication and contains a few exceptions. One clause allows the ISP to view private e-mail if the sender is suspected of attempting to damage the internet system or attempting to harm another user. Another clause allows the ISP to reveal information from a message if the sender or recipient allows to its disclosure. Finally, information containing personal information may also be revealed for a court order or law enforcement's subpoena.

EMPLOYEES AND EMPLOYERS INTERNET REGULATIONS

When considering the rights between employees and employers regarding internet privacy and protection at a company, different states have their own laws. Connecticut and Delaware both have laws that state an employer must create a written notice or electronic message that provides understanding that they will regulate the internet traffic. By doing so, this relates to the employees that the employer will be searching and monitoring emails and internet usage. Delaware charges \$100 for a violation where Connecticut charges \$500 for the first violation and then \$1000 for the second. When looking at public employees and employers, California and Colorado created laws that would also create legal ways in which employers controlled internet usage. The law stated that a public company or agency must create a prior message to the employees stating that accounts will

be monitored. Without these laws, employers could access information through employees accounts and use them illegally. In most cases, the employer is allowed to see whatever he or she pleases because of these laws stated both publicly and privately.

OTHER POTENTIAL INTERNET PRIVACY RISKS

- Malware is a term short for “malicious software” and is used to describe software to cause damage to a single computer, server, or computer network whether that is through the use of a virus, trojan horse, spyware, etc.
- Spyware is a piece of software that obtains information from a user’s computer without that user’s consent.
- A web bug is an object embedded into a web page or email and is usually invisible to the user of the website or reader of the email. It allows checking to see if a person has looked at a particular website or read a specific email message.
- Phishing is a criminally fraudulent process of trying to obtain sensitive information such as user names, passwords, credit card or bank information. Phishing is an internet crime in which someone masquerades as a trustworthy entity in some form of electronic communication.
- Pharming is hackers attempt to redirect traffic from a legitimate website to a completely different internet address. Pharming can be conducted by changing the hosts file on a victim’s computer or by exploiting a vulnerability on the DNS server.

- Social engineering
- Malicious proxy server (or other “anonymity” services)

SPECIFIC CASES

JASON FORTUNY AND CRAIGSLIST

In early September 2006, Jason Fortuny, a Seattle-area freelance graphic designer and network administrator, posed as a woman and posted an ad to Craigslist Seattle seeking a casual sexual encounter with men in that area. On September 4, he posted to the wiki website Encyclopædia Dramatica all 178 of the responses, complete with photographs and personal contact details, describing this as the Craigslist Experiment and encouraging others to further identify the respondents. Although some online exposures of personal information have been seen as justified for exposing malfeasance, many commentators on the Fortuny case saw no such justification here. “The men who replied to Fortuny’s posting did not appear to be doing anything illegal, so the outing has no social value other than to prove that someone could ruin lives online,” said law professor Jonathan Zittrain, while *Wired* writer Ryan Singel described Fortuny as “sociopathic”. The Electronic Frontier Foundation indicated that it thought Fortuny might be liable under Washington state law, and that this would depend on whether the information he disclosed was of legitimate public concern. Kurt Opsahl, the EFF’s staff attorney, said “As far as I know, they (the respondents) are not public figures, so it would be challenging to show that this was something of public concern.”

According to Fortuny, two people lost their jobs as a result of his Craigslist Experiment and another “has filed an invasion-of-privacy lawsuit against Fortuny in an Illinois court.” Fortuny did not enter an appearance in the Illinois suit, secure counsel, or answer the complaint after an early amendment. Mr. Fortuny had filed a motion to dismiss, but he filed it with the Circuit Court of Cook County, Illinois, and he did not file proof that he had served the plaintiff. As a result, the court entered a default judgment against Mr. Fortuny and ordered a damages hearing for January 7, 2009. After failing to show up at multiple hearings on damages, Fortuny was ordered to pay \$74,252.56 for violation of the Copyright Act, compensation for Public Disclosure of Private Facts, Intrusion Upon Seclusion, attorneys fees and costs.

USA VS. WARSHAK

The case *United States v. Warshak*, decided December 14, 2010 by the Sixth Circuit Court of Appeals, maintained the idea that an ISP actually is allowed access to private e-mail. However, the government must get hold of a search warrant before obtaining such e-mail. This case dealt with the question of emails hosted on an isolated server. Due to the fact that e-mail is similar to other forms of communication such as telephone calls, e-mail requires the same amount of protection under the 4th amendment.

SEARCH ENGINE DATA AND LAW ENFORCEMENT

Data from major Internet companies, including Yahoo! and MSN (Microsoft), have already been subpoenaed by the

United States and China. AOL even provided a chunk of its own search data online, allowing reporters to track the online behaviour of private individuals. In 2006, a wireless hacker pled guilty when his Google searches were used as evidence against him. The defendant ran a Google search over the network using the following search terms: “how to broadcast interference over wifi 2.4 GHZ,” “interference over wifi 2.4 Ghz,” “wireless networks 2.4 interference,” and “make device interfere wireless network.” While court papers did not describe how the FBI obtained his searches (e.g. through a seized hard-drive or directly from the search-engine), Google has indicated that it can provide search terms to law enforcement if given an Internet address or Web cookie.

US V. ZEIGLER

In the United States many cases discuss whether a private employee (i.e., not a government employee) who stores incriminating evidence in workplace computers is protected by the Fourth Amendment’s reasonable expectation of privacy standard in a criminal proceeding. Most case law holds that employees do not have a reasonable expectation of privacy when it comes to their work related electronic communications. See, e.g. *US v. Simons*, 206 F.3d 392, 398 (4th Cir., Feb. 28, 2000). However, one federal court held that employees can assert that the attorney-client privilege with respect to certain communications on company laptops. See *Curto v. Medical World Comm.*, No. 03CV6327, 2006 U.S. Dist. LEXIS 29387 (E.D.N.Y. May 15, 2006). Another recent federal case discussed this topic. On January 30,

2007, the Ninth Circuit court in *US v. Ziegler*, reversed its earlier August 2006 decision upon a petition for rehearing. In contrast to the earlier decision, the Court acknowledged that an employee has a right to privacy in his workplace computer. However, the Court also found that an employer can consent to any illegal searches and seizures. See *US v. Ziegler*, ___F.3d 1077 (9th Cir. Jan. 30, 2007, No. 05-30177). Cf. *US v. Ziegler*, 456 F.3d 1138 (9th Cir. 2006). In *Ziegler*, an employee had accessed child pornography websites from his workplace. His employer noticed his activities, made copies of the hard drive, and gave the FBI the employee's computer. At his criminal trial, Ziegler filed a motion to suppress the evidence because he argued that the government violated his Fourth Amendment rights. The Ninth Circuit allowed the lower court to admit the child pornography as evidence. After reviewing relevant Supreme Court opinions on a reasonable expectation of privacy, the Court acknowledged that Ziegler had a reasonable expectation of privacy at his office and on his computer. That Court also found that his employer could consent to a government search of the computer and that, therefore, the search did not violate Ziegler's Fourth Amendment rights.

STATE V. REID

The New Jersey Supreme Court has also issued an opinion on the privacy rights of computer users, holding in *State v. Reid* that computer users have a reasonable expectation of privacy concerning the personal information they give to their ISPs. In that case, Shirley Reid was indicted for

computer theft for changing her employer's password and shipping address on its online account with a supplier. The police discovered her identity after serving the ISP, Comcast, with a municipal subpoena not tied to any judicial proceeding. The lower court suppressed the information from Comcast that linked Reid with the crime on grounds that the disclosure violated Reid's constitutional right to be protected from unreasonable search and seizure. The appellate court affirmed, as did the New Jersey Supreme Court, which ruled that ISP subscriber records can only be disclosed to law enforcement upon the issuance of a grand jury subpoena. As a result, New Jersey offers greater privacy rights to computer users than most federal courts. This case also serves as an illustration of how case law on privacy regarding workplace computers is still evolving.

ROBBINS V. LOWER MERION SCHOOL DISTRICT

In *Robbins v. Lower Merion School District* (U.S. Eastern District of Pennsylvania 2010), the federal trial court issued an injunction against the school district after plaintiffs charged two suburban Philadelphia high schools violated the privacy of students and others when they secretly spied on students by surreptitiously and remotely activating webcams embedded in school-issued laptops the students were using at home. The schools admitted to secretly snapping over 66,000 webshots and screenshots, including webcam shots of students in their bedrooms.

CYBER SECURITY STANDARDS

Cyber security standards are security standards which enable organizations to practice safe security techniques to

minimize the number of successful cyber security attacks. These guides provide general outlines as well as specific techniques for implementing cyber security. For certain specific standards, cyber security certification by an accredited body can be obtained. There are many advantages to obtaining certification including the ability to get cyber security insurance.

HISTORY

Cyber security standards have been created recently because sensitive information is now frequently stored on computers that are attached to the Internet. Also many tasks that were once done by hand are carried out by computer; therefore there is a need for Information Assurance (IA) and security. Cyber security is important in order to guard against identity theft. Businesses also have a need for cyber security because they need to protect their trade secrets, proprietary information, and personally identifiable information (PII) of their customers or employees. The government also has the need to secure its information. One of the most widely used security standards today is ISO/IEC 27002 which started in 1995. This standard consists of two basic parts. BS 7799 part 1 and BS 7799 part 2 both of which were created by (British Standards Institute) BSI. Recently this standard has become ISO 27001. The National Institute of Standards and Technology (NIST) has released several special publications addressing cyber security. Three of these special papers are very relevant to cyber security: the 800-12 titled "Computer Security Handbook;" 800-14 titled "Generally Accepted Principles

and Practices for Securing Information Technology;” and the 800-26 titled “Security Self-Assessment Guide for Information Technology Systems”. The International Society of Automation (ISA) developed cyber security standards for industrial automation control systems (IACS) that are broadly applicable across manufacturing industries. The series of ISA industrial cyber security standards are known as ISA-99 and are being expanded to address new areas of concern.

ISO 27002

ISO 27002 incorporates both parts of the BS 7799 standard. Sometimes ISO/IEC 27002 is referred to as BS 7799 part 1 and sometimes it refers to part 1 and part 2. BS 7799 part 1 provides an outline for cyber security policy; whereas BS 7799 part 2 provides a certification. The outline is a high level guide to cyber security. It is most beneficial for an organization to obtain a certification to be recognized as compliant with the standard. The certification once obtained lasts three years and is periodically checked by the BSI to ensure an organization continues to be compliant throughout that three year period. ISO 27001 (ISMS) replaces BS 7799 part 2, but since it is backward compatible any organization working toward BS 7799 part 2 can easily transition to the ISO 27001 certification process. There is also a transitional audit available to make it easier once an organization is BS 7799 part 2-certified for the organization to become ISO 27001-certified. ISO/IEC 27002 states that information security is characterized by integrity, confidentiality, and availability. The ISO/IEC 27002 standard is arranged into eleven control areas; security policy,

organizing information security, asset management, human resources security, physical and environmental security, communication and operations, access controls, information systems acquisition/development/maintenance, incident handling, business continuity management, compliance.

STANDARD OF GOOD PRACTICE

In the 1990s, the Information Security Forum (ISF) published a comprehensive list of best practices for information security, published as the *Standard of Good Practice* (SoGP). The ISF continues to update the SoGP every two years; the latest version was published in February 2007. Originally the *Standard of Good Practice* was a private document available only to ISF members, but the ISF has since made the full document available to the general public at no cost. Among other programmes, the ISF offers its member organizations a comprehensive benchmarking programme based on the SoGP.

NERC

The North American Electric Reliability Corporation (NERC) has created many standards. The most widely recognized is NERC 1300 which is a modification/update of NERC 1200. The newest version of NERC 1300 is called CIP-002-1 through CIP-009-2 (CIP=Critical Infrastructure Protection). These standards are used to secure bulk electric systems although NERC has created standards within other areas. The bulk electric system standards also provide network security administration while still supporting best practice industry processes.

NIST

1. Special publication 800-12 provides a broad overview of computer security and control areas. It also emphasizes the importance of the security controls and ways to implement them. Initially this document was aimed at the federal government although most practices in this document can be applied to the private sector as well. Specifically it was written for those people in the federal government responsible for handling sensitive systems.
2. Special publication 800-14 describes common security principles that are used. It provides a high level description of what should be incorporated within a computer security policy. It describes what can be done to improve existing security as well as how to develop a new security practice. Eight principles and fourteen practices are described within this document.
3. Special publication 800-26 provides advice on how to manage IT security. This document emphasizes the importance of self assessments as well as risk assessments.
4. Special publication 800-37, updated in 2010 provides a new risk approach: “Guide for Applying the Risk Management Framework to Federal Information Systems”
5. Special publication 800-53 rev3, “Guide for Assessing the Security Controls in Federal Information Systems”, updated in August 2009, specifically addresses the 194 security controls that are applied to a system to make it “more secure.”

ISO 15408

This standard develops what is called the “Common Criteria”. It allows many different software applications to be integrated and tested in a secure way.

RFC 2196

RFC 2196 is memorandum published by Internet Engineering Task Force for developing security policies and procedures for information systems connected on the Internet. The RFC 2196 provides a general and broad overview of information security including network security, incident response or security policies. The document is very practical and focusing on day-to-day operations.

ISA-99

ISA99 is the Industrial Automation and Control System Security Committee of the International Society for Automation (ISA). The committee is developing a multi-part series of standards and technical reports on the subject, several of which have been publicly released. Work products from the ISA99 committee are also submitted to IEC as standards and specifications in the IEC 63443 series.

- ISA-99.01.01 (formerly referred to as “Part 1”) (ANSI/ISA 99.00.01) is approved and published.
- ISA-TR99.01.02 is a master glossary of terms used by the committee. This document is still a working draft but the content is available on the committee Wiki site (<http://isa99.isa.org/ISA99%20Wiki/Master%20Glossary.aspx>)
- ISA-99.01.03 identifies a set of compliance metrics

for IACS security. This document is currently under development.

- ISA-99.02.01 (formerly referred to as “Part 2”) (ANSI/ISA 99.02.01-2009) addresses how to establish an IACS security programme. This standard is approved and published. It has also been approved and published by the IEC as IEC 62443-2-1
- ISA-99.02.02 addresses how to operate an IACS security programme. This standard is currently under development.
- ISA-TR99.02.03 is a technical report on the subject of patch management. This report is currently under development.
- ISA-TR99.03.01 () is a technical report on the subject of suitable technologies for IACS security. This report is approved and published.
- ISA-99.03.02 addresses how to define security assurance levels using the zones and conduits concept. This standard is currently under development.
- ISA-99.03.03 defines detailed technical requirements for IACS security. This standard is currently under development.
- ISA-99.03.04 addresses the requirements for the development of secure IACS products and solutions. This standard is currently under development.
- Standards in the ISA-99.04.xx series address detailed technical requirements at the component level. These standards are currently under development.

ISA SECURITY COMPLIANCE INSTITUTE

Related to the work of ISA 99 is the work of the ISA Security Compliance Institute. The ISA Security Compliance Institute (ISCI) has developed compliance test specifications for ISA99 and other control system security standards. They have also created an ANSI accredited certification programme called ISASecure for the certification of industrial automation devices such as programmable logic controllers (PLC), distributed control systems (DCS) and safety instrumented systems (SIS). These types of devices provided automated control of industrial processes such as those found in the oil & gas, chemical, electric utility, manufacturing, food & beverage and water/wastewater processing industries. There is growing concern from both governments as well as private industry regarding the risk that these systems could be intentionally compromised by “evildoers” such as hackers, disgruntled employees, organized criminals, terrorist organizations or even state-sponsored groups. The recent news about the industrial control system malware known as Stuxnet has heightened concerns about the vulnerability of these systems.

HIGH TECHNOLOGY CRIME INVESTIGATION ASSOCIATION

The High Technology Crime Investigation Association or as known by the abbreviation HTCIA Inc., Roseville, California, is devoted to digital forensics for investigation of crimes. Members of HTCIA Inc. are made up of a professional body of investigators, prosecutors and security professionals. HTCIA is designed to promote, aid encourage

Special Focus on Computer Security Model

and effect the voluntary interchange of data, experience, information, knowledge and ideas about processes, procedures, methods, and techniques relating to investigations and security in advanced technologies, and new technologies introduced into the field of forensic investigation for crime and the law. The HTCIA also promote uniformity in investigative methods, and develop matters of mutual interest. It is also one of the most successful collaborative efforts between law enforcement and private industry working together. The HTCIA serve a common theme within each member state and internationally. This is to foster the growth in knowledge of investigation methods, processes and techniques amongst their Chapter's within their own national capital region in the US and Internationally. The HTCIA hold annual conferences. The HTCIA are primarily intended for Computer forensic analysts, Cybercrime Investigators, Mobile forensic analysts, IT Security, Security Managers, CIOs, Lawyers, Prosecutors, Police officers, Judiciary, and Incident Response specialists. Training opportunities do exist for law enforcement personnel, and investigative professionals. Which is seen today of paramount importance within the organization. The HTCIA has local chapters that sponsor meetings. These meetings attract law enforcement as well as their public sector counterparts and academia. Generally, topics of current interest are presented, allowing members to obtain valuable knowledge. One of the clearest benefits of becoming a member is the ability of members to network & meet other investigators in similar fields. The ability to contact others who have faced the same problems is invaluable. It is today

one of the most respected organizations for professional, in-service training of law officials interested in computers and their role in criminal activity. Some high profile members include Howard Schmidt, Matthew Blake and James Lance.

HISTORY

The HTCIA had its beginnings in the early 1980s when security managers in Silicon Valley saw the need for law enforcement investigators to understand the importance of investigation with high technology and computer crime. In 1984, the Santa Clara District Attorney, Leo Himmelsbach was approached by members of the Santa Clara County Industrial Security Managers Group. Including security manager at Intel and later Sun Microsystems John Callaghan, and Pete Kostner security manager at AMD, they discussed the need for having law enforcement officers trained in the field of high technology crime. This was seen as quiet visionary for its time. Mr. Himmelsbach then applied to the state of California and received a grant from the Office of Criminal Justice Planning Project approved by the California State Assembly, State Assembly Bill 1078 passed into law August 31, 1984, Penal Code Section 13970 (GrntProjSummry) called "SANTA CLARA COUNTY DISTRICT ATTORNEY'S HIGH TECHNOLOGY CRIME PREVENTION PROGRAM" The HTCIA started as part of the 1985 OCJB grant "Project Objectives and Activities. The main objective at that time was to train San Francisco Bay Area investigators and prosecutors in high-technology theft investigation. During that time, law enforcement officers from the Los Angeles area attended DATTA training and wanted to start

their own organization in southern California. A second objective was then conceived “to establish a base that will provide the nucleus for the development of a regional high technology theft prevention effort”(OCJB Grant Contract Feb 1985).

In 1990, the District Attorney’s Theft Technology Association of Santa Carla County affiliated with the other HTCIA chapters in existence in the Silicon Valley Chapter. It was then decided to allow private investigators to become members. From then on in the HTCIA began growing from strength to strength, with national seminars, setting standards for training and guidelines. The HTCIA now cover considerable areas throughout the US and Internationally across the World primarily focusing on High Technology and Crime.

HTCIA CORE VALUES

HTCIA Core Values are defined as follows: I) The HTCIA values the Truth uncovered effective techniques used to uncover that Truth, and within digital information and the so wrongful convictions are avoided! II) The HTCIA values the security of their society and its citizens through the enforcement of our laws and the protection of our infrastructure and economies. III) The HTCIA values the ethical concept of its members and the evidence they expose through investigative procedures and computer forensic best practices including specialized techniques used to gather digital evidence. IV) The HTCIA values the trusted network of forensic and investigative professionals within private and public businesses including law enforcement who share

our values and our vision. V) The HTCIA values the confidentiality of its membership and the information, skills and techniques they share within the association

MEMBERSHIP

I) Membership is open to prosecuting attorneys, and investigators engaged in the investigation of criminal activity associated with computers or technology. II) Senior security specialists, and managerial level professionals engaged in professions covering computers or advanced technology environments III) On October 2008, the international board of directors approved the bylaw provision creating the student membership. The purpose of this membership class is to promote and encourage the study of criminal investigations involving advanced technologies and security by the academic community. This would be composed of students studying in areas such as criminal justice, law enforcement computer science; forensics, corrections, accounting, a minimum Grade Point Average (GPA) is established by the International Executive Committee (IEC). The IEC will establish general application procedures and requirements for Student Members which are not in conflict with these bylaws. Scholarship for Service (SFS) provides scholarships that fully fund the typical costs that students pay for books, tuition, and room and board while attending an approved institution of higher learning. The scholarships are funded through grants awarded by the National Science Foundation (NSF) HTCIA is not affiliated with the SFS.

CASE OF THE YEAR AWARD

Case of the year award The HTCIA also offer a case of the year award to recognize new technology or techniques

which were expended to resolve the case. A synopsis of factors which must be used to evaluate the nominees are either the case was international, national, or regional in scope, it resolved a particularly violent offense & it established an important legal precedent

TRAINING EXAMPLE

Presentations may include discussing the benefits for using live computer forensic investigation techniques, and outline the situations where these techniques may be most appropriate; an example would be the ability to capture encryption passwords. Members or attendees would be introduced to the components of a live computer forensic investigation, shown tools for identifying the machine state to help mitigate the “trojan defense”.

4

Capability-Based Security

Capability-based security is a concept in the design of secure computing systems, one of the existing security models. A capability (known in some systems as a key) is a communicable, unforgeable token of authority. It refers to a value that references an object along with an associated set of access rights. A user programme on a capability-based operating system must use a capability to access an object. Capability-based security refers to the principle of designing user programmes such that they directly share capabilities with each other according to the principle of least privilege, and to the operating system infrastructure necessary to make such transactions efficient and secure. Although most operating systems implement a facility which resembles capabilities, they typically do not provide enough support to allow for the exchange of capabilities among possibly mutually untrusting entities to be the primary

means of granting and distributing access rights throughout the system. A capability-based system, in contrast, is designed with that goal in mind. Capabilities as discussed in this article should not be confused with POSIX 1e/2c “Capabilities”. The latter are coarse-grained privileges that cannot be transferred between processes.

CAPABILITIES AND CAPABILITY-BASED SECURITY

Capabilities achieve their objective of improving system security by being used in place of forgeable references. A forgeable reference (for example, a path name) identifies an object, but does not specify which access rights are appropriate for that object and the user programme which holds that reference. Consequently, any attempt to access the referenced object must be validated by the operating system, typically via the use of an access control list (ACL). Instead, in a system with capabilities, the mere fact that a user programme possesses that capability entitles it to use the referenced object in accordance with the rights that are specified by that capability. In theory, a system with capabilities removes the need for any access control list or similar mechanism by giving all entities all and only the capabilities they will actually need. A capability is typically implemented as a privileged data structure that consists of a section that specifies access rights, and a section that uniquely identifies the object to be accessed. In practice, it is used much like a file descriptor in a traditional operating system, but to access every object on the system. Capabilities are typically stored by the operating system in a list, with some mechanism in place to prevent the programme from

directly modifying the contents of the capability (so as to forge access rights or change the object it points to). Some systems have also been based on capability-based addressing (hardware support for capabilities), such as Plessey System 250. Programmes possessing capabilities can perform functions on them, such as passing them on to other programmes, converting them to a less-privileged version, or deleting them. The operating system must ensure that only specific operations can occur to the capabilities in the system, in order to maintain the integrity of the security policy.

INTRODUCTION TO CAPABILITY-BASED SECURITY

(The following introduction assumes some basic knowledge of Unix systems.) A capability is defined to be a protected object reference which, by virtue of its possession by a user process, grants that process the capability (hence the name) to interact with an object in certain ways. Those ways might include reading data associated with an object, modifying the object, executing the data in the object as a process, and other conceivable access rights. The capability logically consists of a reference that uniquely identifies a particular object and a set of one or more of these rights. Suppose that, in a user process's memory space, there exists the following string:

`/etc/passwd`

Although this identifies a unique object on the system, it does not specify access rights and hence is not a capability. Suppose there is instead the following two values:

`/etc/passwd`

`O_RDWR`

This identifies an object along with a set of access rights. It, however, is still not a capability because the user process's *possession* of these values says nothing about whether that access would actually be legitimate. Now suppose that the user programme successfully executes the following statement:

```
int fd = open("/etc/passwd", O_RDWR);
```

The variable `fd` now contains the index of a file descriptor in the process's file descriptor table. This file descriptor is a capability. Its existence in the process's file descriptor table is sufficient to know that the process does indeed have legitimate access to the object. A key feature of this arrangement is that the file descriptor table is in kernel memory and cannot be directly manipulated by the user programme.

SHARING OF CAPABILITIES BETWEEN PROCESSES

In traditional operating systems, programmes often communicate with each other and with storage using references like those in the first two examples. Path names are often passed as command-line parameters, sent via sockets, and stored on disk. These references are not capabilities, and must be validated before they can be used. In these systems, a central question is “on whose *authority* is a given reference to be evaluated?” This becomes a critical issue especially for processes which must act on behalf of

two different authority-bearing entities. They become susceptible to a programming error known as the confused deputy problem, very frequently resulting in a security hole. In a capability-based system, the capabilities themselves are passed between processes and storage using a mechanism that is known by the operating system to maintain the integrity of those capabilities.

Although many operating systems implement facilities very similar to capabilities through the use of file descriptors or file handles — for example, in UNIX, file descriptors can be discarded (closed), inherited by child processes, and even sent to other processes via sockets — there are several obstacles that prevent all of the benefits of a capability-based addressing system from being realized in a traditional operating system environment. Chief among these obstacles is the fact that entities which might hold capabilities (such as processes and files) cannot be made persistent in such a way that maintains the integrity of the secure information that a capability represents. The operating system cannot trust a user programme to read back a capability and not tamper with the object reference or the access rights, and has no built-in facilities to control such tampering. Consequently, when a programme wishes to regain access to an object that is referenced on disk, the operating system must have some way of validating that access request, and an access control list or similar mechanism is mandated.

One novel approach to solving this problem involves the use of an orthogonally persistent operating system. (This was realised in the Flex machine. See Ten15). In such a system, there is no need for entities to be discarded and

their capabilities be invalidated, and hence require an ACL-like mechanism to restore those capabilities at a later time. The operating system maintains the integrity and security of the capabilities contained within all storage, both volatile and nonvolatile, at all times; in part by performing all serialization tasks by itself, rather than requiring user programmes to do so, as is the case in most operating systems. Because user programmes are relieved of this responsibility, there is no need to trust them to reproduce only legal capabilities, nor to validate requests for access using an access control mechanism.

POSIX CAPABILITIES

POSIX draft 1003.1e specifies a concept of permissions called “capabilities”. However POSIX capabilities differ from capabilities in this article — POSIX capability is not associated with any object — a process having CAP_NET_BIND_SERVICE capability can listen on any TCP port under 1024.

RESEARCH AND COMMERCIAL SYSTEMS

- Tahoe-LAFS - Open Source capability-based filesystem
- KeyKOS
 - o EROS - The Extremely Reliable Operating System
 - KeyKOS successor
 - CapROS - EROS successor, project to further develop EROS code base for commercial use
 - Coyotos - EROS successor, for research
- kaneton
- Cambridge CAP computer

- Carnegie Mellon University C.mmp with Hydra (operating system)
- Carnegie Mellon University CM* with StarOS
- IBM System/38 and AS/400
- Intel iAPX 432
- Plessey System 250
- Symbian
- Flex
- L4 microkernel - Open Kernel Labs - OKL4 and NICTA - seL4, TU-Dresden - Fiasco.OC
- Amoeba distributed operating system

CLOUD COMPUTING SECURITY

Cloud computing security (sometimes referred to simply as “cloud security”) is an evolving sub-domain of computer security, network security, and, more broadly, information security. It refers to a broad set of policies, technologies, and controls deployed to protect data, applications, and the associated infrastructure of cloud computing.

SECURITY ISSUES ASSOCIATED WITH THE CLOUD

There are a number of security issues/concerns associated with cloud computing but these issues fall into two broad categories: Security issues faced by cloud providers (organizations providing Software-, Platform-, or Infrastructure-as-a-Service via the cloud) and security issues faced by their customers. In most cases, the provider must ensure that their infrastructure is secure and that their clients’ data and applications are protected while the

customer must ensure that the provider has taken the proper security measures to protect their information.

DIMENSIONS OF CLOUD SECURITY

While cloud security concerns can be grouped into any number of dimensions (Gartner names seven while the Cloud Security Alliance identifies fifteen areas of concern) these dimensions have been aggregated into three general areas: Security and Privacy, Compliance, and Legal or Contractual Issues.

SECURITY AND PRIVACY

In order to ensure that data is secure (that it cannot be accessed by unauthorized users or simply lost) and that data privacy is maintained, cloud providers attend to the following areas:

DATA PROTECTION

To be considered protected, data from one customer must be properly segregated from that of another; it must be stored securely when “at rest” and it must be able to move securely from one location to another. Cloud providers have systems in place to prevent data leaks or access by third parties. Proper separation of duties should ensure that auditing and/or monitoring cannot be defeated, even by privileged users at the cloud provider.

IDENTITY MANAGEMENT

Every enterprise will have its own identity management system to control access to information and computing

resources. Cloud providers either integrate the customer's identity management system into their own infrastructure, using federation or SSO technology, or provide an identity management solution of their own.

PHYSICAL AND PERSONNEL SECURITY

Providers ensure that physical machines are adequately secure and that access to these machines as well as all relevant customer data is not only restricted but that access is documented.

AVAILABILITY

Cloud providers assure customers that they will have regular and predictable access to their data and applications.

APPLICATION SECURITY

Cloud providers ensure that applications available as a service via the cloud are secure by implementing testing and acceptance procedures for outsourced or packaged application code. It also requires application security measures (application-level firewalls) be in place in the production environment.

PRIVACY

Finally, providers ensure that all critical data (credit card numbers, for example) are masked and that only authorized users have access to data in its entirety. Moreover, digital identities and credentials must be protected as should any data that the provider collects or produces about customer activity in the cloud.

COMPLIANCE

Numerous regulations pertain to the storage and use of data, including Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), the Sarbanes-Oxley Act, among others. Many of these regulations require regular reporting and audit trails. Cloud providers must enable their customers to comply appropriately with these regulations.

BUSINESS CONTINUITY AND DATA RECOVERY

Cloud providers have business continuity and data recovery plans in place to ensure that service can be maintained in case of a disaster or an emergency and that any data lost will be recovered. These plans are shared with and reviewed by their customers.

LOGS AND AUDIT TRAILS

In addition to producing logs and audit trails, cloud providers work with their customers to ensure that these logs and audit trails are properly secured, maintained for as long as the customer requires, and are accessible for the purposes of forensic investigation (e.g., eDiscovery).

UNIQUE COMPLIANCE REQUIREMENTS

In addition to the requirements to which customers are subject, the data centers maintained by cloud providers may also be subject to compliance requirements.

LEGAL AND CONTRACTUAL ISSUES

Aside from the security and compliance issues enumerated above, cloud providers and their customers will

negotiate terms around liability (stipulating how incidents involving data loss or compromise will be resolved, for example), intellectual property, and end-of-service (when data and applications are ultimately returned to the customer

PUBLIC RECORDS

Legal issues may also include records-keeping requirements in the public sector, where many agencies are required by law to retain and make available electronic records in a specific fashion. This may be determined by legislation, or law may require agencies to conform to the rules and practices set by a records-keeping agency. Public agencies using cloud computing and storage must take these concerns into account.

SECURE CODING

History has proven that software defects, bugs and logic flaws are consistently the primary cause of commonly exploited software vulnerabilities. Through the analysis of thousands of reported vulnerabilities, security professionals have discovered that most vulnerabilities stem from a relatively small number of common software programming errors. By identifying the insecure coding practices that lead to these errors and educating developers on secure alternatives, organizations can take proactive steps to help significantly reduce or eliminate vulnerabilities in software before deployment.

ACCESS CONTROL LIST

An access control list (ACL), with respect to a computer file system, is a list of permissions attached to an object.

An ACL specifies which users or system processes are granted access to objects, as well as what operations are allowed on given objects. Each entry in a typical ACL specifies a subject and an operation. For instance, if a file has an ACL that contains (Alice, delete), this would give Alice permission to delete the file.

ACL-BASED SECURITY MODELS

When a subject requests an operation on an object in an ACL-based security model the operating system first checks the ACL for an applicable entry to decide whether the requested operation is authorized. A key issue in the definition of any ACL-based security model is determining how access control lists are edited, namely which users and processes are granted ACL-modification access. ACL models may be applied to collections of objects as well as to individual entities within the system's hierarchy.

FILESYSTEM ACLS

A Filesystem ACL is a data structure (usually a table) containing entries that specify individual user or group rights to specific system objects such as programmes, processes, or files. These entries are known as access control entries (ACEs) in the Microsoft Windows NT, OpenVMS, Unix-like, and Mac OS X operating systems. Each accessible object contains an identifier to its ACL. The privileges or permissions determine specific access rights, such as whether a user can read from, write to, or execute an object. In some implementations an ACE can control whether or not a user, or group of users, may alter the ACL on an

object. Most of the Unix and Unix-like operating systems (e.g. Linux, BSD, or Solaris) support so called POSIX.1e ACLs, based on an early POSIX draft that was abandoned. Many of them, for example AIX, FreeBSD, Mac OS X beginning with version 10.4 (“Tiger”), or Solaris with ZFS filesystem, support NFSv4 ACLs, which are part of the NFSv4 standard. There are two (experimental) implementations of NFSv4 ACLs for Linux. NFSv4 ACLs support for Ext3 filesystem and recent Richacl, which brings NFSv4 ACLs support for Ext4 filesystem.

NETWORKING ACLS

On some types of proprietary computer hardware, an Access Control List refers to rules that are applied to port numbers or network daemon names that are available on a host or other layer 3, each with a list of hosts and/or networks permitted to use the service. Both individual servers as well as routers can have network ACLs. Access control lists can generally be configured to control both inbound and outbound traffic, and in this context they are similar to firewalls.

5

Computer Information Security

Information security means protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, perusal, inspection, recording or destruction. The terms information security, computer security and information assurance are frequently incorrectly used interchangeably. These fields are interrelated often and share the common goals of protecting the confidentiality, integrity and availability of information; however, there are some subtle differences between them. These differences lie primarily in the approach to the subject, the methodologies used, and the areas of concentration. Information security is concerned with the confidentiality, integrity and availability of data regardless of the form the data may take: electronic, print, or other forms. Computer security can focus on ensuring the availability and correct operation of a computer system without concern for the

information stored or processed by the computer. Governments, military, corporations, financial institutions, hospitals, and private businesses amass a great deal of confidential information about their employees, customers, products, research, and financial status. Most of this information is now collected, processed and stored on electronic computers and transmitted across networks to other computers. Should confidential information about a business' customers or finances or new product line fall into the hands of a competitor, such a breach of security could lead to lost business, law suits or even bankruptcy of the business. Protecting confidential information is a business requirement, and in many cases also an ethical and legal requirement. For the individual, information security has a significant effect on privacy, which is viewed very differently in different cultures. The field of information security has grown and evolved significantly in recent years. There are many ways of gaining entry into the field as a career. It offers many areas for specialization including: securing network(s) and allied infrastructure, securing applications and databases, security testing, information systems auditing, business continuity planning and digital forensics science, etc. This article presents a general overview of information security and its core concepts.

HISTORY

Since the early days of writing, heads of state and military commanders understood that it was necessary to provide some mechanism to protect the confidentiality of written correspondence and to have some means of detecting

tampering. Julius Caesar is credited with the invention of the Caesar cipher ca. 50 B.C., which was created in order to prevent his secret messages from being read should a message fall into the wrong hands. World War II brought about many advancements in information security and marked the beginning of the professional field of information security. The end of the 20th century and early years of the 21st century saw rapid advancements in telecommunications, computing hardware and software, and data encryption. The availability of smaller, more powerful and less expensive computing equipment made electronic data processing within the reach of small business and the home user. These computers quickly became interconnected through a network generically called the Internet or World Wide Web. The rapid growth and widespread use of electronic data processing and electronic business conducted through the Internet, along with numerous occurrences of international terrorism, fueled the need for better methods of protecting the computers and the information they store, process and transmit. The academic disciplines of computer security, information security and information assurance emerged along with numerous professional organizations – all sharing the common goals of ensuring the security and reliability of information systems.

BASIC PRINCIPLES

KEY CONCEPTS

For over twenty years, information security has held confidentiality, integrity and availability (known as the CIA

triad) to be the core principles of information security. There is continuous debate about extending this classic trio. Other principles such as Accountability have sometimes been proposed for addition – it has been pointed out that issues such as Non-Repudiation do not fit well within the three core concepts, and as regulation of computer systems has increased (particularly amongst the Western nations) Legality is becoming a key consideration for practical security installations. In 2002, Donn Parker proposed an alternative model for the classic CIA triad that he called the six atomic elements of information. The elements are confidentiality, possession, integrity, authenticity, availability, and utility. The merits of the Parkerian hexad are a subject of debate amongst security professionals.

CONFIDENTIALITY

Confidentiality is the term used to prevent the disclosure of information to unauthorized individuals or systems. For example, a credit card transaction on the Internet requires the credit card number to be transmitted from the buyer to the merchant and from the merchant to a transaction processing network. The system attempts to enforce confidentiality by encrypting the card number during transmission, by limiting the places where it might appear (in databases, log files, backups, printed receipts, and so on), and by restricting access to the places where it is stored. If an unauthorized party obtains the card number in any way, a breach of confidentiality has occurred. Breaches of confidentiality take many forms. Permitting someone to look over your shoulder at your computer screen while you

have confidential data displayed on it could be a breach of confidentiality. If a laptop computer containing sensitive information about a company's employees is stolen or sold, it could result in a breach of confidentiality. Giving out confidential information over the telephone is a breach of confidentiality if the caller is not authorized to have the information. Confidentiality is necessary (but not sufficient) for maintaining the privacy of the people whose personal information a system holds.

INTEGRITY

In information security, integrity means that data cannot be modified undetectably. This is not the same thing as referential integrity in databases, although it can be viewed as a special case of Consistency as understood in the classic ACID model of transaction processing. Integrity is violated when a message is actively modified in transit. Information security systems typically provide message integrity in addition to data confidentiality.

AVAILABILITY

For any information system to serve its purpose, the information must be available when it is needed. This means that the computing systems used to store and process the information, the security controls used to protect it, and the communication channels used to access it must be functioning correctly. High availability systems aim to remain available at all times, preventing service disruptions due to power outages, hardware failures, and system upgrades. Ensuring availability also involves preventing denial-of-service attacks.

AUTHENTICITY

In computing, e-Business and information security it is necessary to ensure that the data, transactions, communications or documents (electronic or physical) are genuine. It is also important for authenticity to validate that both parties involved are who they claim they are.

NON-REPUDIATION

In law, non-repudiation implies one's intention to fulfill their obligations to a contract. It also implies that one party of a transaction cannot deny having received a transaction nor can the other party deny having sent a transaction. Electronic commerce uses technology such as digital signatures and encryption to establish authenticity and non-repudiation.

RISK MANAGEMENT

A comprehensive treatment of the topic of risk management is beyond the scope of this article. However, a useful definition of risk management will be provided as well as some basic terminology and a commonly used process for risk management. The CISA Review Manual 2006 provides the following definition of risk management: *“Risk management is the process of identifying vulnerabilities and threats to the information resources used by an organization in achieving business objectives, and deciding what countermeasures, if any, to take in reducing risk to an acceptable level, based on the value of the information resource to the organization.”* There are two things in this definition that may need some clarification. First, the *process* of risk

management is an ongoing iterative process. It must be repeated indefinitely. The business environment is constantly changing and new threats and vulnerability emerge every day. Second, the choice of countermeasure (computer)s (controls) used to manage risks must strike a balance between productivity, cost, effectiveness of the countermeasure, and the value of the informational asset being protected.

Risk is the likelihood that something bad will happen that causes harm to an informational asset (or the loss of the asset). A vulnerability is a weakness that could be used to endanger or cause harm to an informational asset. A threat is anything (man made or act of nature) that has the potential to cause harm. The likelihood that a threat will use a vulnerability to cause harm creates a risk. When a threat does use a vulnerability to inflict harm, it has an impact. In the context of information security, the impact is a loss of availability, integrity, and confidentiality, and possibly other losses (lost income, loss of life, loss of real property). It should be pointed out that it is not possible to identify all risks, nor is it possible to eliminate all risk. The remaining risk is called *residual risk*. A risk assessment is carried out by a team of people who have knowledge of specific areas of the business. Membership of the team may vary over time as different parts of the business are assessed. The assessment may use a subjective qualitative analysis based on informed opinion, or where reliable dollar figures and historical information is available, the analysis may use quantitative analysis. The research has shown that the most vulnerable point in most information systems is the

human user, operator, designer, or other human The ISO/IEC 27002:2005 Code of practice for information security management recommends the following be examined during a risk assessment:

- security policy,
- organization of information security,
- asset management,
- human resources security,
- physical and environmental security,
- communications and operations management,
- access control,
- information systems acquisition, development and maintenance,
- information security incident management,
- business continuity management, and
- regulatory compliance.

In broad terms, the risk management process consists of:

1. Identification of assets and estimating their value. Include: people, buildings, hardware, software, data (electronic, print, other), supplies.
2. Conduct a threat assessment. Include: Acts of nature, acts of war, accidents, malicious acts originating from inside or outside the organization.
3. Conduct a vulnerability assessment, and for each vulnerability, calculate the probability that it will be exploited. Evaluate policies, procedures, standards, training, physical security, quality control, technical security.

4. Calculate the impact that each threat would have on each asset. Use qualitative analysis or quantitative analysis.
5. Identify, select and implement appropriate controls. Provide a proportional response. Consider productivity, cost effectiveness, and value of the asset.
6. Evaluate the effectiveness of the control measures. Ensure the controls provide the required cost effective protection without discernible loss of productivity.

For any given risk, Executive Management can choose to accept the risk based upon the relative low value of the asset, the relative low frequency of occurrence, and the relative low impact on the business. Or, leadership may choose to mitigate the risk by selecting and implementing appropriate control measures to reduce the risk. In some cases, the risk can be transferred to another business by buying insurance or out-sourcing to another business. The reality of some risks may be disputed. In such cases leadership may choose to deny the risk. This is itself a potential risk.

CONTROLS

When Management chooses to mitigate a risk, they will do so by implementing one or more of three different types of controls.

ADMINISTRATIVE

Administrative controls (also called procedural controls) consist of approved written policies, procedures, standards and guidelines. Administrative controls form the framework

for running the business and managing people. They inform people on how the business is to be run and how day to day operations are to be conducted. Laws and regulations created by government bodies are also a type of administrative control because they inform the business. Some industry sectors have policies, procedures, standards and guidelines that must be followed – the Payment Card Industry (PCI) Data Security Standard required by Visa and Master Card is such an example. Other examples of administrative controls include the corporate security policy, password policy, hiring policies, and disciplinary policies. Administrative controls form the basis for the selection and implementation of logical and physical controls. Logical and physical controls are manifestations of administrative controls. Administrative controls are of paramount importance.

LOGICAL

Logical controls (also called technical controls) use software and data to monitor and control access to information and computing systems. For example: passwords, network and host based firewalls, network intrusion detection systems, access control lists, and data encryption are logical controls. An important logical control that is frequently overlooked is the principle of least privilege. The principle of least privilege requires that an individual, programme or system process is not granted any more access privileges than are necessary to perform the task. A blatant example of the failure to adhere to the principle of least privilege is logging into Windows as user

Administrator to read Email and surf the Web. Violations of this principle can also occur when an individual collects additional access privileges over time. This happens when employees' job duties change, or they are promoted to a new position, or they transfer to another department. The access privileges required by their new duties are frequently added onto their already existing access privileges which may no longer be necessary or appropriate.

PHYSICAL

Physical controls monitor and control the environment of the work place and computing facilities. They also monitor and control access to and from such facilities. For example: doors, locks, heating and air conditioning, smoke and fire alarms, fire suppression systems, cameras, barricades, fencing, security guards, cable locks, etc. Separating the network and work place into functional areas are also physical controls. An important physical control that is frequently overlooked is the separation of duties. Separation of duties ensures that an individual can not complete a critical task by himself. For example: an employee who submits a request for reimbursement should not also be able to authorize payment or print the check. An applications programmer should not also be the server administrator or the database administrator – these roles and responsibilities must be separated from one another.

DEFENSE IN DEPTH

Information security must protect information throughout the life span of the information, from the initial creation of

the information on through to the final disposal of the information. The information must be protected while in motion and while at rest. During its lifetime, information may pass through many different information processing systems and through many different parts of information processing systems. There are many different ways the information and information systems can be threatened. To fully protect the information during its lifetime, each component of the information processing system must have its own protection mechanisms.

The building up, layering on and overlapping of security measures is called defense in depth. The strength of any system is no greater than its weakest link. Using a defence in depth strategy, should one defensive measure fail there are other defensive measures in place that continue to provide protection.

Recall the earlier discussion about administrative controls, logical controls, and physical controls. The three types of controls can be used to form the basis upon which to build a defense-in-depth strategy. With this approach, defense-in-depth can be conceptualized as three distinct layers or planes laid one on top of the other. Additional insight into defense-in- depth can be gained by thinking of it as forming the layers of an onion, with data at the core of the onion, people the next outer layer of the onion, and network security, host-based security and application security forming the outermost layers of the onion. Both perspectives are equally valid and each provides valuable insight into the implementation of a good defense-in-depth strategy.

SECURITY CLASSIFICATION FOR INFORMATION

An important aspect of information security and risk management is recognizing the value of information and defining appropriate procedures and protection requirements for the information. Not all information is equal and so not all information requires the same degree of protection. This requires information to be assigned a security classification. The first step in information classification is to identify a member of senior management as the owner of the particular information to be classified. Next, develop a classification policy. The policy should describe the different classification labels, define the criteria for information to be assigned a particular label, and list the required security controls for each classification. Some factors that influence which classification information should be assigned include how much value that information has to the organization, how old the information is and whether or not the information has become obsolete. Laws and other regulatory requirements are also important considerations when classifying information. The type of information security classification labels selected and used will depend on the nature of the organisation, with examples being:

- In the business sector, labels such as: Public, Sensitive, Private, Confidential.
- In the government sector, labels such as: Unclassified, Sensitive But Unclassified, Restricted, Confidential, Secret, Top Secret and their non-English equivalents.
- In cross-sectoral formations, the Traffic Light Protocol, which consists of: White, Green, Amber and Red.

All employees in the organization, as well as business partners, must be trained on the classification schema and understand the required security controls and handling procedures for each classification. The classification a particular information asset has been assigned should be reviewed periodically to ensure the classification is still appropriate for the information and to ensure the security controls required by the classification are in place.

ACCESS CONTROL

Access to protected information must be restricted to people who are authorized to access the information. The computer programmes, and in many cases the computers that process the information, must also be authorized. This requires that mechanisms be in place to control the access to protected information. The sophistication of the access control mechanisms should be in parity with the value of the information being protected – the more sensitive or valuable the information the stronger the control mechanisms need to be. The foundation on which access control mechanisms are built start with identification and authentication. Identification is an assertion of who someone is or what something is. If a person makes the statement “*Hello, my name is John Doe*” they are making a claim of who they are. However, their claim may or may not be true. Before John Doe can be granted access to protected information it will be necessary to verify that the person claiming to be John Doe really is John Doe. Authentication is the act of verifying a claim of identity. When John Doe goes into a bank to make a withdrawal, he tells the bank

teller he is John Doe (a claim of identity). The bank teller asks to see a photo ID, so he hands the teller his driver's license. The bank teller checks the license to make sure it has John Doe printed on it and compares the photograph on the license against the person claiming to be John Doe. If the photo and name match the person, then the teller has authenticated that John Doe is who he claimed to be.

There are three different types of information that can be used for authentication: something you know, something you have, or something you are. Examples of *something you know* include such things as a PIN, a password, or your mother's maiden name. Examples of *something you have* include a driver's license or a magnetic swipe card. *Something you are* refers to biometrics. Examples of biometrics include palm prints, finger prints, voice prints and retina (eye) scans. Strong authentication requires providing information from two of the three different types of authentication information. For example, something you know plus something you have. This is called two factor authentication. On computer systems in use today, the Username is the most common form of identification and the Password is the most common form of authentication. Usernames and passwords have served their purpose but in our modern world they are no longer adequate. Usernames and passwords are slowly being replaced with more sophisticated authentication mechanisms. After a person, programme or computer has successfully been identified and authenticated then it must be determined what informational resources they are permitted to access and what actions they will be allowed to perform (run, view, create, delete, or change).

This is called authorization. Authorization to access information and other computing services begins with administrative policies and procedures. The policies prescribe what information and computing services can be accessed, by whom, and under what conditions. The access control mechanisms are then configured to enforce these policies. Different computing systems are equipped with different kinds of access control mechanisms - some may even offer a choice of different access control mechanisms. The access control mechanism a system offers will be based upon one of three approaches to access control or it may be derived from a combination of the three approaches. The non-discretionary approach consolidates all access control under a centralized administration. The access to information and other resources is usually based on the individual's function (role) in the organization or the tasks the individual must perform. The discretionary approach gives the creator or owner of the information resource the ability to control access to those resources. In the Mandatory access control approach, access is granted or denied basing upon the security classification assigned to the information resource.

Examples of common access control mechanisms in use today include Role-based access control available in many advanced Database Management Systems, simple file permissions provided in the UNIX and Windows operating systems, Group Policy Objects provided in Windows network systems, Kerberos, RADIUS, TACACS, and the simple access lists used in many firewalls and routers. To be effective, policies and other security controls must be enforceable and upheld. Effective policies ensure that people are held

accountable for their actions. All failed and successful authentication attempts must be logged, and all access to information must leave some type of audit trail.

CRYPTOGRAPHY

Information security uses cryptography to transform usable information into a form that renders it unusable by anyone other than an authorized user; this process is called encryption. Information that has been encrypted (rendered unusable) can be transformed back into its original usable form by an authorized user, who possesses the cryptographic key, through the process of decryption. Cryptography is used in information security to protect information from unauthorized or accidental disclosure while the information is in transit (either electronically or physically) and while information is in storage. Cryptography provides information security with other useful applications as well including improved authentication methods, message digests, digital signatures, non-repudiation, and encrypted network communications. Older less secure application such as telnet and ftp are slowly being replaced with more secure applications such as ssh that use encrypted network communications. Wireless communications can be encrypted using protocols such as WPA/WPA2 or the older (and less secure) WEP. Wired communications (such as ITU-T G.hn) are secured using AES for encryption and X.1035 for authentication and key exchange. Software applications such as GnuPG or PGP can be used to encrypt data files and Email. Cryptography can introduce security problems when it is not implemented correctly. Cryptographic solutions

need to be implemented using industry accepted solutions that have undergone rigorous peer review by independent experts in cryptography. The length and strength of the encryption key is also an important consideration. A key that is weak or too short will produce weak encryption. The keys used for encryption and decryption must be protected with the same degree of rigor as any other confidential information. They must be protected from unauthorized disclosure and destruction and they must be available when needed. PKI solutions address many of the problems that surround key management.

PROCESS

The terms reasonable and prudent person, due care and due diligence have been used in the fields of Finance, Securities, and Law for many years. In recent years these terms have found their way into the fields of computing and information security. U.S.A. Federal Sentencing Guidelines now make it possible to hold corporate officers liable for failing to exercise due care and due diligence in the management of their information systems. In the business world, stockholders, customers, business partners and governments have the expectation that corporate officers will run the business in accordance with accepted business practices and in compliance with laws and other regulatory requirements. This is often described as the “reasonable and prudent person” rule. A prudent person takes due care to ensure that everything necessary is done to operate the business by sound business principles and in a legal ethical manner. A prudent person is also diligent (mindful, attentive,

and ongoing) in their due care of the business. In the field of Information Security, Harris offers the following definitions of due care and due diligence:

“Due care are steps that are taken to show that a company has taken responsibility for the activities that take place within the corporation and has taken the necessary steps to help protect the company, its resources, and employees.” And, [Due diligence are the] *“continual activities that make sure the protection mechanisms are continually maintained and operational.”* Attention should be made to two important points in these definitions. First, in due care, steps are taken to *show* - this means that the steps can be verified, measured, or even produce tangible artifacts. Second, in due diligence, there are continual activities - this means that people are actually doing things to monitor and maintain the protection mechanisms, and these activities are ongoing.

SECURITY GOVERNANCE

The Software Engineering Institute at Carnegie Mellon University, in a publication titled “Governing for Enterprise Security (GES)”, defines characteristics of effective security governance. These include:

- An enterprise-wide issue
- Leaders are accountable
- Viewed as a business requirement
- Risk-based

- Roles, responsibilities, and segregation of duties defined
- Addressed and enforced in policy
- Adequate resources committed
- Staff aware and trained
- A development life cycle requirement
- Planned, managed, measurable, and measured
- Reviewed and audited

INCIDENT RESPONSE PLANS

1 to 3 paragraphs (non technical) that discuss:

- Selecting team members
- Define roles, responsibilities and lines of authority
- Define a security incident
- Define a reportable incident
- Training
- Detection
- Classification
- Escalation
- Containment
- Eradication
- Documentation

CHANGE MANAGEMENT

Change management is a formal process for directing and controlling alterations to the information processing environment. This includes alterations to desktop computers, the network, servers and software. The objectives of change

management are to reduce the risks posed by changes to the information processing environment and improve the stability and reliability of the processing environment as changes are made. It is not the objective of change management to prevent or hinder necessary changes from being implemented. Any change to the information processing environment introduces an element of risk. Even apparently simple changes can have unexpected effects. One of Managements many responsibilities is the management of risk. Change management is a tool for managing the risks introduced by changes to the information processing environment. Part of the change management process ensures that changes are not implemented at inopportune times when they may disrupt critical business processes or interfere with other changes being implemented.

Not every change needs to be managed. Some kinds of changes are a part of the everyday routine of information processing and adhere to a predefined procedure, which reduces the overall level of risk to the processing environment. Creating a new user account or deploying a new desktop computer are examples of changes that do not generally require change management. However, relocating user file shares, or upgrading the Email server pose a much higher level of risk to the processing environment and are not a normal everyday activity. The critical first steps in change management are (a) defining change (and communicating that definition) and (b) defining the scope of the change system. Change management is usually overseen by a Change Review Board composed of representatives from key business areas, security,

networking, systems administrators, Database administration, applications development, desktop support and the help desk. The tasks of the Change Review Board can be facilitated with the use of automated work flow application. The responsibility of the Change Review Board is to ensure the organizations documented change management procedures are followed. The change management process is as follows:

- Requested: Anyone can request a change. The person making the change request may or may not be the same person that performs the analysis or implements the change. When a request for change is received, it may undergo a preliminary review to determine if the requested change is compatible with the organizations business model and practices, and to determine the amount of resources needed to implement the change.
- Approved: Management runs the business and controls the allocation of resources therefore, Management must approve requests for changes and assign a priority for every change. Management might choose to reject a change request if the change is not compatible with the business model, industry standards or best practices. Management might also choose to reject a change request if the change requires more resources than can be allocated for the change.
- Planned: Planning a change involves discovering the scope and impact of the proposed change; analyzing the complexity of the change; allocation of resources

and, developing, testing and documenting both implementation and backout plans. Need to define the criteria on which a decision to back out will be made.

- **Tested:** Every change must be tested in a safe test environment, which closely reflects the actual production environment, before the change is applied to the production environment. The backout plan must also be tested.
- **Scheduled:** Part of the change review board's responsibility is to assist in the scheduling of changes by reviewing the proposed implementation date for potential conflicts with other scheduled changes or critical business activities.
- **Communicated:** Once a change has been scheduled it must be communicated. The communication is to give others the opportunity to remind the change review board about other changes or critical business activities that might have been overlooked when scheduling the change. The communication also serves to make the Help Desk and users aware that a change is about to occur. Another responsibility of the change review board is to ensure that scheduled changes have been properly communicated to those who will be affected by the change or otherwise have an interest in the change.
- **Implemented:** At the appointed date and time, the changes must be implemented. Part of the planning process was to develop an implementation plan, testing plan and, a back out plan. If the

implementation of the change should fail or, the post implementation testing fails or, other “drop dead” criteria have been met, the back out plan should be implemented.

- Documented: All changes must be documented. The documentation includes the initial request for change, its approval, the priority assigned to it, the implementation, testing and back out plans, the results of the change review board critique, the date/time the change was implemented, who implemented it, and whether the change was implemented successfully, failed or postponed.
- Post change review: The change review board should hold a post implementation review of changes. It is particularly important to review failed and backed out changes. The review board should try to understand the problems that were encountered, and look for areas for improvement.

Change management procedures that are simple to follow and easy to use can greatly reduce the overall risks created when changes are made to the information processing environment.

Good change management procedures improve the overall quality and success of changes as they are implemented. This is accomplished through planning, peer review, documentation and communication. ISO/IEC 20000, The Visible OPS Handbook: Implementing ITIL in 4 Practical and Auditable Steps (Full book summary), and Information Technology Infrastructure Library all provide valuable

guidance on implementing an efficient and effective change management programme information security.

BUSINESS CONTINUITY

Business continuity is the mechanism by which an organization continues to operate its critical business units, during planned or unplanned disruptions that affect normal business operations, by invoking planned and managed procedures. Unlike what most people think business continuity is not necessarily an IT system or process, simply because it is about the business. Today disasters or disruptions to business are a reality. Whether the disaster is natural or man-made (the TIME magazine has a website on the top 10), it affects normal life and so business. So why is planning so important? Let us face reality that “all businesses recover”, whether they planned for recovery or not, simply because business is about earning money for survival. The planning is merely getting better prepared to face it, knowing fully well that the best plans may fail. Planning helps to reduce cost of recovery, operational overheads and most importantly sail through some smaller ones effortlessly. For businesses to create effective plans they need to focus upon the following key questions. Most of these are common knowledge, and anyone can do a BCP.

1. Should a disaster strike, what are the first few things that I should do? Should I call people to find if they are OK or call up the bank to figure out my money is safe? This is Emergency Response. Emergency Response services help take the first hit when the disaster strikes and if the disaster is serious enough

the Emergency Response teams need to quickly get a Crisis Management team in place.

2. What parts of my business should I recover first? The one that brings me most money or the one where I spend the most, or the one that will ensure I shall be able to get sustained future growth? The identified sections are the critical business units. There is no magic bullet here, no one answer satisfies all. Businesses need to find answers that meet business requirements.
3. How soon should I target to recover my critical business units? In BCP technical jargon this is called Recovery Time Objective, or RTO. This objective will define what costs the business will need to spend to recover from a disruption. For example, it is cheaper to recover a business in 1 day than in 1 hour.
4. What all do I need to recover the business? IT, machinery, records...food, water, people...So many aspects to dwell upon. The cost factor becomes clearer now...Business leaders need to drive business continuity. Hold on. My IT manager spent \$200000 last month and created a DRP (Disaster Recovery Plan), whatever happened to that? a DRP is about continuing an IT system, and is one of the sections of a comprehensive Business Continuity Plan. Look below for more on this.
5. And where do I recover my business from... Will the business center give me space to work, or would it be flooded by many people queuing up for the same reasons that I am.

6. But once I do recover from the disaster and work in reduced production capacity, since my main operational sites are unavailable, how long can this go on. How long can I do without my original sites, systems, people? this defines the amount of business resilience a business may have.
7. Now that I know how to recover my business. How do I make sure my plan works? Most BCP pundits would recommend testing the plan at least once a year, reviewing it for adequacy and rewriting or updating the plans either annually or when businesses change.

DISASTER RECOVERY PLANNING

While a business continuity plan (BCP) takes a broad approach to dealing with organizational-wide effects of a disaster, a disaster recovery plan (DRP), which is a subset of the business continuity plan, is instead focused on taking the necessary steps to resume normal business operations as quickly as possible. A disaster recovery plan is executed immediately after the disaster occurs and details what steps are to be taken in order to recover critical information technology infrastructure.

LAWS AND REGULATIONS

Below is a partial listing of European, United Kingdom, Canadian and USA governmental laws and regulations that have, or will have, a significant effect on data processing and information security. Important industry sector regulations have also been included when they have a significant impact on information security.

Special Focus on Computer Security Model

- UK Data Protection Act 1998 makes new provisions for the regulation of the processing of information relating to individuals, including the obtaining, holding, use or disclosure of such information. The European Union Data Protection Directive (EUDPD) requires that all EU member must adopt national regulations to standardize the protection of data privacy for citizens throughout the EU.
- The Computer Misuse Act 1990 is an Act of the UK Parliament making computer crime (e.g. cracking - sometimes incorrectly referred to as hacking) a criminal offence. The Act has become a model upon which several other countries including Canada and the Republic of Ireland have drawn inspiration when subsequently drafting their own information security laws.
- EU Data Retention laws requires Internet service providers and phone companies to keep data on every electronic message sent and phone call made for between six months and two years.
- The Family Educational Rights and Privacy Act (FERPA) (20 U.S.C. § 1232 g; 34 CFR Part 99) is a USA Federal law that protects the privacy of student education records. The law applies to all schools that receive funds under an applicable programme of the U.S. Department of Education. Generally, schools must have written permission from the parent or eligible student in order to release any information from a student's education record.
- Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires the adoption of national

Special Focus on Computer Security Model

standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. And, it requires health care providers, insurance providers and employers to safeguard the security and privacy of health data.

- Gramm-Leach-Bliley Act of 1999 (GLBA), also known as the Financial Services Modernization Act of 1999, protects the privacy and security of private financial information that financial institutions collect, hold, and process.
- Sarbanes-Oxley Act of 2002 (SOX). Section 404 of the act requires publicly traded companies to assess the effectiveness of their internal controls for financial reporting in annual reports they submit at the end of each fiscal year. Chief information officers are responsible for the security, accuracy and the reliability of the systems that manage and report the financial data. The act also requires publicly traded companies to engage independent auditors who must attest to, and report on, the validity of their assessments.
- Payment Card Industry Data Security Standard (PCI DSS) establishes comprehensive requirements for enhancing payment account data security. It was developed by the founding payment brands of the PCI Security Standards Council, including American Express, Discover Financial Services, JCB, MasterCard Worldwide and Visa International, to help facilitate the broad adoption of consistent data security measures on a global basis. The PCI DSS is a

multifaceted security standard that includes requirements for security management, policies, procedures, network architecture, software design and other critical protective measures.

- State Security Breach Notification Laws (California and many others) require businesses, nonprofits, and state institutions to notify consumers when unencrypted “personal information” may have been compromised, lost, or stolen.
- Personal Information Protection and Electronics Document Act (PIPEDA) – An Act to support and promote electronic commerce by protecting personal information that is collected, used or disclosed in certain circumstances, by providing for the use of electronic means to communicate or record information or transactions and by amending the Canada Evidence Act, the Statutory Instruments Act and the Statute Revision Act.

SOURCES OF STANDARDS

International Organization for Standardization (ISO) is a consortium of national standards institutes from 157 countries, coordinated through a secretariat in Geneva, Switzerland. ISO is the world’s largest developer of standards. ISO 15443: “Information technology - Security techniques - A framework for IT security assurance”, ISO/IEC 27002: “Information technology - Security techniques - Code of practice for information security management”, ISO-20000: “Information technology - Service management”, and ISO/IEC27001: “Information technology - Security techniques -

Information security management systems - Requirements” are of particular interest to information security professionals.

The USA National Institute of Standards and Technology (NIST) is a non-regulatory federal agency within the U.S. Department of Commerce. The NIST Computer Security Division develops standards, metrics, tests and validation programmes as well as publishes standards and guidelines to increase secure IT planning, implementation, management and operation. NIST is also the custodian of the USA Federal Information Processing Standard publications (FIPS). The Internet Society is a professional membership society with more than 100 organization and over 20,000 individual members in over 180 countries. It provides leadership in addressing issues that confront the future of the Internet, and is the organization home for the groups responsible for Internet infrastructure standards, including the Internet Engineering Task Force (IETF) and the Internet Architecture Board (IAB). The ISOC hosts the Requests for Comments (RFCs) which includes the Official Internet Protocol Standards and the RFC-2196 Site Security Handbook.

The Information Security Forum is a global nonprofit organization of several hundred leading organizations in financial services, manufacturing, telecommunications, consumer goods, government, and other areas. It undertakes research into information security practices and offers advice in its biannual Standard of Good Practice and more detailed advisories for members. The IT Baseline Protection Catalogs, or IT-Grundschutz Catalogs, (“IT Baseline Protection Manual” before 2005) are a collection of documents from the German

Federal Office for Security in Information Technology (FSI), useful for detecting and combating security-relevant weak points in the IT environment (“IT cluster”). The collection encompasses over 3000 pages with the introduction and catalogs.

PROFESSIONALISM

Information security professionalism is the set of knowledge that people working in Information security and similar fields (Information Assurance and Computer security) should have and eventually demonstrate through certifications from well respected organizations. It also encompasses the education process required to accomplish different tasks in these fields. Information technology adoption is always increasing and spread to vital infrastructure for civil and military organizations. Everybody can get involved in the Cyberwar. It is crucial that a nation can have skilled professional to defend its vital interests.

CONCLUSION

Information security is the ongoing process of exercising due care and due diligence to protect information, and information systems, from unauthorized access, use, disclosure, destruction, modification, or disruption or distribution. The never ending process of information security involves ongoing training, assessment, protection, monitoring & detection, incident response & repair, documentation, and review. This makes information security an indispensable part of all the business operations across different domains.

6

Data Security Model

Data security is the means of ensuring that data is kept safe from corruption and that access to it is suitably controlled. Thus data security helps to ensure privacy. It also helps in protecting personal data.

DATA SECURITY TECHNOLOGIES

DISK ENCRYPTION

Disk encryption refers to encryption technology that encrypts data on a hard disk drive. Disk encryption typically takes form in either software or hardware. Disk encryption is often referred to as on-the-fly encryption (“OTFE”) or transparent encryption.

HARDWARE BASED MECHANISMS FOR PROTECTING DATA

Software based security solutions encrypt the data to prevent data from being stolen. However, a malicious

programme or a hacker may corrupt the data in order to make it unrecoverable or unusable. Similarly, encrypted operating systems can be corrupted by a malicious programme or a hacker, making the system unusable. Hardware-based security solutions can prevent read and write access to data and hence offers very strong protection against tampering and unauthorized access. Hardware based or assisted computer security offers an alternative to software-only computer security. Security tokens such as those using PKCS#11 may be more secure due to the physical access required in order to be compromised. Access is enabled only when the token is connected and correct PIN is entered. However, dongles can be used by anyone who can gain physical access to it. Newer technologies in hardware based security solves this problem offering fool proof security for data.

Working of Hardware based security: A hardware device allows a user to login, logout and to set different privilege levels by doing manual actions. The device uses biometric technology to prevent malicious users from logging in, logging out, and changing privilege levels. The current state of a user of the device is read by controllers in peripheral devices such as harddisks. Illegal access by a malicious user or a malicious programme is interrupted based on the current state of a user by harddisk and DVD controllers making illegal access to data impossible. Hardware based access control is more secure than protection provided by the operating systems as operating systems are vulnerable to malicious attacks by viruses and hackers. The data on harddisks can be corrupted after a malicious access is

obtained. With hardware based protection, software cannot manipulate the user privilege levels, it is impossible for a hacker or a malicious programme to gain access to secure data protected by hardware or perform unauthorized privileged operations. The hardware protects the operating system image and file system privileges from being tampered. Therefore, a completely secure system can be created using a combination of hardware based security and secure system administration policies.

BACKUPS

Backups are used to ensure data which is lost can be recovered

DATA MASKING

Data Masking of structured data is the process of obscuring (masking) specific data within a database table or cell to ensure that data security is maintained and sensitive information is not exposed to unauthorized personnel. This may include masking the data from users (for example so banking customer representatives can only see the last 4 digits of a customers national identity number), developers (who need real production data to test new software releases but should not be able to see sensitive financial data), outsourcing vendors, etc.

DATA ERASURE

Data erasure is a method of software-based overwriting that completely destroys all electronic data residing on a hard drive or other digital media to ensure that no sensitive data is leaked when an asset is retired or reused.

INTERNATIONAL LAWS AND STANDARDS

INTERNATIONAL LAWS

In the UK, the Data Protection Act is used to ensure that personal data is accessible to those whom it concerns, and provides redress to individuals if there are inaccuracies. This is particularly important to ensure individuals are treated fairly, for example for credit checking purposes. The Data Protection Act states that only individuals and companies with legitimate and lawful reasons can process personal information and cannot be shared.

INTERNATIONAL STANDARDS

The International Standard ISO/IEC 17799 covers data security under the topic of information security, and one of its cardinal principles is that all stored information, i.e. data, should be owned so that it is clear whose responsibility it is to protect and control access to that data. The Trusted Computing Group is an organization that helps standardize computing security technologies.

DATABASE MODEL

A '*database model*' is the theoretical foundation of a database and fundamentally determines in which manner data can be stored, organized and manipulated in a database system. It thereby defines the infrastructure offered by a particular database system. The most popular example of a database model is the relational model.

OVERVIEW

A database model is a theory or specification describing how a database is structured and used. Several such models have been suggested. Common models include:

- Hierarchical model
- Network model
- Relational model
- Entity-relationship
- Object-relational model
- Object model

A data model is not just a way of structuring data: it also defines a set of operations that can be performed on the data. The relational model, for example, defines operations such as select, project, and join. Although these operations may not be explicit in a particular query language, they provide the foundation on which a query language is built.

MODELS

Various techniques are used to model data structure. Most database systems are built around one particular data model, although it is increasingly common for products to offer support for more than one model. For any one logical model various physical implementations may be possible, and most products will offer the user some level of control in tuning the physical implementation, since the choices that are made have a significant effect on performance. An example of this is the relational model: all serious implementations of the relational model allow the creation of indexes which provide fast access to rows in a table if the values of certain columns are known.

FLAT MODEL

The flat (or table) model consists of a single, two-dimensional array of data elements, where all members of a given column are assumed to be similar values, and all members of a row are assumed to be related to one another. For instance, columns for name and password that might be used as a part of a system security database. Each row would have the specific password associated with an individual user. Columns of the table often have a type associated with them, defining them as character data, date or time information, integers, or floating point numbers. This may not strictly qualify as a data model, as defined above.

HIERARCHICAL MODEL

In a hierarchical model, data is organized into a tree-like structure, implying a single upward link in each record to describe the nesting, and a sort field to keep the records in a particular order in each same-level list. Hierarchical structures were widely used in the early mainframe database management systems, such as the Information Management System (IMS) by IBM, and now describe the structure of XML documents. This structure allows one 1:N relationship between two types of data. This structure is very efficient to describe many relationships in the real world; recipes, table of contents, ordering of paragraphs/verses, any nested and sorted information. However, the hierarchical structure is inefficient for certain database operations when a full path (as opposed to upward link and sort field) is not also included for each record. Parent-child relationship: Child

may only have one parent but a parent can have multiple children. Parents and children are tied together by links called “pointers”. A parent will have a list of pointers to each of their children.

NETWORK MODEL

The network model (defined by the CODASYL specification) organizes data using two fundamental constructs, called *records* and *sets*. Records contain fields (which may be organized hierarchically, as in the programming language COBOL). Sets (not to be confused with mathematical sets) define one-to-many relationships between records: one owner, many members. A record may be an owner in any number of sets, and a member in any number of sets.

The network model is a variation on the hierarchical model, to the extent that it is built on the concept of multiple branches (lower-level structures) emanating from one or more nodes (higher-level structures), while the model differs from the hierarchical model in that branches can be connected to multiple nodes. The network model is able to represent redundancy in data more efficiently than in the hierarchical model.

The operations of the network model are navigational in style: a programme maintains a current position, and navigates from one record to another by following the relationships in which the record participates. Records can also be located by supplying key values. Although it is not an essential feature of the model, network databases generally implement the set relationships by means of pointers that directly address the location of a record on

disk. This gives excellent retrieval performance, at the expense of operations such as database loading and reorganization.

Most object databases use the navigational concept to provide fast navigation across networks of objects, generally using object identifiers as “smart” pointers to related objects. Objectivity/DB, for instance, implements named 1:1, 1:many, many:1 and many:many named relationships that can cross databases. Many object databases also support SQL, combining the strengths of both models.

RELATIONAL MODEL

The relational model was introduced by E.F. Codd in 1970 as a way to make database management systems more independent of any particular application. It is a mathematical model defined in terms of predicate logic and set theory. The products that are generally referred to as relational databases in fact implement a model that is only an approximation to the mathematical model defined by Codd. Three key terms are used extensively in relational database models: *relations*, *attributes*, and *domains*. A relation is a table with columns and rows. The named columns of the relation are called attributes, and the domain is the set of values the attributes are allowed to take. The basic data structure of the relational model is the table, where information about a particular entity (say, an employee) is represented in rows (also called tuples) and columns. Thus, the “relation” in “relational database” refers to the various tables in the database; a relation is a set of tuples. The columns enumerate the various attributes of

the entity (the employee's name, address or phone number, for example), and a row is an actual instance of the entity (a specific employee) that is represented by the relation. As a result, each tuple of the employee table represents various attributes of a single employee.

All relations (and, thus, tables) in a relational database have to adhere to some basic rules to qualify as relations. First, the ordering of columns is immaterial in a table. Second, there can't be identical tuples or rows in a table. And third, each tuple will contain a single value for each of its attributes. A relational database contains multiple tables, each similar to the one in the "flat" database model. One of the strengths of the relational model is that, in principle, any value occurring in two different records (belonging to the same table or to different tables), implies a relationship among those two records. Yet, in order to enforce explicit integrity constraints, relationships between records in tables can also be defined explicitly, by identifying or non-identifying parent-child relationships characterized by assigning cardinality (1:1, (0)1:M, M:M). Tables can also have a designated single attribute or a set of attributes that can act as a "key", which can be used to uniquely identify each tuple in the table. A key that can be used to uniquely identify a row in a table is called a primary key. Keys are commonly used to join or combine data from two or more tables. For example, an *Employee* table may contain a column named *Location* which contains a value that matches the key of a *Location* table. Keys are also critical in the creation of indexes, which facilitate fast retrieval of data from large tables.

Any column can be a key, or multiple columns can be grouped together into a compound key. It is not necessary to define all the keys in advance; a column can be used as a key even if it was not originally intended to be one. A key that has an external, real-world meaning (such as a person's name, a book's ISBN, or a car's serial number) is sometimes called a "natural" key. If no natural key is suitable (think of the many people named *Brown*), an arbitrary or surrogate key can be assigned (such as by giving employees ID numbers). In practice, most databases have both generated and natural keys, because generated keys can be used internally to create links between rows that cannot break, while natural keys can be used, less reliably, for searches and for integration with other databases. (For example, records in two independently developed databases could be matched up by social security number, except when the social security numbers are incorrect, missing, or have changed.)

DIMENSIONAL MODEL

The dimensional model is a specialized adaptation of the relational model used to represent data in data warehouses in a way that data can be easily summarized using OLAP queries. In the dimensional model, a database consists of a single large table of facts that are described using dimensions and measures. A dimension provides the context of a fact (such as who participated, when and where it happened, and its type) and is used in queries to group related facts together. Dimensions tend to be discrete and are often hierarchical; for example, the location might include

the building, state, and country. A measure is a quantity describing the fact, such as revenue. It's important that measures can be meaningfully aggregated - for example, the revenue from different locations can be added together. In an OLAP query, dimensions are chosen and the facts are grouped and added together to create a summary. The dimensional model is often implemented on top of the relational model using a star schema, consisting of one table containing the facts and surrounding tables containing the dimensions. Particularly complicated dimensions might be represented using multiple tables, resulting in a snowflake schema. A data warehouse can contain multiple star schemas that share dimension tables, allowing them to be used together. Coming up with a standard set of dimensions is an important part of dimensional modeling.

OBJECTIONAL DATABASE MODELS

In recent years, the object-oriented paradigm has been applied to database technology, creating a new programming model known as object databases. These databases attempt to bring the database world and the application programming world closer together, in particular by ensuring that the database uses the same type system as the application programme. This aims to avoid the overhead (sometimes referred to as the *impedance mismatch*) of converting information between its representation in the database (for example as rows in tables) and its representation in the application programme (typically as objects). At the same time, object databases attempt to introduce the key ideas of object programming, such as encapsulation and

polymorphism, into the world of databases. A variety of these ways have been tried for storing objects in a database. Some products have approached the problem from the application programming end, by making the objects manipulated by the programme persistent.

This also typically requires the addition of some kind of query language, since conventional programming languages do not have the ability to find objects based on their information content. Others have attacked the problem from the database end, by defining an object-oriented data model for the database, and defining a database programming language that allows full programming capabilities as well as traditional query facilities. Object databases suffered because of a lack of standardization: although standards were defined by ODMG, they were never implemented well enough to ensure interoperability between products. Nevertheless, object databases have been used successfully in many applications: usually specialized applications such as engineering databases or molecular biology databases rather than mainstream commercial data processing. However, object database ideas were picked up by the relational vendors and influenced extensions made to these products and indeed to the SQL language.

OBJECT DATABASE

An object database (also object-oriented database) is a database model in which information is represented in the form of objects as used in object-oriented programming. Object databases are a niche field within the broader database management system (DBMS) market dominated

by relational database management systems. Object databases have been considered since the early 1980s and 1990s but they have made little impact on mainstream commercial data processing, though there is some usage in specialized areas.

OVERVIEW

When database capabilities are combined with object-oriented programming language capabilities, the result is an object-oriented database management system (OODBMS). Today's trend in programming languages is to utilize objects, thereby making OODBMS ideal for object-oriented programmers because they can develop the product, store them as objects, and can replicate or modify existing objects to make new objects within the OODBMS. Information today includes not only data but video, audio, graphs, and photos which are considered complex data types. Relational DBMS are not natively capable of supporting these complex data types. By being integrated with the programming language, the programmer can maintain consistency within one environment because both the OODBMS and the programming language will use the same model of representation. Relational DBMS projects using complex data types would have to be divided into two separate tasks: the database model and the application.

As the usage of web-based technology increases with the implementation of Intranets and extranets, companies have a vested interest in OODBMS to display their complex data. Using a DBMS that has been specifically designed to store data as objects gives an advantage to those companies that

are geared towards multimedia presentation or organizations that utilize computer-aided design (CAD). Some object-oriented databases are designed to work well with object-oriented programming languages such as Ruby, Python, Perl, Java, C#, Visual Basic .NET, C++, Objective-C and Smalltalk; others have their own programming languages. OODBMSs use exactly the same model as object-oriented programming languages.

HISTORY

Object database management systems grew out of research during the early to mid-1970s into having intrinsic database management support for graph-structured objects. The term “object-oriented database system” first appeared around 1985. Notable research projects included Encore-Ob/Server (Brown University), EXODUS (University of Wisconsin–Madison), IRIS (Hewlett-Packard), ODE (Bell Labs), ORION (Microelectronics and Computer Technology Corporation or MCC), Vodak (GMD-IPSI), and Zeitgeist (Texas Instruments). The ORION project had more published papers than any of the other efforts. Won Kim of MCC compiled the best of those papers in a book published by The MIT Press. Early commercial products included Gemstone (Servio Logic, name changed to GemStone Systems), Gbase (Graphael), and Vbase (Ontologic). The early to mid-1990s saw additional commercial products enter the market. These included ITASCA (Itasca Systems), Jasmine (Fujitsu, marketed by Computer Associates), Matisse (Matisse Software), Objectivity/DB (Objectivity, Inc.), ObjectStore (Progress Software, acquired from eXcelon which was

originally Object Design), ONTOS (Ontos, Inc., name changed from Ontologic), O₂ (O₂ Technology, merged with several companies, acquired by Informix, which was in turn acquired by IBM), POET (now FastObjects from Versant which acquired Poet Software), Versant Object Database (Versant Corporation), VOSS (Logic Arts) and JADE (Jade Software Corporation).

Some of these products remain on the market and have been joined by new open source and commercial products such as InterSystems CACHE. Object database management systems added the concept of persistence to object programming languages. The early commercial products were integrated with various languages: GemStone (Smalltalk), Gbase (LISP), Vbase (COP) and VOSS (Virtual Object Storage System for Smalltalk). For much of the 1990s, C++ dominated the commercial object database management market. Vendors added Java in the late 1990s and more recently, C#. Starting in 2004, object databases have seen a second growth period when open source object databases emerged that were widely affordable and easy to use, because they are entirely written in OOP languages like Smalltalk, Java or C#, such as db4o (db4objects), DTS/S1 from Obsidian Dynamics and Perst (McObject), available under dual open source and commercial licensing.

TIMELINE

- 1985 – Term Object Database first introduced
- 1988
 - o Versant Corporation started (as Object Sciences Corp)

- o Objectivity, Inc. founded
- Early 1990s
 - o Gemstone (Smalltalk)
 - o GBase (LISP)
 - o VBase (O2- ONTOS – INFORMIX)
 - o Objectivity/DB launched
- Mid 1990's
 - o Versant Object Database
 - o ObjectStore
 - o Poet
 - o Jade
 - o Matisse
- 2000's
 - o Cache'
 - o db4o project started by Carl Rosenberger
 - o ObjectDB for Java
- 2001
 - o IBM acquires Informix (Illustra) integrates with DB2
 - o db4o shipped to first pilot customer
- 2004 - db4o's commercial launch as db4objects, Inc.
- 2008 - db4o acquired by Versant Corporation

ADOPTION OF OBJECT DATABASES

Object databases based on persistent programming acquired a niche in application areas such as engineering and spatial databases, telecommunications, and scientific areas such as high energy physics and molecular biology. They have made little impact on mainstream commercial

data processing, though there is some usage in specialized areas of financial services. It is also worth noting that object databases held the record for the World's largest database (being the first to hold over 1000 terabytes at Stanford Linear Accelerator Center) and the highest ingest rate ever recorded for a commercial database at over one Terabyte per hour. Another group of object databases focuses on embedded use in devices, packaged software, and real-time systems.

TECHNICAL FEATURES

Most object databases also offer some kind of query language, allowing objects to be found by a more declarative programming approach. It is in the area of object query languages, and the integration of the query and navigational interfaces, that the biggest differences between products are found. An attempt at standardization was made by the ODMG with the Object Query Language, OQL. Access to data can be faster because joins are often not needed (as in a tabular implementation of a relational database). This is because an object can be retrieved directly without a search, by following pointers. (It could, however, be argued that "joining" is a higher-level abstraction of pointer following.) Another area of variation between products is in the way that the schema of a database is defined. A general characteristic, however, is that the programming language and the database schema use the same type definitions. Multimedia applications are facilitated because the class methods associated with the data are responsible for its correct interpretation.

Many object databases, for example VOSS, offer support for versioning. An object can be viewed as the set of all its versions. Also, object versions can be treated as objects in their own right. Some object databases also provide systematic support for triggers and constraints which are the basis of active databases. The efficiency of such a database is also greatly improved in areas which demand massive amounts of data about one item. For example, a banking institution could get the user's account information and provide them efficiently with extensive information such as transactions, account information entries etc. The Big O Notation for such a database paradigm drops from $O(n)$ to $O(1)$, greatly increasing efficiency in these specific cases.

STANDARDS

The Object Data Management Group (ODMG) was a consortium of object database and object-relational mapping vendors, members of the academic community, and interested parties. Its goal was to create a set of specifications that would allow for portable applications that store objects in database management systems. It published several versions of its specification. The last release was ODMG 3.0. By 2001, most of the major object database and object-relational mapping vendors claimed conformance to the ODMG Java Language Binding. Compliance to the other components of the specification was mixed.

In 2001, the ODMG Java Language Binding was submitted to the Java Community Process as a basis for the Java Data Objects specification. The ODMG member companies then decided to concentrate their efforts on the Java Data Objects

specification. As a result, the ODMG disbanded in 2001. Many object database ideas were also absorbed into SQL:1999 and have been implemented in varying degrees in object-relational database products. In 2005 Cook, Rai, and Rosenberger proposed to drop all standardization efforts to introduce additional object-oriented query APIs but rather use the OO programming language itself, i.e., Java and .NET, to express queries. As a result, Native Queries emerged.

Similarly, Microsoft announced Language Integrated Query (LINQ) and DLINQ, an implementation of LINQ, in September 2005, to provide close, language-integrated database query capabilities with its programming languages C# and VB.NET 9. In February 2006, the Object Management Group (OMG) announced that they had been granted the right to develop new specifications based on the ODMG 3.0 specification and the formation of the Object Database Technology Working Group (ODBT WG). The ODBT WG planned to create a set of standards that would incorporate advances in object database technology (e.g., replication), data management (e.g., spatial indexing), and data formats (e.g., XML) and to include new features into these standards that support domains where object databases are being adopted (e.g., real-time systems). The work of the ODBT WG was suspended in March 2009 when, subsequent to the economic turmoil in late 2008, the ODB vendors involved in this effort decided to focus their resources elsewhere. In January 2007 the World Wide Web Consortium gave final recommendation status to the XQuery language. XQuery uses XML as its data model. Some of the ideas developed originally for object databases found their way into XQuery,

but XQuery is not intrinsically object-oriented. Because of the popularity of XML, XQuery engines compete with object databases as a vehicle for storage of data that is too complex or variable to hold conveniently in a relational database.

COMPARISON WITH RDBMSS

An object database stores complex data and relationships between data directly, without mapping to relational rows and columns, and this makes them suitable for applications dealing with very complex data. Objects have a many to many relationship and are accessed by the use of pointers. Pointers are linked to objects to establish relationships. Another benefit of OODBMS is that it can be programmed with small procedural differences without affecting the entire system. This is most helpful for those organizations that have data relationships that are not entirely clear or need to change these relations to satisfy the new business requirements.