CompTIA CASP+ Master Cheat Sheet

1. TPM and vTPM

Trusted Platform Module (TPM). A TPM is a cryptographic processor which behaves much like a **Hardware Security Module (HSM).** Usually seen on newer laptops. **TPM** contains: **true random number generator, key generator, hash generator and secure key store.**

Virtual Trusted Platform **Module** (vTPM). **A vTPM** is a software module that performs the function of a TPM in a virtualized environment. The TPM is a specialized **chip** on an endpoint device that stores RSA encryption keys specific to the host system for hardware authentication. Each TPM **chip** contains an RSA key pair called the Endorsement Key (EIS). The vTPM makes secure storage and cryptographic **functions** available to operating systems and applications running in virtual machines.

NOTE: BitLocker disk encryption normally requires a TPM on Windows. Microsoft's EFS encryption can never use a TPM**. The** new "device encryption" feature on Windows 10 and 8.1 also requires a modern TPM, which is why it's only enabled on new hardware.

2. SELinux Mandatory Access Control (MAC)

MAC evolved out of the multi-level **security (AILS)** program at NSA and Do D. An NSA research project called SE Linux **added a Mandatory** Access Control architecture to the **Linux Kernel,** which was merged into **the mainline** version of Linux in August 2003[1]. **NOTE: Red** Hat Enterprise Linux version **4 (and** later versions) come with a SE Linux-enabled kernel.

3. Storage Area Network (SAN)

A storage area **network (SAN) is a** network **which provides access to** consolidated, block level data **storage. SANS are primarily used to enhance** storage devices, such as disk arrays, tape libraries, and optical jukeboxes, accessible to servers so **that the** devices appear to the operating system as locally attached devices. A SAN typically has its own network of storage devices **that** are generally not accessible through the local area network (LAN) by other devices.

4. Multiple 0/5 encryption on a single SDD A solid-state drive (SSD) is a nonvolatile storage device that stores persistent data on solid-state flash memory. Solid-state drives actually aren't hard drives in the traditional sense of the term, as there are no moving parts involved. A traditional hard disk drive (1-IDD) consists of a spinning disk with a read/write head on a mechanical arm called an actuator. An SSD, on the other hand, has an array of semiconductor memory organized as a disk drive, using integrated circuits (ICs) rather than magnetic or optical storage media. An SSD may also be referred to as a solid-state disk3.

Windows 7: The first issue here is that the **only versions of Windows 7** that support Bitlocker are **Enterprise and Ultimate.** If you have Windows 7 Pro, you're out of luck. You may also find issues on trying to boot the Windows 7 installer on a PC with UEFI, in which case you may have to drop back to legacy mode to work round this.

Unlike Windows 8, if the PC doesn't have a **Trusted Platform Module, the** only way to unlock the Bitlocker drive on boot is to use a USB key. Without a **TPIVI,** you cannot use a password for **Bitlocker in Windows 7.**

If the PC is using **Windows 7 Ultimate or Enterprise,** or Windows 8 Pro or Enterprise, you **can use Bitocker, which comes with** these versions of Windows. Out you have to know what you're **doing.**

BitLocker works best in an environment where a professional **IT** department serves users who may not know what the word ***encrypt*** **means.** You can set it up so that the user doesn't even know that the drive is encrypted. When they log into Windows with their password, they get access to the encrypted files. If they log into another account, or boot with another OS , the files are unreadable .

What's more, if you need to reinstall Windows, or restore the files from a backup, you'll need a special digital key that's created when you encrypt the drive. That key has to be stored elsewhere and someone has to know where **to** find it.

5. TOCTOU attacks

h software development, **time of check to time of use (TOCTTOU** or **TOCTOU, pronounced " *TOCK too")*** is a class of software bug caused by changes in a system between the *chechingof* a condition (such as a security credential) and the *used* **the** results of that check. **This** is one **example** of a race condition.[6]

Time of check (TOO) — When the resource is inspected. For example, all data from **the browser is considered litainted" because a malicious user may have manipulated it. If** the data passes a validation function, then the taint may be removed and the data permitted entry deeper into the app. For example, the app checks whether an email address is well-formed or text contains <script> tags.

Time of use (T U) — When the app performs an operation on the resource. For example, inserting the data into a SQL statement or inserting text into a web page. Weaknesses occur when the app assumes the state of the resource has not changed since the last check-, vulnerabilities occur when the state change relates to a security contro1[7].

CWE-357: Time-of-check Time-of-use (TOCTOU) Race Condition

Description Summary

The software checks the state of a resource before using that resource, but the resource's state can change between the check and the **use in a way that invalidates the results of the check.** This can cause the software to perform invalid actions when the resource is in an unexpected state.

Extended Description

This weakness can be security-relevant when an attacker can influence the state of the resource between check and use. This can happen with shared resources such as files, memory, or e)den variables in multithreaded programs B,

CAPEC-29: Leveraging Time-of-Check and Time-of-Use (TOCTOU) Race Conditions

This attack targets a race condition occurring between **the time of check** (stale) for a resource and the time of use of a resource. The typical exam pre is the **file access.** The attacker can !swage a file access race condition by **"running the race",** meaning that he would modify the resource between the first time the target program accesses the file and the time the target program uses the file. During that period of time, the attacker could do something such as **replace the file** and cause an escalation of privilege

6. Network File System (NFS) and Common Internet File System (CIFS)

What is the difference between NFS and CIFS? Can you explain when you should use CIFS vs NFS?

NFS is the "Network File System" for Unix and Linux operating systems. It allows dies to be shared transparently between servers, desktops, laptops etc. It is a client/sender application that allows a user to view, store and update files on a remote computer as though they were on their own computer, Using NFS, the user or a system administrator can mount all or a portion of a file system.

CIFS is the "Common Internet File System" used by Windows operating systems for file sharing. C1FS uses the client/server programming model. A client program makes a request of a sender program (usually in another computer) for access to a file or to pass a message to a program that runs in the server computer, The sender takes the requested action and returns a response, CIFS is a public or open variation of the Sender Message Block Protocol (SMB) developed and used by Microsoft, and it uses the TCP/IP/Protocol

NFS and CIFS are the primary He systems used in NAS. Comparing CIFS vs. NFS, CIFS tends to be a bit more "chatty" in its communications. This may require file protocol optimization oar a wide area network.

7. Storage Area Network (SAN) protocols

Internet Small Computer System Interface (iSCSI). iSCSI (I.P.-based protocol) works on top of the Transport Control Protocol (TCP) and allows the SCSI command to be sent end-to-end over local-area networks (LANs), wide-area networks (INANs) or the Internet. IBM developed iSCSI as a proof of concept in 1998, and presented the first draft of the iSCSI standard to the Internet Engineering Task Force (IETF) in 2000. The protocol was ratified in 2003.

The major server vendors offer iSCS I as a connectivity option, to provide virtual machines with block-level access to storage volumes, without the need to deploy high-performance Fibre Channel hardware. iSCS I has been proven capable of supporting enterprise-class applications assuming that the specific solution can reach the desired level of performance and scalability and that management is simple enough not to negatively affect the total cost of ownership (·IC0)[1o].

Fibre Channel over Ethernet (FCoE). FCoe is a standards-based protocol that natively maps Fibre Channel to Ethernet for transport in a lossless Ethernet LAN. FCoE allows the consolidation of LAN and Fibre Channel SAN traffic over a single switching infrastructure in the data center.

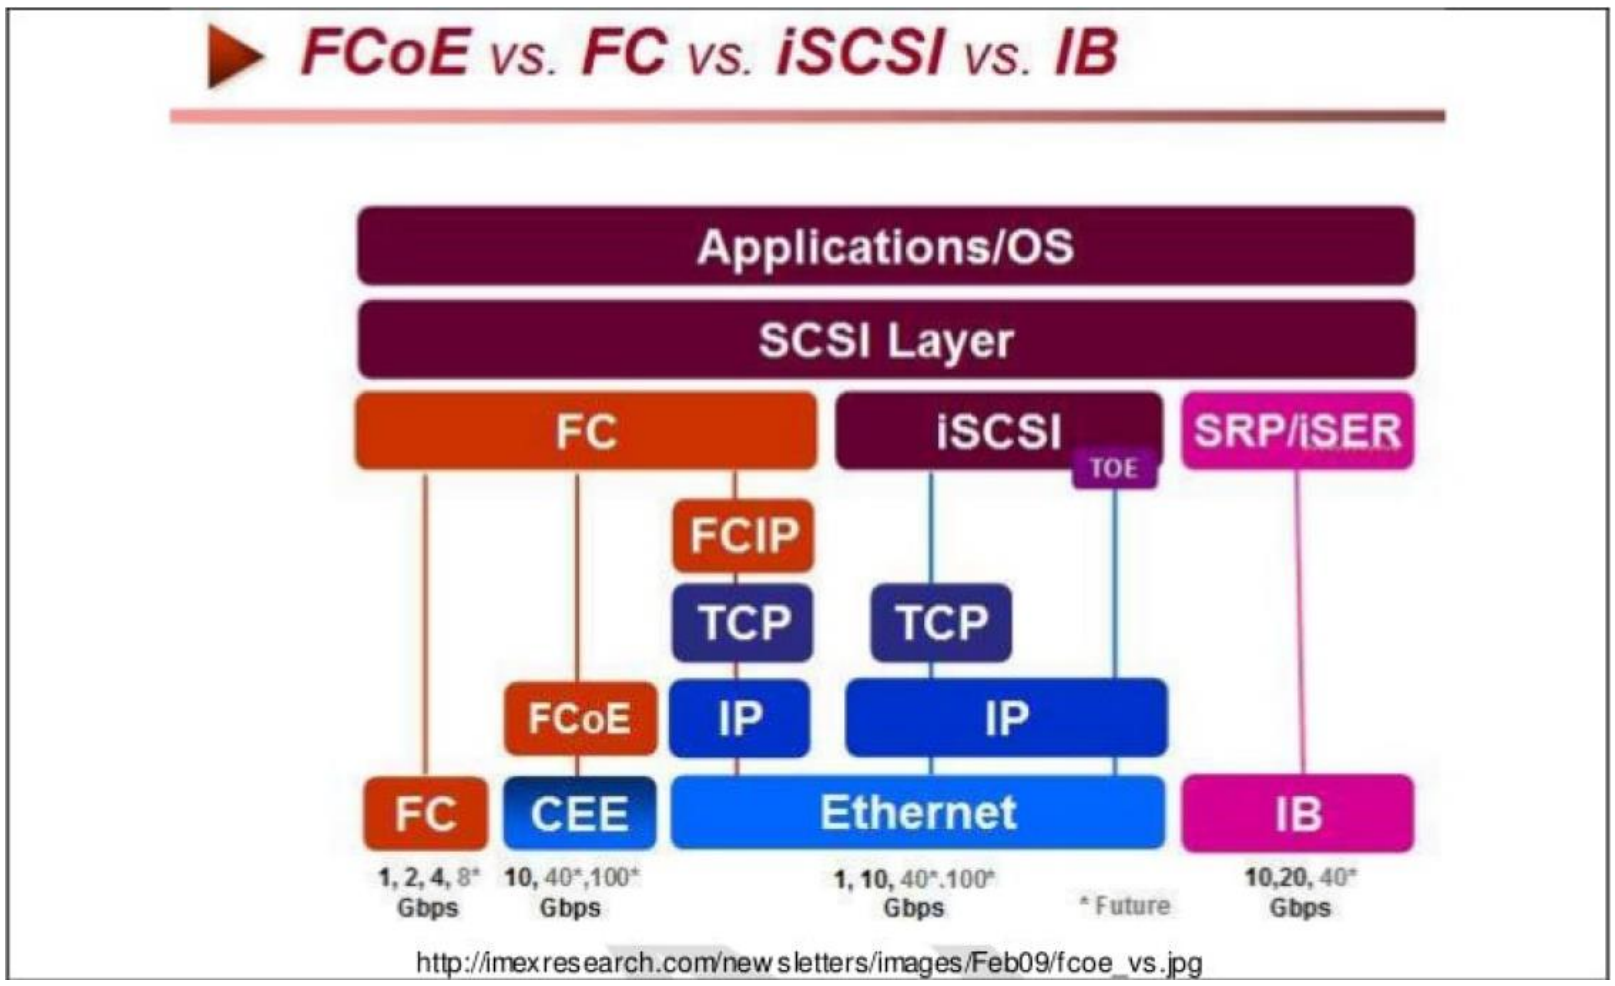

http://imexresearch.com/newsletters/images/Feb09/fcoe_vs.jpg

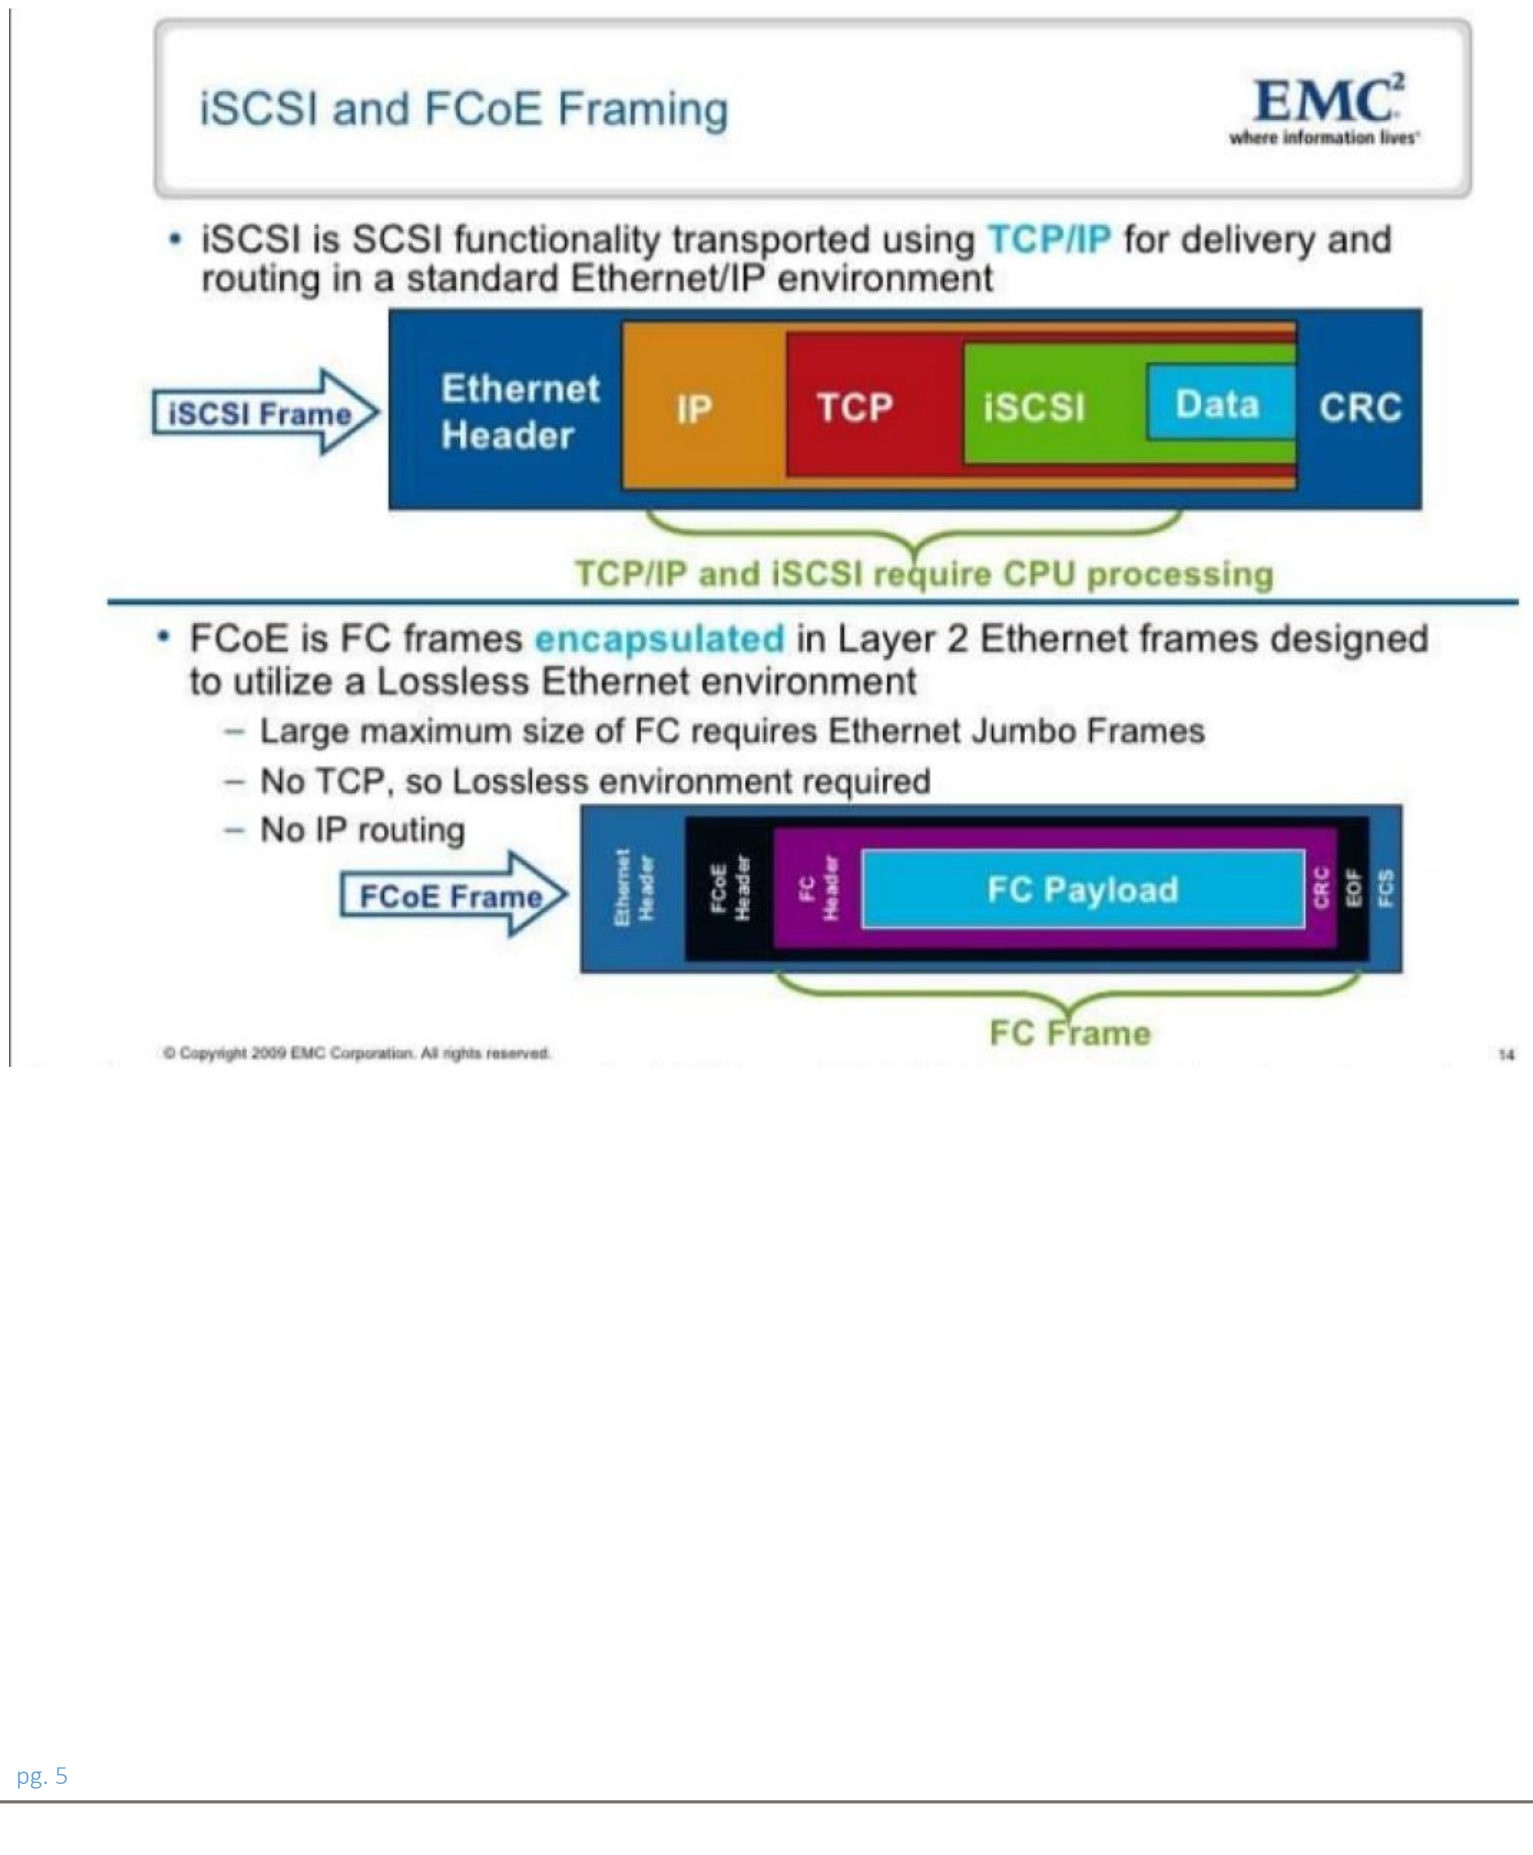

iSCSI, NFS, FC, and FCoE Basics

iSCSI means you map your storage over TCP/IP. You typically put in dedicated Ethernet network cards and a separate network switch. Each server and each storage device has its own IP address(es), and you connect by specifying an IP address where your drive lives. In Windows, each drive shows up in Computer Manager as a hard drive, and you format it. This is called block storage.

NFS means you access a file share like \\MyFileServerName\MyShareName, and you put files on it. In Windows, this is a mapped network drive. You access folders and files there, but you don't see the network mapped drive in Computer Manager as a local drive letter. You don't get exclusive access to NFS drives. You don't need a separate network cable for NFS – you just access your file shares over whatever network you want.

Fibre Channel is a lot like iSCSI, except it uses fiberoptic cables instead of Ethernet cables. It's a separate dedicated network just for storage, so you don't have to worry as much about performance contention – although you do still have to worry.

Fibre Channel Over Ethernet (FCoE) runs the FC protocol over Ethernet cables, specifically 10Gb Ethernet. This gained niche popularity because you can use just one network (10Gb Ethernet) for both regular network traffic and storage network traffic rather than having one set of switches for fiber and one set for Ethernet. https://www.brentozar.com/archive/2012/05/storage-protocol-basics-iscsi-nfs-fibre-channel-fcoe/

8. NAS vs. SANS

Network-attached storage (NAS) is a file-level computer data storage server connected to a computer network providing **data access** to a heterogeneous group of clients. NA S is specialized for serving files either by its hardware, software, or configuration. It is often manufactured as a computer appliance — a purpose-built specialized computer. NAS systems are networked appliances which contain one or more storage drives, often arranged into logical, redundant storage containers or **RAID.**

Network-attached storage removes the responsibility of file **serving from other** servers on the **network.** They typically provide access to files using network file **sharing protocols such** as NFS, SMB/CIFS, **or AFP.** From the mid-1990s, NAS devices began **gaining popularity** as a convenient method of sharing files among multiple computers. Potential benefits of dedicated network-attached storage, compared to general-purpose servers also serving files, include faster data access, easier administration, and simple configuration."

NAS allows files to be retrieved across a computer network. It includes a dedicated hardware device (often called the *head)* that connects to a local area network (usually via Ethernet). This NAS "server" authenticates clients and manages file operations in much the same manner as traditional file servers, through well-established network protocols like NFS (Network File Service) and CIFS/SMB (Common Internet File System /System Message Block).

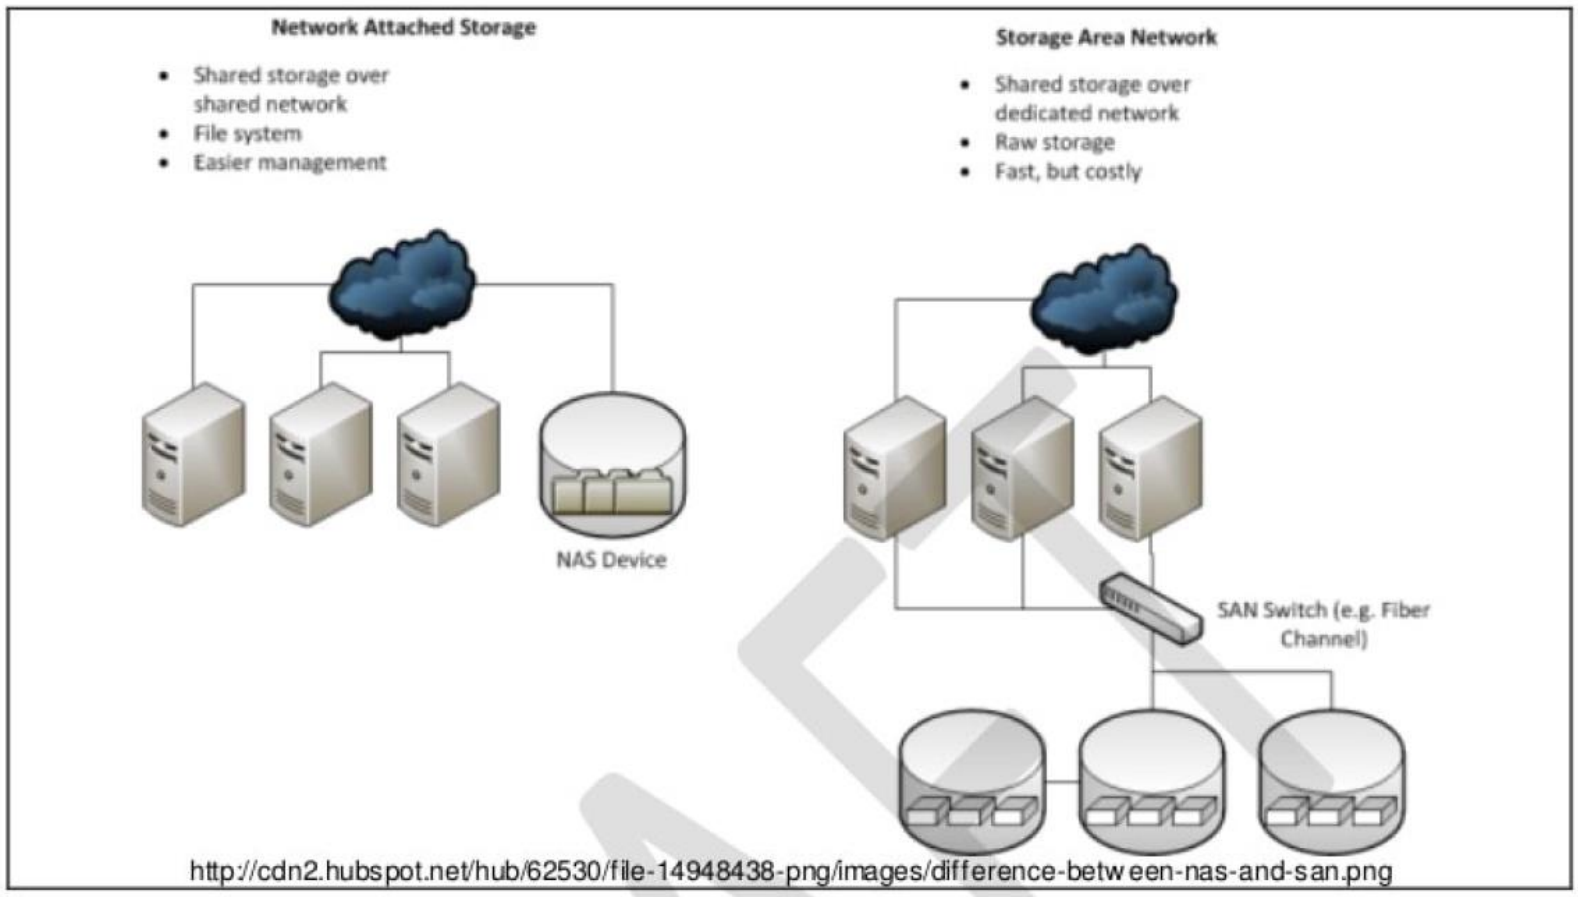

As Internet technologies like TCP/IP and Ethernet have proliferated worldwide, some SAN products are making the transition from Fibre Channel to the same IF -based approach NAS uses. Also, with the rapid improvements in disk storage technology, today's NAS devices now offer capacities and performance that once were only possible with SAN. These two industry factors have led to a partial convergence of NAS and SAN approaches to network storage.[13]

Note about FCoE: Data is transmitted over this channel in an unencrypted fashion (as is true for NAS). Therefore, this channel can be susceptible to hackers. FCoE is NOT mutable, it is operated over Ethernet. The outset of 10 GIbIpis over Ethernet made FCoE possible. FCoE uses FCP, the Fibre Channel Protocol encapsulated in an Ethernet header. FCoE abandons IF protocol and IF addresses (e.g. 192.168.1.xx), instead employing a new 802.3 Ethertype. The approach eliminates TCP altogether, replacing it with a hardware flow control scheme which guarantees packet delivery and can match performance metrics found on native Fiber channel networks. NICs are called Converged Network Adapters (CAN).

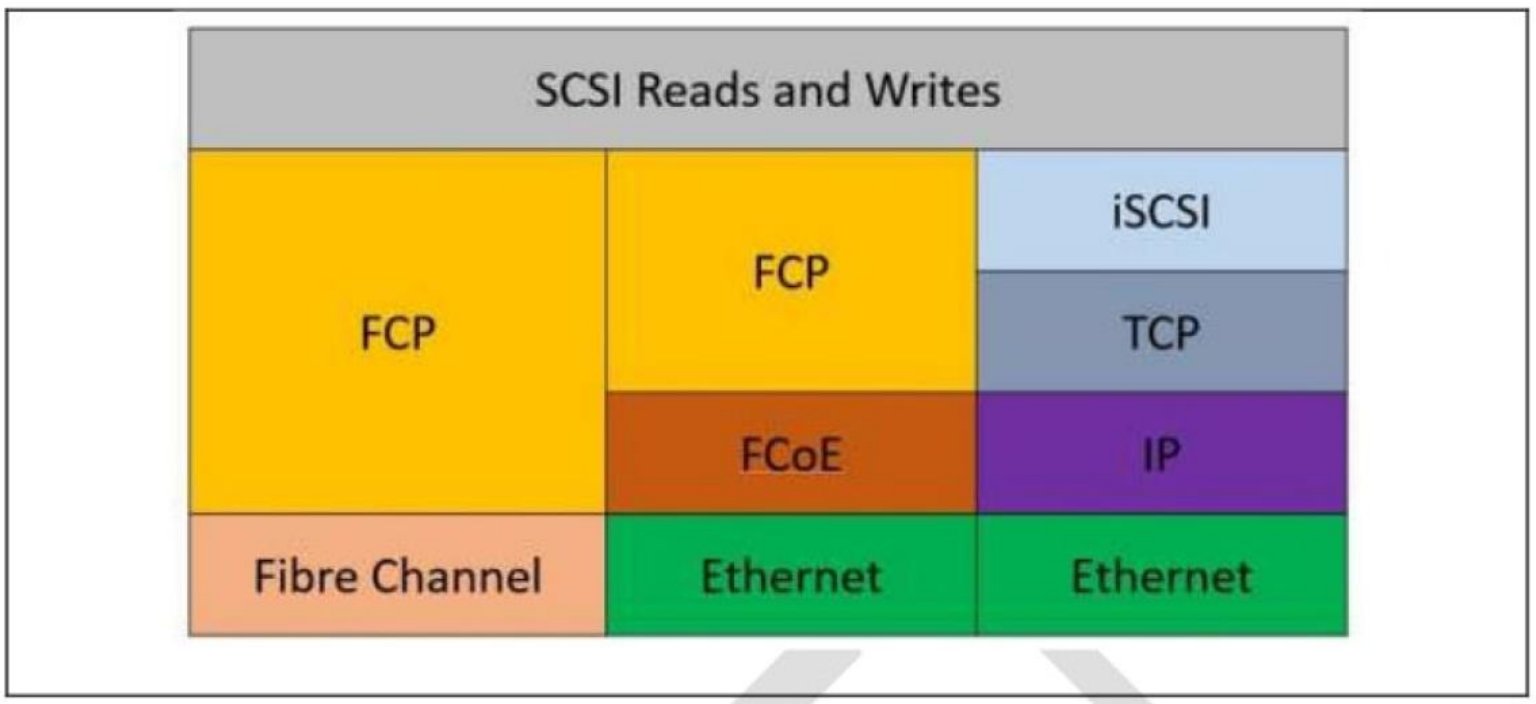
SCSI Reads and Writes
FCP
FCP
iSCSI
TCP
FCoE
IP
Fibre Channel
Ethernet
Ethernet

9. Dynamic Disk Pools vs. RAID

Dynamic Disk Pooling (DDP) dynamically distributes data, spare capacity, and protection information across a pool of disk drives. DDP improves the time and performance of traditional RAID arrays.

Because RAID cannot keep up with increasing disk capacities, DDP was created to be more versatile by providing better rebuild times. In an RAID array when a drive fails, the remaining drives are read, parity recomputed, and the result is written to the spare drive. This is done from the initial Logical block of the array to the last block in the array. This operation is time consuming because all data needs to be recomputed from the beginning of the array to the end of the array, and degrades performance, because although there are parallel reads, there is one single write to the spare drive. Thus, this single write becomes a bottleneck in the system.

In DDP, a disk pool is a set of drives that are logically grouped together in the storage subsystem, where data is distributed across all drives in the pool. The drives in each disk pool must be of the same drive type and drive media type, and they must be similar in size. Unlike RAID, there is no specific spare drive, rather, all drives have spare space that is reserved. When a drive fails, the remaining drives are read, the missing data is recomputed, and the result is written to multiple drives in their spare space. This operation is done on the pieces of data that are missing. The result is parallel reads and parallel writes, which significantly speeds up the rebuild time after a single drive failure.

Both RAID and DDP are techniques for striping data and parity information across a set of disks to provide fault tolerance, but how they operate to attain this goal is different.

10. Microsoft Group Policies

Group Policy, in part, controls what users can and cannot do on a computer system: for example, to enforce a password complexity policy that prevents users from choosing an overly simple password, to allow or prevent unidentified users from remote computers to connect to a network share, to block access to the Windows Task Manager or to restrict access to certain folders. A set of such configurations is called a Group Policy Object (GPO).

As part of Microsoft's IntelliMirrortechnologies, Group Policy aims to reduce the cost of supporting users. IntelliMirror technologies relate to the management of disconnected machines or roaming users and include *roaming* user profiles *folder redirection,* and *offline files)*

Group Policy Objects are processed in the following order (from top to bottom)

1. **Local - Any settings in the computer's local policy. Prior to Windows Vista, there was only one local group policy stored per computer. Windows Vista and later Windows versions allow individual group policies per user accounts.2**
2. **Site Any Group Policies associated with the *Active Directory site* in which the computer resides. (An Active Directory site is a logical grouping of computers, intended to facilitate management of those computers based on their physical proximity.) If multiple policies are linked to a site, they are processed in the order set by the administrator.**
3. **Domain - Any Group Policies associated with the Windows domain in which the computer resides. If multiple policies are linked to a domain, they are processed in the order set by the administrator.**
4. **Organizational Unit - Group policies assigned to the *Active Directory organizational unit (OU)* in which the computer or user are placed. (OUs are logical units that help organizing and managing a group of users, computers or**

other Active Directory objects.) If multiple policies are linked to an OU, they are processed in the order set by the administrator.

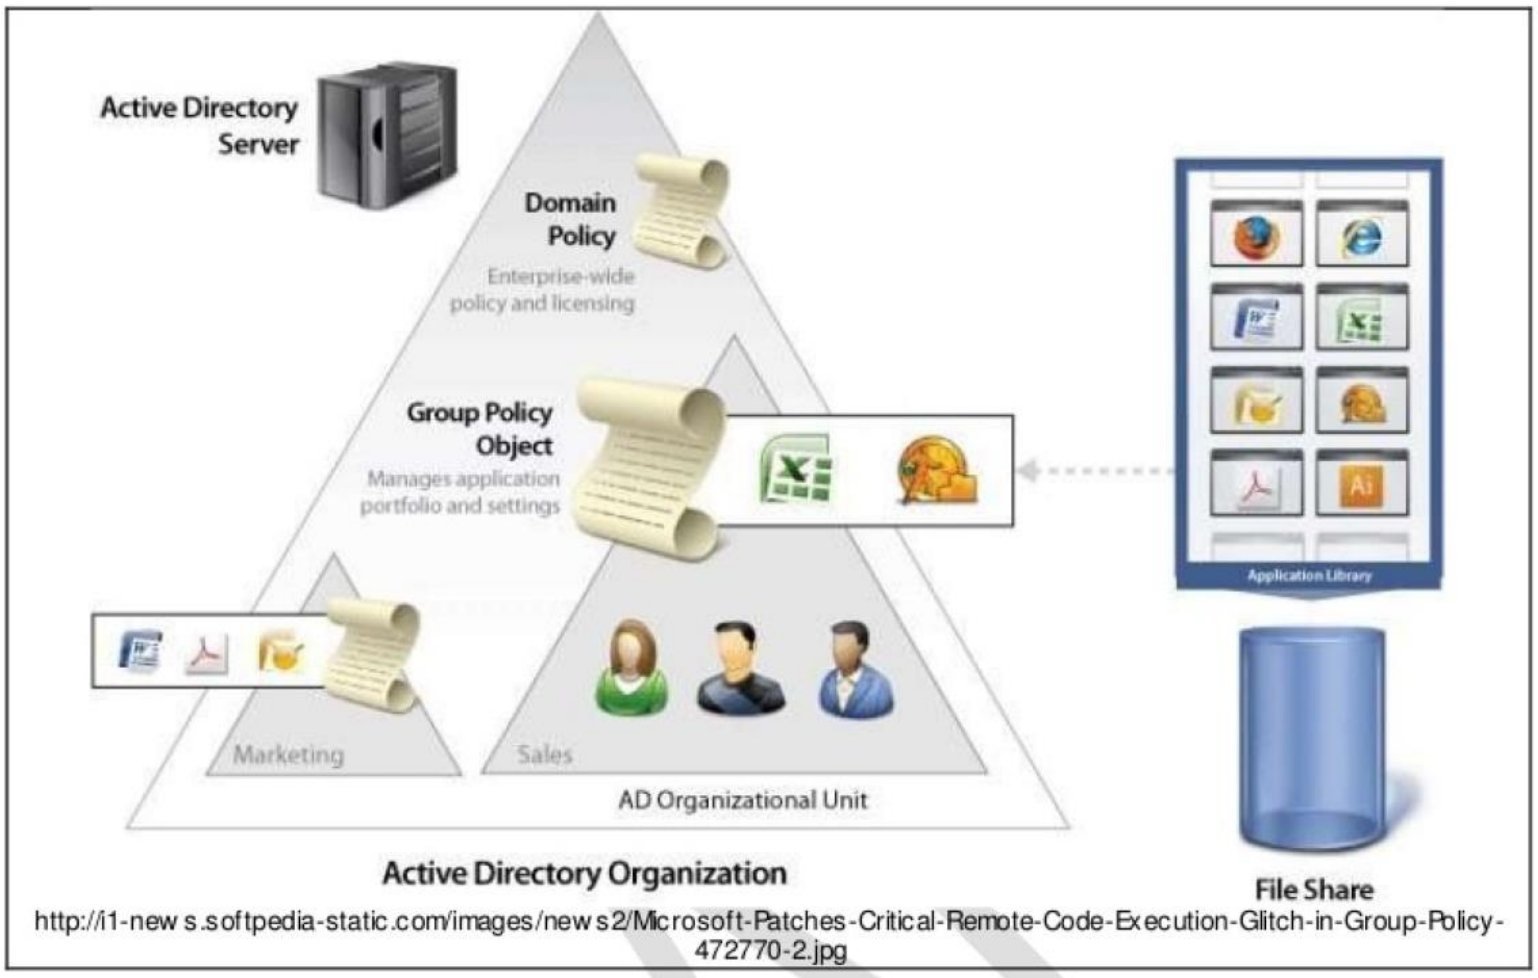

http://i1-news.softpedia-static.com/images/news2/Microsoft-Patches-Critical-Remote-Code-Execution-Glitch-in-Group-Policy-472770-2.jpg

WMI filters

You can use WMI filters to add a decision on when to apply a given group policy. This can be very useful when users or computers are located in a relatively flat structure instead of specific Oils, for example. Filters can also help when you need to apply certain **policies based on server roles, operating** system **version, network configuration,**

or other criteria. Windows evaluates these filters in the **following order of overall Group Policy Processing:**

1. **Policies in hierarchy are located.**
2. **WM/** ***Filters are checked.***
3. **Security settings** are **checked.**
4. **Finally, once everything has 'passed', a policy is applied.**

So we find all the policies that exist in the user/computer's Local, Site, Domain, and 011 hierarchy. Then we determine if the WM I filter evaluates as TRUE. Then we verify that the user/computer has Read and Apply Group permissions for **the GPO. This means that WM I filters are still less efficient than hierarchical linking, but can definitely use filters to make decisions in a non-hierarchical Active** Directory **design's.**

11. Secure Boot

Secure soot is a security standard developed by members of the PC industry to help make sure that your PC boots using only software that is trusted by the PC manufacturer.

When the PC starts, the firmware checks the signature of each piece of boot software, including firmware drivers (Option ROMs) and the operating system. If the signatures are good, the PC boots, and the firmware gives control to the operating system.

The following versions of Windows support Secure Boot: Windows 8.1, Windows Server 2012 R2, Windows RT 8.1, Windows 8, Windows Server 2012, and Windows RT[1]U.

BIOS

Basic Input/ Output System **(BIOS) and** also known as the System BIOS, ROM BIOS or PC BIOS) is a type of firmware used to perform hardware initialization during the booting process (power-on startup) on IBM PC compatible computers, and to provide runtime services for operating systems and programs.[

BIOS works by reading the first sector of the hard drive which has the next device's address to initialize or code to execute. BIOS also selects the boot device that needs to be initialized for starting the operating system. Since BIBS has been in use since the very beginning, it still works in 16-bit mode, limiting the amount of code that can be read and executed from the firmware ROM[17].

UEFI

Unified Extensible Firmware Interface (UEFI) is a specification for a software program that connects a computer's firmware to its operating system (OS). UEFI is expected to eventually replace DOS.

Like BIOS, UEFI is installed at the time of manufacturing and is the first program that runs when a computer is turned on. It checks to see what hardware components the computing device has, wakes the components up and hands them over to the operating system. The new specification addresses several limitations of DOS, including

restrictions on hard disk partition size and the amount of time BIOS takes **to** perform its tasks.

Because UEFI is programmable, original equipment manufacturer (OEM) developers can add applications and drivers, allowing UE FI to function as a lightweight operating system[18].

12. SABSA Architecture

The SABSA Architecture model **(Sherwood Applied Business Security Architecture)** may be mentioned in the same context as the HI ST (National institute of Standards and Technology) Special Publication 800-53 and the CobiT framework **(Control Objectives for Information and Related Technologies). Enterprise Architecture** (EA): the

approach to align business requirernerrts and strategy with the I.T. and security investments made.

SABSA maps business requirements to architectural requirements.

SABSA is a model and a methodology for developing risk-driven enterprise information security architectures and for delivering security infrastructure solutions that support critical business initiatives. The primary characteristic of the SABSA model is that everything must be derived from an analysis of the business requirements for security, especially those in which security has an enabling function through which new business opportunities can be developed and exploited.

The process analyzes the business requirements **at the** outset, **and creates a chain of traceabilit**y **through the strategy and concept, design, implementation, and ongoing 'manage and measure' phases of the lifecycle to ensure that the business mandate is preserved, Framework tools created from practical** experience **further support the whole methodology,**

SABSA defines 85 common attributes; with seven main categories such as: **user, management, operational, risk, technical, legal and business.**

Attribute Profile

Business Attribute	Business Driver	Business Attribute Definition	Measurement Approach	Metric	Performance Target
User Attributes					
Accessible	5	Information to which the user is entitled to gain access should be easily found and accessed by that user.	Search tree depth necessary to find the information	Soft	
Accurate	7	The information provided to users should be accurate within a range that has been preagreed upon as being applicable to the service being delivered.	Acceptance testing on key data to demonstrate compliance with design rules	Hard	
Anonymous	4	For certain specialized types of service, the anonymity of the user should be protected.	Rigorous proof of system functionality Red team review	Hard Soft	
Consistent	23, 41	The way in which log-in, navigation, and target services are presented to the user should be consistent across different times, locations, and channels of access.	Conformance with design style guides Red team review	Soft	
Current	7	Information provided to users should be current and kept up to date, within a range that has been preagreed upon as being applicable for the service being delivered.	Refresh rates at the data source and replication of source and replication of refreshed data to the destination.	Hard	

13. Risk Terms

RISK: Risk is the likelihood or probability that an event will occur that will cause a realization of a threat. Key words: occurrence , event, realization.

THREAT: Threat is the potential of a risk — in that the threat is an agent that COULD cause potential damage to the enterprise. A threat could be a condition, circumstance, environment, etc.; such as: rain storm, civil unrest, bomb threat, etc.

TYPES OF THREATS: Usually defined in three classes: natural, man-made, technological driven. Or, restated: natural, technology, human-caused.

VULNERABILITY: A vulnerability is a weakness in the enterprise (hardware, software, procedural, etc.) that MAY be exploited by a threat agent. Once exploited, the vulnerability now is a realized risk that is damaging the organization.

ASSET IDENTIFUICATION: Normally an asset inventory. This can include hardware, software, intellectual **property, etc. Essentially anything of value** that can **be damaged.**

INFORMATION CLASSIFICATION: This process attempts to create priorities as to how data is protected by creating categories (see TOP SECRET, SECRET, CONFIDENTIAL for example). Labeling information as to how it should be classified can help align needed resources to information protection.

http://acriafrica.com/images/risk.gif

14. Placement of Security Appliances Firewalls

STATELESS

Stateless fire alts watch network traffic, and restrict or block packets based on **source and destination addresses or other static values.** They are not 'aware' of traffic patterns or data flows. A stateless firewall uses simple rule-sets that do not account for the possibility that a packet might be
received by the firewall 'pretending' to be something you asked for.

STATEFUL

Stateful firewalls can watch **traffic streams from end to end.** They are aware of communication paths and can implement various IF Security (IPsec) functions such as tunnels and encryption. In technical terms, this means that stateful firewalls can tell what stage a TCP connection is in (open, open sent, synchronized, synchronization acknowledge or established), it **can tell if the MTU** has changed, whether packets have fragmented etc.