An Information Security Overview

By Adam Firestone

An Information Security Overview

By Adam Firestone

1 Introduction

Modern information security doctrine emphasizes multiple concentric protective rings creating a multilayered defensive perimeter. This concept, known as defense in depth, is based on the premise that if a single security mechanism fails, there will be a second (or third, or fourth) already deployed to defeat an attack. In many ways, it's an admission that the information security scales are weighted in the attackers' favor. There are many potential attackers, a broad attack surface[1] and an almost limitless number of attack methodologies. As a result, there is no single way to protect a system or network, and defense in depth is intended to increase the likelihood that even if one ring is penetrated, a subsequent barrier will stop the attack.[2]

Unfortunately, despite an increasing awareness of information security issues and the proliferation of network security technologies, the total number of reported breaches continues to grow at an alarming rate. One report from the Identity Theft Resource Center indicates a 29% rise in the number of reported[3] breaches between 2016 and 2017.[4] Another report, from security hardware manufacturer Gemalto, indicates that nearly two *billion* data records were stolen in the first half of 2017 alone.[5] The messages in these statistics are both stark and clear:

1. While perimeter defenses are important, they do not reliably prevent successful attacks.
2. The dominant motivation for the attacks was the theft of sensitive information.

Put bluntly, the payoff for the attacker is not the successful breach – it's the successful exfiltration of exploitable data. The breach is nothing. It's the data that's everything.

This white paper explores data, or, more accurately, information security concepts and mechanisms. While it is explicitly not intended to serve as a product guide, certain products may be cited as examples of a particular functionality.[6] It breaks with conventional information security paradigms in that it assumes that all perimeters are permeable and that only defenses that render data, individual devices or both inaccessible to unauthorized parties will provide the necessary defenses and deterrent effects to keep an organization's information secure. Additionally, this paper's primary intended audience comprises senior executives and operational decision makers, as opposed to information technologists and security practitioners.

1. The attack surface is the totality of the different points or potential vulnerabilities at which an unauthorized user can try to breach a network or system environment to enter data, extract data or disrupt operations.

2. https://www.sans.org/reading-room/whitepapers/basics/defense-in-depth-525, Page 1

3. Not all breaches are reported, despite regulatory requirements, so actual numbers are invariably higher.

4. https://www.idtheftcenter.org/Data-Breaches/itrc-and-the-2017-mid-year-data-breach-report

5. http://breachlevelindex.com/assets/Breach-Level-Index-Report-H1-2017-Gemalto.pdf

6. Nothing in this paper should be construed as particular product recommendation, and the reader is encouraged to carefully assess his or her organization's needs and to use formal, structured decision analysis techniques to identify and select the product best meeting those needs.

Technologists and practitioners may find it of use, but its focus and contents are deliberately broad and tailored for the larger community of those who use technology rather than those who implement and deploy it.

2 Information Security: Why Do You Care?

In the 21st century, information is central to the conduct of business. This fact cuts across all verticals, markets and sectors and applies regardless of whether an organization is part of industry, academia or government. In the end, it's good information security that defends and ensures an organization's competitive advantage (e.g., the Apple culture of secrecy[7]), good reputation (e.g., Target and the loss of 40 million customer credit and debit card numbers[8]) and operational continuity (it takes an average of 66 days to recover from a significant data loss[9]). Simply put, an organization cares about securing its information because unauthorized disclosure of that information can existentially impact the organization.

Information security impacts both a company AND its products. For example, a SIA member company that manufactures video surveillance equipment has both internal operational and customer information that must be protected. However, in addition to this, the company needs to protect the information and access inherent to its products. A networked camera (or microphone, or physical access control mechanism for that matter) may be integrated with back-end or cloud-based control systems. Failure to appreciate and understand both the business benefits and the information security risks (and requirements) attendant to connected products can create security issues just as grave as leaving endpoints or primary networks inadequately protected.

How does an organization prioritize which information to protect? To begin, key stakeholders might be asked to identify information, the sudden loss of which would prevent the company from operating normally, or at all. Typically, this is information that supports critical business processes such as revenue generation, accounting, logistics, customer service and regulatory compliance. Loss or theft of such information could result in lost sales and customers; financial, regulatory or criminal sanctions; and/or reputational damage.

Next, the organization ranks the information in order of importance, based on criteria such as:

- the information's impact on revenue and productivity;
- whether there are regulations or laws speaking to the degree of protection required;
- its criticality to continued operations;
- the maximum amount of time that can be allowed to pass between the onset of a disrupting event resulting in data loss and a maximum allowable threshold (i.e., the "Recovery Point Objective" or RPO); and

7. https://theoutline.com/post/1766/leaked-recording-inside-apple-s-global-war-on-leakers

8. https://www.nytimes.com/2017/05/23/business/target-security-breach-settlement.html

9. https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=SEL03130WWEN&

- the duration of time within which and a specified service level to which an information process must be restored in order to avoid unacceptable consequences (i.e., the "Recovery Time Objective" or RTO)[10].

3 An Information Security Primer

This section is intended to provide familiarity with information and data security terminology and concepts. It is emphatically not designed to make the reader an expert, but rather to provide a starting point from which individuals and organization leaders can launch useful conversations and have effective communication with their information security providers and professionals.

3.1 Identity and Access Management

Identity and access management (IAM or IdAM) are the security processes and procedures that enable the right people to access the right resources at the right times for the right reasons. IdAM addresses the fundamental need to ensure both appropriate and assured access to proprietary resources across increasingly complex and diverse environments, and, critically, to meet rigorous legislative, regulatory and compliance requirements. IdAM is a reflection of the organization's business needs and requirements and requires both business skills and technical expertise.[11]

3.1.1 Authentication

Authentication is the mechanism by which an information system securely identifies its users, answering the following questions:

- Who is the user?
- Is the user really who he/she claims to be?[12]

Traditionally, a user is authenticated in of three ways[13], based on what are known as the factors of authentication:

- Knowledge factor: Something the user knows (e.g., a password, partial password, pass phrase, personal identification number (PIN), challenge response or security question).
- Possession factor: Something the user has (e.g., wrist band, ID card, security token, cell phone with built-in hardware token, software token or cell phone holding a software token).
- Inherence factors: Something the user is or does (e.g., fingerprint, retinal pattern, DNA sequence, signature, face, voice, unique bio-electric signals or other biometric identifier).

10. RPO designates the variable amount of data that will be lost or will have to be re-entered during network downtime. RTO designates the amount of "real time" that can pass before the disruption begins to seriously and unacceptably impede the flow of normal business operations.

11. https://www.gartner.com/it-glossary/identity-and-access-management-iam/

12. https://stackoverflow.com/questions/6556522/authentication-versus-authorization

13. There are solutions that offer as many as five factors of authentication, but typically, multi-factor authentication is limited to the three factors mentioned here.

Recently, additional authentication factors have been made available in some products. These include:

- Location factor: Determining that the device from which the user is attempting to authenticate is within a certain defined area (i.e., geofencing), or, in some cases, within a predefined distance of an external locating device such as a proximity beacon.
- Machine inherence factor: Something physically unique and permanent about the device from which the user is attempting to authenticate, such as a burned in processor board or processor serial numbers.

Research has determined that for a positive authentication, elements from at least two, and preferably three or more, factors should be verified.[14]

3.1.2 Authorization

Authorization is the mechanism by which a system determines what level of access a particular authenticated user should have to resources controlled by the system. For example, a database system might be designed to provide some people with the ability to retrieve information from a database but not the ability to modify that information, while concurrently giving other individuals the ability to change information. Authorization systems provide answers to the questions:

- Is user X authorized to access resource R?
- Is user X authorized to perform operation P?
- Is user X authorized to perform operation P on resource R?

Authentication and authorization are related mechanisms; authorization systems depend on secure authentication systems to ensure that users are who they claim to be, and thus prevent unauthorized users from gaining access to secured resources.[15]

3.2 Confidentiality

The ISO 27001 standard describes confidentiality as "the property, that information is not made available or disclosed to unauthorized individuals, entities or processes." Confidentiality guarantees are usually enforced by a number of mechanisms including authentication and authorization (as mentioned above), encryption and the related notion of key management and data life cycle management.

3.2.1 Encryption

The word "encryption" is derived from the Greek word kryptos, meaning "hidden." In keeping with that, the goal of cryptography[16] is to hide a message's meaning, not its presence. An encrypted message is one that has been scrambled according to a predetermined protocol

14. https://en.wikipedia.org/wiki/Authentication

15. https://web.archive.org/web/20121014105355/http://www.duke.edu/~rob/kerberos/authvauth.html

16. Cryptography is the practice and study of techniques for secure communication in the presence of third parties called adversaries or eavesdroppers.

that has been agreed upon by the sender and recipient. With encryption, if an eavesdropper intercepts the message and lacks knowledge of the scrambling protocol, it remains unreadable. Modern cryptography is based on the notion that the cryptographic algorithm (which does the scrambling) can and should be made public, whereas the key, which determines how the scrambling is to be applied, is kept secret. This is because it's assumed, as American information theorist Claude Shannon put it, "The enemy knows the system," and by making the algorithm widely available, a broad vulnerability analysis can be conducted.

This perspective was definitively stated by Auguste Kerckhoffs von Nieuwenhof in 1883 in what came to be known as Kerckhoffs' Principle:

> *The security of a cryptosystem must not depend on keeping secret the crypto-algorithm. The security depends only on keeping secret the key.*[17]

Cryptographic algorithms can be divided into three types:

- symmetric cipher algorithms, in which the sender and recipient share a common key that is required to both encrypt and decrypt the message;
- asymmetric cipher algorithms, where each party has both a public key that is used to encrypt the message, and a private key that is used to decrypt the message; and
- cryptographic hash algorithms, which are one-way functions that produce unique output of a fixed length.

3.2.2 Key Management

The most effective cryptosystems are known as One-Time-Pads (OTP) and rely on completely random keys of great length that are used for only one message. OTPs are so effective that they are said to be "information theoretically secure;" that is, unbreakable. Unfortunately, the difficulty in generating and managing huge numbers of long, random keys is significant. That difficulty, however, is dwarfed by the difficulty of securely distributing those keys across hostile or insecure environments.

Key management deals with the generation, exchange, storage, use, crypto-shredding (destruction) and replacement of keys. It includes cryptographic protocol design, key servers, user procedures and other relevant protocols. It applies to keys at the user level, either between users or systems. It is critical to the security of a cryptosystem. It is the more challenging side of cryptography in a sense that it involves implementation issues such as system policy, user training, organizational and departmental interactions and coordination between all of these elements, in contrast to pure mathematical practices that can be automated.

3.2.3 Data Life Cycle Management

While information may remain useful to an organization for some finite period, it may remain sensitive for a significantly longer period. As a result, it's important to securely manage information throughout its life cycle. An example data, or information, life cycle management process is described below:

17. Singh, Simon. *The Code Book*. New York: Anchor Books, 1999. Amazon Kindle Digital Edition, Location 347

The data life cycle begins with data capture, which is the act of creating data values that do not yet exist and have never existed within the enterprise. Data can be captured by acquiring information that has been produced by an organization outside the enterprise, by having human actors or devices generate new data, or capturing signal or sensor information.

Once data has been captured, it is processed in ways that include movement, integration, cleansing, enrichment, changed data capture and extract-transform-load. It may then be synthesized or analyzed to derive other information, or directly used in support of the enterprise. Data values may also be sent outside of the enterprise, this is called data publication.

After many rounds of usage and publication, the information's end of life nears. At this point the information is copied to an environment where it is stored in case it is needed again in an active production environment, and is then removed from active production environments. When it is certain that the information is no longer of value, it is purged.[18]

3.3 Integrity

Information integrity is the maintenance and assurance of the accuracy and consistency of information over its entire life cycle. Information is said to maintain integrity if it is recorded exactly as intended and upon later retrieval, is the same as it was when it was originally recorded. Integrity means that unintentional or unauthorized changes to information are prevented.

For files or other blocks of digital information, integrity can be provided by mechanisms such as hashing, digital signatures and checksums. Database integrity is typically enforced through a series of integrity constraints or rules. Three types of integrity constraints are an inherent part of the relational data model: entity integrity, referential integrity and domain integrity.

Historically, integrity has been enforced at a single point, such as a database or a file store. If that point was breached or corrupted, the integrity guarantees were lost. Blockchain, a distributed ledger technology most commonly associated with the Bitcoin cryptocurrency, offers great promise with respect to data integrity. Blockchain maintains a complete record of every transaction impacting a piece of information stored within it, and distributes that record to every participant in the overall scheme. As a result, data stored in a blockchain mechanism does not have a single home that can be corrupted. Additionally, blockchain relies on cryptographic hashing to create an immutable signature for each transaction (or set of transactions) – the "block," upon which the signature of every following block depends, so changes are readily apparent. Finally, blockchain introduces a distributed consensus model where no change is permitted to propagate until it is approved by a majority of participants.

18. https://www.bloomberg.com/professional/blog/7-phases-of-a-data-life-cycle/

3.4 Authenticity

With respect to information security, authenticity refers to the notion that the party being communicated with is really who they claim to be. Many techniques are used to ensure this, such as pre-shared keys, public/private key algorithms, and, recently quantum cryptography.[19]

3.5 Availability

For information to be useful, it must be available when required. 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 function properly. Systems requiring high availability assurance must design for unexpected disruptions such as power outages, hardware failures, anticipated downtime associated with system upgrades and malicious activity such as denial-of-service attacks.

Ensuring availability takes place at both design time and runtime. Redundancy, monitoring and control mechanisms are planned and engineered at design time. System health and status checks are continual processes during runtime. Availability is ensured through detection of fault conditions, avoidance of storage problems, review of file system structures, monitoring file system usage and defining disk and tape resources by sizing storage components. Activities specific to availability assurance include:

- Backup/Restore: Backup procedures ensure that critical information is securely saved on a consistent basis. Once saved, data must be readily available through restoration procedures that enable users to request retrieval of any critical data. Copies of backed-up information should be securely stored off-site as well as on-site. Options for backup and archiving have increased significantly with the proliferation of cloud-based solutions.
- Archiving stores critical data on a long-term basis in an efficient manner. While backup data is often readily and locally available, archived data may be saved on tape and stored off-site.
- Storage management ensures that available storage media is available as needed and used efficiently. Continual review of stored data and resources balances availability and cost constraints.
- Database management includes managing the database environment from initial product selection to daily operations monitoring. As databases become more distributed, tools supporting centralized monitoring and sound data architectures are required. These tools detect data errors, conflicts and potential resource issues.[20]

3.6 Non-Repudiation

Non-repudiation in information security refers to guarantees that the author of a statement will not be able to successfully challenge his or her authorship of the statement or validity of an associated contract execution. It can be achieved through a service that provides proof of the integrity and origin of the information and high levels of assurance with respect to authenticity.

19. https://security.stackexchange.com/questions/148173/authenticity-confidentiality-integrity-general-questions

20. http://www.business-esolutions.com/drav.htm

Data integrity can be proven though the use of a cryptographic hash function (e.g., SHA-2), demonstrating that the likelihood of data being undetectably changed is extremely low. Additionally, authenticated encryption modes of operation, such as Galois Counter Mode, can be used to provide non-repudiation guarantees.[21]

Typically, information's origin is established using digital certificates, as discussed above. However, certificates and newer techniques, such as blockchain, that hold great promise with respect to information integrity remain reliant on current asymmetric cryptography technologies. According to the National Institute of Standards and Technology (NIST), these technologies will become vulnerable to quantum computing attacks within 20 years. Fortunately, great progress is being made in post-quantum cryptography. For more information see NIST Internal Report 8105.[22]

4 Information Security Imperatives and Definitions

Once the initial set of critical information is defined, the discovery process continues by examining where the company's critical information resides, categorizing it accordingly and determining which information security imperatives apply.

4.1 Information Types

An organization's critical information can be divided into four broad categories:

- Internal information comprises an organization's operating information. This can include employee information that may be personally identifiable information (PII[23]), or with respect to medical benefits, protected health information (PHI[24]). It can also include an organization's intellectual property, operating policies and procedures, plans, business goals, competitive intelligence or any other information that the organization uses to remain viable and competitive. In the case of some organizations, this may include public trust or national security information.
- Customer information includes any data that the organization maintains about its customers. This can range from order histories to market assessments to shipping and home addresses and financial information (e.g., credit card and bank account numbers). Depending on the type of organization, customer PII and PHI may be held as well.
- Partner, service and tool information includes any of the organization's critical data that is held by a third party. This is especially relevant as more organizations adopt cloud-based services. For example, the financial management and human capital management service Workday has over 1,800 customers, many of whom are Fortune

21. https://en.wikipedia.org/wiki/Non-repudiation

22. https://csrc.nist.gov/publications/detail/nistir/8105/final

23. PII, as used in information security and privacy laws, is information that can be used on its own or with other information to identify, contact, or locate a single person, or to identify an individual in context.

24. PHI under U.S. law is any information about health status, provision of health care or payment for health care that is created or collected, and can be linked to a specific individual.

500 companies.[25] This means that an organization that uses Workday's services has deposited critical information outside its direct control and relies on Workday's information security apparatus to maintain the confidentiality of its PII and PHI.

- Information required to be safeguarded by law for national security reasons, including that deemed for official use only, controlled unclassified information or classified information.

Additionally, there is a fifth category – device information – that is of interest to SIA members. Device information may include proprietary data or trade secrets built into firmware or embedded software, or authentication credentials. While specifics concerning device design and architecture are beyond the scope of this white paper, the endpoint and network protection principles described below remain applicable.

Regardless of how the information is categorized, the unifying theme is that unauthorized disclosure can be expected to have serious negative consequences for the organization, up to and including permanent closure. In some cases, such consequences are driven by the market's reaction to the disclosure, and in others they are driven by regulatory consequences resulting from non-compliance.

4.2 Elements of Information Security

Information security is often described as comprising three discrete capabilities or attributes: Confidentiality, integrity and availability (CIA). Collectively, these are referred to as the CIA triad, and each capability represents a key information security goal:

- **Confidentiality** is maintained when unauthorized access to information is prevented. Confidentiality mechanisms range from physical access controls (e.g., guards, gates, guns) to system and network authentication and authorization tools, to data encryption mechanisms. Confidentiality includes means for protecting privacy and proprietary information.
- **Integrity** is maintained when unauthorized or improper modification or destruction is prevented. Integrity assures users of information accuracy and can include assurances of authenticity[26] and nonrepudiation.[27] Integrity mechanisms include tools such as cryptographic hash algorithms, message authentication codes, digital signatures and digital certificates.
- **Availability** is maintained when information is accessible, in a timely and reliable manner, to authorized users. With respect to information security, availability means defending against unauthorized data destruction (through means such as redundant storage or backups) or defeating denial of service attacks by using tools such as IP throttling and flexible scaling.

25. https://diginomica.com/2017/10/19/workday-rising-and-workday-then-now/

26. Authenticity refers to the validity of claims as to information's origin or authorship.

27. Non-repudiation refers to a state in which the author of a communication or message cannot later validly claim not to have originated the communication or message.

There's no simple answer as to which is most important. Different types of organizational data require different combinations of resource allocations to achieve the proper CIA mix. For example, when considering what type of information security resources to allocate to information that is operationally critical but that can be easily and rapidly recreated, it may make sense to emphasize confidentiality over integrity or availability.

5 Information Security States

Organizational information exists in one of three states. These are commonly referred to as data-at-rest (DAR), data-in-transit (DIT) and data-in-use (DIU). Ideally, an organization's information is protected in each of these states. Information security in each state requires the use of specific tools, protocols and processes.

5.1 DAR

DAR refers to inactive data that is physically stored in any digital form. Storage mechanisms may range from databases to data warehouses to applications such as spreadsheets and word processing documents, network archives, physical tapes, off-site backups and mobile devices. There is no specific temporal definition for DAR. It can refer to static, unchanging information such as historical archives, information that is subject to occasional change, such as reference tables, or information that is regularly used, such as that stored in active databases, but that is not currently in use. DAR is any data that is stored on media and not being moved across a network or used by an active application.

5.2 DIT

DIT (or data in motion) refers to information actively moving from one location to another, such as across the internet or over a local area network. Examples of DIT include information exchanged between a browser and a server, email, instant messaging or any information exchanged in an online manner. DIT can comprise any information of any format and is subject to interception or eavesdropping.

5.3 DIU

DIU refers to information actively being used for computation that is temporarily stored in a volatile digital state, such as in random access memory (RAM), central processing unit (CPU) caches or CPU registers. It can contain sensitive data including digital certificates, encryption keys, intellectual property (e.g., software algorithms, design data) and PII. DIU compromised can enable the compromise of both DIT and DAR. Unfortunately, while research is ongoing, there are few tools to protect DIU.[28]

28. Full memory encryption, CPU-based key storage, enclaves and homomorphic encryption are current solutions, though most of these are not available for general purpose use yet.

6 Information Security by Location

In addition to discrete states, an organization's information exists in a finite set of location types. Information protection mechanisms are tailored for specific location types, including endpoints, local network storage and cloud-based storage.

An **endpoint** is a device connected to the local area network (LAN), wide area network (WAN) or the internet that can receive and send communications back and forth across the network. It was originally meant to indicate networking equipment such as hubs, routers or switches or a host such as a workstation or a server. Today, the term has expanded to represent any device on the periphery of the network, whether inside or outside a firewall. Such devices include laptops, tablets and mobile phones, or any other means by which people connect to the central network. Endpoints are proliferating due to the explosion in the number of mobile devices in the workforce. It's estimated that more than 30 percent of organizational data now exists outside the organizational firewall on mobile endpoint devices.[29] Endpoints also include any smart or networked devices, such as cameras, microphones or access control mechanisms. Such endpoints, or any other non-human actor that has a digital identity, are often referred to as "Non-Person Entities" (NPE). Note that some definitions specify that NPEs are the digital certificates used to identify non-human actors to a public key infrastructure.[30]

Local network storage may refer to a file server, network attached storage (NAS) or a storage array network (SAN). A file server provides shared storage for computer files (such as documents, spreadsheets, presentations, images or video) that can be accessed by workstations that are able to reach the computer providing the access through a network. File servers and NAS devices offer similar capabilities, but file servers usually have more powerful hardware and greater functionality than a NAS device.[31] A SAN provides access to consolidated, block-level data storage. (Block storage is analogous to space on a hard drive.) However, all three generally refer to storage capabilities that are local, or at least proprietary, with respect to a single organization across either a LAN or WAN.

Cloud storage refers to storage mechanisms that use securely shared (multitenanted), remotely sited and externally owned and managed hardware to provide storage capabilities on demand as a service. There are three types of cloud data storage, and they mirror the local network storage types:

- Object Storage: Object storage offers great scalability and metadata characteristics that can be used to flexibly describe the information being stored. Object storage solutions like Amazon Web Services' Simple Storage Service (S3) are used to build modern

29. https://www.druva.com/blog/simple-definition-endpoint/

30. NPE is the credential granted to an authorized device or software application as part of PKI functionality. NPE certificates are issued to devices by a request process that ensures ownership and use of those devices in accordance with guidance and directives that control an organization's IT platforms and use of platforms. The Department of Defense has an NPE initiative to "remove anonymity for devices on DoD networks" starting with "workstations, domain controllers and Web servers." See: https://www.jerichosystems.com/technology/glossaryterms/non_person_entity.html

31. A file server and NAS device might both allow for access control with respect to files and folders, but a file server usually has more security configuration options and more granular access controls than a NAS device.

applications from scratch that require scale and flexibility, and can also be used to import existing data stores for analytics, backup or archive.

- File Storage: This is provided by a shared NAS resource and is intended to support use cases like large content repositories, development environments, media stores or user home directories.
- Block Storage: Cloud block storage is intended to support enterprise applications like databases that require dedicated, low latency storage and is directly analogous to a SAN.

6.1 Endpoint Security

Endpoints manage both DAR and DIT, and, due to the many (and often unpredictable) tasks they are required to perform, present one of the most significant data security environments. Fortunately, there are a wide array of commercial-off-the-shelf (COTS) products available supporting multifaceted requirements of endpoint security. It's worth noting that many products are versatile and cover many endpoint security use cases. The following sections provide an overview of those use cases:

6.1.1 DAR

6.1.1.1 Full Disk Encryption (FDE)

FDE is technology that secures an endpoint computing device by encrypting all the data at rest on its non-volatile storage. (Non-volatile storage may be a hard disk drive, a solid state drive or, for mobile devices, either the onboard flash memory or a removable flash memory device such as a micro-SD card.) This is *anything* stored on the device, including end-user files, application settings and application and operating system (OS) executables.

FDE is intended to protect information in the event that a device is lost, stolen or otherwise physically accessed without authorization. Any organization of any size – to include individual consumers and sole proprietorships – with sensitive DAR to protect will benefit from using full disk encryption software. As long as the device is not in a booted state, FDE provides significant security benefits. As FDE doesn't encrypt data in use, it is often used alongside other storage-encryption types, such as virtual disk encryption, volume encryption and file encryption.[32]

Typically, FDE technology uses a mechanism called on-the-fly-encryption (OTFE). Also known as real-time encryption or transparent encryption, OTFE automatically encrypts data as it is read from or written to non-volatile storage. Put another way, all DAR on the device is always encrypted, regardless of whether the device is powered down, powered but idle or in use. As an authorized user uses the device, application and operating system data requests are shunted through an OTFE engine that is loaded on startup and resides in the device's volatile memory.

The OTFE engine and the encryption key(s) are protected through a mechanism that requires the user to authenticate prior to accessing any information. The OTFE engine stores the relevant encryption keys, also in volatile memory. When information is read from the non-volatile storage

32. http://searchsecurity.techtarget.com/feature/The-top-full-disk-encryption-products-on-the-market-today

in its encrypted state, it is decrypted in volatile memory and then delivered to the application that requested it. Similarly, when an application or operating system writes information to non-volatile memory, it is routed through the OTFE engine, where it is encrypted prior to storage on the disk. When the device is shut down, and typically when a user logs out or, in some cases, locks the device, the encryption keys in volatile memory are destroyed and can only be accessed by a validated user authentication. Thus, the device's information remains secure from unauthorized access.

With OTFE, information is accessible immediately after the encryption key is provided, and the entire non-volatile storage volume is typically mounted as if it were a physical drive, making the secured information as accessible as though it was unencrypted, and typically with very low additional latency. No data stored on an encrypted volume can be read (decrypted) without a valid authentication, and the entire file system within the volume is encrypted (including file names, folder names, file contents, and other meta-data). OTFE usually requires the use of device drivers to enable the encryption process. Administrative access is normally required to install OTFE drivers, but the encrypted devices/volumes can typically be used by normal users.[33]

OTFE products aren't difficult to find. They are available as native operating system components (e.g., Bitlocker, which comes with Microsoft Windows 10 Professional, Enterprise and Education[34]) and as third-party products in both open-source and proprietary formats.
When considering FDE solutions, both individuals and enterprises are urged to develop an understanding of their unique environmental requirements and to compare them to the features offered by different COTS products. FDE feature sets include:

- **Central management:** Whether the product supports enterprise management features such as key recovery, central deployment and the use of hardware security modules or enterprise key managers to support networked pre-boot authentication.
- **Hidden containers:** Whether hidden containers (an encrypted container (A) within another encrypted container (B) so the existence of container A cannot be established) can be created for deniable encryption.
- **Pre-boot authentication:** Whether authentication can be required before booting the computer, thus allowing boot disk encryption.
- **Single sign-on:** Whether credentials provided during pre-boot authentication will automatically log the user in to the host operating system, thus preventing password fatigue and reducing the need to remember multiple passwords.
- **Custom authentication:** Whether custom authentication mechanisms can be implemented with third-party applications.
- **Multiple keys:** Whether an encrypted disk can have more than one active key.
- **Passphrase strengthening:** Whether key stretching is used with plain text passwords to frustrate dictionary attacks, usually using the password-based key derivation function 2.
- **Hardware acceleration:** Whether dedicated cryptographic accelerator expansion cards can be used.

33. https://ipfs.io/ipfs/QmXoypizjW3WknFiJnKLwHCnL72vedxjQkDDP1mXWo6uco/wiki/On-the-fly_encryption.html

34. https://docs.microsoft.com/en-us/windows/device-security/bitlocker/bitlocker-overview

- **Trusted platform module (TPM):** Whether the implementation can use a TPM cryptoprocessor.
- **Filesystems:** What filesystems are supported.
- **Multifactor authentication:** Whether optional security tokens (hardware security modules, such as tokens and smart cards) are supported.[35]

6.1.1.2 Partition Encryption

Disk partitions are logically separate areas on a hard disk drive or other computer storage medium. On a Windows computer, this is often expressed in terms of different drive letters (e.g., C:, D:, K:). The partitions are independently managed by the operating system. Partitioning is usually done before any files or directories have been created and divides a drive's total available storage into different pieces (the partitions). Once a partition is created, it can then be formatted so that it can be used to store information on a computer.[36]

Partition encryption software usually works best on basic disks. (Basic disks are the storage types most often used with Windows. The term basic disk refers to a disk that contains partitions, such as primary partitions and logical drives, and these in turn are usually formatted with a file system to become a volume for file storage.[37]) Partition encryption allows for greater flexibility than FDE as the user is free to open (i.e., authenticate to) different encrypted partitions independently.

Partition encryption products also generally use OTFE. Partition encryption is generally available as a feature of a disk management or encryption solution, not as a standalone capability.

6.1.1.3 Volume Encryption

A volume (also called a "logical drive") is a single storage area with a single file system, usually (though not necessarily) resident on a single partition of a hard disk. A volume can be different from a physical disk drive, but it can still be accessed with an operating system's logical interface. However, a volume differs from a partition. For example, an operating system can recognize a partition without recognizing any volume associated with it. A volume also can be contained within a single file, and this scheme is commonly employed for CD/DVD images and volume encryption mechanisms.[38] Since the advent of Windows NT, Microsoft Windows has enabled users to create multi-partition volumes by combining several partitions (potentially stored on different physical hard drives) into a large single "partition" called a volume.

Volume encryption treats the volume as a single portion of data. A volume is always in one of two states: If the user has not properly authenticated (i.e., provided an encryption key), the whole volume is locked/encrypted. If the user successfully authenticates and opens the volume, everything stored in the volume, regardless of physical location, becomes accessible.

35. https://en.wikipedia.org/wiki/Comparison_of_disk_encryption_software

36. https://en.wikipedia.org/wiki/Disk_partitioning

37. https://msdn.microsoft.com/en-us/library/windows/desktop/aa363785(v=vs.85).aspx

38. https://en.wikipedia.org/wiki/Volume_(computing)

Generally, it is more intuitive to manage a volume than to work separately with every physical drive.[39]

Volume encryption is generally available as a feature of a disk management or encryption solution, not as a standalone capability.

6.1.1.4 Virtual Drive Encryption

Many endpoint information security products intended for desktop and notebook computers running standard operating systems such as Windows, Mac OS X or Linux use dedicated virtual hard drives. Typically, these virtual drives use a file format such as Microsoft's VHD that enables a single file to represent a virtual hard disk drive (HDD). Such files can contain what is found on a physical HDD, such as disk partitions, file systems, files or folders. They are addressable (on Windows systems) with drive letters and can be managed as any other drive.

Encrypted virtual drives usually operate as part of an OTFE scheme under which any data written to them is always encrypted and remains so until decrypted in volatile memory. Virtual drive encryption can be found in both single capability tools and as part of a full solution disk management and encryption solution.

6.1.2 DIT

6.1.2.1 Online Data Exchange

Endpoints are frequently used as mechanisms to exchange information over the internet. Such exchanges include accessing websites, using cloud services and any other communication that employs a web protocol. Most online data exchanges employ a browser as the user's front end, or client. The browser initiates a connection to the website or desired service. It's here that the DIT security for online activity is created. Modern browsers such as Google's Chrome, Mozilla's Firefox, Otello's Opera, Vivaldi Technologies' eponymous Vivaldi and Apple's Safari support secure connection mechanisms such as transport layer security (TLS). Users can determine whether a browser connection is secure by looking at the URL address bar. Secure connections will use the https prefix instead of the insecure http prefix, and often a secure address is shown with an icon such as a closed padlock or other indicators such as green text. Connections that do not show these indicators are to be avoided as the user has no assurance that they have connected to the right site/server, or that the site/server

TLS Handshake Protocol Overview

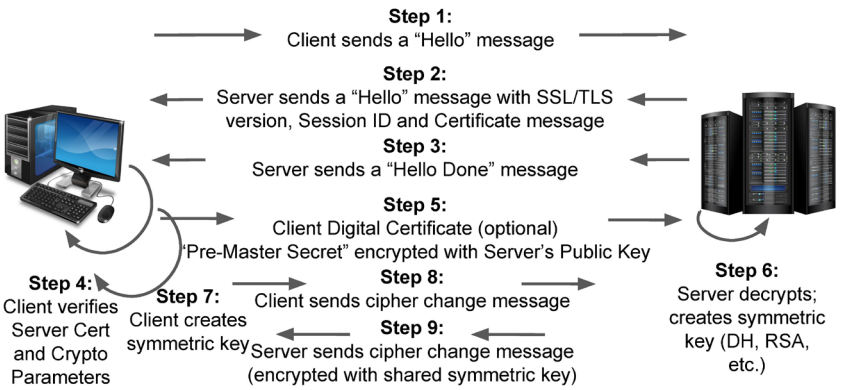

39. https://www.jetico.com/file-downloads/web_help/bcve3/html/01_introduction/02_what_is_ve.htm

to which they have connected is legitimate. Critically, information sent across an insecure connection is sent unencrypted, in the clear, and can be intercepted and read by anyone as it travels across the internet.

TLS' security[40] is the result of a complex, but rapid, interaction between the browser and the server. To initiate a secure connection, the user's browser sends a message requesting a secure session and indicating the cryptographic parameters it can support. The server responds with a message agreeing to use a certain combination of the browser's supported parameters. Included in this message is the server's digital certificate, which provides assurances as to the server's identity and authenticity. The browser verifies the legitimacy of the server's certificate and then the two generate a shared, but secret encryption key that will be used to secure information sent between the two for the duration of the session.

6.1.2.2 Wireless Connections - WiFi

Endpoints frequently access the internet over WiFi connections. Unless the connection is secured, it's trivially easy for an attacker to literally read the communications between a wireless access point or router and an endpoint. Over time, various wireless security protocols were developed to protect wireless networks. Among the first of these was wired equivalent privacy (WEP). WEP is a security mechanism for IEEE 802.11[41] wireless networks introduced in 1997. Since its inception WEP has been **deprecated as insecure**[42] in favor of WiFi Protected Access (WPA2 and WPA2-Enterprise).

Wireless security protocols are intended to prevent unauthorized connections to a wireless network and encrypt data as it is being transmitted over the air. It's worth noting that wireless networks are generally not as secure as wired networks. In an effort to make wireless networking easy to use, the default configuration for early wireless components provided for easy connection but no security. Much of this has been addressed, but the omnidirectional broadcast inherent to wireless networks, intended to be received by every device within range, remains.

Additional detail on wireless security protocols is provided below:

- **WEP:** The original encryption protocol developed for wireless networks. WEP was designed to provide a level of security comparable to wired networks. However, WEP has many well-known security flaws (e.g., the use of a relatively weak, 40-bit encryption key and the flawed RC4 stream cipher), is difficult to configure and is easily broken. No modern wireless network should be using WEP.
- **WPA:** Introduced as an interim security enhancement over WEP while the 802.11i wireless security standard was being developed. Most current WPA implementations use a pre-shared key (PSK, commonly referred to as WPA Personal), and the Temporal Key Integrity Protocol (TKIP, pronounced tee-kip) for encryption. WPA Personal has one

40. TLS provides confidentiality, integrity and authentication between communicating applications.

41. http://ieeexplore.ieee.org/document/7920364/

42. No wireless network in 2018 or beyond should be using WEP for its security.

set of keys for all users whereas WPA Enterprise, which uses an authentication server to generate keys or certificates, has a discrete encryption key for each user.

- **WPA version 2 (WPA2):** This protocol is based on the 802.11i wireless security standard, which was finalized in 2004. The most significant enhancement to WPA2 over WPA is the use of the advanced encryption standard for encryption.

6.1.2.3 Internal Network Connectivity

There are many mechanisms for securing information moving across an organization's internal network. One option is to implement internet protocol security (IPsec). IPsec is a suite of protocols developed to ensure the confidentiality, integrity and authenticity of data communications over an internet protocol (IP) network. It is also used to implement virtual private networks (VPN)[43]. Unlike TLS, which is imposed on a per-session basis, IPsec is omnipresent and transparent to the user. The IPsec administrator can provide protections to all incoming and outgoing data.

IPsec is a powerful network security tool, offering authentication, confidentiality, integrity, access control, protection against replay attacks and some protection against traffic flow analysis. Additionally, it is extremely flexible. It can control whether data packets are protected by confidentiality or message integrity (or both)[44] and how much of the data packet is protected by these assurances[45]. All of this capability comes at a price. IPsec is extremely complex. While it offers great security, many options and a lot of flexibility, these same attributes make it hard to use, and use correctly.

An alternative to IPsec is the use of dedicated network encryptors. Whereas IPsec parameters are configured at the router level operating at the network's datalink layer, network encryptors are dedicated devices generally operating at the network layer (although some network encryptor products can operate at the datalink layer). They offer confidentiality, integrity and authenticity guarantees along with a high throughput and a reduced administrative burden. All of this capability comes at a cost, however; network encryptors are expensive.

6.1.2.4 Cellular/Mobile Communications

The good news for mobile phone communications is that both Global System for Mobile Communications (GSM) and Code Division Multiple Access (CDMA)[46] mobile communication

43. VPNs allow users to securely access a private network and share data remotely through public networks. Much like a firewall protects data on a computer, VPNs protect it online. And while a VPN is technically a WAN (Wide Area Network), the front end retains the same functionality, security, and appearance as it would on the private network. See: https://gizmodo.com/5990192/vpns-what-they-do-how-they-work-and-why-youre-dumb-for-not-using-one

44. Encapsulated Security Protection (ESP) gives both confidentiality and message integrity, whereas Authentication Header (AH) provides only message integrity.

45. In Transport Mode, only the payload of the IP packet is encrypted or authenticated. Since the IP header is neither modified nor encrypted, routing is intact. In Tunnel Mode the entire IP packet is encrypted and authenticated and encapsulated into a new IP packet with a new IP header. Tunnel mode is used to create VPNs for network-to-network communications, host-to-network communications (e.g. remote user access) and host-to-host communications (e.g. private chat).

46. CDMA is primarily used in the U.S., but there is a substantial GSM presence as well. GSM phones will always use SIM cards while only some CDMA phones, notably those on the Sprint and Verizon LTE networks, use SIM cards. See: https://www.pcmag.com/article2/0,2817,2407896,00.asp

standards support encrypted communications from mobile devices. The bad news is that mobile phone voice call security is notoriously weak.

GSM phones use the A5 stream cipher family to encrypt calls:

- **A5/1:** Primarily used in Europe and the United States, A5/1 was developed in 1987. By 2014, it was estimated that some 7.2 billion GSM device users relied on A5/1 for their voice call confidentiality. A number of attacks on A5/1 have been published, and it is believed that national agencies are able to decrypt A5/1 messages at will.[47]
- **A5/2:** Developed in 1989 (and re-engineered in 1999), A5/2 was a deliberate weakening of the algorithm for certain export regions in Asia. A5/2 was cryptanalyzed in the same month it was published and demonstrated to be so weak that it can be broken by commodity equipment in real time. On July 1, 2006, the GSM Association declared that GSM mobile phones will not support A5/2 any longer, due to its weakness and the fact that A5/1 was deemed mandatory by the 3GPP association. In July 2007, the 3GPP association approved a change request to prohibit the implementation of A5/2 in any new mobile phones.[48]
- **A5/3 (also known as Kasumi):** A5/3 is a block cipher derived from the MISTY1 cipher developed by Mitsubishi Electric Corporation and is the successor to the A5/1 cipher used in GSM phones. Unfortunately, it didn't fare much better than A5/1 in terms of security. In 2010, Orr Dunkelman, Nathan Keller and Adi Shamir published a cryptanalysis that recovered the full A5/3 key. Worse, the computational requirements were low enough that a 2010 vintage computer powered by Intel Core 2 Duo was able to complete the attack in less than two hours.[49] [50]

CDMA phones use the Cellular Message Encryption Algorithm (CMEA), or one of its derivatives, for encrypting voice communications. CMEA is designed to encrypt the control channel, rather than the voice data. In 1997, a group of cryptographers published attacks on the cipher showing it had several weaknesses that give it a trivial effective key length of a 24-bit to 32-bit cipher. CMEA and its improved successor, CMEA-I, should be considered to be insecure.[51]

Fortunately, there is an alternative.

6.1.2.5 Secure Chat

The period between 2015 and early 2018 saw the emergence and proliferation of a class of smartphone apps designed to ensure communications privacy. In many cases they support real-time chat, voice and file transfer over standard mobile data and WiFi networks. In general, these apps are intuitive and easy to use and eliminate the need for the user to understand or, in most cases, be aware of the encryption mechanism.

47. https://yro.slashdot.org/story/13/12/14/0148251/nsa-able-to-crack-a51-cellphone-crypto

48. https://en.wikipedia.org/wiki/A5/2

49. https://en.wikipedia.org/wiki/KASUMI

50. https://eprint.iacr.org/2010/013.pdf

51. https://en.wikipedia.org/wiki/Cellular_Message_Encryption_Algorithm

These encrypted messaging apps all have a feature called end-to-end encryption. End-to-end encryption means that the message is encrypted on the sender's device and sent to the recipient where it is decrypted. The mechanism prevents eavesdroppers, even if they are from the company that hosts the service, from being able to recover messages. It also means that the service provider does not have the ability to disclose the messages even if their servers were seized pursuant to a subpoena or breached by a malicious actor. Examples of end-to-end encrypted messaging apps include:

- Signal
- Cyphr
- Pryvate
- Wickr
- Silence
- Viber
- Voxer
- Threema
- Chat Secure

As an example of the security guarantees offered by modern secure chat apps, the Signal protocol provides confidentiality, integrity, authentication, participant consistency, destination validation, forward secrecy, backward secrecy (aka future secrecy), causality preservation, message unlinkability, message repudiation, participation repudiation and asynchronicity. (It does not provide anonymity preservation and requires servers for the relaying of messages and storing of public key material.)[52]

6.1.2.6 Secure Email

Email, with the possible exception of SMS messaging, may be the most ubiquitous electronic communications mechanism available today. It is asynchronous and ubiquitous, spans almost every platform imaginable and is reliable for both message and file communications. As a result of its centrality to both personal and public life, email security is largely a solved problem. Unfortunately, the implementation of email security solutions has lagged their availability, with some email providers making design decisions to hamper users' efforts to ensure the confidentiality of email communications.[53] The two primary means of securing email are through the use of clients that make use of digital certificates and secure email services.

Digital certificates were designed as a means to provide assurance that the recipient of a message secured using asymmetric (or "public key") cryptography was, in fact who he or

52. https://en.wikipedia.org/wiki/Double_Ratchet_Algorithm

53. For much of its existence, the price for Google's Gmail service was that Google's advertising mechanism "read" each email that traversed its servers so as to better craft targeted advertisements. As a result, none of the official Gmail clients (e.g., web or mobile) supported security or encryption. As of mid-2017, Google no longer reads Gmail messages (see: https://www.npr.org/sections/thetwo-way/2017/06/26/534451513/google-says-it-will-no-longer-read-users-emails-to-sell-targeted-ads), but the clients still do not support security mechanisms. Google claims that messages are secured via TLS between a Gmail client and Google's servers. However, the messages remain unsecured when they reach Google and when they are transmitted to a non-Gmail recipient.

she claimed to be. Public key cryptography solves the problem of securely sharing secrets (i.e., encryption keys) across large distances (i.e., those large enough to frustrate face-to-face interaction). It does this by using two distinct encryption keys: A **public key**, which can (and arguably should) be distributed to anyone in the world without compromising the message's confidentiality, and a **private key**, which is kept confidential by its owner. What one key encrypts, the other decrypts. Thus, if Alice wants to send Bob an email that only he can read, Alice would use Bob's public key to encrypt the email, secure in the knowledge that only Bob's private key is able to decrypt the email. Once it is encrypted with Bob's public key, Alice cannot decrypt the message.

When Bob receives the message, he uses his private key to decrypt the message. Importantly, there is an inverse relationship between the keys; what one encrypts, the other decrypts. Consequently, messages encrypted with a private key can be decrypted with a public key and vice versa.

Public Key Encryption

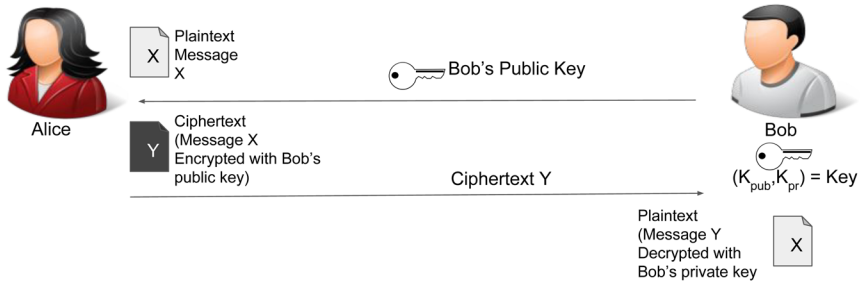

All Alice and Bob have to exchange in order to secure their messages, or in this case, emails, are their respective public keys. Unfortunately, Alice and Bob are still left with a problem. Alice wants her public key to be distributed as widely as possible so that she can communicate securely with as many people as possible. However, how do the public key's recipients know they're getting Alice's public key, and not a public key from Mallory (the malicious actor), pretending to be Alice? If they were to encrypt messages using Mallory's spurious key, Mallory could read every message intended for Alice.

The answer is found in digital certificates.[54] A digital certificate is a short text file that includes its owner's public key and identifies (among other things) the certificate's issuer and/or creator and the certificate owner.[55] Additionally, the certificate contains a representation of the identifying information that has been hashed[56] and then digitally signed with the *certificate issuer's* **private** key. The recipient of a digital certificate examines its identification portion to make sure it's relevant, hashes the identification portion and then uses the certificate issuer's **public** key to decrypt the encrypted/digitally signed portion of the certificate. If the recipient's calculated hash value matches the hash value of the decrypted digitally signed hash, then the

54. At least, the answer today is found in digital certificates. Emerging technologies such as blockchain may one day render the X.509 (or public key) infrastructure relied upon by digital certificates obsolete, or at least provide a viable alternative. Such technology is not widely available as of the time of publication of this white paper in early 2018.

55. The widely used X.509 standard includes a number of other data elements.

56. Hashing means that the data has been run through a hash function. A hash function maps data of arbitrary size to data of fixed size such that the result, for all possible inputs, is always the same length. In a cryptographic hash function, it is statistically impossible to recreate the original input from the hash value output. Because of this property, hash functions are used to verify data integrity. If two data have the same hash, they are identical.

recipient knows that the certificate is valid and that the information provided with it (e.g., the public key) can be trusted. The network of certificate issuers, or **certificate authorities,** that supports this distributed trust mechanism is called a "public key infrastructure" or PKI.

Digital certificates provide a means of protecting (or merely validating) emails when they are used in conjunction with email clients that have functionality to store and use them, generally by supporting secure multipurpose internet mail extensions (S/MIME) functionality. Such clients include (but are very much not limited to) Microsoft Outlook (both thick client and online versions), Mozilla Thunderbird and eM Client.

Digital certificates can be purchased in bulk for an organization and allocated to employees. Such certificates usually provide both public and private keys, and therefore support email confidentiality, and they support digital signatures, message integrity, authenticity and non-repudiation assurances as well.

Digital certificates can also be used as an authentication mechanism for NPEs and networked security devices that connect to back end and

Digital Certificate Chain of Trust

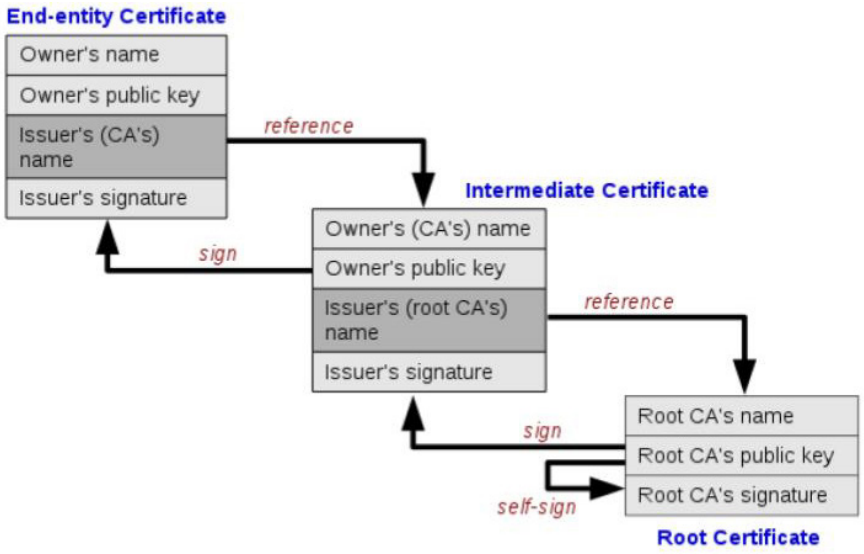

Public Key Certificate

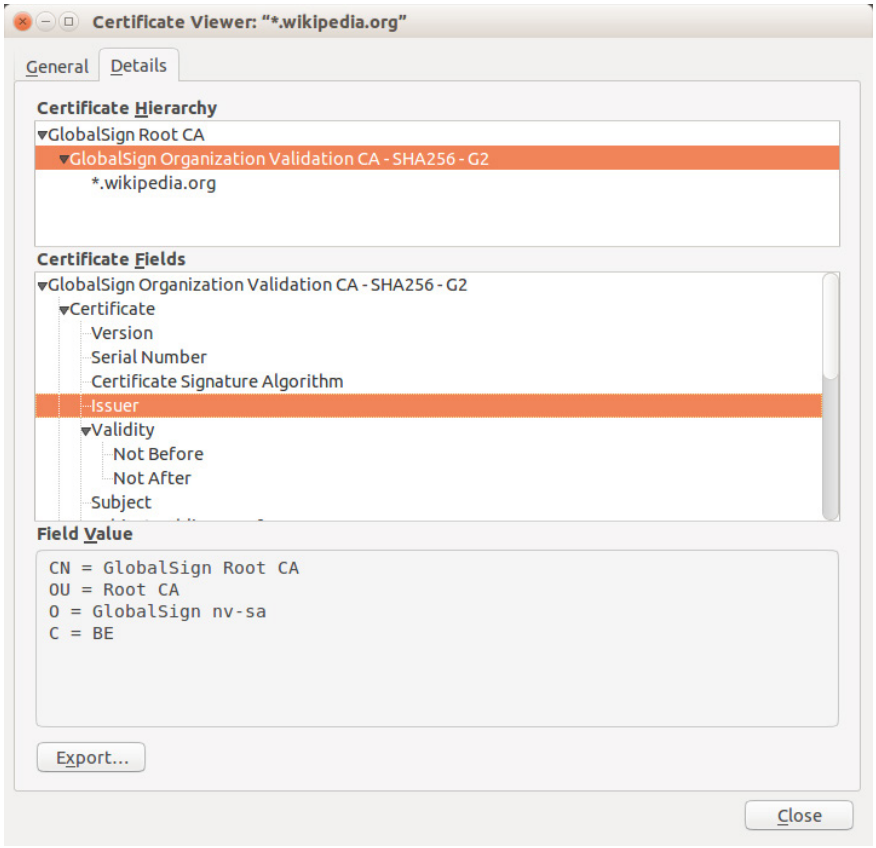

cloud management systems. Certificate-based authentication can be used for all endpoints – users, machine, devices and even the growing Internet of Things. In such cases, certificates are stored locally on the NPE device and generally require no additional management once installed at the time of manufacture.

In this case, certificates provide for mutual authentication, where both parties involved in the communication positively identify. For example, a remote microphone must prove its identity to a cloud-based command and control server and the server must prove its identity to the microphone, before a connection can be made. Certificate-based authentication is flexible and can support fine grained access control, ensuring that access is granted only to approved NPEs, preventing the entry of unauthorized NPEs or rogue machines.[57]

Unfortunately, there are hurdles to digital certificate adoption. On many email clients, the setup process is complex and can be challenging. There's also a cost issue. Depending on the degree of identity verification and certificate lifespan required, individual certificates can be expensive. Finally, there's a data ownership issue. Unless an enterprise invests in some mechanism to securely store the private keys associated with employee certificates, encrypted emails may not be recoverable when an employee leaves the organization.

Still, for individual use, or when the primary concern is confidentiality, certificates can provide a powerful information security tool. For example, individual certificates are available for free at InstantSSL.com.

For individuals and entities interested in securing their email communications without the overhead of managing S/MIME (digital) certificates, there are services that provide true client-side encryption as well as the administrative functions required by enterprise. Client-side encryption means that an email is fully encrypted before it leaves the email client (whether that client is a browser or a dedicated email application). Many of these services do not require email recipients to use the service and also provide controls over functionality such as forwarding, message time-to-live and offer data loss prevention capability.

6.2 Network and Cloud Storage

Many of the endpoint security controls discussed above are also applicable to information stored in shared network or cloud locations:

- Individual files can be encrypted on the endpoint prior to upload.
- Encrypted volumes can be created on network or cloud locations, ensuring that all information stored within them is protected.
- Cloud storage mechanisms can have encryption enabled such that all data stored within them is transparently encrypted at rest.

Additionally, access control plays a significant part in providing information security on remote storage. Sensitive information stores should not be publicly exposed. Role-based and other fine-grained access control mechanisms are available on all major networking (e.g., Microsoft Windows domains) and cloud platforms. For example, on Amazon Web Services' (AWS') Simple

57. https://www.globalsign.com/en/blog/what-is-certificate-based-authentication/

Storage Service (S3), all storage resources (e.g., buckets, the AWS equivalent to a Windows folder), objects, and related subresources such as life cycle website configurations are private by default and only visible to the resource owner. The resource owner can optionally grant access permissions to others. S3 access policy options can be attached to resources (buckets and objects) or to users. In S3, these mechanisms can be flexibly combined to manage access to information stored in S3.[58]

For non-object data stored remotely, AWS' Elastic Block Storage (EBS) service provides useful illustrations. EBS ties its access control to that of the virtual machine controlling the storage asset, which in turn leverages the overall AWS IAM. IAM enables AWS customers to:

- create users and groups;
- assign unique security credentials to each user;
- control each user's permissions to perform tasks using AWS resources;
- enable users in another AWS account to share the customer's AWS resources;
- create roles for the AWS account and define the users or services that can assume them; and
- use existing identities to grant permissions to perform tasks using AWS resources.[59]

With respect to encryption, encrypted EBS volumes can be created that secure the following types of data:

- data at rest inside the volume;
- all data moving between the volume and the virtual machines using it;
- all virtual machine snapshots created from the volume; and
- all volumes created from those snapshots.

EBS was designed to handle encryption and decryption transparently, requiring no additional action from AWS customer applications.[60]

6.3 Information Security by Location Recap

6.3.1 Endpoints

Endpoints manage both DAR and DIT.

6.3.1.1 DAR

- FDE is technology that secures an endpoint computing device by encrypting all the data at rest on its non-volatile storage. FDE is intended to protect information in the event that a device is lost, stolen or otherwise physically accessed without authorization. Typically, FDE technology uses OTFE.

58. https://docs.aws.amazon.com/AmazonS3/latest/dev/s3-access-control.html

59. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UsingIAM.html

60. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

- Partition encryption allows for greater flexibility than full disk encryption as the user is free to open (i.e., authenticate to) different encrypted partitions independently. Partition encryption products also generally use OTFE. Partition encryption is generally available as a feature of a disk management or encryption solution, not as a standalone capability.
- Volume encryption treats a disk volume as a single portion of data. Generally, it is more intuitive to manage a volume than work separately with every physical drive. Volume encryption is generally available as a feature of a disk management or encryption solution, not as a standalone capability.
- Virtual drive encryption tools operate as part of an OTFE scheme under which any data written to them is always encrypted and remains so until decrypted in volatile memory. Virtual drive encryption can be found in both single capability tools and as part of a full solution disk management and encryption solution.

6.3.1.2 DIT

- Online data exchange is generally done via a browser employing TLS, and users must be educated as to the indicators of a secure browsing session.
- WiFi is inherently insecure and users should only connect to known WiFi networks over authenticated connections.
- Enterprise networks should secure information moving across an organization's internal network using a mechanism such as IPsec or a network encryptor. External connections to the network should be controlled and secured using a mechanism such as a VPN.
- The mechanisms for preserving the confidentiality of voice calls over a mobile phone connection are broken.
- Secure chat provides a viable, secure alternative for mobile communications.
- Email can be secured using certificates or through a secure email service. Both have costs and relative advantages and disadvantages.

6.3.2 Network and Cloud Storage

- Many of the endpoint security controls discussed above are also applicable to information stored in shared network or cloud locations.
- Individual files should be encrypted on the endpoint prior to upload.
- Encrypted volumes can be created on network or cloud locations, ensuring that all information stored within them is protected.
- Cloud storage mechanisms can have encryption enabled such that all data stored within them is transparently encrypted at rest.
- Access control plays a significant part in providing information security on remote storage.
- Cloud services provide extensive and flexible access control and encryption options that should be exploited by administrators.

7 Roles and Responsibilities

Information security impacts everyone from individuals to enterprises and requires thought and understanding at all levels. It can no longer be relegated to a dedicated security team.

The conversation must be expanded and awareness raised as to how everyone plays a role in protecting sensitive information. The dissemination of information about how to secure and safeguard personal data and devices will create habits that extend into the workplace. When possible, provide security awareness training to educate employees not only on phishing, social engineering, ransomware, how to identify attacks and how to report possible attacks, but also how to use their devices more effectively.

Ensure that individuals, enterprises and leadership view information security as a long-term priority and that everyone supports the development of a holistic information security approach. Information security is a blend of technology, processes and people working together.

8 Conclusion

This white paper has broadly explored data security concepts and mechanisms. It bears repeating that all electronic perimeters are permeable and that only defenses that render data, individual devices or both inaccessible to unauthorized parties will provide the necessary defenses and deterrent effects to keep an organization's information secure. Additionally, while this paper's primary intended readers are senior executives and operational decision makers, this information is equally applicable to information technologists and security practitioners.