Issue 3 | May | 2024

Terminus Shield

ULTIMATE SECURITY IS WHAT YOU DESERVE

Our goal is to increase the level of awareness of the audience about the latest topics on information security, so with sending your feedback and comments, help us to do it better.

info@terminusys.com

From the Editor

The convergence of blockchain and cybersecurity presents a paradigm shift in combating malware and ensuring data integrity. While malware remains a persistent threat, blockchain's decentralized, transparent, and secure architecture offers promising avenues for fortifying our digital ecosystems against evolving cyber risks. By leveraging the inherent strengths of blockchain technologies, organizations can fortify their defenses and safeguard critical assets in an increasingly interconnected world.

In this editorial, we've explored the intricate relationship between blockchain and malware, underscoring the potential of blockchain as a transformative force in modern cybersecurity.

As we navigate this complex landscape, harnessing the synergy between innovation and vigilance will be pivotal in safeguarding digital assets and preserving trust in our interconnected society.

In the realm of cybersecurity, the dynamic interplay between emerging technologies like blockchain and persistent threats such as malware continues to unfold. Blockchain, heralded for its security and transparency, presents a formidable barrier against tampering and fraud. On the other hand, malware, an enduring challenge, exploits vulnerabilities to compromise systems and steal data. Understanding these concepts is key to navigating the evolving landscape of Cybersecurity.

Blockchain, at its core, is a decentralized, distributed digital ledger that records transactions across multiple computers in a secure and transparent manner. It ensures trust through its foundational attributes:

- Distributed and Sustainable: The ledger is shared among participants, eliminating dependence on a single controlling entity.
- Secure, Private, and Indelible: Utilizes cryptography to prevent unauthorized access and tampering, ensuring transactions are immutable.
- Transparent and Auditable: Participants have access to the same records, allowing for validation without intermediaries.
- Consensus-Based and Transactional: Requires agreement among network participants for transaction validity.

In contrast, malware represents a persistent threat to cybersecurity. Malicious software is designed to infiltrate systems, steal data, or disrupt operations. One common vector for malware delivery is through email attachments, leveraging user interactions to execute malicious payloads.

What is BlockChain?

Introduction

Blockchain builds trust through the following five attributes:

- Distributed and sustainable: The ledger is shared, updated with every transaction, and selectively replicated among participants in near real time. Because it's not owned or controlled by any single organization, the Blockchain platform's continued existence isn't dependent on any individual entity.
- Secure, private, and indelible: Permissions and cryptography prevent unauthorized access to the network and ensure that participants are who they claim to be. Privacy is maintained through cryptographic techniques and/or data partitioning techniques to give participants selective visibility into the ledger; both transactions and the identity of transacting parties can be masked. After conditions are agreed to, participants can't tamper with a record of the transaction; errors can be reversed only with new transactions.
- Transparent and auditable: Because participants in a transaction have access to the same records, they can validate transactions and verify identities or ownership without the need for third-party intermediaries. Transactions are time-stamped and can be verified in near real time.
- Consensus-based and transactional: All relevant network participants must agree that a transaction is valid. This is achieved through the use of consensus algorithms. Each Blockchain network can establish the conditions under which a transaction or asset exchange can occur.
- Orchestrated and flexible: Because business rules and smart contracts (that execute based on one or more conditions) can be built into the platform, Blockchain business networks can evolve as they mature to support end-to-end business processes and a wide range of activities.

Distributed and sustainable: The ledger is shared, updated with every transaction, and selectively replicated among participants in near real time. Because it's not owned or controlled by any single organization, the Blockchain platform's continued existence isn't dependent on any individual entity.
Secure, private, and indelible: Permissions and cryptography prevent unauthorized access to the network and ensure that participants are who they claim to be. Privacy is maintained through cryptographic techniques and/or data partitioning techniques to give participants selective visibility into the ledger; both transactions and the identity of transacting parties can be masked. After conditions are agreed to, participants can't tamper with a record of the transaction; errors can be reversed only with new transactions.
Transparent and auditable: Because participants in a transaction have access to the same records, they can validate transactions and verify identities or ownership without the need for third-party intermediaries. Transactions are time-stamped and can be verified in near real time.
Consensus-based and transactional: All relevant network participants must agree that a transaction is valid. This is achieved through the use of consensus algorithms. Each Blockchain network can establish the conditions under which a transaction or asset exchange can occur.
Orchestrated and flexible: Because business rules and smart contracts (that execute based on one or more conditions) can be built into the platform, Blockchain business networks can evolve as they mature to support end-to-end business processes and a wide range of activities.

Blockchain concepts and structure

Here, you get a glimpse of how Blockchain stores transactions in a way that prevents recorded transactions from being changed. You discover the four concepts that form the foundation of a Blockchain for business, and you meet the network participants and find out about the various roles they play. A blockchain is a decentralized, distributed and public digital ledger that is used to record transactions across many computers so that the record cannot be altered retroactively without the alteration of all subsequent blocks and the collusion of the network. This allows the participants to verify and audit transactions inexpensively. A blockchain database is managed autonomously using a peer-to-peer network and a distributed timestamping server. They are authenticated by mass collaboration powered by collective self-interests.

The result is a robust workflow where participants' uncertainty regarding data security is marginal. The use of a blockchain removes the characteristic of infinite reproducibility from a digital asset. It confirms that each unit of value was transferred only once, solving the long-standing problem of double spending. Blockchains have been described as a value-exchange protocol. This blockchain-based exchange of value can be completed more quickly, more safely and more cheaply than with traditional systems. A blockchain can assign title rights because, when properly set up to detail the exchange agreement, it provides a record that compels offer and acceptance. Here I explain some basic concepts used in Blockchain:

Blocks

Blocks hold batches of valid transactions that are hashed and encoded into a Merkle tree. Each block includes the cryptographic hash of the prior block in the blockchain, linking the two. The linked blocks form a chain. This iterative process confirms the integrity of the previous block, all the way back to the original genesis block.

Sometimes separate blocks can be produced concurrently, creating a temporary fork. In addition to a secure hash-based history, any blockchain has a specified algorithm for scoring different versions of the history so that one with a higher value can be selected over others. Blocks not selected for inclusion in the chain are called orphan blocks. Peers supporting the database have different versions of the history from time to time. They keep only the highest-scoring version of the database known to them. Whenever a peer receives a higher-scoring version (usually the old version with a single new block added) they extend or overwrite their own database and retransmit the improvement to their peers. There is never an absolute guarantee that any particular entry will remain in the best version of the history forever. Because blockchains are typically built to add the score of new blocks onto old blocks and because there are incentives to work only on extending with new blocks rather than overwriting old blocks, the probability of an entry becoming superseded goes down exponentially as more blocks are built on top of it, eventually becoming very low. For example, in a blockchain using the proof-of-work system, the chain with the most cumulative proof-of-work is always considered the valid one by the network. There are a number of methods that can be used to demonstrate a sufficient level of computation. Within a blockchain the computation is carried out redundantly rather than in the traditional segregated and parallel manner.

Block time

The block time is the average time it takes for the network to generate one extra block in the blockchain. Some blockchains create a new block as frequently as every five seconds. By the time of block completion, the included data becomes verifiable. In cryptocurrency, this is practically when the money transaction takes place, so a shorter block time means faster transactions. The block time for Ethereum is set to between 14 and 15 seconds, while for bitcoin it is 10 minutes.

Hard forks

A hard fork is a rule change such that the software validating according to the old rules will see the blocks produced according to the new rules as invalid. In case of a hard fork, all nodes meant to work in accordance with the new rules need to upgrade their software.

If one group of nodes continues to use the old software while the other nodes use the new software, a split can occur. For example, Ethereum has hard-forked to "make whole" the investors in The DAO, which had been hacked by exploiting a vulnerability in its code. In this case, the fork resulted in a split creating Ethereum and Ethereum Classic chains. In 2014 the Nxt community was asked to consider a hard fork that would have led to a rollback of the blockchain records to mitigate the effects of a theft of 50 million NXT from a major cryptocurrency exchange. The hard fork proposal was rejected, and some of the funds were recovered after negotiations and ransom payment. Alternatively, to prevent a permanent split, a majority of nodes using the new software may return to the old rules, as was the case of bitcoin split on 12 March 2013.

Decentralization

By storing data across its peer-to-peer network, the blockchain eliminates a number of risks that come with data being held centrally. The decentralized blockchain may use ad-hocmessage passing and distributed networking. Peer-to-peer blockchain networks lack centralized points of vulnerability that computer crackers can exploit; likewise, it has no central point of failure. Blockchain security methods include the use of public-key cryptography. A public key (a long, random-looking string of numbers) is an address on the blockchain. Value tokens sent across the network are recorded as belonging to that address. A private key is like a password that gives its owner access to their digital assets or the means to otherwise interact with the various capabilities that blockchains now support. Data stored on the blockchain is generally considered incorruptible.

While centralized data is more easily controlled, information and data manipulation are possible. By decentralizing data on an accessible ledger, public blockchains make block-level data transparent to everyone involved. Every node in a decentralized system has a copy of the blockchain. Data quality is maintained by massive database replication and computational trust. No centralized "official" copy exists and no user is "trusted" more than any other. Transactions are broadcast to the network using software. Messages are delivered on a best-effort basis. Mining nodes validate transactions, add them to the block they are building, and then broadcast the completed block to other nodes. Blockchains use various time-stamping schemes, such as proof-of-work, to serialize changes. Alternate consensus methods include proof-of-stake. Growth of a decentralized blockchain is accompanied by the risk of node centralization because the computer resources required to process larger amounts of data become more expensive.

Openness

Open blockchains are more user-friendly than some traditional ownership records, which, while open to the public, still require physical access to view. Because all early blockchains were permissionless, controversy has arisen over the blockchain definition. An issue in this ongoing debate is whether a private system with verifiers tasked and authorized (permissioned) by a central authority should be considered a blockchain. Proponents of permissioned or private chains argue that the term "blockchain" may be applied to any data structure that batches data into time-stamped blocks. These blockchains serve as a distributed version of multiversion concurrency control (MVCC) in databases. Just as MVCC prevents two transactions from concurrently modifying a single object in a database, blockchains prevent two transactions from spending the same single output in a blockchain. Opponents say that permissioned systems resemble traditional corporate databases, not supporting decentralized data verification, and that such systems are not hardened against operator tampering and revision. Nikolai Hampton of Computerworld said that "many in-house blockchain solutions will be nothing more than cumbersome databases," and "without a clear security model, proprietary blockchains should be eyed with suspicion." Business analysts Don Tapscott and Alex Tapscottdefine blockchain as a distributed ledger or database open to anyone.

Permissionless

The great advantage to an open, permissionless, or public, blockchain network is that guarding against bad actors is not required and no access control is needed. This means that applications can be added to the network without the approval or trust of others, using the blockchain as a transport layer. Bitcoin and other cryptocurrencies currently secure their blockchain by requiring new entries to include a proof of work. To prolong the blockchain, bitcoin uses Hashcash puzzles. While Hashcash was designed in 1997 by Adam Back, the original idea was first proposed by Cynthia Dwork and Moni Naor and Eli Ponyatovski in their 1992 paper "Pricing via Processing or Combatting Junk Mail".

Financial companies have not prioritised decentralized blockchains. In 2016, venture capital investment for blockchain-related projects was weakening in the USA but increasing in China. Bitcoin and many other cryptocurrencies use open (public) blockchains. As of April 2018, bitcoin has the highest market capitalization.

Permissioned (private) blockchain

Permissioned blockchains use an access control layer to govern who has access to the network. In contrast to public blockchain networks, validators on private blockchain networks are vetted by the network owner. They do not rely on anonymous nodes to validate transactions nor do they benefit from the network effect. Permissioned blockchains can also go by the name of 'consortium' or 'hybrid' blockchains.

The New York Times noted in both 2016 and 2017 that many corporations are using blockchain networks "with private blockchains, independent of the public system."

Remote Desktop Protocol (RDP)

Restrict access via firewalls, RDP gateway, VPNs
Use strong passwords and 2FA
Limit users who can log in using RDP
Set an account lockout policy

Now that we've covered a few basic steps toward making malware delivery more difficult, let's turn our attention back to the challenge of email attachments you can't practically filter out. What happens when a malicious attachment successfully makes its way into a user's inbox and that user gets fooled into opening it? Well, it's not game over yet. In the vast majority of cases, that attachment isn't going to do anything inherently malicious on its own. Its main role is to serve as a dropper that retrieves the primary malware payload from a C2 server or compromised website. Luckily, that means there's still time to nip this attack in the bud.

Here's a simple diagram to illustrate that point. It depicts one of the most common infection scenarios — a spam email with a Word document attachment, weaponized with a macro designed to launch PowerShell and download a malicious payload. As you can see, there are multiple opportunities for disrupting the process, even after a user makes the unfortunate mistake of opening the attachment.

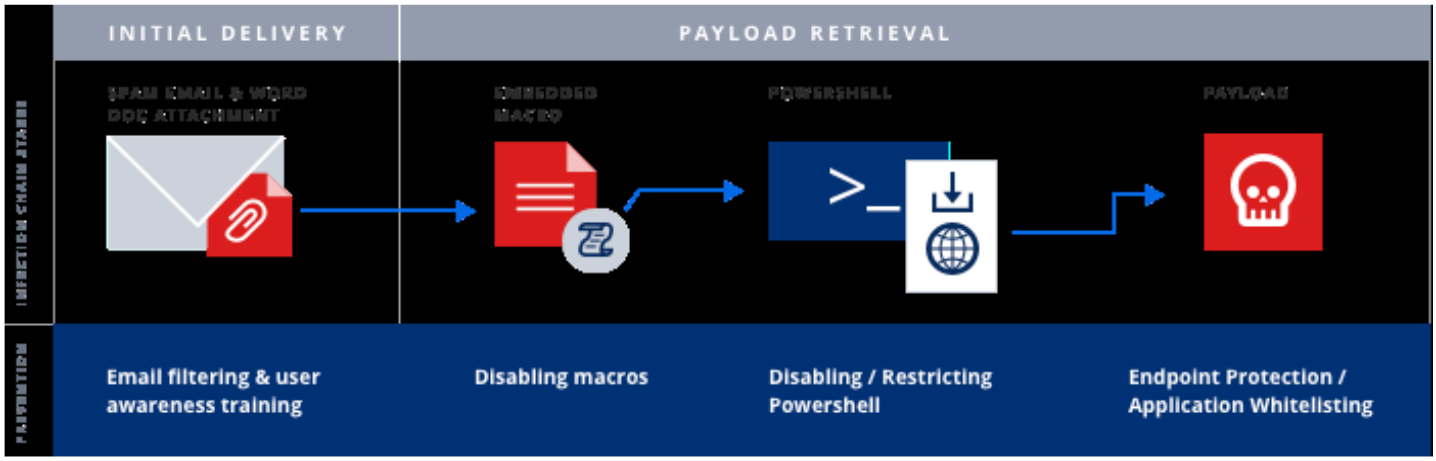

Abusing Microsoft Office

There are few legitimate applications attackers love abusing more than Microsoft Office programs. Why? A couple of reasons...

1) Office files are almost universally accepted

Attackers are able to capitalize on the fact that using and sharing Office documents is baked into most users' day-to-day work. Rather than trying to trick users into downloading suspicious executable files that many AV programs would block anyway, instead they can deliver the types of familiar-looking documents, reports, invoices, spreadsheets, etc. that users expect to receive at work.

2) Office files are easy to weaponized

Microsoft has gone to great lengths to build powerful capabilities into Office programs designed to make them more useful. Unfortunately, those capabilities are just as attractive to attackers as they are to users, and because Microsoft considers these capabilities as features rather than vulnerabilities, the company has rarely issued patches or fixes in response to abuse. Let's take a look at some of the most commonly-abused capabilities.

There a few methods for attackers to abuse Microsoft Office to retrieve payloads

METHOD#1: MACROS

Macro abuse has been around for years, having experienced its first heyday back in the late 1990s and early 2000s. Over the past few years, there's been a major revival, with a large percentage of spam campaigns utilizing Word documents and embedded macros to download malware.

For attackers, part of the draw of macros is how simple they are to build and configure, but they have their drawbacks, too. For one thing, macros are disabled by default, so utilizing them requires tricking a user into enabling them.

Macros also don't provide attackers with all the functionality they need to accomplish their goals, so they're typically used to send a request out to a command-and-control server or compromised website to grab a second-stage payload.

METHOD#2: OBJECT LINKING AND EMBEDDING (OLE)

Microsoft developed OLE to give users the ability to link to and add data from other applications inside Office documents, or even embed one type of document inside another. Attackers haven't been shy about abusing this capability, often using it to trick users into inadvertently launching embedded scripts (typically Visual Basic or JavaScript) that in turn download a second-stage payload.

Like macros, one downside from the attacker's perspective is OLE-embedded objects require user interaction. First, the user has to interact with the object (often disguised as a file icon), then respond to a warning prompt by confirming that they do want to open it.

METHOD#3: DYNAMIC DATA EXCHANGE (DDE)

DDE is actually an older feature that was superseded by OLE, and it provided similar capabilities by allowing Office programs to automatically load data from other Office programs. Attackers latched onto that functionality, as well, abusing DDE to launch code to the command line, instead.

Once again, the good news is DDE requires user interaction. When a user opens a document with DDE fields they will receive a warning notifying them that the document contains links that may refer to other files. To continue, the user then has to confirm that they do want to update the document with data from the linked files.

How to prevent abusing Microsoft Office

Preventing macro abuse

Here are a few options for protecting your organization from macro malware, listed from most to least demanding:

- **Disable macros across the entire organization:** While it might be a natural conclusion to jump to, simply disabling macros entirely isn't always a practical option. Some legitimate software and business processes rely on macros for functionality.
- **Whitelisting:** What about only allowing authorized macros? Maybe. Macros do have a signature format that can support allowing only digitally-signed macros to run, but that's difficult to maintain from an IT perspective, especially as an organization grows.
- **Disable macros in certain scenarios:** Microsoft ramped up macro protection in its Office 2016 suite, providing admins with more granular controls. One helpful option is the ability to block macros in high-risk situations only, such as when a user is attempting to enable macros in a document downloaded from the Internet.

Discussion

Table 1 shows some of the recent cyberattacks and the ETIF for each of the attacks. The ETIF ranged from as low as three (3) hours (University of Maryland) to as long as one (1) year (Sony Pictures & the US Government's Office of Personnel Management). The actual cyberattack date and time, as mentioned before, is determined by forensic analysis. The forensic analysis process is not standard across all industries and as such the data in the public domain is not easily comparable. As of today, there is no agreed definition of what activities would characterize the start of a cyberattack. For example, in some cases, certain viruses are dormant in the system before becoming malignant and in other cases there are multiple attacks over a period of time. Corporations are concerned about confidentiality, competitive pressures, litigation, image, etc., and are therefore reluctant to disclose the specific details of a cyberattack including the date of the attack and the date of the identification. Given the lack of standardization in publicly available data on start of a cyberattack, ts, and on the identification of a cyberattack ti, it is currently difficult to consistently and specifically calculate ETIF and hence cyber resiliency

Table 1

Calculated ETIF based on recent cyberattacks

No	Company/ Organization	Cyberattack Date	Cyberattack Identified Date	Elapsed Time to Identify Failure (ETIF)
1	Premera Blue Cross[11]	May-14	29-Jan-15	> 7 months
2	Anthem Blue Cross[12]	Dec-14	29-Jan-15	> 1 month
3	Sony Pictures[6]	as early as Nov 2013	24-Nov-14	1 year
4	Staples[13]	as early as July 20, 2014	20-Oct-14	> 3 months
5	Home Depot[14]	as early as April 2014	2-Sep-14	> 4 month
6	JP Morgan[15]	as early as mid-June 2014	Mid-August 2014	> 2 months
7	Community Health[16]	Apr-14	1-Jul-14	3 months
8	Toys R Us[17]	28-Jan-15	30-Jan-15	3 days
9	University of Maryland[18]	17-Feb-15	17-Feb-15	3 hours
10	Office of Personnel Management (OPM)[19]	May-14	1-May-15	1 Year

If skilled companies in the IDS calculate and report ETIF instead of the corporations experiencing the cyberattack, it would enable standardization of the forensic process and development of the tools necessary to clearly define, measure and calculate the start of a cyberattack. By reporting ETIF instead of the start of cyberattack, t_s, and the identification of cyberattack, t_i, corporations may feel less exposed to litigation and negative publicity.

Publishing ETIF and reducing the value of ETIF could spur competition in IDS space to develop advanced algorithms to identify anomalies in the quality of service data. However, for ETIF to be meaningful and to be an effective metric for improving cyber resiliency, clear definitions for t_i and t_e, and agreed upon processes and tools need to be established to measure t_i and t_e.

The discussion until now has been focused on the loss of resiliency of an information system. However, an ideal state would be to have no loss of resiliency. That is, the threat of a failure is identified early enough before it becomes an attack. Such an outcome is possible if there was a mechanism to disseminate information about the adverse events and the threats among various entities on a real-time basis. Ideally, if an entity that experienced failure could share the system's vulnerability with other entities, analysis of such shared failure event would be critical in detecting the presence of a threat or recover quickly from the failure. We define this metric as Elapsed Time to Identify Threat (ETIT). The effectiveness of such a metric can be measured by its ability to move t_i (time to identify failure) closer to t_s (start time of the failure). ETIT is critical in changing limits on the loss and recovery functions and thus impacting the quality of service. If there is an ability for early identification of the threat that is causing the failure, then the overall time to recovery and hence the loss of resiliency could be reduced.

It is conceivable that different information systems could have the same loss of resiliency (i.e., same area of the triangle) but with different slopes for loss and recovery functions. This is illustrated in Figure 5 where System 1 and System 2 have the same loss of resiliency but System 1 has much steeper loss and recovery functions and it fully recovers to perform its full intended function much sooner than System 2. From a user/customer perspective, System 1 would be preferred in the sense that it recovered to perform its fully intended function in a much shorter duration than System 2. Therefore, in addition to the loss of resiliency, the slope of the loss and recovery functions would become important. The steeper the slope for the loss and recovery functions and shorter the duration between ts and ti (i.e., ETIF), the smaller the loss of resiliency, the slope of the loss and recovery functions would become important. The steeper the slope for the loss and recovery functions and shorter the duration between ts and ti (i.e.,ETIF), the smaller the loss of resiliency.

Conclusions and Future Wor

A modified information system resiliency model was presented and introduced two new variables: Elapsed Time to Identify Failure (ETIF) and Elapsed Time to Identify Threat (ETIT). The model qualitatively demonstrated that the lower the ETIF, the smaller the loss of resiliency for an information system. Challenges in calculating ETIF in recent cyber-attacks showed difficulty in measuring the start of the cyberattack (t_s) and the identification of cyberattack (t_i).

By transferring the responsibility of calculation and publication of ETIF from the companies that experience cyberattacks to the companies that are in IDS space could bring standardization and continuous improvement in measuring and reporting ETIF. In addition, it could encourage competition in the IDS space and would bring innovative and superior algorithms to find anomalies in the data quality thus improving cyber resiliency.

With the growth in cloud computing, information about the operational assets is being stored away from the local servers. Although the benefits of cloud computing are obvious, such decoupling could result in poor operational resiliency if the information asset is compromised. Therefore, to keep the operational resiliency unaffected, it is essential to bolster information asset resiliency in the cloud. Hence, a technical framework, along with appropriate regulatory framework, needs to be created to enable the measurement and reporting of ETIF and ETIT. The Cybersecurity Act of 201516 passed by the US Congress, aims at setting a regulatory framework for inter-government and private company to government exchanges. NIST has released a special publication outlining the building blocks of an incident exchange program. Future work would entail architecting a framework that allows a set of cooperative systems to aggregate, summarize, and utilize ETIF and ETIT to improve resiliency.