Earning Trust in the 21st Century

Cloud Security Alliance - DC (CSA-DC)
Research

CLOUD SECURITY ALLIANCE OFFICIAL CHAPTER

D.C.
Chapter

CSA

CSA cloud security alliance®

Acknowledgments

CSA-DC Research Chair:

Dr. Mari Spina

Authors:

Simran Sakraney

Contributors:

Justin Agostine
Al Aresenault
Bob Klankurkan
Juanita Koilpillai

Special Thanks:

Todd Edison
Frank Guanco
Dan Hiestand
Anil Karmel

CSA Staff:

AnnMarie Ulskey (Graphic Design)

Table of Contents

1.0 Background

In today's interconnected and technology reliant world, the expectation of trust and need to trust is growing. [1] Today's trust-based solutions may become non-viable in the future. [2] [3] As use of the cloud grows, we are experiencing a shift in resource allocation from on-premise to off-premise systems. As systems move to cloud hosted environments, the loss of control over the access network becomes a concern. Today's trust-based solutions typically start at the network level. If a user has access to a network, they are typically trusted to have access to some or all of the resources, data, and systems on that network.

But, when networks are unknown and untrusted, how is trust acquired? Zero Trust (ZT) architectures seek to provide access control techniques that assume the network is not trustworthy. One of the approaches suggested by industry is the use of trust scores. Like a credit score, a cyber trust score could be used to assess the risk potential associated with allowing any given user access to systems and information. But how would a trust score be calculated? Current approaches smack of a violation of privacy where the right to gain access is issued only by agreeing to be monitored.

This paper addresses the technical, social, policy, and regulatory issues associated with creating trust frameworks in a Zero Trust world. Industry and government are called to solve issues in ways that continue to protect the right to a users' privacy.

Traditional Domain-Based Trust Systems

Transitive Trust is a two-way relationship automatically created between two domains in a forest. For example, transitive trust may allow the resource domain to trust the account domain through a chain of trust relationships—and even between intermediate domains. [22] Inter-Domain Trust occurs when a domain provides another trusting domain with its user's security access token. The trusting domain may use the token to determine if the user has the necessary permissions to access its resources. For example, in operation, a user logs into the first trusted domain and then opens a file in the second trusted domain without logging into the second domain. [22]

Authentication is the process an entity undergoes to prove its identity to a second entity—which is often a system the former entity is attempting to access. [4] Authentication may occur between any system, such as a computer program and an end user (human), a computer system, a piece of hardware, or a mobile device. Credentials authenticate users and are considered proof of identity. There are different credentials—including Public Key Infrastructure (PKI)—which use digital certificates, passwords, pins, and even biometrics (such as fingerprints and iris scans). [4]

Trust today is achieved between networks and domains using Single Sign On (SSO), Federation, and Kerberos protocols. SSO is a session and user authentication service that permits end users to enter one set of login credentials (such as a name and password) to access multiple applications. SSO allows for a user's identity to provide access across numerous service providers. [4]

Federated identity management (FIM) is an arrangement between multiple enterprises that allows subscribers to use the same identification data to obtain access to all networks in the group. Federation works by establishing common standards and protocols to manage and map user identities between Identity Providers (IdPs) across organizations (and security domains). [4] When two domains are federated, a user can authenticate to one domain and then access resources in the other domain without performing a separate login process due to trust relationships previously established between the domains via digital signatures, encryption, and PKI. For FIM to work efficiently, the partners must develop mutual trust. [4]

Identity and access management (IAM) is a framework of business processes, policies, and technologies that enable digital identity management. Federated Identity Management is a sub-discipline of IAM. Identity, Credential, and Access Management (ICAM) include IAM. Federation allows SSO without requiring passwords because the federation server knows the username for a person in each application and presents that application with a token via Security Assertion Markup Language (SAML), OpenID, WS-Trust, WS-Federation, and OAuth. The token implies that "this person is an accepted user of domain\johndoe or johndoe@example.com." As a result, only one password is required for the user to login to multiple systems. [4]

Kerberos is a prominent example of a trusted SSO system and functions as a computer-network authentication protocol. Kerberos uses tickets that allow nodes—communicating over a non-secure network—to prove their identity to one another in a secure manner. Kerberos has the foundation of symmetric key cryptography and requires a trusted third party. For certain phases of authentication, Kerberos may use public-key cryptography. [4]

Zero Trust Obviates the Network

Technology advancements in cloud computing applications provide end users with data access from anywhere and have spurred a dramatic migration to the cloud. The transition from traditional, domain-based network security to the cloud has forever altered the rules. These previously established network frameworks have lost form and function and are quickly becoming obsolete.

Furthermore, a surge in the "bring your own device" (BYOD) trend is only hastening this movement.

The ZT model arose in response to insufficient domain-based perimeter security approaches that could not deter hackers from breaching corporate firewalls. [21][22]

A Zero Trust Architecture (ZTA), as defined by the National Institute of Standards and Technology, provides a collection of concepts, ideas, and component relationships (architectures) designed to eliminate the uncertainty in enforcing accurate access decisions in information systems and services. [22]

Every actor, system, or service requires thorough verification. The Zero Trust model conducts more due diligence than perimeter-based security did in the past. Zero Trust allows enterprises to define internal trust boundaries to granularly control traffic flow, enable secure network access, and implement continuous network monitoring. Authentication and associated data determine whether a user, machine, or application is trustworthy and granted access to a trust boundary.

Zero trust establishes trust zones where resources with similar trust levels and functionality operate alongside one another. This process minimizes pathways and malicious threats. Zero trust encourages maximum security and controls for the highest level of network visibility, threat detection, and compliance reporting. Zero trust solutions may include a combination of multi-factor authentication, strong encryption, file system permissions, and analytics enablement for threat detection and prevention. Zero trust also supports the principle of least privilege. [24]

However, today's Zero Trust Architectures (ZTA) may not be sufficient to make trust determinations. When introduced, Google's BeyondCorp alliance proposed the trust or inference engine for devices accessing its network. Such systems need appropriate information from which to make trust assessments. [24] Trust scoring approaches have been posited, but questions remain. What measures should be considered? And how will such data impact system use and utility?

2.0 Solution Landscape

Trust Scoring Models for Entities

FICO® created an Enterprise Risk Suite, a pilot program that allows businesses to access their FICO® Enterprise Security Score. The score derives from machine learning models that forecast a future breach event's likelihood by analyzing key risk indicators—including IT system health and hygiene, network infrastructure, and software and services. These current and historical data signal behaviors are compared to past actions of organizations that have (and have not) suffered a material data breach. This worldwide program gives any company the ability to assess their cybersecurity posture before they are evaluated by other organizations in their supply chain, free of charge. [26]

Other vendors are attempting to create enterprise cyber trust scores as well. RiskSense, a vulnerability management firm, offers a service that scans agency assets, applications, databases, and networks for vulnerabilities and assigns a cybersecurity risk score (much like a credit score). [18] According to Anand Paturi, vice president of RiskSense, the "RiskSense-RS3" score is essentially a prediction of the level of peril the organization is subject to, based on its vulnerabilities and contextual data. [18] The results have been positive because of the financial credit score system's widespread popularity.

For example, all participating agencies in Arizona currently have scores of 700, and RiskSense is trying to elevate this result to 725. These scores prioritize patching and give the state's security experts an easier way to communicate a system's security posture with leadership and other collaborators. [26] [27]

Another player in the marketplace is Sift. Machine learning powers Sift's Digital Trust Platform, which claims to automatically enhance digital interactions in real time based on individual risk scores that are developed by predicting user intent. [2] Sift protects businesses from all vectors of fraud and abuse, including payment fraud, account takeover, fake accounts, and abusive user-generated content.

According to Geoff Huang, Sift's vice president, "The objective of Sift is to make the best predictions of outcomes for behaviors on the internet." [2] Sift's methods help determine online user trustworthiness in real time and enable participating businesses to enter new markets previously inaccessible due to security concerns. Many global brands such as Twitter, Airbnb, Yelp, Indeed, Zillow, and Wayfair, use

Sift. Additionally, Sift allows validated businesses to publicly display a Sift trust and security "badge" to indicate participation in the Sift program. This program strengthens the digital trust network by connecting users with trustworthy businesses, which creates a "win-win" situation.

According to a recent article in the Wall Street Journal, Sift uses more than 16,000 signals to inform the Sift score, "which is a rating that ranges from 1 to 100." [16] The Sift score can flag and/or block devices, credit cards, and accounts owned by any entity, whether human or robot. This score is like a credit score, but for overall trustworthiness. [2] There is no way to find out your Sift score. Companies use products, such as Sift and SecureAuth, to detect bots and determine who to subject to additional screening (such as requesting a user to upload a form of identification). [2]

Cyber risk scores are often based on characteristics relative to specific industries, cyber supply chains, vulnerabilities, network connectivity, and interaction efforts, among other variables. [6] Insurance companies use scoring systems to provide an empirical foundation for issuing policies and pricing insurance premiums, commonly known as actuarial tables. This framework equips high-risk companies with clear benchmarks and incentives to purchase insurance and prioritize internal investments to mitigate specific cyber-risk areas.

Admittedly, many of these scores are curated by machine learning models with known flaws. Many security rating companies use a combination of data points collected or purchased from public and private sources—in conjunction with proprietary algorithms—to articulate an organization's security effectiveness into a quantifiable measure or score. Ratings rely upon harvesting accurate data from various sources in dynamic environments—a challenging process that can potentially produce inaccurate, biased, irrelevant, or incomplete results.

Trust Scoring for Individuals

In 2014, Beijing established a "Plan for Establishing a Social Credit System" as a nationwide, regulated project. In 2015, the Central Bank of China granted licenses to 12 pilot projects, including China's biggest tech giant, Alibaba—which released Sesame Credit, a private credit scoring and loyalty program system. By 2016, the government established 31 pilot projects in addition to Sesame Credit, including Unified Social Credit Numbers, the Honest Shanghai app, and Tencent Credit. [29] The apps are currently voluntary and give users a public credit score based upon citizens' social, financial, personal, and behavioral habits.

According to the Chinese government, these apps were implemented to judge citizen trustworthiness and encourage more honesty. The apps track credit histories, financial purchases, and even behavior towards friends and family. This compiled data feeds algorithms, which create individual scores that are adjusted based on positive or negative actions. [29] Scores determine whether individuals are placed on a "blacklist" or a "red list," Beijing's version of a "white list." Both lists are readily available on a website called China Credit.

In China, the exact methodology behind social credit scores remains secret—much like financial credit score tabulations in the United States. [29] However, it is important to note this may be due to censorship and a general lack of transparency in Beijing. Rachel Botsman, author of *Who Can You Trust? How Technology Brought Us Together—and Why It Could Drive Us Apart*, gathered information from Chinese and Western media sources, summarized on the next page.

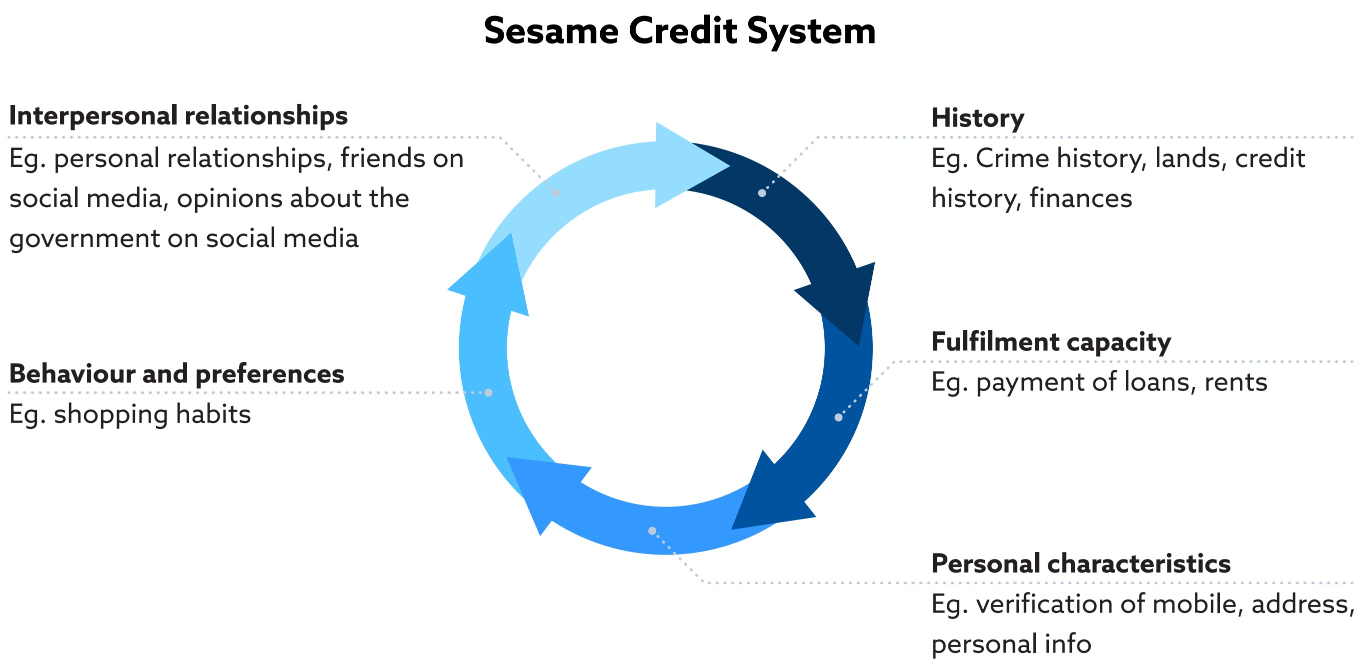

Sesame Credit System

Interpersonal relationships
Eg. personal relationships, friends on social media, opinions about the government on social media

History
Eg. Crime history, lands, credit history, finances

Behaviour and preferences
Eg. shopping habits

Fulfilment capacity
Eg. payment of loans, rents

Personal characteristics
Eg. verification of mobile, address, personal info

Figure 1: Sesame Credit System [5]

The diagram above reveals the five metrics in the Sesame Credit system: history, fulfillment capacity, personal characteristics, behavior and preferences, and interpersonal relationships. [26]

Could the U.S. adopt a scoring framework that merges the current FICO® system and the Chinese social credit system? An initial implementation could be launched as a cyber-centric pilot program, similar to a FICO® score. A FICO® score represents the probability a customer will default on a loan or other form of credit. [26] [27] The new scoring system might assign individuals a numeric value based on their cyber-attack susceptibility—making cyber trustworthiness equally as crucial as creditworthiness. Furthermore, individual cyber trust scores could be associated with a unique identification, such as social security numbers.

Consumers reap benefits from financial credit scores because they can grant access to financial instruments not otherwise available. High credit scores reward past positive behavior with lower rates on future loans or greater access to credit. However, identity theft, abuse, or creditor error can generate false and misleading scores. Individuals must keep track of all purchases, not spend outside their means, pay off balances in a timely fashion, and monitor accounts for fraud to prevent access roadblocks.

Cyber trust scores may infuse various elements (similar to the FICO® score) to calculate a result, including personal internet usage, individual propensity to interact with malicious sites or files, security hygiene of devices—and perhaps even the integrity of their blogs or level of social media interactions. Results would concurrently measure an individual's cyber trustworthiness and cyber risk.

While such a scoring solution may be technologically viable, gathering the necessary data causes privacy concerns. Should a person's online habits be monitored as a means of granting access to systems and data? Would access only be achievable by "opting-in" and keeping one's "surfing nose" and desktop "clean." Would the internet remain a place for discovery and free speech? Could the internet become a place of censure? Data collection solutions will drive answers to these inquiries.

Privacy Technologies for Trust Scoring

Although privacy is primarily based upon a legal construct, privacy-enhancing technologies (PETs) may also protect user identities. Systems could be developed to analyze data through artificial intelligence (AI) via pattern analysis, which would use a pattern-based query focused on a specific, uniquely identifiable individual (or individuals). [13] In 2017, Apple Inc. started a massive experiment with new privacy technology to build products that understand users without invasively monitoring their activities. Anonymizing technologies have also been considered, including "one-way hashing" concepts, to scramble data.

Differential privacy obscures data so that true user identities associated with data is unknowable. An interface system covers analyzed data by adding measurable amounts of statistical noise, making it sufficiently difficult for someone without proper authorization to connect specific data to a particular user. In this case, privacy issues may be avoidable. The following excerpt from The Journal of Big Data [12] explains how differential privacy works and how it would solve potential cyber trust score privacy issues.

Figure 2 shows the differential privacy mechanism. In this figure, the analyst sends a query to an intermediate piece of software, the Privacy Guard. The Privacy Guard assesses the privacy impact of the query using a special algorithm. Then, the database receives the Guard's query and obtains a clear answer based on data that has not been distorted in any way. The Guard then adds the appropriate amount of "noise" (scaled to the privacy impact) to ensure confidentiality for individuals whose information is in the database. Finally, the analyst receives a modified response. [12]

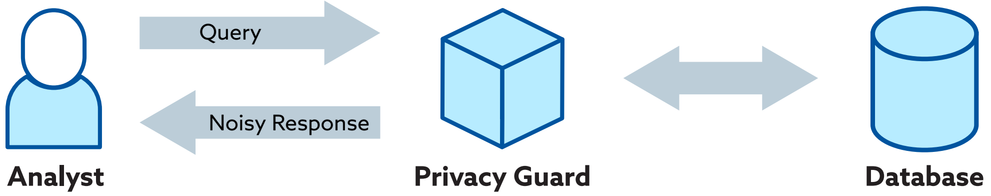

Figure 2: Differential Privacy Mechanisms

Additionally, immutable audits can keep logs of who accesses data and when data is accessed and altered. Audit records can also be protected from alteration.

However, differential privacy and AI-based PETs are not perfect solutions, as they don't have a legal construct for protection and can suffer from entropy balance and bias. Furthermore, privacy guards would be required in front of every transaction database and log store to solve all privacy concerns associated with cyber worthiness scoring. This scenario would necessitate negotiated agreements between system owners and scoring companies. Finally, this full framework would require policing—a daunting task that demands robust policy and process infrastructure to authorize privacy guards and assure fair access. [12]

3.0 Implications

Technology

The cybersecurity industry recognizes the need to make risk-based decisions for controlling access to IT resources, such as networks, applications, and data. For example, the National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207 discusses the development and use of a "Trust Algorithm" in (ZTA). [19] Technologies, such as those associated with software-defined perimeter (SDP) and conditional access control, can be enabled to grant access by devices and users to IT resources. [13] [15] Software-defined perimeter solutions may employ a policy decision point (PDP) or controller that may evolve to act as a trust broker that ingests various contextual data feeds. This process would generate a trust score for use in making fine-grained access control decisions. The MITRE Common Weakness Scoring System (CWSS™) may offer a means for computing trust scores for systems. [6]

According to research conducted by Tufts University and Mastercard in 2017 across 60 countries, people across the globe are simultaneously becoming dependent on—and distrustful of—technology. [7] This paradox is a global phenomenon and accurately describes the current landscape of cloud adoption and migration. In many ways, industries move to the cloud with good intentions only to discover gaps in business processes and skills that create security vulnerabilities. As cyber scoring approaches are devised, there must be a recognition of potential shortcomings that adoption may bring. An immediate concern relates to the algorithms necessary to make cyber trust scoring possible. Algorithms should not artificially hinder access. AI algorithms may have a bias because data used to train AI has a bias. The industry must develop a consensus on what information individuals monitoring and analyzing data can see and determine who monitors and aids machine learning algorithms to ensure objectivity, accuracy, and sensibility (if possible). For example, concerns have been raised about personal digital assistants (such as Amazon's Alexa or Google's Assistant) and the people who analyze user commands given to the machines. In theory, these commands make the devices more intelligent and well-cultured.

Social

In recent decades, the Chinese government determined the country's rapid shift from an agrarian to a postindustrial society necessitated an established framework for unknown parties to determine each other's trustworthiness. Hence, the social credit system was born. In China, the system encourages behavior that is deemed legal and acceptable by the Chinese government. The social credit systems give points for acts of kindness, such as making charitable donations or taking an elder to the doctor.

Unfortunately, China has an authoritarian history, and critics fear the country is attempting to use this system to maintain control and oppressively modify citizen behavior (which is uncomfortably reminiscent of George Orwell's novel *1984*, and "Nose Dive", an episode from Netflix's *Black Mirror* series).

For example, Chinese citizens are punished harshly for bad behavior the government discourages under the social credit system. [14] Punishments may include travel bans, public shaming, slower internet connectivity, and rejection from high visibility jobs and/or higher quality education.

Furthermore, social credit network participants may face penalties, reduced scores, and negative lifestyle impacts if others in their social network display improper behavior, are punished, or actively disregard government regulations. Perhaps most daunting: punishments occur outside of the legal system without a presumption of innocence, a jury, or a trial.

Like the Chinese social credit score, the cyber trust score system would generate a "blacklist" accessible to prospective private and governmental employers. This list would brand blacklisted individuals as ineligible from accepting employment or contracts related to national security or senior industry level positions to diminish risks of attack to vital economic and governmental entities. [14]

Despite this similarity to the Chinese system, the cyber trust score system could potentially impact society positively because individuals would be more mindful of how they interact with other people through connected devices. Theoretically, people would take more precautions when posting content online and more concerned about what would be tracked. As a result, cybercrime would have a less detrimental effect on the economy, society, and individual livelihoods because cybercriminals would be dissuaded from engaging in risky behavior.

Positive and negative personal scores may help citizens better understand how their cyber habits reflect on them as members of society and if they contribute to the common good.

One major hurdle for the cyber trust industry to overcome is determining how to encourage individuals to achieve and maintain a high cyber trust score.

The current financial credit score system penalizes those who do meet specific score requirements, making it increasingly difficult for such individuals to make large purchases, for example. Similarly, the cyber trust score system would discipline individuals who fail to meet a predefined baseline score. The U.S. will not be the sole arbiter of punishment for system participants, so it is up to the industry to determine punitive alternatives—other than forcing individuals off the grid—if the cyber trust score system is implemented.

Policy

Cyber trust score implementation in the United States would result in massive policy implications. Consider the credit score industry: currently, private companies unaffiliated with the government generate consumer credit scores. However, government agencies allow corporations to consider additional factors to augment credit scores when determining crediting decisions. For example, in 2019, the New York State Department of Financial Services announced that life insurance companies could base premiums on what they find in consumer social media posts. [7]

More examples: a company called Patronscan sells a kiosk, desktop, and handheld system designed to protect bar and restaurant owners from objectionable customers. Patronscan maintains a list of these problematic customers (such as people previously removed from other premises due to fighting, sexual assault, drugs, theft, and other malicious behavior) and produces a public list shared among all Patronscan customers. This functionality is similar to the blacklist the Chinese government utilizes and is often used by Patronscan customers for access control. [7] Airbnb, a Patronscan customer and a major travel accommodation provider, has more than six million listings in its system.

Patronscan bragged that a ban from Airbnb could limit an individual's travel options—an assertion backed up by Airbnb's policy that it can disable customer accounts for life for any reason it chooses. Furthermore, Airbnb reserves the right not to tell customers the reason for the deactivation. [7]

An Uber policy, announced in May 2019, enforced similar restrictions on customers using the service. Under the policy, customers with average ratings "significantly below average" may face a ban. Another popular application, WhatsApp, also bans users if it distrusts them. While the application has a relatively small footing in the U.S., it is an essential communication tool for many users worldwide. [7]

The cyber trust score system would likely be governed similarly to the examples cited—but perhaps with an appeal process. Individuals with low cyber trust scores could face bans from certain internet privileges.

If trends continue, most misdemeanors (and even some felonies) could face punishment by corporations instead of governments in the future. In essence, Silicon Valley—and not Washington, D.C.—would deliver justice. [9] This slippery slope from democracy toward corporatocracy provides clear-eyed reasoning for why checks and balances are needed to continuously validate and question consumer loyalties and dependencies on technologies—and firms that create them. The National Institute of Standards and Technology offers a framework in NISTIR 8149 for defining identity trust that may be valuable in this context where legality meets technology. [21]

Regulatory

Recently, there has been more pressure to tailor data protection laws to the innumerable circumstances in which they are applied. In response, more than 50 organizations have joined forces as a consortium to create the "Principles of Fair and Accurate Security Ratings." These principles inject transparency into the methodologies, appeal and dispute resolution processes, and advocate that ratings should be empirical, data-driven, or based on an expert opinion. [23] Furthermore, the principles address the significant change process, the nuances of commercial agreements, and how ratings will be appropriately protected. Finally, the principles urge companies to not publicize an individual organization's rating or provide third parties with sensitive or confidential information on rated organizations that could lead directly to system compromise.

The European Union General Data Protection Regulation (GDPR) provides an example of data protection regulation adapting to technology's challenges and giving data subjects increased control over personal data while allowing necessary flexibility in its implementation.

The bottom line is that individual cyber trust scores will raise privacy concerns and violations, regardless of any implemented system guardrails. The development of cyber trust ratings for individuals can become very complicated and uncomfortable because technology rating systems are typically invasive. There is always the risk that citizens will feel distrusted and uncomfortable under surveillance, thus diminishing their trust in government or industry.

4.0 Recommendation

Today, society is at a “trust crossroads” as traditional domain-based approaches are obviated by cloud adoption and (ZTA). For this new reality to succeed, cyber trust scores for individuals and entities may become the next measure for granting access to systems and data. Is it possible a whole new industry for measuring cyber trust could evolve? Would it be modeled after the credit scoring industry created by Equifax, Experian, and TransUnion?

Technologists are responsible for solving difficult problems. Cyber trust scoring could have significant social, policy, and regulatory implications, and today’s technology solutions may not align or cooperate with this potential shift. Monitoring user behavior raises many privacy concerns. Automation can cause bias, and differential privacy techniques require extensive inter-connectivity and complicated obfuscation algorithms. [12]

Accordingly, industry and government stakeholders should work diligently to address critical gaps before cyber trust scoring in America becomes a reality. Current industry collaborations should be recognized and built upon to include government initiatives and partnerships, including:

- National Strategy for Trusted Identities in Cyberspace (NSTIC) [25]
- Georgia Research Tech Institute NSTIC Trustmark Pilot
- European Commission Strategy on Trust Services and Electronic Identification (eID)

Further topic research should address the creation of a framework to provide industry guidance. The framework should cover privacy protection, input measures, effective computation, and implementation phasing over time. Industry and federal government participants should collaborate to this end.

5.0 Bibliography

[1]. Agre, Philip E., Rotenberg, Marc. “Technology and Privacy: The New Landscape.” https://pages.gseis.ucla.edu/faculty/agre/landscape.html 1997. MIT Press.

[2]. Brieger, Kelly. “Sift Science Launches Digital Trust Platform™–Holistic Approach Boosts Customer Loyalty and Retention.” https://www.econotimes.com/Sift-Science-Launches-Digital-Trust-Platform--Holistic-Approach-Boosts-Customer-Loyalty-and-Retention-982321 1 Nov. 2017. *EconoTimes*.

[3]. Britz, J. J. “TECHNOLOGY AS A THREAT TO PRIVACY: Ethical Challenges to the Information Profession.” http://web.simmons.edu/~chen/nit/NIT›96/96-025-Britz.html Department of Information Science University of Pretoria.

[4]. Broeckelmann, Robert. “Authentication vs. Federation vs. SSO.” https://medium.com/@robert.broeckelmann/authentication-vs-federation-vs-sso-9586b06b1380 24 Sept. 2017.

[5] Botsman, Rachel. “Who Can You Trust?: How Technology Brought Us Together–and Why It Could Drive Us Apart.” https://books.google.com/books/about/Who_Can_You_Trust.html?id=7cGWDgAAQBAJ 5 October 2017.

[6] Coley, Steven. “Common Weakness Scoring System (CWSS™).” https://cwe.mitre.org/cwss/cwss_v1.0.1.html 5 Sept. 2014. *Common Weakness Enumeration*.

[7] Chaturvedi, Shankar Ravi. Chakravorti, Bhaskar. “DIGITAL PLANET 2017: HOW COMPETITIVENESS AND TRUST IN DIGITAL ECONOMIES VARY ACROSS THE WORLD.” https://sites.tufts.edu/digitalplanet/files/2020/03/Digital_Planet_2017_FINAL.pdf July 2017. *The Fletcher School, Tufts University*.

[8]. Day, Jamison. “Cyber Threat Scoring Is Not Risk Assessment.” https://www.lookingglasscyber.com/blog/tech-corner/cyber-threat-scoring-not-risk-assessment/ 28 Nov. 2017. Looking Glass Cyber Solutions Inc.

[9]. Elgan, Mike “Uh-oh: Silicon Valley is building a Chinese-style social credit system.” https://www.fastcompany.com/90394048/uh-oh-silicon-valley-is-building-a-chinese-style-social-credit-system 26 Aug. 2019. *Fast Company*.

[10]. Huang, Geoff. “Sift Scores: Growing and Protecting Businesses.” https://blog.sift.com/2019/sift-score 1 Aug. 2019. *Sift Blog*.

[11]. IDG Communications. “2020 Cloud Computing Study.” https://www.idg.com/tools-for-marketers/2020-cloud-computing-study/ 8 June 2020. IDG.

[12] Jain, Priyank. Gyanchandani, Manasi. Khare, Nilay. “Differential privacy: its technological Prescriptive using big data.” https://journalofbigdata.springeropen.com/articles/10.1186/s40537-018-0124-9 13 Apr. 2018. *Journal of Big Data, SpringerOpen*.

[13]. Koilpillai, Juanita, Murray, Nya Allison. "Software Defined Perimeter (SDP) and Zero Trust." https://cloudsecurityalliance.org/artifacts/software-defined-perimeter-and-zero-trust/ 27 May 2020. *Cloud Security Alliance*.

[14]. Ma, Alexandra. "China has started ranking citizens with a creepy 'social credit' system—here's what you can do wrong, and the embarrassing, demeaning ways they can punish you." https://www.businessinsider.com/china-social-credit-system-punishments-and-rewards-explained-2018-4 29 Oct. 2018. *Business Insider*.

[15] Microsoft Inc. "What is Conditional Access?" https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/overview 16 Oct. 2020.

[16]. Mims, Christopher. "The Secret Trust Scores Companies Use to Judge Us All." https://www.wsj.com/articles/the-secret-trust-scores-companies-use-to-judge-us-all-11554523206?mod=hp_listb_pos1 6 Apr. 2019. *The Wall Street Journal*.

[17]. Reiter, Andrew. "THE ROLE OF TRUST IN TECHNOLOGY." https://www.shield.ai/content/2018/12/19/the-role-of-trust-in-technology 19 Dec. 2018. *Shield AI*.

[18]. Rosenzweig, Paul. Dempsey, James. "Technologies That Can Protect Privacy as Information Is Shared to Combat Terrorism." https://www.heritage.org/homeland-security/report/technologies-can-protect-privacy-information-shared-combat-terrorism 26 May 2004. *The Heritage Foundation*.

[19]. Rose, Scott. Borchert, Oliver. Mitchell, Stu. Connelly, Sean. "Zero Trust Architecture." https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf Aug. 2020. *NIST Special Publication 800-207*.

[20] Rouse, Margaret. "What is identity and access management? Guide to IAM." https://searchsecurity.techtarget.com/definition/identity-access-management-IAM-system#:~:text=Identity%20and%20access%20management%20(IAM,critical%20information%20within%20their%20organizations Search Security Tech Target.

[21] Temoshok, David. Abruzzi, Christine. "Developing Trust Frameworks to Support Identity Federations." https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8149.pdf Jan. 2018. National Institute of Standards and Technology.

[22]. "Transitive Trust." https://www.ibm.com/support/knowledgecenter/SSLTBW_2.4.0/com.ibm.zos.v2r4.euvfa00/euv2b3_Transitive_trust.htm *IBM Knowledge Center.*

[23]. "Principles of Fair and Accurate Security Ratings." https://www.uschamber.com/issue-brief/principles-fair-and-accurate-security-ratings 20 June 2017. U.S. Chamber of Commerce.

[24]. Saltonstall, Max. "Google BeyondCorp with Max Saltonstall." https://softwareengineeringdaily.com/2018/02/09/google-beyondcorp-with-max-saltonstall/ 9 Feb. 2018. *Software Engineering Daily*.

[25]. "GTRI NSTIC Trustmark Pilot." https://trustmark.gtri.gatech.edu/ The National Institute of Standards and Technology.

[26]. Waterman, Shaun. "Insurance regulators pitched on FICO-style score for cybersecurity." https://www.cyberscoop.com/insurance-regulators-pitched-fico-style-score-cybersecurity/
7 Aug. 2017. *CyberScoop*.

[27]. "FICO Offers Free Cybersecurity Ratings to Companies Worldwide." https://www.fico.com/en/newsroom/fico-offers-free-cybersecurity-ratings-to-companies-worldwide 27 June 2018. *FICO*.

[28]. Webster, Teri. "Apple quietly implements 'trust scores' based on users' phone data." https://www.theblaze.com/news/2018/09/21/apple-quietly-implements-trust-scores-based-on-users-phone-data 21 Sept. 2018. *the Blaze*.

[29]. de Froidmont, Cassandre. Donati, Ludovica. Steffen, Miriam. Zaidi, Aiman. "China's Social Credit System–Analysis of a Socio-technical Controversy." https://mastertsinghua.wordpress.com/2018/04/30/chinas-social-credit-system-analysis-of-a-socio-technical-controversy-by-cassandre-de-froidmont-ludovica-donati-miriam-steffen-and-aiman-zaidi/ 30 Apr. 2018. *Master Tsinghua Data Science and Expertise*.