CUTTER Business Technology Journal

Management, Innovation, Transformation

Vol. 33, No. 5, 2020

'TRUST'

Is IT the Problem or the Solution?

Claude Baudoin
Guest Editor

CUTTER Business Technology Journal

As business models for creating value continue to shift, new business strategies are constantly emerging and digital innovation has become an ongoing imperative. *Cutter Business Technology Journal* delivers a comprehensive treatment of these strategies to help your organization address and capitalize on the opportunities of this digital age.

Cutter Business Technology Journal is unlike academic journals. Each monthly issue, led by an expert Guest Editor, includes five to seven substantial articles, case studies, research findings, and/or experience-based opinion pieces that provide innovative ideas and solutions to the challenges business technology professionals face right now — and prepares them for those they might face tomorrow. *Cutter Business Technology Journal* doesn't water down or delay its content with lengthy peer reviews. Written by internationally known thought leaders, academics, and practitioners — you can be certain you're getting the uncensored perspectives of global experts.

You'll benefit from strategic insight on how the latest movements in digital innovation and transformation, artificial intelligence/machine learning, Internet of Things, blockchain, analytics, and cloud, to name a few, are changing the business landscape for both new and established organizations and how cutting-edge approaches in technology leadership, enterprise agility, software engineering, and business architecture can help your organization optimize its performance and transition to these new business models.

As a subscriber, you'll also receive the *Cutter Business Technology Advisor* — a weekly bulletin featuring industry updates delivered straight to your inbox. Armed with expert insight, data, and advice, you'll be able to leverage the latest business management thinking to achieve your organization's goals.

No other journal brings together so many thought leaders or lets them speak so bluntly — bringing you frank, honest accounts of what works, what doesn't, and why. Subscribers have even referred to *Cutter Business Technology Journal* as a consultancy in print and likened each month's issue to the impassioned discussions they participate in at the end of a day at a conference!

Get the best in thought leadership and keep pace with the technologies and business models that will give you a competitive edge — subscribe to *Cutter Business Technology Journal* today!

Founding Editor: Ed Yourdon
Publisher: Karen Fine Coburn
Group Publisher: Christine Generali
Managing Editor: Cindy Swain
Copy Editors: Jennifer Flaxman, Tara K. Meads
Production Editor: Linda Dias
Client Services: service@cutter.com

***Cutter Business Technology Journal®** is published monthly by Cutter Consortium, an Arthur D. Little company, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA • Tel: +1 781 648 8700 • Fax: +1 781 648 8707 • Email: cbtjeditorial@cutter.com • Website: www.cutter.com • Twitter: @cuttertweets • Facebook: Cutter Consortium. ISSN: 2475-3718 (print); 2475-3742 (online).*

Subscription rates are US $485 a year in North America, US $585 elsewhere, payable to Cutter Consortium. Reprints, bulk purchases, past issues, and multiple subscription and site license rates are available on request.

☐ Start my print subscription to *Cutter Business Technology Journal* ($485/year; US $585 outside North America).

Name _____ Title _____

Company Address _____

City _____ State/Province _____ ZIP/Postal Code _____

Email (Be sure to include for weekly *Cutter Business Technology Advisor*) _____

Fax to +1 781 648 8707, call +1 781 648 8700, or send email to service@cutter.com.
Mail to Cutter Consortium, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA.

Request Online License Subscription Rates

For subscription rates for online licenses, email or call: **sales@cutter.com** or **+1 781 648 8700.**

CUTTER CONSORTIUM
Access to the Experts

Opening Statement

by Claude Baudoin, Guest Editor

We are living though complex and worrying times. Much of the world, as I write this, is dealing with a pandemic that has upset our daily lives in a way that is unparalleled in recent history, except in countries that were literally, not just metaphorically, at war or undergoing a revolution. At the same time, old demons of our societies, never completely slayed, have reemerged to haunt us: populism, racism, intolerance….

As we fight through this and eventually come out on the other side to whatever will be the "new normal," a word has increasingly made its way into the daily discourse of business and technology leaders, as well as of politicians (at least those who know how to spell it): *trustworthiness*. Some things are noticed only when we miss them; trustworthiness is one of those. It is interesting to delve into how we got to this situation; the role that information technology has been playing in the erosion of trustworthiness; and how it, like many double-edged tools, might help solve the very problems it has helped create.

Our notion of trustworthiness has been evolving for several centuries. This evolution has accelerated recently, and this is clearly attributable to technology. "Trust" initially referred only to trust in people. Then we had to trust certain institutions — rulers, governments, banks — to do what they had promised to do for us. What is paper money if not a testament to the trustworthiness of a central bank? The technology of the First Industrial Revolution added to that meaning, when the trustworthiness of the machine designer and supplier became important. Still, in a world of tangible objects, inspecting a machine before putting it into service was relatively easy. When the automobile appeared, it was a whimsical, inconsistent machine whose driver needed to have some of the skills of a mechanic to be able to trust that driver and machine would reach the intended destination. Yet although an untrustworthy car could leave you stranded by the roadside, it was unlikely to start a conflict or change the result of an election.

Interestingly, early forms of communication were also fraught with trustworthiness issues. Much of that was rooted in the ambiguities of human language, and the multiplicity of them. The trustworthiness of a translation was a huge issue, and at least one war started in that manner.[1]

Enter IT. For several decades, all was well, because IT did not touch us personally all that much. It was easy enough to verify that our computer-generated paycheck showed the right numbers, and that was about it. Isaac Asimov certainly brought up the trustworthiness of robots in his writings and, in 1942, even invented laws about robotics. But despite the philosophical import of his prose, his ideas remained simply fiction. It is only when email, the World Wide Web, and the computerization and automation of a growing number of activities became part of our lives that the issue of whom and what we can trust exploded. That shift into uncertainty has only grown since then.

To a large extent, we naturally mistrust what we do not understand. The fact that software developers use voluminous and arcane code to automate things — and that even other professionals have a hard time deciphering what the code means and validating that it cannot produce ill effects — is enough to explain the loss of trustworthiness. This is well illustrated by the old joke of the real-time software engineers who, upon boarding an early Airbus plane, hear the captain announce, "Today, this will be an entirely fly-by-wire experience." Upon which they hastily disembark. So even without considering any malicious intent, complex systems already stretch our ability to trust them. Is the code bug-free? Certainly, we know that beyond a few hundred or perhaps a few thousand lines, no code is completely bug-free and that real-world systems contain millions of lines. Has the supplier performed all the necessary tests? This is an almost impossible feat, as so many combinations of conditions would have to be tested, and the requirements against which a system is tested can be riddled with ambiguity. Was this measurement supposed to be in imperial or metric units (as

was the cause of the Mars Climate Orbiter failure[2])? Losing a spacecraft is one thing, but what if an entire country was plunged into a blackout because of untrustworthy data in the control systems of the national electric grid?

In the 20 years since the Mars Orbiter crashed, we have become, individually and collectively, dependent on information systems at a much deeper level. Instead of a few reasonably balanced news channels, we now have a choice of hundreds of electronic "echo chambers" in which we only receive news selected according to the biases of like-minded people. We do not know how many of the posts we read on Facebook or Twitter were generated by Russian bots. The majority of us who live in democratic countries are no longer sure that the often abysmally insecure technology of electronic voting machines has not been subverted. We know that this stash of money from an unknown Nigerian official is fake, but what about this offer for a shopping coupon at my pharmacy? Did this politician really say these things, or is it a deep fake? And speaking of "fake," how do I know what is real or not, now that some people call "fake news" anything that does not agree with their opinions, while you can almost be sure that what they praise as newsworthy (mostly because it flatters them) is almost surely false?

Artificial intelligence is now adding yet another twist to this story: How do we know why a neural network denied a loan application or confused the face of a person with that of a terrorist? Do we know whether such misidentification occurs at the same frequency for people of different ethnicities? Who wrote this software, and which data set did they use to train it? Is placing a human in the loop likely to improve or degrade the trustworthiness of the system? A loan officer might look at a strange rejection recommendation, question it, and redo some calculations by hand, but he might also be prejudiced against certain applicants. A human driver might override the controls of an autonomous car that is going to run over a misclassified pedestrian, but we also know that many airplane crashes were caused by pilots ignoring the warnings of their cockpit instruments. Who (or what) should be trusted more?

I addressed some of the above points at greater length a few months ago in a *Cutter Business Technology Journal* (*CBTJ*) article I called "Trustworthiness: A Mouthful That Shouldn't Leave a Bad Taste."[3] But it seemed too important a subject, in these times of uncertainty and confusion, to leave it at that without seeking the opinion of a broader panel of experts. Hence this issue of *CBTJ*, for which we asked the question: "Is IT the problem or the solution?" In other words, while IT has created the conditions, the products, and the insecure protocols that permit the problems listed above, can IT also be used to counter these threats? For example:

- Internet protocols and telephone caller ID could be updated to prevent spoofing.
- Voting systems could be designed with redundancies and paper trails to permit verification of the counts, without creating a risk of vote buying (which is one of the unintended consequences of primitive paper trail systems).
- "Provenance and pedigree" standards might be used to create a tamper-proof trail of where news items, photographs, videos, data sets, or software come from.
- Internet of Things (IoT) sensors could be required to "sign in" to the networks using cryptography methods and to only send data in encrypted form to protect critical infrastructure and assets from industrial spying or hacker attacks.

Clearly, technology solutions are not the only things we need to restore the needed sense of trust in systems and information, especially when it comes to news and social media. At a minimum, it seems that we require a new regime of checks and balances that covers a whole range of qualities, such as *reliability, resiliency, visibility of provenance, safety, security, privacy, and absence of bias*. These checks and balances must be put in place by well-trained and ethical-minded humans, working for organizations that respect and protect their independence. The all-too-common practice of dismissing the warnings of someone who says "we haven't tested this system enough, so we shouldn't release it" should be

Upcoming Topics

Beyond Automation: AI, ML, RPA
San Murugesan

How Technology and Business Leaders Help in a Crisis
Steve Andriole

Proactive Risk Management
Tom Teixeira

banned. We need the guarantors of correct systems design to attain the same level of training and professional certification as other experts, to maintain the same unblemished records and accountability as other esteemed professionals, and to hold the same exalted level in our society as judges hold — and usually deserve.

In This Issue

This issue's contributors have addressed the question of trustworthiness from a variety of angles. Each article offers a significant contribution to the challenge of restoring and maintaining trust.

In our first article, Cutter Consortium Senior Consultant and frequent contributor Paul Clermont uses his well-known "straight talking" style to paint a clear picture of the "broad scope of threats" we are facing. He uses anthropological analogies to explain the "circles of trust" we use in deciding what to believe. Clermont doesn't shy from the potential conflicts among information transparency, privacy, and intellectual property. He then looks at the proper role of governments in creating the frameworks and standards that can help improve trustworthiness.

Next, Philippe Flichy tells us that there are three complementary facets we need to consider, particularly in an industrial environment: trusting the data, a challenge made more difficult by the emergence of IoT, digital transformation, and cyberattacks; trusting the tools, for example, the machine learning algorithms whose innards are, almost by design, largely inscrutable; and trusting the people, given the pandemic-era new work practices.

David Tayouri then brings us the perspective of the Israeli defense environment, justly famous for its leadership in cybersecurity. For Tayouri, the combination of biometrics, asymmetric cryptography (think "PKI"), and blockchain can help construct a strong authentication and authorization environment, which is crucial to, in his words, "reconstruct virtual trust."

Following along the same "technology as the solution" line of thought but with the added twist of putting a human in the loop, a team of eight coauthors led by Greek academic Panagiotis Monachelis proposes to combine peer-to-peer decentralized networks and blockchain technology to address the challenge of misinformation in social media. The authors provide a detailed description of an architecture, embodied in their research project called EUNOMIA, that allows end users to review posts and feed a secure voting system.

Finally, Robert A. Martin addresses in the last article the complete ecosystem involved in the procurement of products and services. What does it mean to trust that what you buy, and the organizations that sell to you, meet all the conditions required to merit your trust? Martin describes the elements of a system of trust for supply chain security that is currently under development and is based on collecting information from a wide community of procurement departments and standards organizations.

Even if all the ideas presented by this issue's authors are implemented, serious challenges will remain. One is the tension between trust and anonymity, when the latter is required; in particular, to protect whistleblowers or opponents of authoritarian regimes. The other is the fact that society and its actors (politicians, media, product or service suppliers, and consumers) do not change as quickly as the technology. Levels of trust that have been destroyed in just a few years may take decades to rebuild. But we can be thankful to our authors for pointing us toward several useful building blocks of the solution.

References

[1]The "Ems Dispatch" started the Franco-Prussian War of 1870-1871, at least in part, because the translation of the German word "Adjutant" into the French "adjudant" reinforced a sense of insult among the French government and populace; see Wikipedia's "Ems Dispatch."

[2]See Wikipedia's "Mars Climate Orbiter: Cause of Failure."

[3]Baudoin, Claude. "Trustworthiness: A Mouthful That Shouldn't Leave a Bad Taste." *Cutter Business Technology Journal*, Vol. 33, No. 1, 2020.

Claude Baudoin is a Senior Consultant with Cutter Consortium's Business & Enterprise Architecture and Data Analytics & Digital Technologies practices. He has been owner and Principal Consultant of cébé IT and Knowledge Management, a boutique consulting practice, since 2009. Prior to that, Mr. Baudoin spent 35 years in IT and software management roles in industry, including 26 years at Schlumberger, the global oilfield services company, in France and the US. Since 2015, he has served as an energy domain advisor to the OMG and the Industrial Internet Consortium (IIC) and has sat on the steering committee of the Cloud Standards Customer Council. Mr. Baudoin has published two books on software engineering, Méthodes de Programmation *and* Realizing the Object-Oriented Lifecycle, *and numerous papers and conference presentations. He holds two patents related to IT infrastructure and security. Mr. Baudoin's undergraduate engineering degree is from École Polytechnique in Paris, and he holds a master of science degree in computer science from Stanford University. He can be reached at cbaudoin@cutter.com.*

REACHING FOR NEW HEIGHTS

Trust Our Technology? Hmmm. Let's Think.

by Paul Clermont

Gun rights activists like to tell us that "guns don't kill people; people kill people." They're right, but only in a literal — and trivial — sense. If nobody was motivated to kill, nobody would be murdered, no matter how many guns were out there, but real people do have motives, and guns provide a far more easy and convenient means to do the job than knives or blunt instruments. Likewise, we could say that computers don't cheat people or robots don't kill or maim them, and we would be literally, but still meaninglessly, correct.

For IT to be trustworthy, we must be able to trust that designers are:

- Competent to ensure the technology works as legitimately intended
- Diligent in minimizing opportunities for hackers' intrusion

Circles of Trust

First, a couple of definitions from *Merriam-Webster*:

1. **Trust** — "assured reliance on the character, ability, strength, or truth of someone or something"
2. **Trustworthy** — "worthy of confidence, dependable"

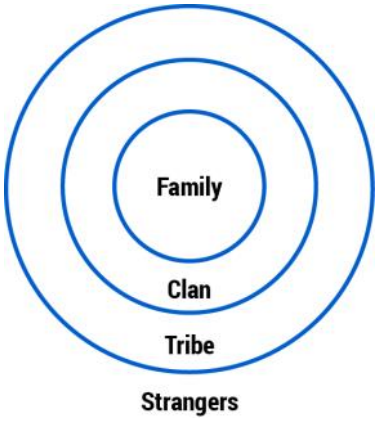

Figure 1 — Prehistoric circles of trust.

Trusting People

Trust is an essential part of being human. Without the ability to trust, our lives would be far more nasty, brutish, and short than they ever were back in the Stone Age. The challenge facing humans is to assess correctly the trustworthiness of people and the institutions with which we deal. Fortunately, the ability to make some relatively safe assumptions has helped us (see Figure 1).

It's certainly not foolproof, but our family is, on average, more trustworthy than our clan, which is more trustworthy than the rest of our tribe, which is more trustworthy than everyone else. Over the millennia, those who trusted too much in the benign intentions of unknown people tended to disappear from the gene pool; thus, while xenophobia is neither good nor nice, it has been baked into our DNA.

Trusting Institutions

As our species moved on from caves, we formed various "institutions" upon which we depended for transactions — exchanging goods and services — and information. These institutions needed to build and maintain a reputation for trustworthiness, or they would eventually fail. To do this, they had to devise and institute various controls to ensure that they were not (detectably!) cheating their customers or being cheated by customers, suppliers, or employees.[1] As the scope of institutions increased, these controls had to become ever-more sophisticated, the circle of trust expanding to include those who designed and administered the controls that prevent and detect cheating.[2] The banking system is a prime example of institutions built on trust (see sidebar "First Bank and Trust").

The advent of automation required another circle of trust: the people designing and building the *computer systems* that kept the books and implemented the controls (see Figure 2.) The resulting standardization of methods and calculations greatly reduced the incidence of random errors that are inevitable when data entry and calculation are manual. Automation enabled the handling of far more complexity and options. This

progression has led us to today's world in which we can trust a company such as Fidelity Investments with our complete personal portfolios, including retirement funds often worth millions, with no more evidence of our portable wealth than a monthly statement received online. Viewed in the perspective of centuries, that's pretty impressive! It's a major step beyond when wealth was evidenced by paper bonds or stock certificates, which themselves required a level of trust far beyond the days when only precious metals or jewels were good enough to warrant trust.

Today, with artificial intelligence (AI), we need to be able to trust the integrity and competence of IT designers and builders even more. It's not merely just about telling the computer correctly what to do but rather telling it how to learn to recognize and correctly handle new situations, including detecting suspect and spurious input. For example, in a recent experiment, a slightly modified 35-mph-speed-limit sign (see Figure 3)[3] was read by a passing Tesla as 85, so the car started to accelerate accordingly, lacking the context to recognize something as fake — something that would never have fooled any sentient human.

Traditional computing may seem like black boxes to nontechnical users (as it, in some cases, should be to minimize gaming the system), but its creators understand it. AI algorithms that learn are "black boxes on steroids"; they may learn wrongly and it's up to their creators to minimize this scenario. For example, some automotive GPSs know the speed limit based on car location, enabling a cross-check of what the AI sees (or thinks it sees) on a sign.

Trusting Information

Building the trustworthiness of information (other than that for our accounts) has not been as advanced as Fidelity's customer accounting. That's not a condemnation; it's a recognition that information is inherently squishy. Facts are often hard to discern clearly, and their interpretation is subject to human limitations like bias or lack of experience-based intuition to spot something fishy. There is abundant room for selective fact gathering and tendentious presentation, and there's no arithmetic audit trail to reality as there is (or should be) in our accounts. While we are long past the era of taking claims of miracles and dragon sightings at face value, not even icons like *The New York Times* or the BBC are immune to mistaken reporting and interpretation. Obviously, there is a spectrum of trustworthiness; a story in *The New York Times* deserves and receives

First Bank and Trust

We deposit our money, trusting that we'll be able to withdraw it whenever we want, which requires that the bank both maintain a reserve and make loans to trustworthy borrowers who will almost always repay them. Rumors of a troubled loan portfolio can provoke a "run on the bank" in which most depositors end up stiffed, amplifying the economic downturn that compromised the portfolio in the first place, as during the Great Depression. One New Deal reform was *deposit insurance*, which banks bought from the government and were then subject to periodic inspections to ensure prudent practices. In other words, your trust moved from the individual bank to the government — a far safer bet — and the approach has worked well for almost 90 years. Unfortunately, a parallel deposit insurance scheme for US savings and loan associations was not up to monitoring their charges' use (and misuse) of vast new powers granted them around 1980. Failures were legion, and by 1990, the insurance fund ran out of money and taxpayers ended up bailing it out — not a great regulatory success, but it maintained the trust of millions of depositors in the overall system.

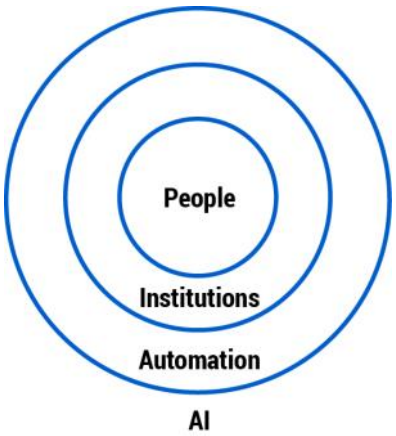

Figure 2 — 21st-century circles of trust.

Figure 3 — When is a 3 not a 3? (Source: Whitwam.)

infinitely more credibility than one in a supermarket tabloid.

Unfortunately, we seem to be going backward in this area in recent years, and technology is partly to blame. The Internet and social networks have given powerful megaphones to cranks, conspiracy theorists, and certifiable lunatics who flood cyberspace with junk information (e.g., the Pope endorsing Donald Trump for president[4]) designed to misinform or otherwise sow fear, uncertainty, and doubt on behalf of some institution, ideology, or mischief maker. Journalists have curatorial power but working to deadline (and not wanting to get scooped) imposes practical limits. Publications that cover current events from a bit of distance (e.g., *The New Yorker*) employ fact checkers for every assertion, noting when the person or institution being checked has not responded.

The Internet and social networks have given powerful megaphones to cranks, conspiracy theorists, and certifiable lunatics who flood cyberspace with junk information.

Broad Scope of Threats: Why We Cannot Assume Trustworthiness

This section addresses the myriad reasons that things can go wrong in such a way as to erode the general public's trust in technology, its purveyors, and the organizations that use it.

At the Public Level

Threats from institutions and mischief makers include:

- **Damage to critical infrastructure.** For example, our dependence on the power grid is fundamental. When it's controlled by a network of computers connected to the Internet, as it needs to be to assure systemic resilience, a hacker or saboteur, even a continent away can damage it, as actually happened in Ukraine a few years ago.[5]
- **Theft of personal data.** There is a constant stream of news (e.g., Experian, Starwood Hotels, the US Office of Personnel Management) in which our personal data has been accessed and stolen.
- **Breach of privacy.** What we write on paper can be burned or shredded, but every keystroke is potentially discoverable.
- **Hacked voting systems.** Voting machines connected online with no consistent security standards and no paper audit trail for recounts can be subtly but effectively hacked, not just to throw an election but to destroy confidence in the whole voting process.
- **Defective products.** The Boeing 737 MAX 8 has been grounded for a year and counting due to fatal crashes caused by a computer control that didn't cross-check a spurious input or allow the flight crew to override the automatic action that misinformation caused.[6]

Threats from governments include:

- Fudged statistics leading to bad macroeconomic decisions by central banks and microeconomic decisions by companies and investors
- Fudged census data leading to misallocation of public resources
- Voting processes and technology vulnerable to manipulation from inside the government
- Unethical algorithms, whether intentional or accidental, defeating the intent of laws requiring equal rights or access
- Insufficient regulation to ensure safety in use of automation and AI (e.g., the regulatory failure with the Boeing 737 MAX 8 that has tarnished the US Federal Aviation Administration's worldwide reputation as the gold standard of air safety assurance)
- Thwarting of privacy laws, as agencies of governments have done in the name of national security

At the Enterprise Level

External threats include:

- Measurable losses from hacker-enabled theft of goods or services by "customers"
- Less measurable losses[7] from hacker-enabled theft of proprietary and customer data[8] and intellectual property (IP)
- Impersonation by scammers who damage the enterprise's credibility

- Malware that affects customers (e.g., a distributed denial of service bringing down operations)

Internal threats include:

- Dishonest and disgruntled employees, especially those who know how internal systems work
- Incompetent employees designing and managing controls and building their associated IT in such a way that it's too easy for threats to materialize and go undetected for too long
- Error-prone or easily fudged data-capture procedures

Beneficial "threats" include:

- Hackers unmotivated by personal gain or desire to create mischief who break in as a challenging hobby to reveal previously unknown vulnerabilities

At the Personal Level

In general, personal threats include:

- Being scammed by impersonators of trusted individuals or organizations
- Being impersonated by scammers or other mischief makers to victimize people who trust you
- Having identity and personal data stolen
- Having your devices invaded by malware that degrades performance, corrupts applications or data, or demands payment of a ransom to return to the status quo

When procuring goods and services, threats include:

- Misleading information
- Important information made deliberately hard to find
- Error-prone processes
- Poorly designed record keeping
- Lack of audit trails to sort out problems
- Unethical algorithms that reflect biases or unfair and shoddy but not technically illegal business practices
- Voluminous "Terms & Conditions" seemingly designed to not be read, with the intent of evading responsibility for problems

When obtaining general info, threats include:

- Deliberately false or misleading info (e.g., deep fake photos and sound recordings)
- Insufficient/nonexistent curation and fact checking of information available online

What We Should Know & Do

Guiding Premises

- Lost trustworthiness is very difficult to regain; it can take years, if ever.
- Absolute trustworthiness is an unattainable goal, just like eliminating crime and accidents, yet we must maximize our diligence in getting as close as possible.
- Cutting corners will breed fear and cynicism about the whole computer-based world and jeopardize the good things accomplished to date.

Technical Challenges

Cybercrime, like all crime, requires a potential perpetrator (*who*), the confluence of motive (*why*), the means (*what* and *how*), and the opportunity (*when* and *where*). There will always be people with motives to commit crime, so technologists must address the following:

- **Prevention** — limiting means and opportunity:
  - In IT, minimizing vulnerabilities to intrusion
  - Minimizing risk from adding complexity ("cool" features) and interconnectedness simply because we can, without a compelling purpose[9]
  - Ensuring that purchased software components, particularly open source, are traceable and have been thoroughly scanned and tested for vulnerabilities[10]
- **Detection** — having the logic to recognize something wrong and deal with it:
  - Catching intrusions and anomalies, including potentially spurious inputs and deep fakes
  - Isolating and containing the damage
  - Ensuring that machine learning "training" does not introduce unintentional bias into algorithms

(e.g., the way a too limited spectrum of inputs has resulted in face recognition whose effectiveness varies by race)

  - Checking algorithms in use to uncover "adverse effects" that may have materialized despite best efforts to avoid them

- **Solution** — identifying and neutralizing the perpetrator(s):

  - Audit trails to ID the source of intrusion

Continuous learning about how better to anticipate, prevent, detect, and solve cybercrimes is a necessity — "Fool me once, shame on you; Fool me twice, shame on me."

Audit trail techniques that ferret out the source of intentional misinformation can also lead to the outing of valuable whistleblowers, exposing them to possibly lethal reprisals.

Conundrums

Efforts to ensure trustworthiness can lead to situations where worthy goals conflict, as in the following examples.

Audit Trails, Good and Bad

Audit trail techniques that ferret out the source of intentional misinformation can also lead to the outing of valuable whistleblowers, exposing them to possibly lethal reprisals. Perhaps a specialized warrant could be designed that would allow governments to penetrate walls of secrecy only when there is a specific crime to be investigated (i.e., habeas corpus). Without having such a warrant presented, ISPs and others in the chain could not be required to comply. (This might have solved the standoff between the US government and Apple over breaking into the iPhone of one of the San Bernardino, California, terrorists.)

IP: Transparency vs. Proprietary Rights

On the one hand, it seems obvious that a "black box" defense against a claim of a wrongful decision made by a purchased algorithm should not be allowed. If one uses a proprietary algorithm, one should be responsible for what it does, unless it did something the vendor said explicitly it would not do. That means transparency about how it works and why it works that way. Someone in the organization (or a very limited group[11]) that uses the algorithm needs to be cognizant of how it works and sophisticated enough to envision how it could make unacceptable mistakes that would require modification or tuning. A purchased algorithm that can't be modified or tuned is a pig in a poke — never a safe choice.

On the other hand, developers of sophisticated algorithms should not have to fear that their great ideas can be easily copied. Some specialized form of nondisclosure agreement (NDA) may be required for cognizant employees of the user. Alternatively, public or private entities with appropriate NDAs could certify an algorithm, taking responsibility if they missed the possibility of something the user of the algorithm explicitly didn't want to have happen or a violation of generally accepted ethical rules for algorithms.

Of course, computer scientists can correctly argue that explainable AI is not as powerful as AI that is not constrained by a need to explain itself, but that may be more power than we should allow it when the AI is applied to decisions that can affect whole lives. It's hard to imagine trustworthiness taking a worse hit than a news story about inexplicable actions that harmed real people and that were not the result of human error.

The Role of Governments

There is a rapidly emerging consensus, even among some big names in tech (e.g., Tim Cook of Apple, Satya Nadella of Microsoft,[12] Eric Schmidt, formerly of Google), that the libertarian approach to tech has run its course and become counterproductive as more people lose trust. Because it has shown no serious propensity to self-regulate, as some once naively hoped, tech needs help from governments to save itself. This help can take a number of forms.

Antitrust

Companies can be broken up, horizontally or vertically, to increase competition. The recent antitrust "doctrine" that government action is justified only by evidence of consumers having to overpay may be relevant for commodity industries, which tech, other than basic devices (e.g., PCs, mobile phones), surely isn't. However, that doctrine is questionable when it comes to

amassing market power (e.g., Facebook and Google in online advertising) or dominating a tech area in order to steer innovation in a desired direction (e.g., big companies gobbling up startups for insane prices, possibly to bury an innovative technical approach before a competitor buys it).

Privacy Guarantees

Europe is leading the way with its General Data Protection Regulation (GDPR), hardly a surprise given that tens of millions of Europeans remember life under intrusive and unaccountable governments.

Legal Standards and Certification

Part of being trustworthy is being able to boast of some kind of certification by impartial experts that algorithms and the practices around their use meet standards, implying the need to create standards. This is not just a matter for technicians alone; there are roles for experts in the business, as well as lawyers, ethicists, and consumer advocates. We need standards involving the following:

- Ensuring algorithms that "make" critical life-changing decisions (as in criminal justice, home and small business loans, admission to higher education, hiring into tenure/partner track positions) are subject to external review to certify lack of bias and the boundaries defining close situations that require human intervention to resolve
- Record-keeping standards for algorithm usage and how close decisions were resolved
- Professional certifications for experts in cybersecurity and ethical algorithms

Standards for Voting Technology

Paper ballots, once the norm, could be counted and recounted. Keeping the process honest was a matter of standard procedures and physical security. Mechanical voting machines introduced the possibility of physical malfunction. Computer-based machines added the possibility of hacking within the precinct. Connection to the Internet added the possibility of external hacking. Meanwhile, at least in the US, responsibility for voting technology is devolved to the state or even local level, as in the 2000 presidential election Florida fiasco.[13] Confidence in the integrity and auditability of the voting process lies at the heart of a functioning democracy. It's too important to leave to the unguided discretion of often unsophisticated buyers of technology, and self-certification by vendors is essentially meaningless, as the Boeing 737 MAX 8 situation shows.

Fact-Checking Requirements

The 1964 *New York Times Co. v. Sullivan*[14] court decision regarding a libel suit used the phrase "wanton disregard for the truth," which describes a lot of what's on the Internet. Note that the defendant was *The New York Times,* not the reporter in its employ, implying that the enterprise distributing the information bears responsibility for what it distributes. Operators of social networks have resisted taking responsibility for the veracity of the often-questionable information they put out. Curating and policing it may be expensive and difficult, but it would not be unreasonable to demand it by law, given the extraordinary profit margins in the business. In fact, it is required by law in much of Europe, where platform operators must proactively review what's posted in order to filter out noxious material, such as neo-Nazi messages.

A High-Stakes Challenge

Since the dawn of the computer age, investments in automation have been enormous, and society has benefited tremendously in terms of product and service quality at acceptable cost and with a level of convenience in the business of daily life that could not have been imagined even as recently as the early 1990s. The expanding circles of trust — justified and largely well maintained — have been as critical as the technology that enabled these good things to happen. If that trust is lost or even notably chipped away, it cannot be recovered simply by upgrading to a new generation of technology.

When it comes to accepting products and services, people's feelings and emotions matter. So far, tech has profited immensely from generally positive attitudes on the part of the public, but clear signs of skepticism and cynicism have appeared as a result of eroding trust. That trust must be restored and maintained before a tipping point is reached. There's far too much on the table to risk in the name of quarterly profits, libertarian philosophies, breaking things as part of moving fast, or oversize egos.

Most people are reasonable enough to understand that absolute trustworthiness will not and cannot ever be

guaranteed, but that does not reduce the imperative of constantly driving toward that goal. Falling too far short breeds disgruntlement and cynicism about any institution and leader, and much of the world is heading in that direction to the detriment of all. How far is "too far"? We should not want to find that out empirically!

Statements of principle and intention are fine, but they must be accompanied by visible good faith efforts made credible by accepting delays or earnings hits in pursuit of getting it right (i.e., putting one's money where one's mouth is).[15]

References

[1]Some of the oldest specimens of written language, when translated, turned out to be records of transactions.

[2]Double-entry bookkeeping, for instance, was invented in 14th-century Florence, Italy.

[3]Whitwam, Ryan. "Researchers Tape Speed Limit Sign to Make Teslas Accelerate to 85 MPH." ExtremeTech, 19 February 2020.

[4]Evan, Dan. "Pope Francis Shocks World, Endorses Donald Trump for President." Snopes, 10 July 2016.

[5]See Wikipedia's "December 2015 Ukraine power grid cyberattack."

[6] German, Kent. "Boeing 737 Max Could Stay Grounded Until Late 2020." CNET, 30 April 2020.

[7]Although, an IBM study estimates the total cost of a data breach at $150 per record. This includes mitigation, lost revenue, etc. See: "Cost of a Data Breach Report." IBM, 2019.

[8]Sometimes made very measurable by massive fines imposed.

[9]Under what bizarre circumstances might we want our household clothes dryer connected to our Wi-Fi and the Internet? But our recently purchased drier can be. Good grief!

[10]The National Telecommunications and Information Administration (NTIA) in the US Department of Commerce has convened a working group around a software bill of materials to address this issue.

[11]Loan officers in a bank are judged on their loan production, motivating fudges to get the algorithm to approve a questionable application. If officers understand the algorithm in detail, they'll know the right fudges.

[12]Faw, Larissa. "Microsoft's Nadella: 'Privacy Is a Human Right.'" MediaPost, 24 May 2018.

[13]See Wikipedia's "2000 United States presidential election recount in Florida."

[14]See Wikipedia's *"New York Times Co. v. Sullivan."*

[15]The US auto industry had a long-held ethic that the release date of a new model (Job 1 in industry parlance) was sacred, regardless of the readiness of the product, and canny consumers never bought a significantly changed model during its first few months on sale. In the mid-1980s, Ford used the slogan "Quality is Job 1" and put its money where its mouth was by delaying the introduction of the much-ballyhooed new Taurus model by six months from its original date.

Paul Clermont is a Senior Consultant with Cutter Consortium's Business Technology & Digital Transformation Strategies practice. He has been a consultant in IT strategy, governance, and management for 30 years. His clients have been primarily in the financial and manufacturing industries, as well as the US government. Mr. Clermont's major practice areas include directing, managing, and organizing IT; reengineering business processes to take full advantage of technology; and developing economic models and business plans. He is known for successfully communicating IT issues to general managers in a comprehensible, jargon-free way that frames decisions and describes their consequences in business terms. In his consulting engagements, he follows a pragmatic approach to the specific situation and players at hand and is not wedded to particular models, methodologies, or textbook solutions. Mr. Clermont has spoken and written about the challenges of getting significant and predictable value from IT investments and has taught executive MBA courses on the topic. His undergraduate and graduate education at MIT's Sloan School of Management was heavily oriented toward operations research. He can be reached at pclermont@cutter.com.

DO YOU TRUST ME?

Trust: The Keystone of Digital Transformation

by Philippe Flichy

Although the emphasis on digital transformation has recently given way to dealing with the COVID-19 pandemic, digital transformation is likely to regain the spotlight once we weather the storm we are currently battling. At that point, successful digital transformation will require trust more than ever.

Microsoft CEO Satya Nadella has reportedly said, "Digital transformation is about reimagining how you bring together people, data, and processes to create value and maintain a competitive advantage in a digital-first world." Nadella is proposing that we create a shift in our corporate cultures. Bringing together people, data, and processes requires a fundamental change from the way most companies are working today; one in which data will be central, thanks to the ubiquitous computing power we now enjoy.

Transforming all the data we generate into insights requires many steps. These are the processes to which Nadella refers. For someone to have the confidence level to use the resulting information and insights, these data manipulations must be trusted. Before understanding where and how trust can be established, let's review these steps.

First, we must collect the data, which means identifying the data points that can be reliably collected at a reasonable cost. An increasing number of data points are now collected directly where they are produced, from all kinds of smart devices and sensors — a process referred to as "edge computing." This field has evolved significantly recently, with Syntiant's Neural Decision Processors capable of processing data at the edge with a processor so small that 54 chips can fit on a single penny.[1]

Next, data points need to be aggregated and normalized so they can be validated. This task can now be performed onsite to benefit from high-frequency data, allowing algorithms to take advantage of the best granularity of data available. Transferring this data reliably and securely at the right frequency presents a considerable challenge. Cybersecurity requires edge processing power as well as machine learning (ML) algorithms to distinguish valid traffic from illegitimate traffic. The advent of cube satellites,[2] Wi-Fi 6,[3] and 5G protocol[4] will help to lower the cost of transmission and increase the available bandwidth and coverage areas.

Short-lived data resides at the edge, while longer-lived data, as well as data derived from the field data stream, needs to be transferred centrally. The decision to transmit certain data at certain frequencies from a remote location may depend on the activity being performed. The same level of data granularity is not needed for all data points. The level will depend on what activity is being performed at a given time, as well as other operating states, such as weather conditions or other specific situations (e.g., a recent incident or some level of alert). Algorithms at the place of collection can help correlate near-real-time data feeds to eliminate outliers and create synthetic data when a value is missing or erroneous.

The decision to transmit certain data at certain frequencies from a remote location may depend on the activity being performed.

Nowadays, we can store high-frequency data locally and only transmit data centrally at various useful sample rates depending on the level of detail required for remote analysis. Thus, we keep the full granular data for on-demand requests and asynchronous replication by physically moving hard drives (such as with the import/export offering from Amazon Web Services) or by using spare bandwidth when it is available. The high-frequency data can be used locally to perform the initial validation steps, while compute-intensive mining and algorithm optimization can be performed centrally. The resulting improved algorithm can then be pushed back to the edge units to keep improving the data quality of the edge processing. Once available to subject matter experts (SMEs), they can analyze and correlate the data so it can be shared and transformed into actionable information or, better, reliable insights.

How Can You Trust This Complex Flow of Information?

For the end user to trust data after it has trickled through all these steps, it needs to be openly processed and securely distributed. The aim of data trust is to give people and organizations confidence when enabling access to data in ways that provide them with value in return (directly or indirectly). Data, wherever it is collected, needs to be understood and correlated to other data (some of which may reside in different parts of the company and, therefore, in various types of databases) in order to become information and yield insights, preferably actionable ones. Good correlation depends on good data. If there is no reason to doubt the data received, we tend to trust it. But how do we establish that required level of confidence?

According to the Industrial Internet Consortium, trustworthiness is "the degree of confidence one has that the system performs as expected."[5] Trustworthiness includes attributes such as constancy, congruity, reliability, and integrity, all of which involve various facets of data management. Let's start with: Can we believe the data in light of its context? Does it make sense? A simple illustration comes from a physics perspective: as the temperature of a gas rises, the pressure must rise as well. But the observer should also notice whether the data range being read corresponds to the environment being monitored. Data types and many other attributes can be analyzed to get a sense of whether some data is being distorted.

We may also rely on subjective criteria, such as whether the source of the data has a good reputation. Trust is not always based on scientific notions, as with our physics example, but also on aspects rooted in human psychology. From prehistoric times, humans have had to rely on and trust each other; for example, they have had to trust that the person on guard at night would wake up those sleeping when a threat was detected. Humans have developed a strong sense of belonging to and affiliation with tribes and then, later, groups of all kinds (from church to alumni associations and sports clubs, to name a few). Knowing that someone with the same values, education, or other shared common ground has curated data tends to predispose the user to accept that data as trustworthy. A corollary would be whether the data being received bears official endorsement. Conversely, some people might question whether data coming from or sanctioned by a particular group or institution is neutral and objective. The ultimate determination of trustworthiness may come down to whether the data source is usually reliable.

Once the source of the data is accepted as trustworthy, the next step is to better understand how the data is being handled. Can we verify that the data is from who it is said to be from, that it does indeed come from the source we believe it to come from? To address this concern, we look for traceability. Is it possible there was some transcoding (manual or automated)? Could alterations have been made, intentionally or not? Could someone have made an honest mistake? Can any changes be traced back? Was the data modified by an expert, and if so, who is that person, and on what grounds was the modification made? We find ourselves evaluating whether we can trust the individual/entity or the algorithm selected, and we will want to eventually assess the process used for any alteration. Is there any reason to believe that criminals could have intentionally altered the data feed? What guarantees are there that the data being received is genuine? Can the security processes used to secure the data be trusted? It is very difficult for an end user to have a good sense of the various steps and parameters necessary to ensure a secure transmission. The tendency, once again, is to trust the specialists and the entity's established processes and to believe that they provide the required level of confidence to trust that the data is legitimate.

One emerging technology being used for data exchange that promotes trust is blockchain. A blockchain is a decentralized, distributed, and, often, public digital ledger used to record transactions across many computers while preventing any of those records from being altered retroactively without the modification of all subsequent blocks. Since the records are disseminated on many distributed machines, it is virtually impossible to alter an entry once it has been entered unless an official update is submitted.

To establish confidence, we might also want to know the configuration parameters of a sensor and the associated software used to capture and process the data. We would again evaluate whether the parties involved in those steps are to be trusted. We should record as many as possible of these trustworthiness parameters in metadata. Metadata is a set of data that describes and gives information about other data. For example, metadata can provide information about a sensor's position, maintenance records, and the calibration tables over time to understand the drift of the sensor. Metadata can help users understand the

type of conversion used to convert analog readings to digital values, such as the sampling rates used or the number of bits of the transcoder, to name two obvious examples.

Timestamps can be complicated to reconcile to obtain a proper understanding of what happened at a precise instant. Vendors may not synchronize their tools' clocks across the same network. We therefore need to calculate offsets and may opt to calculate synthetic data to get all the measurements to align at a certain time. All these alterations must be traceable so they can be understood and then, we hope, agreed upon. Traditional code can be audited and confirmed, but we are increasingly relying on ML algorithms, which often are not human readable. The trend recently is for ML to create some level of human-auditable trails. Programmers also tend to include physics-based equations within ML algorithms to prevent aberrations. All these approaches can promote trust in the data.

Can We Create a Culture of Trust?

For any of these processes to be efficient, people must be willing to cooperate and share data and methodologies, which means crossing the boundaries of the typical departmental silos that exist in many companies. Thus, a culture of trust needs to emerge so we can comfortably rely on data and its analysis from other groups.

For data to be successfully shared, efforts must be made to normalize it and access it from various databases — an exercise that requires cooperation across disciplines and across companies. A data model needs to be identified, agreed upon, and implemented. The data then needs to be validated and, next, transferred or pointed to (one instance of the data but many places from where it can be accessed) so it can be aggregated in central repositories, such as data lakes, with data from other sources. Only then can the real digital correlation work start.

To discover valuable correlations, data analysts and SMEs ought to sit together and use each other's expertise to come up with valuable insights. There is no magic in the manipulation of data mining tools. The results obtained are the fruit of intentional cooperative efforts to understand what the data is telling us. As Mark Twain put it so eloquently, "Get your facts first, and then you can distort them as much as you please."[6] Sometimes various types of expertise are required to really distinguish true, meaningful correlations from what are aberrations. Did you know, for example, that the amount of US spending on science, space, and technology correlates with the number of suicides by hanging, strangulation, and suffocation? Spurious correlations[7] can destroy trust very quickly.

As we leverage an increasing number of digital tools, opportunities for miscomprehension of the various data sets and what they really mean increase exponentially. Unfortunately, so does the opportunity to distort reality and make numbers tell a story based on a specific agenda. This opportunity for misunderstanding or distortion has been true since humans started using statistics, but the increasing complexity of the correlation tools we use makes it even more of a problem. Some users may decide to resort to an A-B team approach to cross-validate some of the findings, while others may opt for a more subjective approach by trusting the entities and individuals involved in the analysis. How much we trust emerging tools aimed at validating correlations to do the right counter-analysis will be based on experience with, and comparison of, various validations.

For data to be successfully shared, efforts must be made to normalize it and access it from various databases — an exercise that requires cooperation across disciplines and across companies.

How Do We Access the Data?

Statistics about the volume of data we keep on generating and accumulating are well known; more data has been generated in the last two years than in all the previous years of human history. This brings up an interesting question as to how data is shared within a company, which illuminates yet another aspect of how an organization trusts its members. The more complex an access matrix the company builds, the more costly and difficult it is to enforce. Cross-pollination between groups often results in great discoveries, but for this type of effort to happen spontaneously, the boundaries between entities must be limited. As data is more and more protected at the core, in complement to firewall peripheral defense, company security officers need to identify and establish a simple yet efficient level of access.

This access issue has become even more important as the need for remote access from home has increased during the COVID-19 pandemic, and it is likely to continue after this confinement. The connectivity that employees working from home have with their companies' remote servers can be a huge challenge for some organizations. Although experts suggest that companies that already had deployed a good cloud infrastructure do not have an increased exposure to cyber threats, home network security will become a challenge if we are to continue working from home at any substantial rate, and could pose a threat to data trustworthiness.

Conclusion

Trustworthy data is required to accept the meaningful and useful insights we can derive from that data, and such insights can lead to the shift in corporate culture essential for digital transformation. In conclusion, to broaden the discussion outward again, for digital transformation to take place within an organization, a new culture based on trust must emerge. Leaders should show the way. Leadership is what will foster a culture of trust and make possible the adoption of data-driven tools to make more suitable and efficient decisions. However, the desired transformative shift is such that we need to evolve from trusting leaders to *having* trusted leaders. The level of transparency required dictates a strong reliance on empathy to build true and honest rapport and conversation.

For teams trying to unveil new discoveries resulting from data analysis in a fluid process, trust must first exist among the various participants. Trust can be initiated at a technical expertise level, but long hours digging into data require true interaction and demand a greater level of trust. Trust is based on empathy and getting to understand others and what they feel and think. Empathy leads to kindness and respect.

A culture built on trust is also dependent on clarity. Thriving by improving clarity in all processes is key to increasing efficiency. A clear, well-articulated vision and purpose will allow all actors to align themselves to the common goal. Knowing how to clearly communicate with one another on achieving the expected results will allow all involved to understand how they can contribute, at their level, to the common goal. Trust is also based on humility; being humble means that we are willing to listen to others and recognize we can be wrong without much fuss.

Much research has demonstrated that stress can be a transformative change agent. COVID-19 has certainly brought about much stress. But beyond the imperatives that will result from the current pandemic, the changes it has brought about may be a profound agent in creating the cultural precepts of a more digitally structured — and transformed — world.

References

[1]Nordmark, Jon. "How Amazon Will 'Own' Your 'Digital Me' and Detect a COVID-19 Infection Before You Do; Plus Strategies to Help Your Company Compete." *Iterate to Innovate*/ LinkedIn, 17 March 2020.

[2]A cube satellite is a constellation of low-orbit nanosatellites weighing between one and 10 kilograms.

[3]Wi-Fi 6, also known as 802.11ax, is the next-generation standard in Wi-Fi technology, with speed improvement of about 40% over the 2014 Wi-Fi 5, increased range and battery life, and the ability to handle more devices per router.

[4]5G protocol is the next generation of telecom networks, which can serve communication needs for billions of connected devices at up to 10Gbs — up to 100 times faster than 4G.

[5]"The Industrial Internet of Things Vocabulary Technical Report, Version 2.2." Industrial Internet Consortium, 6 November 2019.

[6]Kipling, Rudyard. "An Interview with Mark Twain." *From Sea to Sea and Other Sketches: Letters of Travel*. Doubleday, 1899.

[7]Vigen, Tyler. *Spurious Correlations*. Hachette Books, 2015.

Philippe Flichy is a digital transformation strategist and advisor at Endeavor Management with a deep knowledge of the energy industry. He is also an independent director and has held senior positions at Schlumberger, Baker Hughes, and Weatherford. Mr. Flichy joined the oil and gas industry after leaving the 2002 Salt Lake Olympics, where he delivered all the games results. He initiated the Digital Transformation Study Group at the Society of Petroleum Engineers and, as an entrepreneur, started his first company while still in college. Mr. Flichy earned a bachelor's degree in management information systems from Boston University. He can be reached via PhilippeFlichy.com, where he publishes blogs related to this article.

IS THAT REALLY YOU?

Reconstructing Virtual Trust

by David Tayouri

The Internet has been around for decades, and issues around security and user authentication are well known and have been much discussed. Although public key infrastructure (PKI) technology was a great advance for secure communication, current authentication methods are lacking.

Case in point: criminals, terrorists, and other illegal actors methodically abuse the openness of the Web, which enables the almost-instant creation of email and social media accounts (with potentially missing or false data) to create fake identities to hide their activities. Consequently, current authentication methods provide a low level of trust, which is particularly problematic when it comes to accessing personal sensitive data in the cloud (e.g., bank accounts and health records), remotely accessing business environments, and so on.

And there's more. A high level of trust in users being who they say they are is ever-more crucial in order to eliminate political interference by foreign actors impersonating local citizens or to prevent sham entities from spreading fake news. This high level of trust is also essential when our kids communicate through the Web, chatting or exchanging pictures with their friends. We want to be absolutely certain they are contacting other children and not imposter pedophiles. Can we be sure of this with the Web's current trust model?

Current Authentication Methods and Drawbacks

Three main types of credentials enable authentication today:

1. Something you know, usually a username and password
2. Something you have, such as a key generator device or cellphone
3. Something you are, such as biometric authentication (e.g., fingerprint, face scan, iris scan, voiceprint)

The use of username and password, despite remaining the most popular authentication method, has several well-known disadvantages, including password cracking, password theft, and reusing a single password for multiple applications. Two-factor authentication improves the security process somewhat but is inconvenient, still requires remembering many passwords, and doesn't prevent the creation of fake identities, since a one-time SIM card can be used for the authentication process.

Biometric authentication is widely known to be the most effective type of authentication because it is extremely difficult to transfer biological features from one user to another. However, the cost of early biometric authentication implementations has made it a less desired option for many projects. Recently, new technologies are making biometric authentication more feasible for a range of use cases, and that is why it is one of the building blocks of my proposed solution to improve authentication and increase trust.

Building Blocks for Solving Web Authentication Drawbacks

In addition to biometric authentication, asymmetric cryptography and blockchain are technologies that can help create an inherently authenticated and secure Web.

Asymmetric Cryptography

Asymmetric cryptography uses pairs of keys: public keys, which may be disseminated widely, and private keys known only to the owner. Effective security only requires keeping the private key private; the public key can be openly distributed without compromising security.

One important issue with asymmetric cryptography is the confidence/proof that a particular public key is authentic (i.e., that it is correct and belongs to the entity claimed and has not been tampered with or replaced by a malicious third party). There are several possible

approaches, including a PKI, in which one or more third parties — known as certificate authorities — certify ownership of key pairs (Transport Layer Security [TLS] relies upon this), and a "Web of Trust" that decentralizes authentication by using individual endorsements of the link between user and public key (Pretty Good Privacy [PGP] and Domain Name System [DNS] lookup use this approach).

Blockchain Technology

By design, a blockchain is resistant to modification of the data. Once recorded, the data in any given block cannot be altered retroactively without alteration of all subsequent blocks, which requires a consensus of the network majority. Although blockchain records are not unalterable, most consider blockchain secure by design. Recording credentials in a blockchain database assures that the credentials are authentic and were not altered by a third party.

Strengthening Authentication and Improving Trust

To reconstruct virtual trust, it's advantageous to strengthen current authentication methods with the addition of biometric authentication, asymmetric cryptography, and blockchain. We can achieve this by establishing a trusted Web layer, *in which users are uniquely identified and strongly authenticated*. This trusted Web layer can be accomplished by creating a database of trusted users. The database records will include a person's biometric ID, such as fingerprint, iris, or 3D digital face photo (or a combination thereof) and the credentials for cloud service providers (CSPs). The credentials will be encrypted and passed directly, automatically, and securely to each CSP for the authentication process. The use of blockchain as the database ensures the security of the recorded authentications and eliminates the need for a managing authority. Figure 1 depicts the high-level architecture of my proposed solution.

The records in the database of trusted users should be unalterable to avoid fraud and forgery, and the database should be decentralized to avoid a single point of failure and the need for a management entity. Blockchain meets these database property requirements because its records cannot be modified or deleted.

Users who want to change their password for a CSP create a new record. The authentication application will always present a user's latest record for a given CSP. In the event of a user's death or if he or she unsubscribes from a CSP, records do not need to be deleted; the relevant records will never be accessed again. From a law enforcement perspective, this database behaves similarly to databases of digital currency transactions; both use blockchain as the infrastructure, and, therefore, only users can access their own records.

The proposed solution has the following essential features that overcome the drawbacks of current authentication methods:

- **Strong biometric identification**, which uniquely identifies a physical person and is much stronger than other identification methods, eliminates the

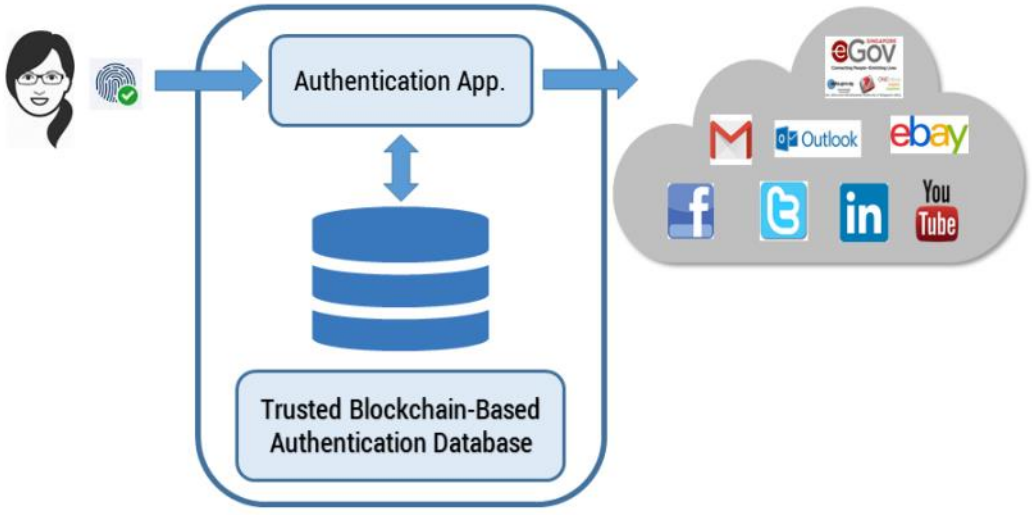

Figure 1 — High-level architecture of proposed solution.

fake identity problem. To avoid the theft of biometric keys, the database limits access to a record to the owner of that record.

- **Strong authentication** provides a high level of trust, a high level of confidence that the person is who he or she claims to be, and assurance that only the intended user can access his or her accounts.
- **A single authentication process** for all of a user's virtual accounts avoids the issue of needing to remember the various credentials for all the different cloud services that a user accesses.

Creating a Record in the Database

When accessing a CSP for the first time, users register their credentials in the database (i.e., they create a new record in the database). To create a new record, the user employs a biometric device (e.g., a fingerprint reader) and a dedicated application that creates a record with the user's biometric data, the name of the CSP, and the user's encrypted credentials for that CSP. The credentials are encrypted with the CSP public key, so only the CSP can decrypt them, using its private key. Validated records are distributed to all the database instances. The process repeats if users decide to change their credentials for a CSP.

The Authentication Process

For the authentication process, the user opens a CSP login page and employs a biometric device and a dedicated application searches the database for the user's biometric data and the relevant CSP. The application then passes to the CSP the latest relevant record, which includes the user's credentials for the CSP, encrypted with the CSP's public key. If a user has changed his or her credentials, the latest relevant record includes the most updated credentials. The CSP uses its private key to decrypt the credentials and to authenticate the user.

To enable this process, there are two prerequisites. First, users must have a biometric device. They may use the fingerprint scanner built into their smartphones or, for face recognition, the cameras in their smartphones, tablets, laptops, or workstations. Second, the CSP must enable logging in with encrypted credentials, which can be an option alongside the regular username-password authentication.

To clarify the process, let's take an example. Let's say Dave wants to access his bank account. At the login page, he selects the authentication option of "global biometric authentication" (the name we'll use for the described solution). He uses his fingerprint as identification. If this is his first time logging in, he is required to enter his credentials to access this bank. These credentials are saved in the database and sent to the bank for authentication. If this is not Dave's first time logging in, his encrypted credentials are retrieved from the database and sent to the bank for authentication.

Since the proposed trusted Web layer requires that users become accustomed to biometric identification and that the CSP provides authentication with encrypted credentials retrieved from the authentication database, the implementation of this proposed solution may need to be done gradually. The proposed approach doesn't eliminate current methods of authentication. Rather, it is an alternative, so users can continue using the existing authentication methods until the new method is fully adopted.

Privacy vs. Security

The openness of the Web and its inherent anonymity are considered advantages in that they enable wide access and privacy, but they can also be viewed as disadvantages, since illegitimate actors abuse these characteristics to create false identities and conceal their (often nefarious, and even criminal) activities. Anonymity can assure privacy, but it comes at a price: a low level of trust as well as a low level of security.

A database of users who are biometrically authenticated can be the solution to this problem. The described authentication process ensures that individuals are who they claim to be, no matter how many profiles they may have. Individuals in this database may still be illegal actors, but once exposed, they cannot create new identities and hide behind them. The use of a physical (biometric) method for identification makes the anonymous creation of a new identity impossible.

Let me clarify that the proposed solution does not eliminate the possibility of creating multiple (distinct) identities (e.g., for private and professional purposes). However, for each virtual identity, the person behind it is uniquely identified and strongly authenticated, maintaining trust.

Note that, as the solution offered here does not suggest discontinuing current authentication methods that do not rely on biometric identification, users can still access the Web anonymously. In situations where a low level of trust is acceptable (e.g., in some social interactions or for restaurant or lodging recommendations) or when people could be at risk if they reveal their identity, anonymity can be preserved.

Summary

The current non-biometric authentication methods have a low level of trust, since they cannot assure that users are who they declare themselves to be. This enables criminals, terrorists, and other illegal actors to abuse the openness of the Web and create forged identities to mask their activities.

While the current level of trust may be satisfactory for some uses, we must secure a higher level of trust when accessing personal sensitive data. Establishing a trusted Web layer, in which users are uniquely identified and thoroughly authenticated, reduces the lack of trust issue. We can achieve trust by creating a decentralized, unalterable database of users who are biometrically authenticated using encrypted credentials.

Combining biometric authentication, asymmetric cryptography, and blockchain technology enables us to improve authentication and reconstruct virtual trust, which is currently low and insufficient. My proposed solution enables more trust when accessing personal sensitive data in the cloud, allowing children to communicate safely on the Internet, and more. Biometric authentication drives less anonymity, but if a CSP requires open and secure biometric authentication (such as with the solution proposed here), we will have more trust in the people accessing that CSP. When I think of my children surfing the Web, I prefer a situation where we all lose some anonymity and privacy to ensure that children are safer. How about you?

David Tayouri is Cyber R&D Manager in the Engineering Directorate, Cyber Division, ELTA Systems, Israel Aerospace Industries (IAI). Mr. Tayouri is among the cyber activity leaders at IAI, managing its Cyber Intelligence Department from 2013-2017, prior to the establishment of the Cyber Division. During the last 17 years, he has been developing intelligence-gathering systems for defense organizations in different layers, mastering the cyber domain in the last nine years by heading cyber technology and business units and developing innovative cyber solutions. Mr. Tayouri has nearly three decades' experience as a technology leader with a keen business understanding as a software developer, team leader, system engineer, project manager, and system architect in various domains. He earned a master of science degree with honors in computer science from Bar-Ilan University, Israel. He can be reached at dtayouri@elta.co.il.

PROPER HYGIENE

How to Trust a Stranger

by Panagiotis Monachelis, Panagiotis Kasnesis, Charalampos Patrikakis, Xing Liang, Ryan Heartfield, George Loukas, Nelson Escravana, and Soulla Louca

The increasing role of social media has led to a corresponding need for methods to assess the trustworthiness of user-generated, online-shared content. This demand emerges from the fact that, while a significant portion of social media posts have proven to be a great source of knowledge and news, other posts have purposefully spread false information. While trustworthiness of some such posts is easily determined following simple rules, others are cleverly designed to deceive even experienced users. At the same time, there is a growing trend toward decentralization of social media platforms. Mastodon[1] is an excellent example of a successful paradigm, and the "Bluesky" project[2] announced in 2019 by Twitter CEO Jack Dorsey is representative of the trend toward social network decentralization.

We begin this article with a presentation of contemporary decentralized networks and enabling technologies and argue that addressing misinformation in these environments must follow a similar path. Rather than relying on third parties for fact-checking, we should exploit the potential end users represent as contributors to the credibility-checking process. We discuss how to successfully implement a human-centric solution to assist social media users in gauging the trustworthiness of information in a manner that complements rather than replaces "information hygiene" guidelines.

Most mainstream online social networks are centralized platforms, each operated by a single authority. This brings about a situation where the data of vast numbers of users is held by, and in some cases subject to the censorship of, a very small number of social network providers. In contrast to mainstream centralized online social networks, decentralized online social networks are based on distributed information management schemes, empowered by trusted servers or peer-to-peer (P2P) systems. A primary motivation of decentralized social media is censorship resistance, often achieved through anonymity. Another motivation is the guarantee of ownership of personal data, which leads to new content monetization models with fewer or no ads. In addition to Mastodon, several decentralized platforms exist, including diaspora*,[3] SocialX,[4] Sola,[5] Minds,[6] Steemit,[7] and DTube,[8] each with its own architecture.

But anonymity comes with a price. Checking the reliability and trustworthiness of information is much harder with anonymous information. To assess the veracity of social network post information, it is useful to identify sources of the information (i.e., when it was first published and in which post), how similar the information is to other published sources, and how all related sources of information may have changed over time. But this is not always an easy task. The increased anonymity and decentralized nature can make it more challenging for the user to follow information hygiene guidelines, such as "check the source" or "check how the information reached you." As a result, we argue that to be in line with the ethos of decentralized social media, any tools developed to address misinformation should be designed with an aim to help the user follow information hygiene guidelines rather than deciding for them what information should or should not be considered trustworthy.

Below, we present a brief overview of two main underlying technologies that help in decentralization while still providing a degree of trust between networked entities: (1) P2P social networks and (2) blockchain technology.

P2P Social Networks

A P2P network uses a distributed architecture featuring networked devices called "peers." Peers interact with each other and share resources. All peers can share files, use instant messaging, make audio and video calls, interact with posts and photos, and join or leave the network at any time. The overall performance of the network increases as the number of peers increases. Peers can also be combined into groups as they communicate, interact, and share bandwidth, allowing the network to continue to function, even if one or more peers disconnects. In a more complex interconnection, a P2P social network consists of small groups of peers,

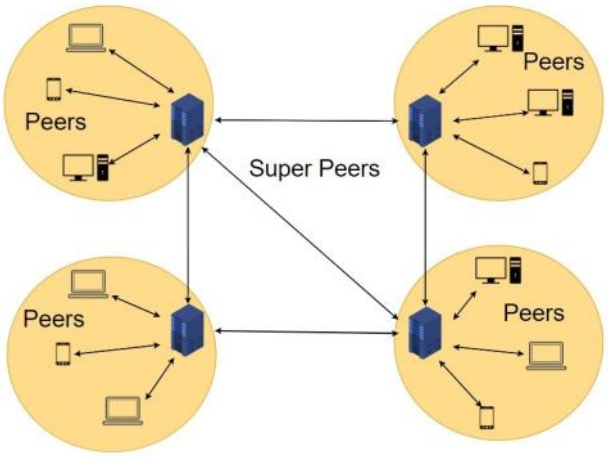

Figure 1 — P2P social network.

where each group has its own "super peer," a device that undertakes the role of mediator with the rest of the network via other super peers (see Figure 1).

In this scheme, each peer connects with a super peer to route its data via the latter. Super peers manage the incoming messages and forward them to other super peers or peers, according to the content. If authentication is needed, it can be implemented between peers and super peers before connecting to the network. A popular file-sharing P2P system is BitTorrent,[9] which users can use to exchange files.

Blockchain Technology

Originally developed to support the Bitcoin cryptocurrency, blockchain technology has proven instrumental in supporting applications related to trust in several other distributed environments as well.[10] In distributed social networks, the technology provides security and privacy through cryptography. In a blockchain structure, nodes (which can be physical or logical entities) interact with each other (e.g., making Bitcoin transactions recorded in the form of a decentralized digital ledger), and participating nodes on the blockchain network have the ability to validate, synchronize, and retain a copy of the ledger. Each user's action enters the network in the form of a block. The blocks connect to each other, creating a chain in which each block (except the first) is connected to the previous with a hash function. Figure 2 provides a high-level overview of a blockchain.

Typically, the blocks in a blockchain reference a specific user or system transaction. However, blocks may also be designed to reference an immutable record of a "state" of information trust or integrity. For example, in the case of social network posts, a blockchain ledger can be considered a reference state, recording votes of trustworthiness that have been anonymously recorded for specific post content. To uphold privacy, however, the "right to be forgotten" must be upheld, whereby blockchain transactions can be deleted or expunged from the ledger. This is somewhat of a paradox for blockchain, as it was designed with the premise that the blockchain is an immutable record of events. To address this architectural limitation, rather than define blocks based on individual posts and their trustworthiness votes, it may be better to define blocks as a representation of the "state" of the P2P database underpinning a decentralized social network platform (e.g., where a "state" block is added periodically after either a predetermined amount of time, a set number of posts has been published, or a prescribed number of trustworthiness votes have been cast). With this approach, if posts are revoked from the social media platform (and therefore the P2P database), it does not affect the immutability of the blockchain because each block references the P2P database "state" of posts and votes cast against them, and not the posts or votes content themselves. Thus, registered actions (and, in this case, a post's trustworthiness votes) can be verified as a record by the most recent committed state, while ensuring

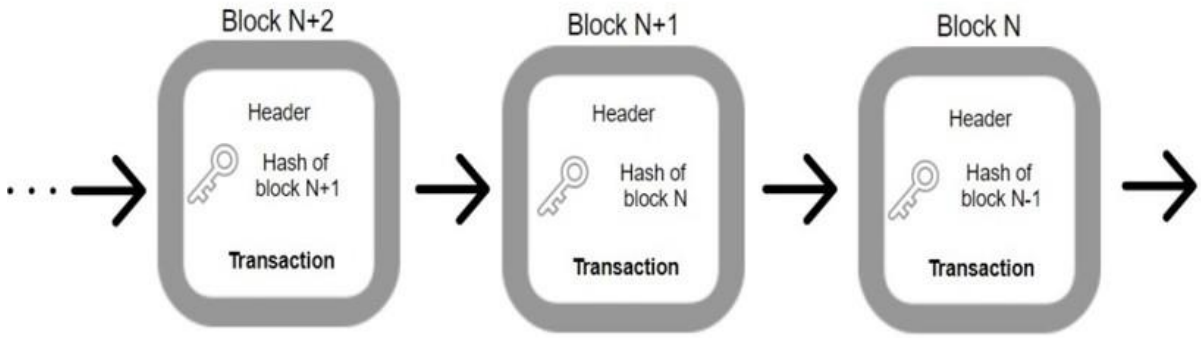

Figure 2 — A chain of blocks.

posts that are revoked (and with them any trustworthiness votes stored in the P2P) do not compromise the integrity of the ledger. In a decentralized social network ecosystem, applying blockchain in this manner is highly attractive, as it provides a single source of intermediary-free "truth" (i.e., verified state) for posts and trustworthiness votes that is resistant to manipulation of any single node in the network.

Assessing Trustworthiness in Social Media

Research in addressing challenges related to information trustworthiness in social media focuses on the development of automated reasoning techniques based on machine learning (ML), where a news article or post is classified as legitimate or "fake news," based on a variety of indicators. However, there is increasing interest in involving the user in the process and exploiting the potential of crowdsourcing. Where algorithms cannot draw accurate conclusions, a crowdsourcing service may improve reliability.[11] An assessment of the validity of news is not delegated to any specific authority but is instead supported by a blockchain system where users (including news experts such as journalists) are responsible for assessing the validity of news. The published news is deployed in a chain of transactions. News beyond a certain level of virality is verified by user validators. In the particular work by Paul Shobon et al.,[12] the validation has two phases: During the first phase, professional journalists validate the news. In the second phase, a random group of users with knowledge of the article's subject validate it. The validators hold different scoring weights, and the final rating of the article is a combination of the validators' scoring. The weights of the validators' scoring change over time as valid votes increase a user's weight while incorrect votes decrease it. Validity is evaluated by the deviation of a validator's score from that of the professional journalists. This system is a step toward decentralizing the process and more actively involving the user.

A completed implementation of a decentralized online social network (DOSN) is PrPl (private-public), which proposes a "personal cloud butler" service as a safe place for all personal data that may be located in different data stores.[13] Users and butlers can be authenticated by a federated identity management system based on the OpenID protocol. The butler can run on a home server or on an external provider. A butler provides a federation of data storage where a butler's data can be stored encrypted in different locations, using an indexing system for more efficient accessibility.

An interesting proposal that puts humans in the center and exploits the potential of humans as trustworthiness sensors is the research project EUNOMIA, which addresses the challenges of misinformation in social media. EUNOMIA actively encourages citizen participation in content verification by voting on content trustworthiness. The goal is for users to take ownership of the problem of disinformation, rather than relying on third-party fact-checkers or computer software. The number of votes appears as one of several indicators that the user can visualize along the information cascade of each post. Currently under development in the context of the EU's Horizon 2020 research framework,[14] by a consortium of academic and business partners (with this article's authors among them), EUNOMIA is a fully decentralized P2P platform.

The goal is for users to take ownership of the problem of disinformation, rather than relying on third-party fact-checkers or computer software.

A blockchain infrastructure supports users, providing a reference and immutable state of the most current trustworthiness votes registered for a post. This "state" integrity methodology enables a participating user to show that he or she is the original source of a piece of information posted online, or to identify the earliest sources of specific information posted in the platform. A user can see information about a post — but not the initial source. At the press of a button, a user can view an information cascade that shows not only where a post originated, but also if and how it has been modified. The user may then determine whether a post is an important news item or just a joke. As different users have different ways of determining for themselves whether a piece of information is or is not trustworthy, users can choose what kind of information related to trustworthiness they want to see. This may include a measure of activity, such as the ratio of followers to following, or other indicators the users suggest or that are identified in the scientific literature.

In the remainder of this article, we describe how components and services in EUNOMIA's architecture enable this capability to assess trustworthiness.

The EUNOMIA Architecture

The EUNOMIA architecture contains three components that ensure its secure, decentralized nature: (1) a P2P network, (2) a blockchain infrastructure, and (3) a security and privacy framework. It also comprises four tools that assist a user in assessing information trustworthiness: (1) a human-as-trust-sensor component, (2) a social media content and context data analysis component, (3) trustworthiness scoring, and (4) a user application in the form of a digital companion.

These tools and components are implemented in a network of interconnected services nodes, making their services available to end users through an application running on user devices (a "digital companion"). The nodes communicate through a P2P network that features a distributed file system protocol.

Figure 3 illustrates these EUNOMIA components and their interconnections. The architecture is based on the use of services nodes (which consist of services related to the rules that determine how data can be created, stored, and updated), authentication, authorization, accounting, and data storage. The services nodes make use of two core components:

1. **Both the adapter layer and decentralized social network** work as an external component used to relate the EUNOMIA accounts to social network accounts and provide verification that a corresponding account exists on the social network. To abstract the P2P service from the mechanisms inherent to the different possible social networks connected to EUNOMIA, an adapter layer is implemented, serving as middleware to the third-party authentication mechanisms of the social networks.

2. **The digital companion,** an application deployed on all types of devices (desktops/tablets and smartphones), features a responsive Web-based and a personal (mobile/wearable) app version, allowing for the active involvement of social media users. The digital companion client runs on the user's device and can communicate with one or more decentralized social networks as well as one of the EUNOMIA services nodes. In the client, the end user can read a post from a decentralized social network, reply or "like" the post, and also view computed EUNOMIA indicators about the post, such as its sentiment score. The user can also create a new post from the client and vote about the trustworthiness of a specific post. The result of the analysis is displayed on the client in the form of interactive visualization graphs. All the queries between the client and the rest of the EUNOMIA services are implemented through the EUNOMIA services interface.

The components in EUNOMIA that comprise the services nodes include:

- **The P2P infrastructure**, which supports the storage and communication service based on the use of a storage server and the corresponding logical storage

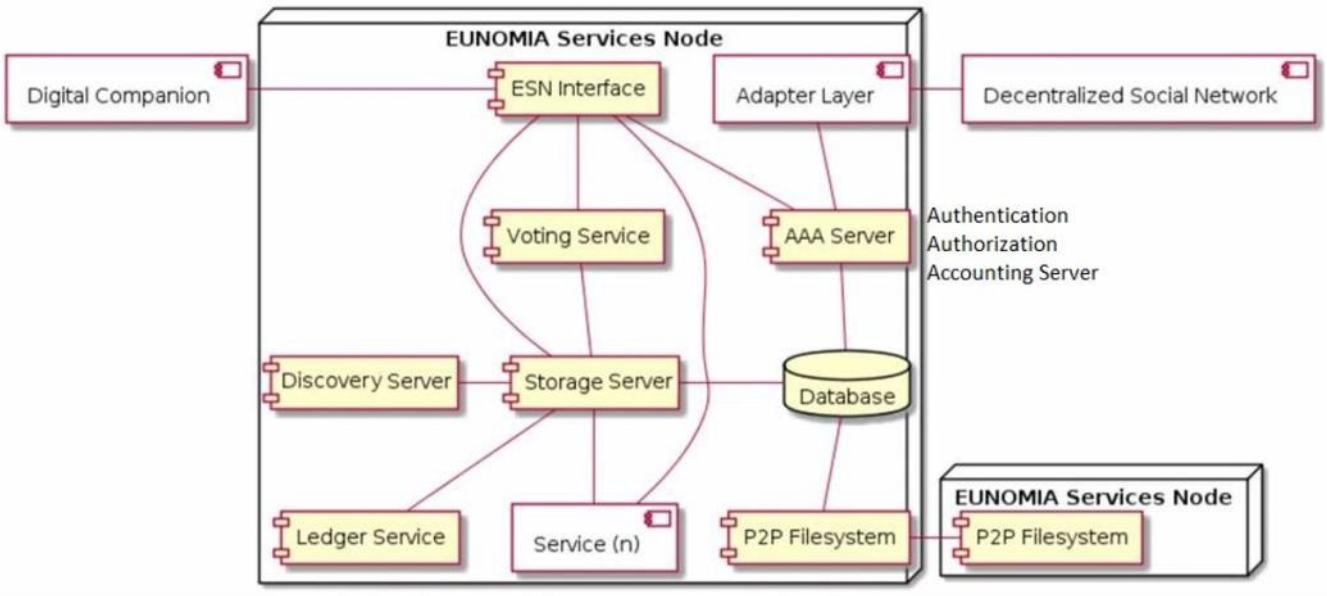

Figure 3 — EUNOMIA component diagram.

service (running on the storage server). The storage server features a distributed database supported by a P2P file system and a ledger service (see below) used to record the content information cascade (i.e., a group of posts containing similar content) as content metadata to ensure transparency and consistency. A P2P server API provides endpoints to enable other components to manipulate data, provided they have been authenticated. For authentication, the requests must be accompanied by access keys that the entity itself generates to certify users in order to accept the request.

- **The blockchain infrastructure,** which consists of two main components:

    1. **The ledger service** consists of the EUNOMIA blockchain network and an API interface, which acts as an abstraction service layer between the blockchain network infrastructure and the rest of the service components (see Figure 4). The ledger service is implemented using the Hyperledger Fabric framework as the underlying blockchain network. The blockchain network infrastructure is comprised of peer nodes and certificate authority (CA) nodes. The authentication of the network is performed using a set of policies and roles over a public key infrastructure (PKI). The orderer nodes participate in the consensus process, and peer nodes act as gateway nodes to the network. Bootstrapping the EUNOMIA's blockchain network requires at least one orderer peer and one CA server. The orderer peer plays a significant role in the transaction flow, organizing transactions into blocks that are then pushed to the committing peers and recorded on the ledger. The CA server is used for the authentication of users.

    2. **The EUNOMIA blockchain API** invokes EUNOMIA's chaincode (i.e., smart contract) through a general-purpose remote procedure calls (gRPC) interface protocol over an SSL connection. The chaincode deployed on the blockchain infrastructure implements and controls the transactions logic and object states. For the EUNOMIA blockchain API to connect to a blockchain peer, it must be authenticated as a blockchain user. The EUNOMIA blockchain API is considered as a user of the EUNOMIA blockchain network and offers an abstraction layer for the other EUNOMIA services node components and the ledger service.

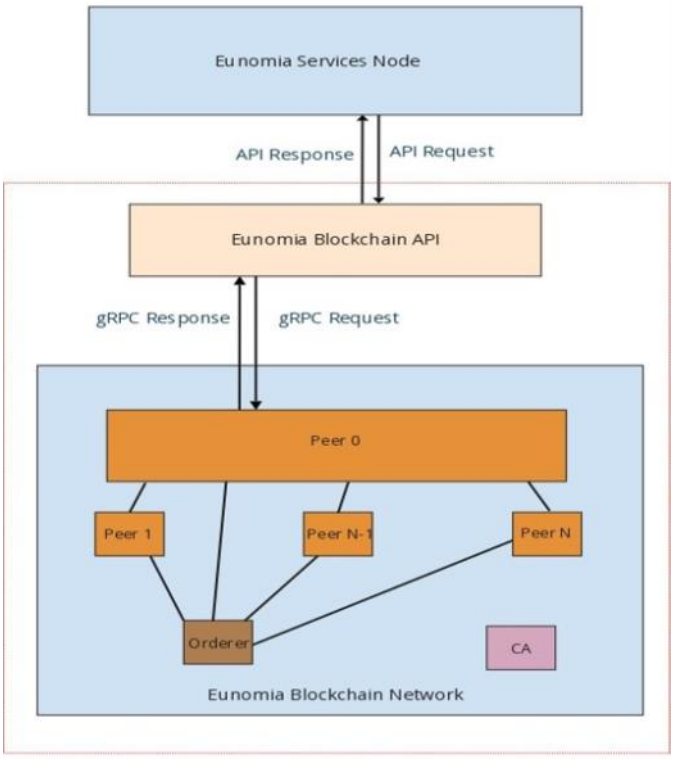

Figure 4 — Ledger service.

EUNOMIA uses three components to ensure security and privacy:

1. **The AAA server,** which is responsible for authentication, authorization, and accounting

2. **The discovery server,** which provides the means to allow the discovery of other services and corresponding metadata, including public keys for sharing sensitive information

3. **The voting server,** through which users can react and express opinions on available content, features a voting mechanism that supports the submission, logging, and querying of votes by users for individual posts and cascades

Services nodes contain one or more services (e.g., sentiment analysis service) to communicate and exchange information, while providing an API abstracting the EUNOMIA's business logic for CRUD (create, read, update, delete) operations.

Trust and Information Cascade Management

In EUNOMIA, an information cascade module performs data collection, preprocessing, feature extraction,

and data aggregation to generate an information cascade (where applicable) of aggregated posts and author trustworthiness indicators. The module also performs aggregated visualization of posts and author data and metadata over the digital companion graphical user interface, allowing assessment and trustworthiness voting.

The automated trustworthiness scoring is an experimental research component of the architecture that requires additional evaluation to determine its ability to help curb the spread of misinformation.

The architecture allows the incorporation of easy-to-understand data aggregation and presentation approaches that take into consideration trustworthiness indicator preferences set by the users in the digital companion. Users may set indicator criteria and specific thresholds by which they assess trustworthiness for a post or information cascade. For example, if the top primary criteria a user defines are "account_age > 1 month," "average_post_sentiment > 0.5," and "follower_count > 50," then they may be presented with the percentage of these criteria that are satisfied. In a binary classification mode, a simple rule-based majority vote can be employed to suggest trustworthiness based on the user's preferences.

Trustworthiness scoring can also be facilitated at the local level by the user's voting. Such a component would require sufficient past voting data by a specific user to learn what indicators matter the most to that user and would accordingly suggest an allocation of indicators, just like a recommendation engine. For each user vote cast on a post via EUNOMIA, the platform collects and measures specific post/information cascade indicators that are recorded locally and used as training data for the EUNOMIA platform to learn the indicator thresholds. Note that any such form of automated trustworthiness scoring is based on the specific user's voting behavior and generates a unique model that applies inference to the selected posts. It is available only for that specific user. The automated trustworthiness scoring is an experimental research component of the architecture that requires additional evaluation to determine its ability to help curb the spread of misinformation.

Conclusion

Purposeful disinformation can have a big impact in public health, public affairs, and several other critical social areas. Privacy and data ownership are also controversial issues that require an effective strategy to address. Several techniques have been proposed for this purpose. In the quest for information trustworthiness, we must remember that the user should be at the center and must not be forgotten or replaced completely by IT tools. A suite of ML algorithms may help in quickly assessing the language, sentiment, or stance of a post — or to compare against databases of items of fake news collected by external entities. But this cannot and should not replace digital literacy and the encouragement of the user to practice critical thinking. In this respect, "offloading" information trustworthiness assessment to a system is no different from trusting a mainstream social network provider or third-party fact-checkers. Instead, an effective direction for addressing information trustworthiness challenges is to develop tools that help users in following information hygiene guidelines rather than effectively replacing them with solely technical solutions.

EUNOMIA has currently finished its first phase of development across all key components and has developed an integration for Mastodon, which is tested in the project's own private Mastodon instance. The next step is to test with social media users. By its completion in November 2021, EUNOMIA will be validated across several use cases, including journalists, social media journalists, and citizen users.

References

[1]Mastodon, 2020.

[2]Twitter CEO Jack Dorsey announced the "Bluesky" project in a tweet on 11 December 2019.

[3]diaspora*, 2020.

[4]SocialX, 2019.

[5]"Sola (SOL)." CryptoTotem, 2020.

[6]Minds, 2020.

[7]Steemit, 2020.

[8]DTube, 2020.

[9]BitTorrent, 2020.

[10]Tama, Bayu Adhi, et al. "A Critical Review of Blockchain and Its Current Applications." *Proceedings of the International Conference on Electrical Engineering and Computer Science (ICECOS)*, IEEE, 2017.

[11]Shovon, Paul, et al. "Fake News Detection in Social Media Using Blockchain." *Proceedings of the 7th International Conference on Smart Computing & Communications (ICSCC)*, IEEE, 2019.

[12]Shovon et al. (see 11).

[13]Seong, Seok-Won, et al. "PrPl: A Decentralized Social Networking Infrastructure." *Proceedings of the 1st ACM Workshop on Mobile Cloud Computing and Services: Social Networks and Beyond*, ACM, June 2010.

[14]"Horizon 2020." European Commission, 2020.

Panagiotis Monachelis is a PhD student in the Electrical and Electronic Engineering Department at the University of West Attica (UniWA), Greece, and a teaching staff member in UniWA's Computer Networks & Services Research (CoNSeRT) Lab. His research focuses on big data mining and visualization over online social networks. Currently, Mr. Monachelis is involved with the EUNOMIA project in the visualization field. Previously, he worked in the domain of microcontrollers at a technical company and IT in the banking sector. Mr. Monachelis earned a bachelor's degree in electronic engineering from the Technological Educational Institution of Athens and a master's degree in networking and data communications from Kingston University, UK, in partnership with Piraeus University of Applied Science, Greece. He can be reached at pmonachelis@uniwa.gr.

Panagiotis Kasnesis is cofounder and R&D Director of THINGENIOUS. He has participated in several European R&D projects as a machine learning engineer and semantic application developer in the Computer Networks & Services Research (CoNSeRT) Lab of the University of West Attica (UniWA), Greece. His research interests include machine/deep learning, Semantic Web technologies, multiagent systems, and the Internet of Things. Dr. Kasnesis has published and presented more than 20 scientific articles in international journals and conferences. He is certified as Instructor and University Ambassador by the NVIDIA Deep Learning Institute in natural language processing and computer vision. Dr. Kasnesis earned a diploma in chemical engineering, a master's of science degree in techno-economic systems, and a PhD in computer science from the National Technical University of Athens, Greece. He can be reached at pkasnesis@uniwa.gr.

Charalampos Z. Patrikakis is a Professor in the Department of Electrical and Electronic Engineering at the University of West Attica (UniWA), Greece. He is also Director of UniWA's Computer Networks & Services Research (CoNSeRT) Lab. Dr. Patrikakis has more than 100 publications in book chapters, international journals, and conferences and has made two contributions to national legislation. He is a senior member of IEEE, Associate Editor responsible for special issues in IEEE's IT Professional*, and counselor of IEEE's student branch at UniWA. Dr. Patrikakis has acted as scientific/technical coordinator in European projects, leading dissemination activities in many of them. He is also a member of many technical program committees and editorial boards of international conferences and journals and has presented results in academic and industry fora. He can be reached at bpatr@uniwa.gr.*

Xing Liang is a postdoctoral Research Fellow in data analytics and machine learning in the IoT and Security (ISEC) research group at the University of Greenwich, UK, with particular interest in information trustworthiness in social media. Currently, she is involved with the EUNOMIA project. Previously, Dr. Liang was a Research Fellow in computer vision at the University of Westminster, UK. She earned a PhD from the University of Bradford, UK, where she participated in two EU-funded projects in the area of mobile satellite communication networks. She can be reached at X.Liang@greenwich.ac.uk.

Ryan Heartfield is a postdoctoral Research Fellow in cybersecurity in the IoT and Security (ISEC) research group at the University of Greenwich, UK. He has led R&D in multiple national and international research projects, focusing on artificial intelligence in cybersecurity and cyber-physical intrusion detection. Currently, Dr. Heartfield is leading work in developing social media information trustworthiness tools for the EUNOMIA project. His research in defining the human-as-a-security-sensor paradigm has gained both academic and industrial recognition. Dr. Heartfield earned a PhD in cybersecurity from the University of Greenwich. He can be reached at R.Heartfield@greenwich.ac.uk.

George Loukas is an Associate Professor and head of the IoT and Security (ISEC) research group at the University of Greenwich, UK. He has led several national and international research projects, including coordination of the EUNOMIA project on information trustworthiness in social media as well as a project on cyber hygiene for artificial intelligence applications in domestic life for the Engineering and Physical Sciences Research Council. Dr. Loukas is the author of Cyber-Physical Attacks*, an Association for Computing Machinery (ACM) Top 10 book in ACM's 2015 annual list of notable books and articles in computing. He is on the editorial board of* Simulation Modelling Practice and Theory *and* The Computer Journal*. Dr. Loukas earned a PhD in network security from Imperial College, UK. He can be reached at G.Loukas@greenwich.ac.uk.*

Nelson Escravana is Director of the Communications & Cybersecurity Unit at INOV, where he is also a board member. Mr. Escravana has over 20 years' professional experience, actively participating in R&D and integration projects as a technical leader and project manager in the fields of penetration testing, technological developments for intrusion detection, forensic auditing, and mobile authentication, and has been a leading security auditor. He has participated in multiple projects and studies for the Portuguese government, European Space Agency (ESA), and the North Atlantic Treaty Organization (NATO). In 2011, Mr. Escravana was responsible for the creation and management of the Institute for Systems and Computer Engineering, Technology and Science (INESC) shared service center (INESC Serviços), which includes INESC CSIRT (cyber incident response group). He earned a degree in computer science engineering and management post-graduation, both from the Technical University of Lisbon, Portugal. He can be reached at nelson.escravana@inov.pt.

Soulla Louca is a Professor in the Department of Management and MIS as well as Director of the Institute for the Future at the University of Nicosia, Cyprus. Dr. Louca was nominated by LATTICE80 in partnership with Miss Kaya as one of the top 100 women in fintech for 2019. With a background in computer science, she is among the first academics to get involved in blockchain technologies. Previously, Dr. Louca headed the Domain Committee for ICT for European Cooperation in Science and Technology (COST) for the 36 COST member states. She earned a PhD in computer science from the Illinois Institute of Technology. She can be reached at louca.s@unic.ac.cy.

TO TRUST, OR NOT TO TRUST ...

The Supply Chain Security System of Trust: A Framework for the Concerns Blocking Trust in Supplies, Suppliers, and Services

by Robert A. Martin

In this article, we examine an effort to develop and validate a process for integrating evidence of organizational, technical, and transactional trustworthiness of supplies, suppliers, and services for decision makers. Leveraging the MITRE community's decades of experience across a widely diverse set of customers and domains, along with the work of colleagues in industry and academia, the supply chain security system of trust (SoT) has four focused objectives. The first objective is to gather and organize a single structured corpus of the concerns surrounding trust with respect to organizations, products/components, and service offerings. Second is to capture the techniques and mechanisms available to gather evidence as to whether concerns with trust are or are not substantiated (e.g., a bill of materials, whether for software or hardware, can provide a streamlined, repeatable, and scalable method for collecting evidence about salient trustworthiness aspects for supplies you use to build and operate your systems). The third objective is to provide a mechanism for winnowing and tailoring the overall SoT to a custom set of concerns and investigative questions that considers the resources of an organization, the significance of the system or service to its operations, and the consequences that could result from failing to fully vet concerns. Finally, the SoT aims to put in place objective scoring mechanisms that an organization can adapt to its priorities, operational sensitivities, and experience with its type of business and partners.

The topic of trust and trustworthiness is one confronting many communities around the world.

Along with these elements of SoT, this article describes how various community efforts — like the Industrial Internet Consortium (IIC)[1] and its focus on trustworthy Industrial Internet of Things (IIoT) systems in operation and Plattform Industrie 4.0[2] with its ideas on trustworthy organizations — and other endeavors to define trustworthy systems can be integrated.

Context for the SoT

The supply chain security SoT is a MITRE community initiative aimed at defining, aligning, and addressing the specific concerns and risks that stand in the way of organizations trusting suppliers, supplies, and services. The SoT effort includes an examination of the issues that can undermine trust in a supplier and/or in the supplies and services it provides. This examination encompasses the multiple contributors to the development and provisioning of systems and services, as well as the suppliers' offerings of services to maintain systems and proffer upgrades/updates.

The topic of trust and trustworthiness is one confronting many communities around the world. At the same time, these communities often reflect distinct perspectives and different purposes. The IIC, for example, has been exploring how to define a trustworthy IIoT system and how to demonstrate or establish that it is trustworthy.[3] Simultaneously, Germany's Plattform Industrie 4.0[4] use of the term "trustworthy" suggests what it means to work with trustworthy partners, and this group collaborates to align with the IIC's work on the trustworthiness of systems.

When we explore what many define as supply chain issues, most organizations tend to focus on whether delivered goods and all associated components and raw materials will be available where and when they are needed — the classic "logistics" approach to supply chain. When industry and government use the terms "trust" and "trustworthiness," however, they focus on risk management, most often in terms of financial and regulatory perspectives.

Parties to transactions do not usually consider yet another source of risk: how to decide to trust those different parts. Moreover, they often do not have a

complete understanding of the different areas of concern and risks that could disrupt a supply chain, or, importantly, the risks that could undermine trust in a system or a supplier. A frequent result is that, when two organizations "talk to each other," they are instead talking past each other because of differences in each party's definition of the word "trust" and accompanying concept of transactional trustworthiness.

As illustrated in Figure 1, the various entities involved in transactions exhibit different intentions, along with individual unique understandings of the meaning of trustworthiness. Real-world examples include parties concerned with creating or using trustworthy commercial-off-the-shelf (COTS) systems that incorporate information and communications technology (ICT), or those interested in having trustworthy high-value items like medical devices with embedded ICT capabilities, or those concerned with having trustworthy surveillance cameras across a large manufacturing complex. These parties undoubtedly often have different intentions, along with individual, unique understandings of the meaning of trustworthiness.

While each of the above roles and types of concerns uses the terms "trust" and "trustworthiness," generally the various entities presume different concepts, scope, and level of veracity underlying those terms. This causes immense confusion and misunderstandings that negatively impact organizational transactions and trust and trust relationships (primarily external, and potentially internally). The SoT harmonizes the vocabulary of trust and trustworthiness while supporting the subtexts and nuances that delineate varied areas of concern and various roles among suppliers and those using their products and services. Without this harmonization by the SoT, each of these unique participants with their unique experiences and perspectives will continue to talk past each other due to their use of differing concepts or scope of issues when they talk about trust and trustworthiness.

Supply Chain Security's Challenging Evolution

From a supply chain security perspective, the least understood risk to most organizations comes from needing to almost be a programmer to use everything from a vehicle, to an elevator, or to today's medical devices, never mind trying to manage a "smart" power grid. Software- and network-enabled ICT has become a key enabler for a large portion of people's homes and businesses, with new and enduring supply chain challenges.

While organizations have always had suppliers, the introduction of the supply chain as a unique source of economic harm and national security risk came in World War I, when "anti-tamper" became both a practice espoused for protecting the goods and services of the munitions industries and combat material and a focused law enforcement effort aimed at defeating

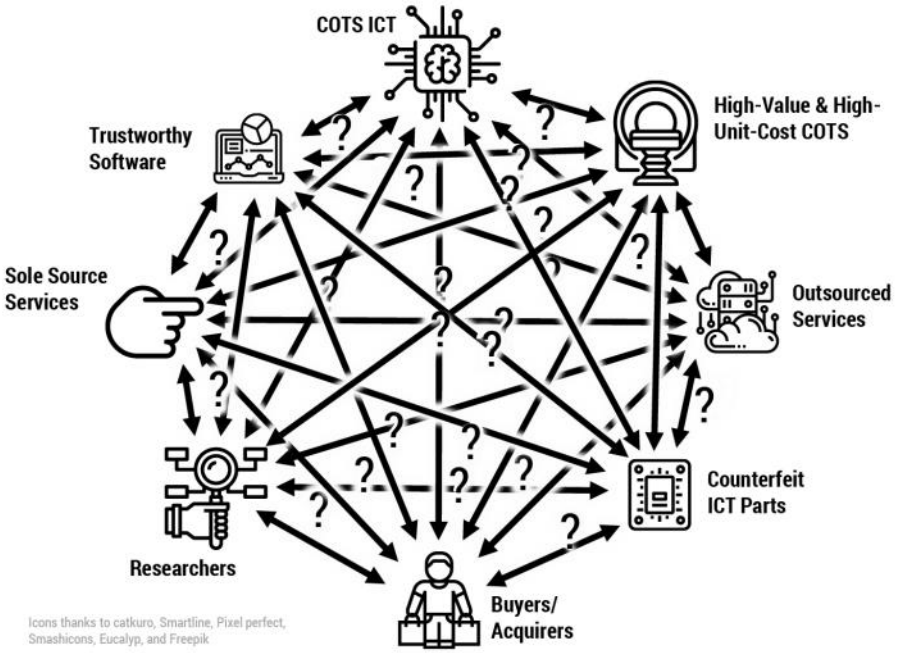

Figure 1 — Confused usage of the terms "trust" and "trustworthiness."

sabotage and counterfeiting arose. This was followed by World War II–era efforts to protect nuclear and other defense industries' production and critical infrastructure products throughout their full lifecycles. Subsequently, we saw the introduction of just-in-time (JIT) manufacturing, where flexibility, predictability, and multi-sourcing considerations came to bear. This created a much-expanded universe of supply chain partners and all manner of supplier qualification issues, including the transitive financial stability risks of those upstream of an organization and the need to trust a supplier to deliver the right quantities on time. If a supplier's supplier is disrupted, the finely tuned JIT supply line might not just falter, but crumble.

In the 2000s, many US federal government practices of supply chain logistics management from the Cold War era were extended into the broader commercial IT marketplace, as those technologies and the efficiencies they brought to business and government started to become key enablers of the information economy. As if this were not enough of a challenge, the computerization of everything gave rise to pervasive cyber threats. For many suppliers serving the US Department of Defense (DoD) with commercial goods, the concept of a "cleared industry partner" became part of their way of life. In parallel, the globalization of ICT manufacturing and markets, coupled with security concerns such as the growth in export licensing of sensitive US technologies, extended the definition of, and areas of concern with, the supply chain. Consequently, visibility and control of highly complex microelectronic component supply chains and their global network of specialized software developers has become a difficult, perhaps impossible, technical management problem.

Naturally, the resulting potential for disruption and harm from the supply chain has become a key concern at the personal, organizational, and societal levels. Such supply chain–rooted risks require specific attention and different perspectives about an organization's practices across the many business functions of both commercial enterprises and government. Figure 2 illustrates these disparate aspects of supply chain security and the variety of organizational elements that need to work together effectively to address them.

Today's organizations must coordinate and foster cooperation across their various business functions to enumerate, measure, and manage their supply chain risks. In particular, techniques addressing the supply chain risks to software and software-enabled ICT components do not easily fall within the traditional

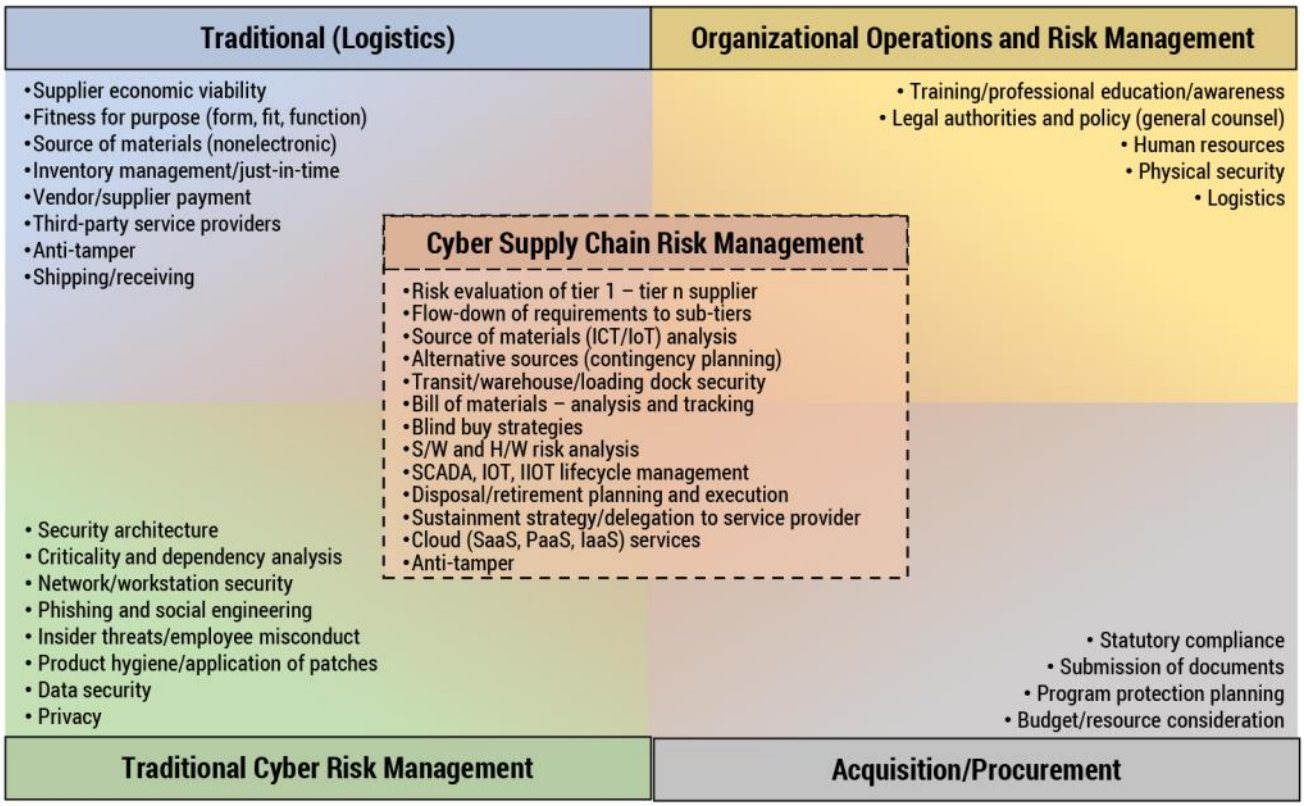

Figure 2 — Elements of practice and multiple aspects of supply chain security.

logistics supply chain practices applied for the protection of physical goods and services in transit, although they share some commonality.

Supply chain security is the superset of these different elements. Depending on the nature of a product, service, or provider, players from one or more of these elements may take the lead in addressing the current set of concerns and risks. However, an organization should be ready to shift focus to other aspects of supply chain security as circumstances change (e.g., the item is now installed and needs servicing), such as shifting focus from acquisition to operations risk management and cybersecurity. The relationship between supplier and customer also inevitably changes over the lifecycle of an item and its role in an organization. An organization may even fill various roles simultaneously, as an acquirer for one item, an integrator for another, and the supplier for a third.

Usually, the discussion of a supply chain will start with discussing immediate suppliers (often referred to as tier 1 suppliers), but when addressing many of the concerns with suppliers, we need to consider the suppliers' suppliers as well (referred to as tier 2), and so forth, all the way to tier n. The realm of supply chain concerns includes both intentional and unintentional acts. So not only do we need to consider someone inserting tainted or counterfeit items, but also someone accidentally introducing poor-quality items due to errors or poor-quality checks of hardware, software, and/or services. A system of trust, therefore, needs to have a pervasive, holistic approach to everything that can be of concern; only then can the SoT be truly effective in supporting the management of relevant items of concern.

Guiding Goals for Building the SoT

As these past efforts have grown, evolved, and bumped into each other, the need for consistency in terminology and harmony in approach have become increasingly important. Today, enterprises of all sizes need to interact and work with a diverse community. Given the diversity of approaches used to manage and consider supply chain issues, the current likelihood that any two organizations will have the same considerations and sets of concerns is limited. Consequently, there will be wasted time and energy as they struggle to understand each other's expectations and abilities to provide and utilize trustworthy products and services that can be trusted.

The SoT's four objectives are as follows:

1. **Gather and organize a single structured corpus of the concerns** surrounding trust with respect to organizations, products/components, and service offerings.
2. **Capture the techniques and mechanisms** available to gather objective evidence as to whether the concerns with trust are or are not substantiated.
3. **Provide a mechanism for winnowing and tailoring the overall SoT** to a custom set of concerns and investigative questions that considers the resources of an organization, the significance of the system or service to its operations, and the consequences that could result from failing to fully vet concerns.
4. **Put in place objective scoring mechanisms** that an organization can adapt to its priorities, operational sensitivities, and experience with its type of business and partners.

A system of trust needs to have a pervasive, holistic approach to everything that can be of concern and needs to be truly effective in supporting our management of all items of concern.

In fulfilling the first objective, one goal of the MITRE SoT is bringing alignment and harmony to the array of concerns organizations — particularly large acquirers of ICT in government and in a country's critical infrastructures — may have with the supplies they receive, the services they utilize, and the suppliers they rely on, so that the entire community has a common holistic set of concerns to consider. Additionally, the SoT seeks to inform users of the *present trust posture* of an entity and its products by means of a static assessment, using an objective set of metrics. The SoT is also constructed to provide a dynamic view supporting a *forward-looking prognosis of the trust posture over time*. This temporal aspect will make a SoT assessment more useful, but also more challenging to validate.

The second and third objectives focus on the process we are using to design and implement tooling to support the operational use of the SoT and assessments against it. The next section discusses the data model created to support these objectives and the piloting being established to validate that the SoT can meet these objectives when applied to real situations.

To fulfill the fourth objective, the SoT needs to:

- Be the basis of a tailorable assessment that focuses on the context of a specific organization and the supply and service items it is examining
- Be able to adjust the importance of the consequences of the risks to focus on the context of the operational capability into which the supplies/services are being integrated
- Be widely applicable to different sizes and types of organizations and domains of businesses and operations
- Be widely adoptable by regulated and unregulated entities, as well as governmental organizations

Acknowledging the vast collection of past and current supply chain security work from and the variety of questions each effort has established, the MITRE effort is both building from the ground up and from the top down.

By meeting these objectives, the SoT will be well positioned to become the generally accepted principle for supply chain security, similar to the generally accepted accounting principles (GAAP) used in all businesses in the US or the globally equivalent international financial reporting standards (IFRS).

The SoT Creation Approach

The SoT is an amalgamation of MITRE's decades of experience in supporting its customers' efforts to address supply chain issues, along with its engagement in various national and international standards efforts on the topic and all the associated published work from MITRE, the standards developing organizations, and others working in various aspects of supply chain security. The SoT permits us to transform that knowledge base into a coherent whole that serves as a taxonomy and implementation model. Acknowledging the vast collection of past and current supply chain security work from across the world and the variety of questions each effort has established, the MITRE effort is both building from the ground up (using individual questions) and from the top down (using supply chain risk taxonomies).

Moreover, the SoT effort is arranging unilateral non-disclosure agreements (NDAs) with a variety of industry organizations. These unilateral NDAs are almost identical to the ones MITRE previously used to create the Common Weakness Enumeration (CWE) body of knowledge.[5] Like other community efforts MITRE has pursued, feedback and engagement with experts and practitioners will be critical in creating something that promotes integration, eases the burden of working with others, and streamlines communications and understanding when multiple parties discuss a topic or share information.

In addition to the community published materials, MITRE's past internal work, and the knowledge shared through the unilateral NDAs, the SoT will also integrate legal-focused content to allow for the SoT to identify areas of concern that an organization might be obligated (or precluded) from investigating.

By collecting the various sets of measures and questions currently in use across the various aspects of the supply chain security community depicted in Figure 2, MITRE is establishing a wide-ranging set of risk-determination questions for input into the SoT. However, rather than just have these as open-ended questions or questions requiring subject matter experts (SMEs) to interpret when used, the SoT will refine them into a series of yes/no questions that embed SME experience into the questions themselves so that less experienced practitioners can use them to identify where and at what level one might trust a supplier, its supplies, or its service offerings.

To illustrate this, consider the different modes of taking a measurement of human blood pressure. If we ask only for the systolic and diastolic measurement, there is an implicit requirement for someone with appropriate training and understanding to interpret the measurements (the evidence) into a finding of "normal," "elevated," "hypertension stage 1," "hypertension stage 2," or "hypertensive crisis." If, however, we ask about these as separate range-based questions that are yes/no–based, the medical expertise is embedded into the questions themselves, for example:

- Is the systolic number less than 120 and the diastolic less than 80?

- Is the systolic number higher than 180 and/or the diastolic higher than 120?

To the extent possible, the supply chain security SoT will follow this paradigm of embedding informed risk expertise into the questions so that the SoT produces objective metrics that can easily be used by those who are not supply chain risk SMEs.

Rather than collect these items into a spreadsheet, word processing document, or some other unstructured form, the SoT will use a knowledge base along with an application tool to support viewing, organizing, and tailoring the content or subset of the content within the SoT. Further, this application tool will support evaluating a supplier, supplies, and services against the tailored SoT subset, as well as support adjusting the contributing weights of specific SoT questions and areas of concern. The SoT tools will focus on assessing the supply chain security concerns captured in the SoT taxonomy against specific supply items, service offerings, and companies. The SoT tools should not be confused with general risk management tools in common use at many organizations.

The SoT Knowledge and Evaluation Model

As part of bringing together the vast past work in supply chain–related risks and concerns and to support reviews, tailoring, and evaluation, the MITRE SoT effort has a knowledge and evaluation model at its core. This model leverages the work within the standards community[6] for representing measures[7] where high-level categories of the taxonomy are broken into lower-level collections of subcategories as appropriate and then into factors and measures for those factors, as shown in Figure 3.

Through the use of this model, the SoT will bring alignment and consistency for the concerns it captures and the risk questions it incorporates, as well as support the contextual relationships between some aspects of the risk measures themselves and how one risk measure may potentially contribute to another. Specific scoring mechanisms, weights, confidence indicators, and measurement prerequisites should provide for wide adoption in diverse industries and organizations.

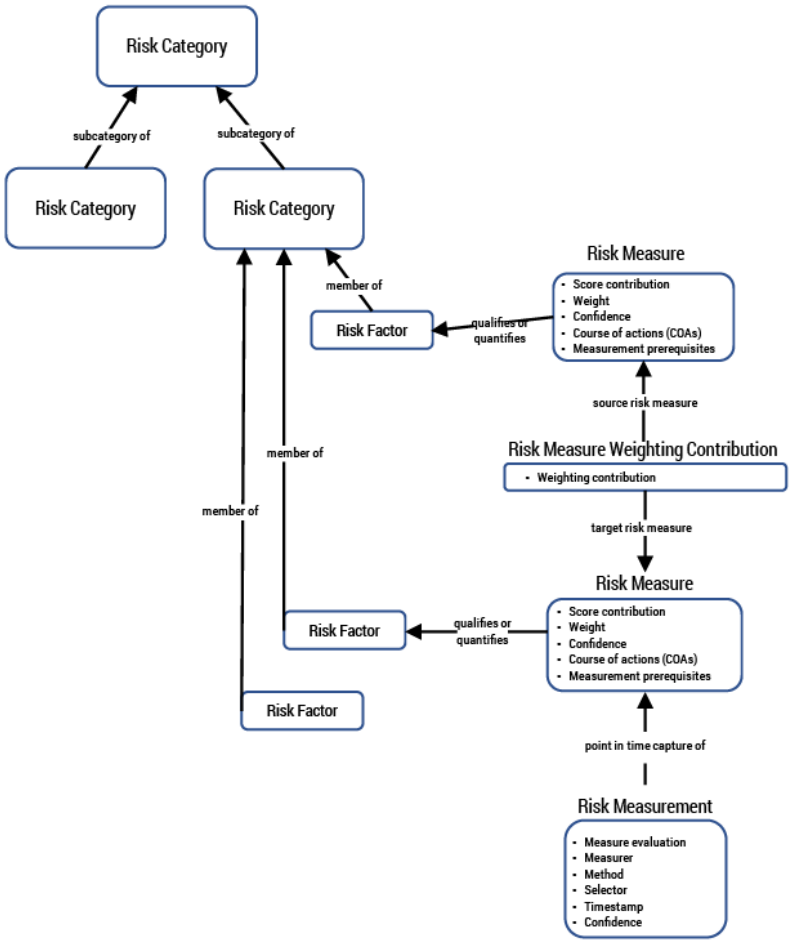

Figure 3 — Initial SoT knowledge and evaluation model.

The SoT Draft Structure

The initial version of SoT tooling mechanics is still in final development at the time of this writing; however, essential aspects of the analytical tool have been developed. The top-down portions of the SoT taxonomy are currently represented by elements defined as "concerns about supply chain security," which are then decomposed into: (1) concerns about suppliers, (2) concerns about supplies, and (3) concerns about services.

Figure 4 illustrates the first few levels of the draft SoT taxonomy as a mind map. This draft taxonomy is a reasonable generic starting point but will continue to evolve as other organizations contribute to it.

The Concept for Applying the SoT

Applying the SoT as an evaluation framework and tool will challenge organizations on both sides of any supply chain transaction to confront deeply embedded characteristics of their corporate behavior and culture. Customers/acquirers will use a variety of information sources to provide themselves with the types and depth of data they need to enumerate risks and build trust. That data may force customers/acquirers to adjust their internal risk management strategies to compensate for whatever they discover is blocking their ability to trust. Product/service providers should pay attention to what public and private information is being collected and shared about themselves, their products and activities, and their own suppliers. Providing evidence that demonstrates trust-building behavior will increase

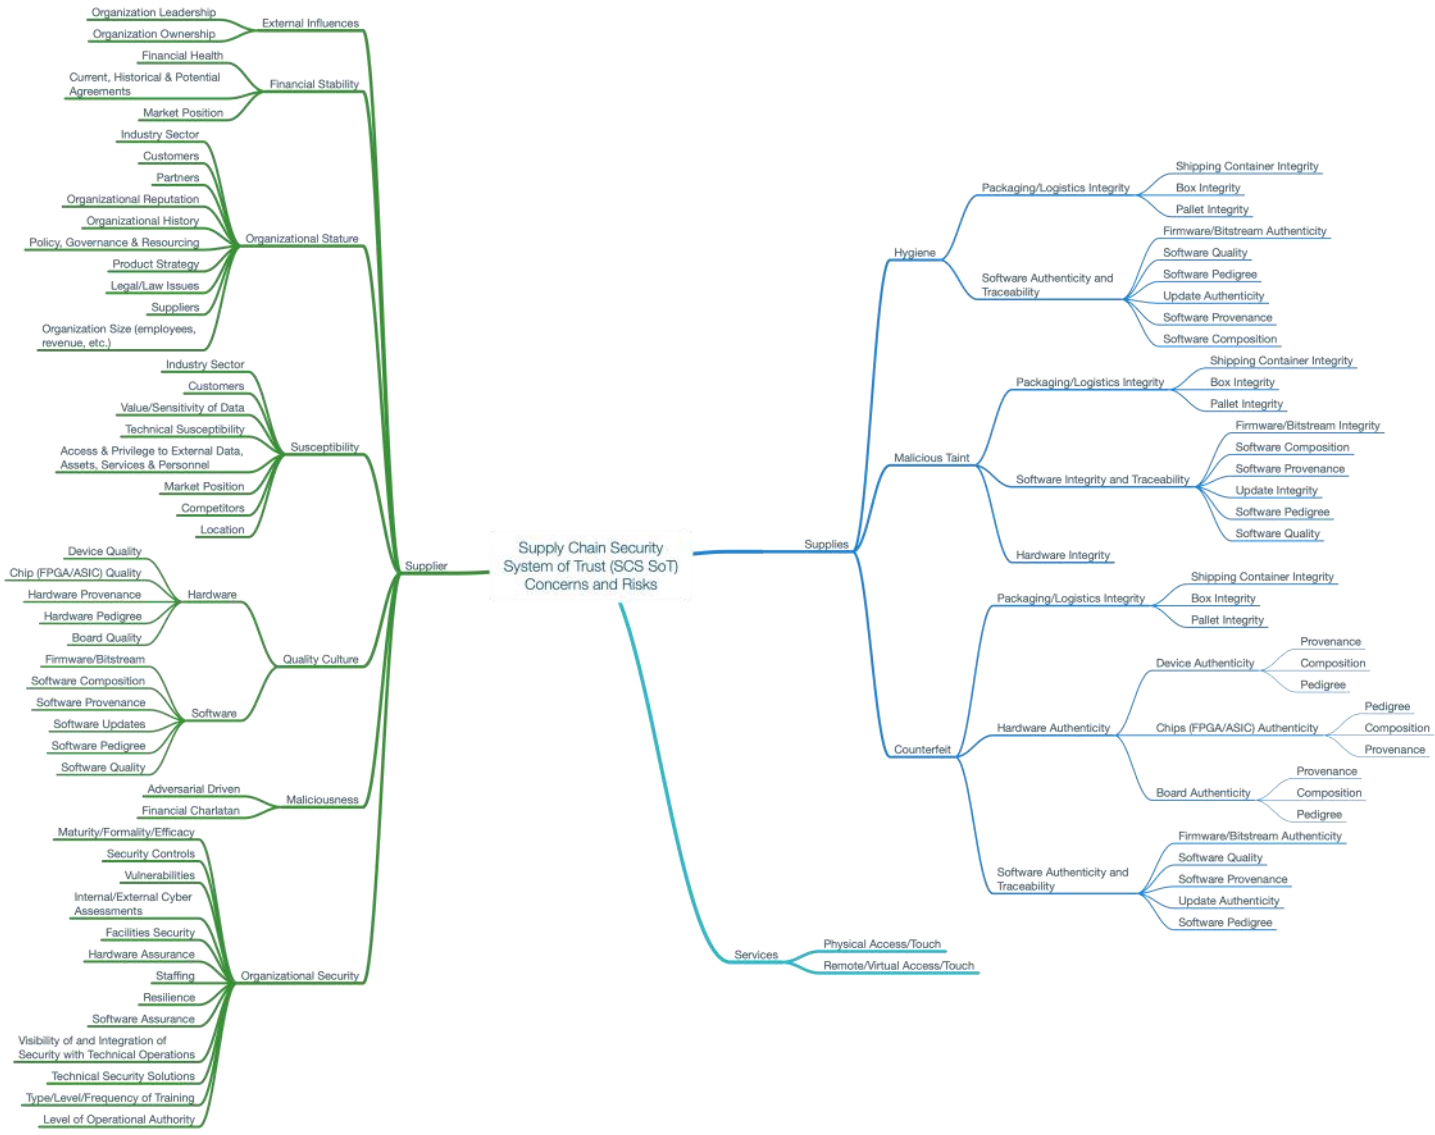

Figure 4 — SoT draft taxonomy of concerns and risk.

the attractiveness of their company as a supplier to those seeking to acquire items using the SoT.

For example, every organization faces a challenge in how to balance leadership goals, employee morale, and consistency of behavior across the organization with the financial incentives the organization offers to achieve those behaviors, goals, and the desired morale level. The outward perception of these characteristics shapes the ability of the organization to radiate trustworthiness to external stakeholders. Many consumers care only about the features and costs of an organization's products and services. But supply chain intermediaries, including those that represent the needs of consumers, such as governmental and industry purchasers, should look beyond these basic, consumer considerations to assess the trustworthiness of a potential supplier. Even those suppliers with whom the acquirer has an existing relationship should be periodically assessed for their continued trustworthiness as an organization, to include their supply chain.

Beyond the obvious factors of features and cost, what remain are the characteristics addressed when applying the SoT. An assessment of trust — and of its counterpart, risk — should be an essential step in the acquisition process. Applying the SoT will assure that all parties to a transaction use a consistent assessment process and vocabulary.

Many of the factors in building trust and managing risk are not quantifiable. Applying the SoT will require the acquirer to compare its own organizational values to the demonstrated or acquired data about its potential supplier's values. Customers seek evidence to trust or not trust organizations and to quantify risks to their own organizations. Applying the SoT, evaluating the resultant set of questions, and utilizing their weighted scoring will provide the foundation of a defensible rationale for an appropriate level of trust by assembling enough data points to distinguish those suppliers most worthy of trust.

Validating the SoT

To fulfill its envisioned role, the SoT needs a robust set of verification and validation (V&V) processes to assure users of its value in driving adoption of trust-building strategies for all ICT supply chain participants. During development, these processes will incrementally and methodically ensure that the indicators of both risk and trust make sense. Using internal, data-driven examination, benchmarking, and testing, as well as external application and critique, the V&V process will ensure the SoT is usable in its initial release and over time.

The Way Ahead

In addition to continuing to build the SoT, socialize the concepts it entails, and gather ideas for additional questions that explore how to make concrete the concerns organizations have with their suppliers, supplies, and services, the SoT effort will continue creating an assessment tool. This tool will allow an organization to tailor the SoT to just those select areas of concerns an organization feels are the most useful for its decision making. Eventually, it will support tuning of the weights and combinatorial mechanisms used to combine the individual answers into an overall trustworthiness finding. Finally, it will enable the exploration of mechanisms used for answering the questions, whether manually, through data analysis tools, or with both quantitative and repeatable qualitative methods.

At the same time that we explore these analytic approaches and methods, we will be identifying which of the questions can be answered with public or private information sources and exploring how to document these possible sources of evidence about those aspects of the SoT.

Many communities exist for the topics being amalgamated under the SoT. Their respective past, current, and future work will need to be explored and integrated, where appropriate, so that the SoT and those efforts can coexist, proffer synergy, and ensure their respective work, foci, and best practices are clearly understood.

Finally, there will be an immediate and evolving need to explain and train individuals and organizations on the use and utility of the SoT. Additionally, many of those currently working in the areas touched by the SoT work will, we hope, evolve their efforts to align with the integrated SoT vocabulary and concepts as we collectively solve the ICT trust and trustworthiness challenge, with a shared taxonomy and methodology and the SoT tooling (more ICT).

A Piloting Plan for the SoT

In view of both the novelty of the SoT concept as an element of supply chain security abilities and the anticipated broad scope of interest from organizations across industry and government, together with the broad applicability of the SoT's concepts, taxonomy, evaluation methodology, and tool set, we have decided to embed the practices of frequent testing, expert evaluations, and interim exercises (both mini/tabletop exercises and full formal simulations) to test assumptions about the structure, functioning, and applicability of the SoT's elements to the full technology lifecycle.

The SoT must be designed as a living system that evolves with changing needs. In addition to these continuing practices of testing and evaluation throughout the creation and maturation process of the SoT, there will also be two or three pilot implementations of multiple weeks to provide further evidence of functionality and to support further editing and restructuring, gap analysis, and overall refinement.

The first assessment pilot will focus on an acquisition of a large operational technology–focused system that is about to conclude. The planned pilot will analyze the procurement artifacts and compare the scope and coverage of issues considered by the just-concluded acquisition with what the SoT considers. The aim is to identify possible after actions for application to the winning solution, as well as to provide ideas for revisions to subsequent acquisition activities.

The second assessment pilot will involve a set of enhancements where a planned set of COTS and open source items are being integrated for use across a large community. The pilot will assess a widely used platform capability from the perspective of supply chain security concerns about the supplier and the software of the capability itself.

The third assessment pilot will address the gamut of cybersecurity and other ICT supply chain security challenges that many organizations face when looking to purchase COTS items. The pilot will assess a widely used commercial capability from the perspective of supply chain security concerns about the supplier, the software of the capability itself, and offering the service to the world.

Finally, a regular cycle of exercises and user evaluations is intended to be a signature element of the SoT's "maturity model" during its future life. Ongoing exploration of applying the SoT to other acquisitions, suppliers, and products/services will occur in other transactional environments, including various intergovernmental and critical infrastructures transactions, as already described. The flexibility and adaptability of the SoT are expected to become key elements favoring SoT incorporation and adoption in these diverse transactional settings.

Acknowledgments

The summary work contained in this article is based on discussions with a number of individuals at MITRE and throughout the industry. The author offers a special thanks for the contributions of the MITRE SoT team, which provided valuable input to this article.

References

[1]Industrial Internet Consortium (IIC), 2020.

[2]Plattform Industrie 4.0, 2020.

[3]IIC (see 1).

[4]Plattform Industrie 4.0 (see 2).

[5]"Common Weakness Enumeration, Version 4.0." MITRE, 2020.

[6]McGarry, John, et al. *Practical Software Measurement: Objective Information for Decision Makers*. Addison-Wesley Professional, 2001.

[7]"ISO/IEC/IEEE 15939:2017, Systems and Software Engineering — Measurement Process." ISO/IEC/IEEE, 2017.

Robert A. Martin is a Senior Principal Engineer in MITRE's Cyber Solutions division. For the past 38 years at MITRE, his efforts have focused on the interplay of enterprise risk management, cybersecurity assessment standardization, critical infrastructure protection, and managing risks from software-based technologies and services. Mr. Martin is a member of the Industrial Internet Consortium (IIC) Steering Committee, the Consortium for IT Software Quality (CISQ) Advisory Board, and the Object Management Group (OMG) Board of Directors. Within IIC, he co-leads the Trustworthiness Group and is a frequent speaker in many public venues. Mr. Martin is an ISC2 Certified Secure Lifecycle Professional; a member of ACM, AFCEA, NDIA, INCOSE, IEEE, and IEEE Computer Society; as well as a contributor to standards in ETSI, SAE, The Open Group, UL, OMG, ISO, and the ITU-T. He earned a bachelor's and master's degree in electrical engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He can be reached at ramartin@mitre.org.

Cutter Consortium Membership
The Ultimate Access to the Experts

Cutter Consortium Membership opens up multiple avenues to interact with Cutter's experts to brainstorm and gain guidance to transform your organization and boost success.

Like everything business technology, one size does not fit all.

That's why we encourage you to choose the Membership that's right for your organization. Whether you choose Digital Transformation & Innovation, Enterprise-wide, Practice-specific, or CIO Membership you'll see a strong return.

Contact us at +1 781 648 8700 or sales@cutter.com to arrange a sample inquiry call with a Cutter expert and see for yourself how quickly your return on Membership can be realized.

What Do You Get from Cutter Membership?

- Get guidance in leveraging new strategies, emerging technologies, and business management practices to enable digital transformation and boost competitive advantage
- Learn how to mine data to create new products and services and improve customer experience
- Get input on how to reduce expenses through more cost-effective strategies
- Gain insights and get ideas on achieving sustainable innovation, successful change management, and prudent risk management
- Get coaching and insights on leadership and team-building practices that boost productivity
- Discover vendor-agnostic advice to ensure unbiased purchasing decisions

Test Drive Cutter Consortium Today

Have a question? A challenge you're trying to overcome? Have you reached a bottleneck and need some expert advice? Find out for yourself how valuable Cutter Membership can be — request a sample inquiry with a Cutter expert now.

Call +1 781 648 8700 or email sales@cutter.com.

CUTTER CONSORTIUM
Access to the Experts

37 Broadway, Suite 1
Arlington, MA 02474 USA
Tel: +1 781 648 8700
Fax: +1 781 648 8707
Web: www.cutter.com
Email: sales@cutter.com

"I have personally been able to leverage Cutter's services since 1999. Among the attributes that differentiate Cutter from other firms, two remain at the top of my list: **Thought Leadership and Real Value**. Thought Leadership is driven by Cutter's experts. Real Value, driven by 'Access to the experts,' pushes us beyond just understanding the options. We can develop relationships with the experts, and tailor the options they present so that they can be quickly and practically executed within our organization, enabling our Business Technology team to continually improve, engage, and contribute to business growth."

— Doug Mikaelian
VP Business Technology,
Dairy Farmers of America

"The information technology revolution continues to change how businesses work and how people work in and with them. Today's demands on software to provide useful, usable, secure, and scalable services stress all aspects of systems development. These include architecture, technology selection, product development, and lifecycle management. Cutter provides comprehensive coverage for these crucial topics — through research reports and customized virtual training delivered globally — helping us to understand needs and solutions that span traditional silos."

— Eric Schoen
Director of Engineering, i2k Connect LLC
Former Chief Software Architect,
Schlumberger

"Thanks to Cutter´s experts, in a short time we have been able to:

- Improve our capabilities, using the contract scorecard for outsourcing
- Gain key insights and practical strategies for responding to challenges we face in building trust and partnerships within the organization and outside our institutional borders.
- Improve our methods for working with people and anticipating the challenges of cultural transformation within our organization.
- Have immediate access to key 'experts' at critical times, who spend quality time with our teams, provide state-of-the-art orientation, and help us implement projects.
- Implement agile, effective, and articulated innovation around a common vision for achieving quality education in Mexico."

— Arturo Cervantes
Director General of Information Systems
for Evaluation Results, National Institute for
Educational Evaluation, Mexico

CUTTER Business Technology Journal

About Cutter Consortium

Cutter Consortium is a unique, global business technology advisory firm dedicated to helping organizations leverage emerging technologies and the latest business management thinking to achieve competitive advantage and mission success. Through its research, training, executive education, and consulting, Cutter Consortium enables digital transformation.

Cutter Consortium helps clients address the spectrum of challenges technology change brings — from disruption of business models and the sustainable innovation, change management, and leadership a new order demands, to the creation, implementation, and optimization of software and systems that power newly holistic enterprise and business unit strategies.

Cutter Consortium pushes the thinking in the field by fostering debate and collaboration among its global community of thought leaders. Coupled with its famously objective "no ties to vendors" policy, Cutter Consortium's *Access to the Experts* approach delivers cutting-edge, objective information and innovative solutions to its clients worldwide.

For more information, visit www.cutter.com or call us at +1 781 648 8700.