Shifting Paradigms: Using Strategic Foresight to Plan for Security Evolution

Heather Vescent
Futurist
puissant@heathervescent.com

Bob Blakley
Citigroup
bob.blakley@citi.com

ABSTRACT

Everyone wants to know the future. Exploring the future is not to make predictions, but to anticipate which futures might happen, so we may make better decisions today. The foresight process is a tool which enables researchers to become more attuned to the future; the foresight process is designed to help anticipate emerging trends, rather than be surprised by dramatic change. This paper identifies possible futures in two ways: extrapolating from the present into the future, and starting from future states to reconstruct how they might be arrived at from the current state. The result is 12 fresh scenarios and 13 new paradigms.

CCS CONCEPTS

• **Security and privacy** → **Economics of security and privacy**; **Social aspects of security and privacy**; **Usability in security and privacy**; • **Social and professional topics** → *Governmental regulations*; • **Computing methodologies** → *Modeling methodologies*;

KEYWORDS

Information Security, Strategic Foresight, Paradigms, Future, Scenarios

ACM Reference Format:
Heather Vescent and Bob Blakley. 2018. Shifting Paradigms:, Using Strategic Foresight to Plan for Security Evolution. In *New Security Paradigms Workshop (NSPW '18), August 28–31, 2018, Windsor, United Kingdom.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3285002.3285013

1 INTRODUCTION

Computer security was described as a "dumpster fire" by some security professionals we interviewed. Others opined that it's in an "abysmal" state that is "often put together with the technology equivalent of duct tape".

The attack surface of computer security is increasing rapidly as we introduce new technologies and further integrate technology into all aspects of our lives, from cars to pacemakers. To date, security has frequently been bolted on rather than designed in from the beginning. As a result, each year we have more hacks and more data breaches than ever.

NSPW '18, August 28–31, 2018, Windsor, United Kingdom

ACM ISBN 978-1-4503-6597-0/18/08…$15.00
https://doi.org/10.1145/3285002.3285013

This research is an initial look at what security could be in 2038—20 years from the date this paper was written. The following questions formed the kernel of our research:

- What is changing in the security world to improve or cause unintended consequences today?
- What has changed about what needs to be secured?
- New technology gives us new power and solves problems of the present, but it opens new problems in the future. What are the security problems of the next 20 years?
- How do we change the thinking of security professionals to better consider problems of the future?

The paper is structured in three parts. Part one sets out our methodology, including the foresight methods we used as well as other futurist research methods that could be used in the security industry. Part one also enumerates our baseline assumptions, including a description of the state of security today and the results of our Future of Computer Security Survey. It sets the stage for part two, which dives into a set of future scenarios and a discussion of the forces that might lead us down the path to each of the scenarios. Each scenario is based on changes to the current state via a common set of variables. We start by identifying the variables, and we explain how adjusting the variables gives rise to each scenario. Part three presents a dozen alternative paradigms, each of which would contrast with today's existing paradigm in specific, identified ways.

2 METHODOLOGY

For this research, we used methods from the discipline of strategic foresight to understand and analyze trends and identify possible futures in the information security space.

This foresight methodology is flexible in that it can be applied to many industry areas alongside traditional research to gain insight. Foresight research supplements traditional research activities and extends the data to explore possible futures. The reason one should spend resources exploring possible futures is not to make predictions, but to identify among a set of possible futures the future which is most desirable, anticipate which futures might happen under the influence of various forces, and gain insight into the action of forces on the futures in order to make better decisions in the present. Insights gained from foresight research can help practitioners more clearly see the operation of forces and more clearly identify trends; this in turn helps practitioners make effective and timely decisions to influence the future, rather than being surprised and forced into change belatedly.

There are a variety of foresight methods in use, including:

- Delphi expert panel [24]
- Casual Layer Analysis (CLA) [13], [12]
- Futures Wheel [21]

- Backcasting [9]
- Character based narratives [22], [8]
- Appreciative Inquiry [CW08] [6]
- Foresight interview protocol [4]

For this research, we used the Foresight Interview Protocol, Appreciative Inquiry, Backcasting, a light version of Casual Layer Analysis, and Scenario Development. We took existing research methods of surveying and interviewing and applied a foresight lens to understand the past, current conditions, and possible futures.

While it is easy to identify problems and go into problem solving mode, it is harder to identify the things that already work well. We pay attention to the squeaky wheel that needs the oil, not the well oiled machine. Appreciative Inquiry is a powerful method used to identifying existing positive characteristics. Once these are identified, we can support activities and decisions that increases the power and influence of these positive characteristics. Rather than identify and solve problems, we identify what is already working and increase what makes it work so well.

Initial Expert Interviews: Five interviews were conducted with security experts using foresight questioning. These questions focused on the current conditions, the impact of certain trends on the future, other changes impacting the space, and speculative questions about security problems and solutions in the future. The results of these interviews were used to create a survey, as well as to guide scenario development.

Survey: We formulated a 20-question survey and solicited participation using Twitter and via direct solicitation. 89 security practitioners completed the survey. The questions included a combination of queries about demographics, current security activities, tools used, attack surfaces, and "keep you up at night" concerns, as well as speculative questions about future security problems, new technologies, and possible solutions. Open ended survey questions gave us insights to be used when creating the scenarios. The results were analyzed and used to guide the development of our scenarios.

Limited Historical Analysis: Since this topic is focused 20 years in the future from 2018, we wondered what the view from 20 years in the past might have been. To get an idea of this, we did a brief historical analysis of the papers from NSPW 1998 and we asked ourselves, what problems from 1998 have we solved? Which ones have changed? Which ones have stayed the same? We used the result of this analysis in our scenario development.

Appreciative Inquiry Interviews: We conducted one-on-one interviews with 11 individuals using a standard set of appreciative inquiry questions customized to the computer security domain. This data was analyzed, with similarities grouped and outliers identified. Appreciative Inquiry focuses the query on what is positive, energetic and already working. Our questions focused on the positive energy of the space, the professionals, existing breakthroughs, the problems already solved, and the dreams for the future of the security industry. An analysis of this data was used to develop the scenarios and identify existing paradigms, key properties of which we "flipped" to identify desirable new paradigms.

Scenario Development: The findings from all of the research activities drove the development of the 12 scenarios. We started with a baseline scenario, which assumes that nothing important changes in information security in the next 20 years, and which serves as the starting point from which changes to infosec-relevant variables are applied to create alternative future scenarios. We identified some scenarios organically from the research and "reverse engineered" the changes to our variables which would lead us into the scenario; we created other scenarios by analyzing what would happen in the future if specific adjustments were made to identified collections of variables.

New Paradigms: When analyzing the Appreciative Inquiry data, several current paradigms were clearly identified. Despite the positive outlook of appreciative inquiry, many interview subjects identified negative aspects of the current paradigm. We took this as an opportunity to identify better paradigms, by flipping the negative properties of the existing paradigm to create new positive paradigms.

We then employed "backcasting" [9] which takes the desired paradigm as the end point in a strategic plan and looks at what actions need to be taken to make that future occur.

3 RESEARCH FINDINGS

3.1 Initial Expert Interview Results

We started our research by asking a selected group experts to describe the current state of computer security. The experts' answers were not too positive; the following descriptions are representative:

- It is a dumpster fire.
- Solved with duct tape and bailing wire.
- Some companies are starting to take information security seriously.
- Sometimes has negative effects, for example, chasing shiny bug bounties.

But the experts hadn't given up hope; they suggested a number of innovations they thought might be productive:

- "All source code is public": One subject proposed that all source code should be readable and public, even if it's not open source for the purposes of reuse.
- "Security lawsuits": Another subject suggested that companies be required to take legal responsibility for the result of their insecure code.
- "Regulation": Another subject proposed that software be regulated and companies be allowed to push only "security verified" code into production.
- "Augmented humanity": Another subject proposed removing the "human" from human computer interaction by adding a layer of human-augmentation AI between software's HCI layer and the human to improve the human's performance in using the technology safely and securely.

A consistent theme in our research interviews was the influence of economic factors on information security. Interview subjects mentioned:

- **The cost of doing security.** What is a business's ROI when it is trying to get a product to market ASAP? Security is not emphasized in the economic trade offs of many companies; rather than doing the right thing security wise, companies risk the bad PR a data breach brings and litigate against security researchers who find vulnerabilities. Corporate incentives for security are often muted because of internal

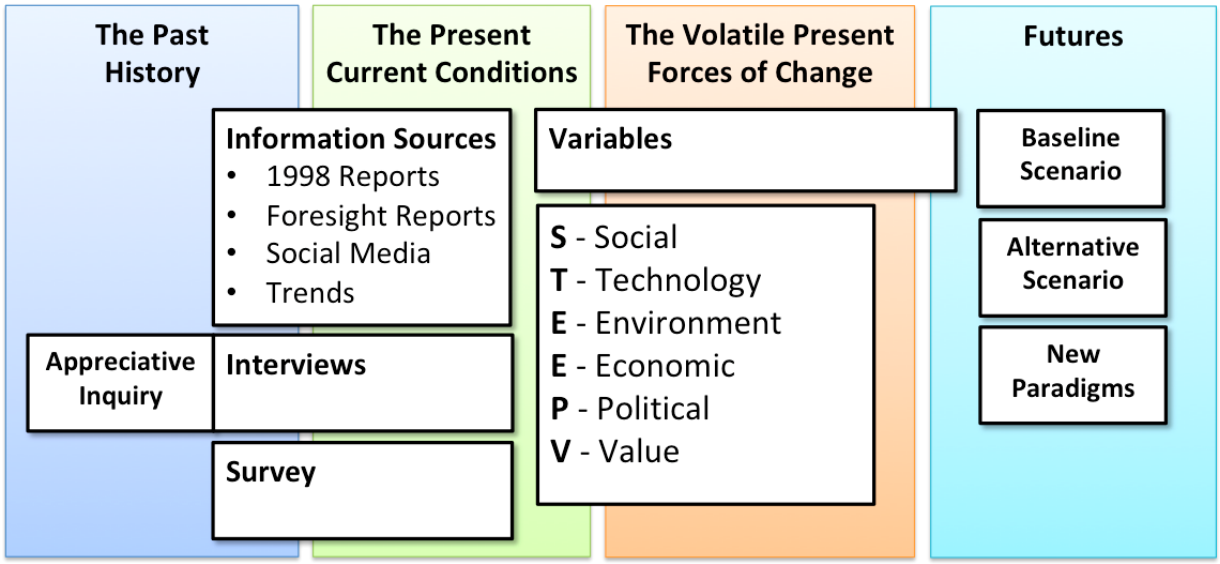

Figure 1: Foresight Research Methodology. Adapted from Dr. Peter C. Bishop, Studies of the Future [4].

reporting structure. But security is a commons from which everyone benefits when done well. Unfortunately, when it is done well, you don't hear about it. Only when it is done poorly does a company get PR.

- **Time rich, money poor.** There is an increasing number of highly educated, but underemployed or unemployed people with security skills. Idle hands are the devil's workshop; people with skills and time, but without money, have incentives to become black hat hackers. This is prevalent in countries that have good education systems but poor job markets, including especially Brazil and the former Eastern Bloc. We wonder if some security problems might be solved by merely having full employment for all educated security experts.
- **Finite vs Infinite Rules.** Building upon the previous point, the security theatre is uneven, with defense bounded by finite rules due to limited resources and attackers utilizing resources (time, money, tenacity) to find vulnerabilities. This is an uneven playing field that economic solutions may change.
- **The economics of computing.** People expect their water heater to last 30 years, but most technology companies don't last 10 years. What is going to happen to the many IoT products that will be built with this throw-away attitude? These products will have security vulnerabilities but will have long lifetimes in the world with minimal, if any user interfaces to upgrade their security. Economic incentives can put pressure on the need to ensure technology lasts and is secure in this constantly changing space.

3.2 Survey Results Summary

After analyzing the expert interviews, we designed a 20-question survey and distributed it to security professionals via Twitter and via direct solicitation. The survey generated 89 responses; this section summarizes the survey results.

Responses represented a broad cross-section of demographic categories, from brand new practitioners and students to professionals near retirement age. About 25 percent of respondents were female, but this is partly selection bias as we deliberately reached out to female practitioners toward the end of the survey period as the initial sample was disproportionately male. We got at least two non-binary respondents. We did not ask about ethnicity or national origin. The distribution of respondents across years of tenure in the industry is very even, with a slight bias toward younger practitioners. We had only two respondents who claimed to be educators or trainers, and only one who claimed to be a researcher.

We included a question about the current state of security; responses to this question were overwhelmingly negative. Only two respondents out of more than 80 gave responses which could be characterized as net positive, and many were strikingly negative.

We asked about top-of mind worries. Unsurprisingly, respondents are worried about a broad range of risks, but the standout finding here was that nation-state attacks were more worrying to respondents than any risk other than data breach.

We asked about what respondents thought the key infosec issues were. In response, our survey participants said that humans are overwhelmingly seen as the key vulnerability.

We asked what the most important new infosec risks are; IoT is the consensus winner for worst new threat vector. AI is starting to be a concern to respondents.

We asked for areas in which respondents are hopeful about the future of infosec. Participants replied that they have significant hope that authentication, and identity, and access management will be solved in the next 20 years.

We asked what problems definitely will not be solved in 20 years. No respondent thinks that vulnerabilities will be substantially reduced or that problems arising from humans will get better in the next 20 years.

We asked what problems respondents would fix if they had a magic wand. They replied that they would fix security education and issues arising from human users.

We asked what new technologies or developments might have positive effects in the next two decades. Respondents replied that AI/ML, new development methods and tools, and increasing diversity in the field are promising trends.

We asked what's likely to get worse. Respondents worry that humans, education, complexity and scale, and incentives are issues that will continue to get worse.

We asked what's changed in the field since the beginning of respondents' careers. They answered the field has gotten more complex, but that it's also now got more publicity and more management focus.

We asked what's changed in individual respondents' roles. They answered that the role now requires broader scope and expertise (though many also said roles are increasingly specialized); respondents also felt that prestige, pay, and visibility to executive management have improved.

We asked what keeps respondents up at night? They answered "IoT and nation-state attacks". And kids. And spiders.

3.3 Appreciative Inquiry Results

After analyzing early survey responses, we began conducting appreciative inquiry interviews to focus on some themes which had emerged from our expert interviews and from the survey, with a view to identify characteristics of the baseline ("nothing much changes") future, and possible characteristics of desirable futures to guide us in scenario development. Findings from Appreciative Inquiry included:

Baseline Future: Things that will stay the same

- The dynamic nature of the industry, there will always be new attacks and new fixes. This industry clearly attracts people who are energized by this dynamic and changing landscape, and the need to learn new things constantly.
- New tools will be developed.
- Users will be users (and they're human).

Alternate Futures: Trends today that could impact the future

- Gamification of tools and SIEM management/response.
- Diversity in gender, race, and cultural methods (consider Asian & Russian attack and defense methods).
- Breakthroughs have not been in technology, they have been in other areas (social, cultural, political).
- Increasing disclosure of vulnerabilities and how they are disclosed (e.g. directly to vendor vs public).

Desired Futures: Things that need to change

- More women and diversity.
- More education. There aren't enough practitioners, so education to prepare new entrants for work in the field is clearly required. This is true for practitioners in all disciplines, from system administrators to SOC analysts and incident responders to software developers. Indeed, as integrated development methodologies like devops become more prevalent, security may become a quality attribute of other jobs rather than a standalone discipline, and education will need to keep up with this evolution. Ongoing continuing education is needed to enable established practitioners to share new information as well as to stay up to date with adversaries' innovation, and network system changes.
- Less litigation targeted at researchers who are sharing vulnerabilities publicly (disclosure). "Hey, we're just trying to help, don't sue us."

3.4 A Retrospective Look at NSPW 1998

As background to help us look forward, we took a look back, specifically in the context of NSPW. Since our target future is 20 years in the future (as of 2018), we looked back 20 years, to 1998, and tried to use the content of NSPW 1998's papers to get an idea of what we did back then that worked, and what didn't work. The point of the retrospective review was to gain insight into whether the "dumpster fire" observed by our experts was a result of a failure of the research community to generate new ideas, or a failure of the practice community to implement new research ideas, or just a consequence of the insurmountable difficulty of the information security problem. We therefore undertook the retrospective with three questions in mind: (1) Is infosec research coming up with workable ideas? (2) Are infosec practitioners coming up with workable ideas which are not being implemented? or (3) Are infosec practitioners implementing workable research ideas, but they aren't solving the problem?

Twelve of NSPW 1998's thirteen papers proposed new models, new metrics, or new mechanisms for securing systems. Six of the twelve proposals - or proposals like them from other sources - saw at least some adoption (a pretty good success rate!). But in spite of this, during the period between 1998 and 2018, vulnerabilities, malware variants, security expenditures, data breaches, and many other measures of attacker opportunity and accomplishment steadily worsened. NSPW 1998 is of course just a small sample of even the academic work on information security - let alone the commercial work - in a single year of the last two decades. But even based on this small sample it's clear that the information security research community isn't running out of ideas, and the information security practitioner community is successfully putting many of those ideas into practice. Nevertheless, the attacker community has grown more successful since 1998. Based on the NSPW 1998 results alone, we lean toward believing that Infosec's status is best described by option (3): the field is coming up with workable ideas, and many of them are being implemented, but they're not solving the problem (perhaps because of the speed at which the problem is getting worse).

Looking back at the NSPW 1998 proceedings [1] in a little more detail, we see the following:

In "Toward A Secure System Engineering Methodology", Salter et. al. [25] propose a methodology for identifying security vulnerabilities, assessing how each vulnerability might be exploited by attackers with known characteristics (based on an adversary model), and judging which vulnerabilities are worth remediating given likelihood of exploitation by identified attackers and cost and complexity of the remediation. Since 1998, several similar approaches have been implemented. OWASP recommends Threat Risk Modeling and identifies several methodologies which can be

used to implement it, including Microsoft STRIDE and DREAD [14] and AS/NZS 4360:2004 (Risk Management) [19]. Mature organizations today use these methodologies to achieve most of the goals Salter et. al. set out in this paper.

In "Security Engineering in an Evolutionary Acquisition Environment", Abrams [2] proposes the integration of the Systems Security Engineering process into software development lifecycles of government agencies via an integration of the Spiral Software Development Lifecycle with Evolutionary Acquisition. Since 1998, software development methodologies have evolved from Spiral to Agile, and Agile methodologies are frequently used in conjunction with API-based environments which enable consumption of security and other services from cloud-hosted service providers. Today's software development lifecycle and security component integration methodologies meet many of the criteria Abrams describes in this paper.

In "An Integrated Framework for Security and Dependability", Jonsson [15] lays out a framework for simultaneously achieving information security and dependability goals, where dependability goals are focused on the system remaining available and working reliably and safely, and security goals are focused on the system resisting attacks. Jonsson observes that security goals and dependability goals sometimes conflict, and describes a unified framework for thinking about how to simultaneously achieve both kinds of goals. Jonsson's methodology hasn't been widely adopted, and dependability and security remain separate disciplines in most commercial enterprises.

In "Meta Objects for Access Control: A Formal Model for Role-Based Principals", Riechmann and Hauck [23] propose the use of a role-based mechanism to solve some problems arising from trying to control access in an object-oriented system with encapsulation. The paper introduces a method of associating roles with object references to prevent unwanted escalation of privilege when objects are passed across encapsulation boundaries between objects operating at different privilege levels. This mechanism hasn't been widely adopted, and managing and enforcing access control in environments with strong encapsulation remains a problem.

In "Evaluating System Integrity", Foley [10] proposes a formal definition of integrity which is intended to capture aspects of segregation of duty, fault-tolerance, and other application dependability properties by considering transactions at the level of an entire enterprise. This formalism hasn't been widely adopted; indeed, formal methods in general are still rarely applied to the development of commercial software.

In his Position Paper, "Prolepsis on The Problem of Trojan-Horse-Based Integrity Attacks", McDermott [18] observes that integrity mechanisms deployed in 1998 were not effective, but argues that other known mechanisms, including replication, session replay, and pre- and post-condition checks could work. Some of these approaches have subsequently been tried in earnest (sandbox solutions for example, implement session replay), but Trojan Horse attacks remain an unsolved problem.

In "Death, Taxes, and Imperfect Software: Surviving the Inevitable", Cowan and Pu [7] propose a notion of security bug tolerance as an alternative to the high-integrity Trusted Computing Base model advocated by the Orange Book, and demonstrates a method for categorizing how security-bug-tolerant a system is. The security bug tolerance metric has not been widely adopted, and bugs and vulnerabilities continue to increase.

In "A Graph-Based System for Network-Vulnerability Analysis", Philips and Swiler [20] propose a method for analyzing networks to identify attack paths along which attackers have a high probability of success in compromising the target network. Since the publication of this paper, MITRE has developed a method of breaking down attacks into phases (the "Kill Chain"), and several commercial vendors have developed products which use network topology, system vulnerability scan information, and asset inventory information to identify attack paths in a way similar to that proposed in the paper. Nevertheless, successful system penetrations continue to occur.

In "Parsimonious Downgrading and Decision Trees Applied to the Inference Problem", Chang and Moskowitz [5] propose a new paradigm for analyzing sensitive data to decide how much of it can be released, and how released data should be sanitized, to avoid giving away specified secrets. Since the publication of this paper specific mechanisms for parsimonious downgrading have been developed and have been shown to be effective to various degrees for the right kinds of datasets and the right use cases.

In "Server-Assisted Cryptography", Beaver [3] proposes a mechanism to share the workload of resource-intensive cryptographic computations across a number of machines operating at relatively low levels of trust. Ideas discussed in this paper, such as Secure Multiparty Computation, have found their way into commercial use.

Greenwald's "Discussion Topic: What is the Old Security Paradigm?" [11] is a retrospective and does not propose new mechanisms.

In "Tolerating Penetrations and Insider Attacks by Requiring Independent Corroboration", Kahn [16] proposes a redundancy mechanism in which agents within a system seek independent corroboration of inputs they receive, and treat the degree of independent corroboration of an input as a metric of its trustworthiness. Simpler versions of independence-based security, including Byzantine agreement, segregation of duty, and maker-checker systems are in commercial use, but Kahn's more complex system has not been adopted.

In "A New Model for Availability in the Face of Self-Propagating Attacks", Lin, Ricciardi, and Marzullo [17] provide a model for measuring how severely a system degrades under attack by propagating malicious code (for example, worms, viruses, or ransomware). This metric wasn't widely adopted, and propagating malware has continued to be an effective vector for compromise of real-world systems.

4 INTRODUCING THE FUTURE

We'd like to emphasize again that in strategic foresight, visions of the future are not predictions. A prediction requires the ability to control or accurately estimate values of all or most variables to ensure high confidence that a prediction comes true - in reality no one controls enough variables to successfully predict what will happen in complex situations.

We begin with a baseline future, which is the current scenario extrapolated if nothing changes in our current conditions of the world. This gives us the baseline from which to explore alternate scenarios. Some of our scenarios are mutually exclusive, while

others can be combined. Some of the scenarios are desirable and positive, while others are downright dystopian and terrifying.

In order to derive and understand our scenarios, we used our background research to identify a set of variables which determine how the scenarios differ from one another, and how the future might change if one or more variables change. The broad categories of factors we heard from our interview subjects were: (i) the number of devices, (ii) the number of vulnerabilities, (iii) the number and sophistication of adversaries, (iv) the number and resourcing of defenders, (v) the severity of consequences of successful attacks, (vi) the likelihood of detection and punishment of attackers, (vii) the effectiveness of security controls, (viii) the economic balance of attacker and defender costs, (ix) who pays the costs of defense and of security failures, and (x) how the law treats security attacks and breaches. As we examined how these factors might evolve in the future, we refined the factors into more granular variables. The variables we chose in the end are:

- Vulnerabilities - how many vulnerabilities exist in the world-wide security-relevant attack surface? This includes both hardware and software vulnerabilities.
- Connected Devices - how many potentially vulnerable devices are connected to the global internet?
- Non-State Adversary Population - how many technically capable malicious attackers exist as potential threat actors?
- Nation-State Adversary Population - how many highly-skilled malicious attackers are employed by national governments to serve as threat actors?
- Defender Population - how many technically capable information security professionals are employed by security vendors, commercial entities, research institutions, and national governments?
- Attack Impact - how much damage can be done by an attacker, given the nature and functionality of connected devices?
- Application Criticality - how critical are the connected devices? Can they cause property damage? Financial loss? Harm to life, health, or safety?
- Attribution Effectiveness - how easy or hard is it to tell what actor caused a specific action (including attacks) in a system?
- Control Effectiveness - how effective are information security products and processes at identifying and remediating vulnerabilities and attacks before they can be used to cause damage?
- Attacker Cost - how much does it cost for a malicious actor to attack a system? Is there risk to a malicious actor's life, safety, or liberty as a result of attempting an attack?
- Vendor Cost - how much does it cost for an infosec vendor to create and sell a security control?
- System Operator Cost - how much does it cost a company or government entity to buy, test, install, and operate a security control?
- End User Cost - how much (in money and time) does it cost an end user to install and use a security control?
- Individual Criminal Sanctions - how severe are criminal penalties for attacking systems?
- Institutional Criminal Sanctions - how severe are criminal sanctions against institutions for failing to prevent attacks, or for hosting or perpetrating attacks?
- Institutional Civil Sanctions - how severe are civil sanctions against institutions for failing to prevent attacks, or for hosting or perpetrating attacks? Is it easy for individuals to initiate civil actions for attacks?
- User Sophistication - how effective are individual end-users at detecting and responding to attacks?
- Data Risk Aggregation - how much sensitive data can be stolen or damaged by attacking a single system or a single entity?
- System Risk Aggregation - how damaging (to economics, life, safety, or other critical properties) is an attack on a single system or a single entity?

5 SCENARIOS

5.1 Scenario 0: The Baseline Future

In the baseline scenario, cat and mouse games in security continue. There are increasing numbers of skilled security experts and hackers, but few jobs, especially in highly educated developing nations (e.g. Brazil, Romania). Companies continue to add security to systems after development or even deployment, so their systems continue to be fixed bailing wire and duct tape style. Many security researchers (including amateurs with free time) find and report vulnerabilities. Some companies offer bug bounties, but many others don't respond and take legal action against security researchers. Congresses and states pass anti-hacking laws, that can be used to prosecute researchers. Some researchers become black hat researchers to continue their research, sharing discovered vulnerabilities – much to the delight of the attackers and sensationalist media – while companies try to hide the vulnerabilities and corporate legal teams spend time tracking down black hat researchers who share vulnerabilities.

Users continue to be relatively unsophisticated and vulnerable. High-profile data breaches and ransomware attacks remain common; nation-state attacks, including on civilian systems, slowly become more common. Diversity among information security professionals increases slowly, but underrepresented populations, including socioeconomically disadvantaged, minority, and geographical populations remain underrepresented - and hence serve as a productive source of black-hat personnel. The attack surface grows larger as older systems pass out of support and as the number of new systems explodes with the growth of IoT. This adds to the complexity not only of the infrastructure but also of the required knowledge base of white-hat practitioners. Criminal and civil penalties for insecure software remain lax, and attribution remains difficult. Cloud failures become increasingly high-impact events as more and more services move into the public cloud. More money is allocated to information security defense, both by venture capitalists and by the commercial market, but this is a mixed blessing, as it drives up the cost disparity between defense and attack (because attack continues to be cheap). New cascade failure modes emerge as infrastructure becomes more complex and interconnected. Privacy and fraud risks to individuals increase steadily.

Table 1: Baseline Future

	Increase	Same	Decrease
Vulnerabilities	*		
Connected Devices	**		
Non-State Adversary Population		*	
Nation-State Adversary Population		*	
Defender Population		*	
Attack Impact	*		
Application Criticality	*		
Attribution Effectiveness		*	
Control Effectiveness		*	
Attacker Cost		*	
Vendor Cost		*	
System Operator Cost		*	
End User Cost		*	
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation	*		
System Risk Aggregation	*		

Table 2: Scenario 1: Brazil

	Increase	**Same**	**Decrease**
Vulnerabilities	***		
Connected Devices	***		
Non-State Adversary Population	*		
Nation-State Adversary Population	*		
Defender Population		*	
Attack Impact	*		
Application Criticality	*		
Attribution Effectiveness		*	
Control Effectiveness		*	
Attacker Cost		*	
Vendor Cost		*	
System Operator Cost	**		
End User Cost	**		
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation	*		
System Risk Aggregation	*		

The matrix above represents the change from the current state which leads to this scenario; a similar matrix will be used to describe each of our scenarios. Each row in the matrix represents one of our variables. The entries in the row indicate how the variable changes in the scenario in question; a star in the appropriate column indicates that the variable's value increases, decreases, or stays the same. Double and triple stars indicate larger changes. In the baseline scenario, the number of devices grows at the rate we expect today, and vulnerabilities and attack impact increase proportionally. More critical applications are gradually moved onto vulnerable systems, more and more data is aggregated into large repositories, and more processing is concentrated into a few large processors (including cloud providers); these trends have all been steady for many years.

5.2 Scenario 1: Brazil

In this scenario, whose title is derived from the Terry Gilliam movie, the bad guys have basically won. Ordinary consumer systems (both computers and IoT "things") still exist, but they frequently malfunction and cause damage in a variety of ways because of malware and human attacks. Privacy intrusions are frequent, and financial fraud is a significantly worse problem than it was in 2018, with Cryptocurrency a particular focus of instability, theft, fraud, and other losses. This scenario arises because of a set of driving changes from the baseline scenario; note that in the matrix below, the driving changes are represented by the shaded cells in the table; other changes are consequences of the driving changes.

Changes from Baseline: In this scenario, vulnerabilities increase dramatically because of the economics of IoT devices, which don't have a profit margin sufficient to sustain large investments in security. The number of connected IoT devices amplifies the new vulnerabilities, and new failure modes arise as a result of emergent properties of the vast number of new, insecure devices interacting with one another. Sensing the weakness, both private-sector and nation state attackers pounce, with the result that most devices are penetrated and operate in a nearly continuous state of compromise, resulting in frequent system failures, frauds, data breaches, and increasingly serious real-world instances of physical harm. People muddle along through the chaos and try to live in a very dysfunctional electronic world.

5.3 Scenario 2: Rollerball

In this scenario, whose name is derived from the title of the Norman Jewison film, all computing is hosted in the cloud by one of a small number of very large providers (Google, Facebook, Apple, Amazon, Weibo, etc…). Security is good, but not perfect, and content is not shared across platforms except when content providers pay large fees to the platform hosts for compatibility. User behavior and content is heavily surveilled, and the fruits of this surveillance are shared with government as part of the price of doing business as a multinational content host. Everyday reliability of computing systems is good, but highly-evolved organized crime syndicates still pull off spectacular heists from time to time. Nation-state activity is kept covert and subtle, as the major players each have their own affiliated large platform providers, and their surveillance partnerships with these providers are sources of national stability.

Changes from Baseline: The expense of security controls, together with the shortage of trained security staff and the increasing frequency and severity of security breaches has led most organizations to give up on in-house security in favor of hosted solutions from large providers with big security budgets and big security staffs. In response to the demand for secure hosted solutions, these providers charge more for security solutions, and devote significant resources to hiring security professionals and building and operating security controls.

5.4 Scenario 3: Minority Report

In this scenario, whose name is inspired by the Philip K. Dick novel and the subsequent film, national governments have followed the

Table 3: Scenario 2: Rollerball

	Increase	Same	Decrease
Vulnerabilities			*
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population	*		
Defender Population	*		
Attack Impact			*
Application Criticality	*		
Attribution Effectiveness	*		
Control Effectiveness		*	
Attacker Cost	**		
Vendor Cost	*		
System Operator Cost	***		
End User Cost			*
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions			*
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation	***		
System Risk Aggregation	***		

Table 4: Scenario 3: Minority Report

	Increase	Same	Decrease
Vulnerabilities		*	
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population	*		
Defender Population	*		
Attack Impact			*
Application Criticality	*		
Attribution Effectiveness	***		
Control Effectiveness	*		
Attacker Cost	***		
Vendor Cost		*	
System Operator Cost		*	
End User Cost			*
Individual Criminal Sanctions	***		
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation		*	

Table 5: Scenario 4: A Canticle for Leibowitz

	Increase	Same	Decrease
Vulnerabilities		*	
Connected Devices			***
Non-State Adversary Population			*
Nation-State Adversary Population			*
Defender Population			*
Attack Impact			***
Application Criticality			***
Attribution Effectiveness		*	
Control Effectiveness			*
Attacker Cost	***		
Vendor Cost			*
System Operator Cost			*
End User Cost			*
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation			***
System Risk Aggregation			***

lead of China and established panoptic surveillance and rigid content control over the Internet within their borders. Strong (but not perfect) attribution has been implemented, and even minor social and legal infractions are reliably detected and severely punished. The Dark Web flourishes, both as a venue for organized crime and as a haven for dissidents and investigative journalists - but it is very dangerous, and a tiny tradecraft error can lead to a life in prison, or worse.

Changes from Baseline: Nation-states, tiring of risks to civilian-owned critical infrastructure and economic harm from cyber fraud, and under pressure from citizens unhappy about frequent breaches and service outages, have stepped in to increase budgets for online crime fighting. They have dramatically increased both surveillance of online activity and criminal penalties for hacking. The dark web has truly become "the FBI's computer".

5.5 Scenario 4: A Canticle for Leibowitz

In this scenario, whose name is inspired by the Walter M. Miller novel, business and government have revehttps://v2.overleaf.com/4856187932bnvzmyhxpqbqrted to manual and paper-based processes, in some cases assisted by special-purpose devices like printers, which are not tied to global or even regional networks. The computer age is over, at least for financially- or safety-critical functions.

Changes from Baseline: After frequent and escalating catastrophes, individuals and businesses have concluded that the electronic world cannot be trusted with anything related to money, safety, security, or privacy. Computers are originally banned by law in some critical infrastructure and health-device applications after large-scale failures causing significant loss of life, and computerization is gradually abandoned in other economically- and safety-sensitive applications.

5.6 Scenario 5: Cold War 2

In this scenario, the leading tech nations have developed significant offensive cyber capabilities, and, realizing the potential for catastrophic damage resulting from Cyber attacks, have established new laws of electronic warfare. These nations have settled into a kind of electronic Cold War, with small players unable to damage larger nations without catastropic consequences, and with the larger players locked into a mutual-destruction standoff.

Changes from Baseline: After years of escalating attacks by criminal gangs, and infrastructure disruptions by nation-states attacking each other in cyberspace, national governments establish vigorous programs to recruit the best hackers into government

Table 6: Scenario 5: Cold War 2

	Increase	Same	Decrease
Vulnerabilities		*	
Connected Devices	***		
Non-State Adversary Population			***
Nation-State Adversary Population	*		
Defender Population	**		
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness	**		
Control Effectiveness	*		
Attacker Cost	***		
Vendor Cost		*	
System Operator Cost			*
End User Cost			*
Individual Criminal Sanctions	***		
Institutional Criminal Sanctions	***		
Institutional Civil Sanctions	***		
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation	*		

Table 7: Scenario 6: Colossus

	Increase	Same	Decrease
Vulnerabilities		*	
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population		*	
Defender Population			*
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness	**		
Control Effectiveness	***		
Attacker Cost	***		
Vendor Cost	**		
System Operator Cost			*
End User Cost			*
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation	*		

service at early ages. Comprehensive online surveillance tools are built which radically shrink the "dark web", and hacking becomes the exclusive domain of nation-states, who hire essentially everyone with relevant technical skills as red-team or blue-team staff, except for occasional lone-wolves, who are often caught and imprisoned.

5.7 Scenario 6: Colossus: The Forbin Project

In this scenario, whose name is based on the Stanley Chase film, information security has been recognized as being too difficult for humans, particularly given the speed at which malicious software such as ransomware can propagate and act. Artificial Intelligence and machine learning have advanced to fill the gap, and the large majority of system defense has been turned over to the machines, which effectively prevent the vast majority of attacks, and quickly ascertain the sources of attacks they can't prevent via machine-learning augmented attribution systems operating on the vast pool of personal and behavioral data available because of the proliferation of sensors in every device in every location. The global internet hosts an equally global AI, which enforces strict security policies, kicks offending devices and users off the network, alerts national authorities to hacking activity, and issues warnings about vendors and programs who produce insecure or malicious software. Human system defenders use augmented AR/VR environments to monitor and investigate incidents more quickly and to find defects in their systems.

End users also benefit from machine augmentation; when a user buys a new device it is automatically examined by AI in the network to ensure that it is secure and properly configured, and before a user can install a new application on a device, the application is automatically scanned for malicious code. The results of these scans are logged and reported; applications and companies are ranked based on their vulnerability scores, data breach losses, and other security performance metrics. All this information is presented to device users before they purchase and activate new devices.

Changes from Baseline: Rapid progress in artificial intelligence combined with the availability of more and more identity analytics and network diagnostic data enables the construction of automated security monitoring and enforcement systems, which become good enough at attribution to quickly identify hackers and locate them physically, and which also become able to reconfigure computers and deploy active countermeasures to stop almost all attacks within milliseconds. Errors and failures continue to plague electronic systems, but widespread and high-impact malicious attacks become very rare, and usually lead to prosecution and imprisonment of the perpetrators. Vendors devote significant resources to securing their systems, because the global AI will quickly detect security weaknesses and issue warnings which will make them unsellable.

5.8 Scenario 7: In Code We Trust

This scenario, like scenario 6, assumes that AI has become much more sophisticated in the security space. But whereas in scenario 6, the security AI is essentially a global policeman residing in the cloud, in this scenario AI is a quality-assurance tool in the vendor's production line, where it eliminates vulnerabilities before they can find their way into deployed products.

Changes from Baseline: Artificial intelligence has been combined with natural language processing and formal methods to produce a programming environment that creates highly reliable programs based on specifications negotiated with humans in dialog with secure software design chatbots; the resulting programs have very few vulnerabilities. The AI tools take over most code review functions, and software development teams abandon the use of tools and libraries which don't support strong, AI-based security assurance. Security becomes truly integrated into the development process in a seamless sec/dev/ops discipline. Hacking becomes so difficult that it's no longer profitable for any organizations other

Table 8: Scenario 7: In Code We Trust

	Increase	Same	Decrease
Vulnerabilities			**
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population		*	
Defender Population		*	
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness		*	
Control Effectiveness	***		
Attacker Cost	***		
Vendor Cost	***		
System Operator Cost	*		
End User Cost	*		
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation		*	

Table 9: Scenario 8: Send Lawyers Guns and Money

	Increase	Same	Decrease
Vulnerabilities			**
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population		*	
Defender Population		*	
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness		*	
Control Effectiveness	***		
Attacker Cost	***		
Vendor Cost	***		
System Operator Cost	**		
End User Cost	**		
Individual Criminal Sanctions	***		
Institutional Criminal Sanctions	***		
Institutional Civil Sanctions	***		
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation		*	

than nation-state intelligence services. Criminals abandon electronic crime and revert to older methods, and the electronic world becomes reliable, safe, and easy to use. All of this requires major investments by software vendors; the cost of this additional tooling is passed along to software customers - especially enterprise customers.

5.9 Scenario 8: Send Lawyers, Guns, and Money

In this scenario, whose name is inspired by the Warren Zevon song, the software market, under pressure from newly-imposed regulatory burdens and a rising flood of customer lawsuits, has shrunken to a few dozen large vendors worldwide, with slow release cycles, strict engineering processes like those seen in the Aerospace sector, and high prices. Software is very reliable and relatively secure, but for many applications it's prohibitively expensive. Some business functions have reverted to paper, and entertainment and personal productivity systems are rigidly separated from critical infrastructure, financial systems, and mechanical control systems.

Changes from Baseline: Governments have regulated the development of software, requiring insurance for failures of software-based systems and imposing strict liability on software vendors for failures leading to financial losses, property damage, injuries, or deaths. Even researchers are not immune from liability; strict regulations govern responsible disclosure of vulnerabilities, which must be reported to centralized authorities and remediated before any publication is allowed.

5.10 Scenario 9: Money, Cash, Hoes

In this scenario, whose name is inspired by the Jay Z song, security defensive work has become so attractive and well-paid that black-hat work is no longer attractive, and people with security-relevant technical skills are almost universally employed in defensive jobs. Governments can still hire offensive hackers, but criminal syndicates have a very hard time competing with white-hat employers.

Changes from Baseline: Employers pay great salaries, offer significant flexibility and benefits, and recruit extremely aggressively, especially in poor countries with good technical talent and in ungoverned spaces.

During the workshop discussion, several questioners noted that in section 3.1 we noted the adverse effects arising from an underdeveloped security job market in specific geographic regions, but did not discuss the adverse effects arising from an underdeveloped security job market for female professionals even in areas with a strong security job market. This is an important point; it would be very unwise to assume that malicious actors will overlook underemployed candidates with security expertise just because they are female. As we observed in section 3.3, our interviews identified more women and diversity as one of the most important desired future states for the security community.

5.11 Scenario 10: I'm From the Government, and I'm Here to Help

In this scenario, the government has created an entity like an FDA or FAA for code; this entity creates and reviews standards and enforces strict adherence to those standards among information technology developers. The result is a mixed bag, with some standards honestly improving the security landscape, but others stifling innovation and others weakened or rendered ineffective through regulatory capture. Coding also loses some of its cutting-edge attractiveness as the Silicon Valley "cowboy culture" fades under the restraints of regulation; there's less innovation and diversity in development tools, and more mandatory code reuse. There's no more MVP bolt-on security: security is a core system capability and embedded in all releases. There's also more professional education, licensure, liability insurance, and other trappings of professionalism. Some

Table 10: Scenario 9: Money, Cash, Hoes

	Increase	Same	Decrease
Vulnerabilities			**
Connected Devices	***		
Non-State Adversary Population			***
Nation-State Adversary Population		*	
Defender Population	***		
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness		*	
Control Effectiveness	*		
Attacker Cost	***		
Vendor Cost	**		
System Operator Cost	**		
End User Cost			*
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation		*	

Table 11: Scenario 10: I'm from the Government

	Increase	Same	Decrease
Vulnerabilities			**
Connected Devices	***		
Non-State Adversary Population	*		
Nation-State Adversary Population		*	
Defender Population		*	
Attack Impact		*	
Application Criticality		*	
Attribution Effectiveness		*	
Control Effectiveness	**		
Attacker Cost	**		
Vendor Cost	***		
System Operator Cost	**		
End User Cost	**		
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions	*		
Institutional Civil Sanctions	***		
User Sophistication		*	
Data Risk Aggregation		*	
System Risk Aggregation		*	

practitioners abandon development of enterprise software for less-regulated fields like entertainment software; others turn to dark-side hacking, though this is more difficult because of improved software quality and reduced vulnerability density.

Changes from Baseline: Governments have heavily regulated information technology to improve security; civil and criminal penalties have been imposed for selling vulnerable software, or software which fails and causes damage. Commercial software development is a licensed profession. Only regulated technology is allowed in production software. Regulation enforces which tools and methodologies can be used to write released software. If a manufacturer releases defective or unauthorized code, there are financial and legal repercussions – not to mention they have to fix it. Even sunset code must be safely disposed of.

5.12 Scenario 11: Everyman is Iron Man

In this scenario, whose title is inspired by the Stan Lee Marvel Comics series, the information security community has devoted massive resources to designing user experiences which support safe and secure use of computer systems - and this effort has succeeded. Systems effectively alert users to dangers, and elicit appropriate behaviors from users when dangers are detected. As a result of user interfaces designed by usable-security experts, and as a result of AI which detects and warns users before they can make critical security mistakes, the human attack surface is dramatically reduced, and attackers are forced to rely more and more on complex exploits targeting relatively rare hardware and software vulnerabilities.

Changes from Baseline: Scenario 6 applied AI augmentation to the network, to make it self-policing. Scenario 7 applied AI to the system development process, to make devices more secure by default. In scenario 11, AI augmentation has been applied to the user, in conjunction with carefully tested usable-security interfaces designed to present the user in real time with comprehensible threat information and actionable response choices to create a kind of Iron-Man suit for cybersecurity awareness and capability - but in a form-factor that makes it affordable and accessible to every technology user. Like a turbocharged version of the browser lock icon, the AI makes every invisible threat visible, and shows the user how to step around the threats and navigate the online world safely.

A workshop participant noted that this scenario is the only one which contemplates an improvement in end-user performance on security tasks. That's deliberate. In our view, human responses to situations of risk change on an evolutionary timescale (because, by definition, bad risk choices get you weeded out of the population). Given this, it seems likelier that improvements in human risk-management behavior will result from augmentation via fast-adapting artificial systems than by improvements in human education and awareness.

6 PARADIGM SHIFTS

During the analysis of our Appreciative Inquiry interview data, many characteristics of current paradigms were identified (note that we selected the old security paradigm characteristics from our raw interview data, which is only summarized, and not every characteristic in the table below appears in our summary). Since this paper is written for the New Security Paradigms Workshop, we added this section, which explores possible alternatives to the current security paradigm (we note that Greenwald's NSPW 1998 paper [11] highlighted the difficulty of defining the current paradigm; this section is in some sense an answer to that complaint).

To create the new paradigms, we took the properties of the paradigm identified by our interview subjects and reversed them to understand possible characteristics of new paradigms. Many of the new paradigms thus envisioned are positive, but not all - for example, a world of war-flying microdrones might be good for auditors and regulators, but it's probably good for attackers too. We applied the "backcasting" method to walk back to the present

Table 12: Scenario 11: Everyman is Iron Man

	Increase	Same	Decrease
Vulnerabilities		*	
Connected Devices	***		
Non-State Adversary Population		*	
Nation-State Adversary Population		*	
Defender Population	**		
Attack Impact			*
Application Criticality		*	
Attribution Effectiveness		*	
Control Effectiveness	*		
Attacker Cost	*		
Vendor Cost		*	
System Operator Cost			*
End User Cost		*	
Individual Criminal Sanctions		*	
Institutional Criminal Sanctions		*	
Institutional Civil Sanctions		*	
User Sophistication	***		
Data Risk Aggregation		*	
System Risk Aggregation		*	

from the new paradigm of the future to understand how some of the characteristics in the table below might manifest into scenarios; where appropriate, we've indicated in the table the scenarios we've developed with the aid of this "backcasting" technique.

Not all the new security paradigm characteristics we list here have been incorporated into our scenarios; in many cases this is because it's not clear (to us) what steps would be required to get from today's state to a new paradigm with the "opposite" characteristic listed in the table below. "Shoot the wounded" is a good example; nature does this, but at great cost. Could a business model in which large numbers of computing devices simply die and are thrown away as a result be economically sustainable? We're not sure, but we think the question is worth asking. In other cases, the new characteristic may be a bad thing (e.g. Non-Human adversaries), and we didn't want to spend time exploring unremittingly dystopian scenarios.

7 CONCLUSIONS

Applying foresight methods to traditional security research resulted in a rich and substantial set of scenarios and new paradigms.

- Most security solutions have been focused on technology, leaving the human/user as the main fault.
- In the future, it's clear humans and technology will work together. Whether human weaknesses are reduced through augmentation with machine learning, through working in a gamified environment, or through automated mitigation of human vulnerabilities, humans are not going away, so we need to be creative about accepting and reducing their attack surfaces.
- The attack landscape may need to change from the adversarial way it is structured today. Today the defense plays by a different set of rules (finite rules and resources) from the attacker (infinite rules, potentially unlimited resources); this asymmetry ensures that the security landscape is chaotic.

Table 13: Old and New Paradigms

Old Security Paradigm	**New Security Paradigm**
Security plays catch-up.	Security is pro-active.
Cat and mouse games.	No more adversaries. (Scenario 9)
User is the weakest link.	User behavior is part of security. User no longer the weakest link. (Scenario 11)
Hard to get funding for security.	Infosec is the cost of doing business. (Scenario 7)
Human is the (primary) adversary.	Non-human adversary. (Bot, AI, Machine Intelligence, Quantum cracking)
Security is doesn't give business benefit.	Security is a (business) advantage.
Centralized data, an irresistible attack surface.	Decentralized Data.
War driving. (pentesting pineapples)	War flying. (drone pineapples)
Passwords.	No more passwords. (e.g. via zero knowledge proofs, biometrics, etc...)
Defend Everything.	Shoot the wounded - let weak systems die.
Expanding attack surface.	Shrinking attack surface. (Scenario 7)
Technology attack surface.	Human attack surface. (Scenarios 6, 7, 11)
Security is a Commercial Enterprise.	Security is a Government Enterprise. (Scenarios 5, 8, 10)

- Solving security isn't just about using technology to solve threat and vulnerability problems. There are economic incentives, education issues, and job market issues. These areas will change regardless of any effort the security industry puts on them - however the security industry can proactively influence trends in these areas toward its own ends. To fully address security issues today and in the future, a broad look at the economics of security could be helpful.
- The future of security isn't limited to what we think is likely - what we came up with in our baseline future. There are many alternate futures and we can already put resources behind the ones we want to create.
- We successfully applied foresight methods to the security industry and came up with interesting, relevant, and useful scenarios. More collaboration between security researchers and futurists may be helpful in identifying strategies and points of leverage in the security market.

Discussion Prompts for NSPW Session During the workshop presentation session, we sought feedback on the following questions:

- Do you think the foresight methodologies are useful applied to security? Would security research benefit from more structured futurist analysis?
- Have we chosen the right set of variables in the scenarios?
- Based on these scenarios, what is the single highest impact variable? For example, if you could can only adjust one variable, which variable would you adjust?
- Are the scenarios believable? Why? Why not?
- Do these scenarios help you deal with security issues today? Could you use them to prioritize your work?
- Do the scenarios, in aggregate, suggest that information security is on the right path or the wrong path?

8 FUTURE WORK

During our research interviews, a number of subjects noted that education, both of users and of IT practitioners, would significantly affect the future of security. Interview subjects suggested a number of possible future developments that would make a difference, ranging from from more in-depth security training for IT practitioners, to better information resources about new and emerging threats and vulnerabilities, to more effective end-user security training, to the establishment of an "education commons" which would contain a wide variety of open-source information resources freely available to the entire world population. We found it difficult to formulate hypotheses about how these proposals would affect the future of security, in part because some past education efforts have not evidently made big differences, and in part because it seems difficult to predict the effects of education on populations including malicious actors. This is an area that could be investigated further.

During the workshop discussion, a participant asked if we intended movement between the scenarios, rather than simply from the current state to a single future scenario, to be possible. Our answer was (and is) that the scenarios aren't presented as alternative roads to a single future, with all the choices made at the beginning. Instead, the scenarios should be thought of as tools for thinking about how variables and combinations of variables influence the state of security as the future evolves. In this context, the scenarios we've presented here are really just a few vertices in a much larger graph depicting how the state of security changes as society changes the underlying variables - and movement "sideways" in the graph should absolutely be analyzed more thoroughly in future work. It might (for example) turn out that getting to a desirable end state requires passing through one or more less desirable intermediate states.

During the workshop, a participant raised the question: If you were to hire a futurist to do a full study, what are the additional things you should do. Foresight research uses specific methods to help see around corners into the future, so a futurist would use a combination of methods similar to the ones we used in this research, and then create scenarios based on the results of that research. The futurist may work with a team and conduct workshops to co-develop and apply scenarios to a specific company, product or future timeframe. The scenarios may be communicated in a variety of mediums - from traditional reports to immersive design futures, like films, artifacts, and other graphically designed elements, to give an experience of the future. The results of the research and scenarios can be used in strategic planning and product development to "future-proof" decisions. A futurist working with other researchers would bring fresh perspectives.

REFERENCES

[1] 1998. Proceedings of the 1998 Workshop on New Security Paradigms. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[2] Marshall D. Abrams. 1998. Security engineering in an evolutionary acquisition environment. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[3] Donald Beaver. 1998. Server-assisted cryptography.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[4] P. Bishop and A. Hines. 2012. *Teaching About the Future* (1 ed.). 0, Vol. 1. Palgrave Macmillan.
[5] LiWu Chang and Ira S. Moskowitz. 1998. Parsimonious downgrading and decision trees applied to the inference problem.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[6] D.L. Cooperrider, D. Whitney, and J.M. Stavros. 2008. *Appreciative Inquiry Handbook* (2 ed.). 1, Vol. 1. Crown Custom Publishing, Brunswick, OH.
[7] Crispin Cowan and Calton Pu. 1998. Death, taxes, and imperfect software: surviving the inevitable. (9 1998).
[8] E. Dragt. 2017. *How to Research Trends* (1 ed.). 1, Vol. 1. BIS Publishers BV.
[9] K.H. Dreborg. 1996. Essence of Backcasting. *Futures* 28, 9 (7 1996), 813–828. An optional note.
[10] Simon N. Foley. 1998. Evaluating system integrity.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[11] Steven J. Greenwald. 1998. Discussion topic: what is the old security paradigm?. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[12] Sohail Inayatullah. 1993. Causal Layered Analysis: Poststructuralism as Method. *Futures* 30, 8 (10 1993), 815–29.
[13] Sohail Inayatullah. 2014. Causal Layered Analysis Defined. *The Futurist* 48, 2 (2 2014).
[14] Michael Dunner Srinath Vasireddy Ray Escamilla J.D. Meier, Alex Mackman and Anandha Murukan. 2003. *Improving Web Application Security: Threats and Countermeasures Roadmap* (1 ed.). Microsoft Corporation.
[15] Erland Jonsson. 1998. An integrated framework for security and dependability.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[16] Clifford Kahn. 1998. Tolerating penetrations and insider attacks by requiring independent corroboration.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[17] Meng-Jang Lin, Aleta M. Ricciardi, and Keith Marzullo. 1998. A new model for availability in the face of self-propagating attacks.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[18] J. McDermott. 1998. Prolepsis on the problem of Trojan-horse based integrity attacks. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[19] Standards Association of Australia. 1995. *AS/NZS 4360: Risk Management*. Standards Association of Australia.
[20] Cynthia Phillips and Laura Painton Swiler. 1998. A graph-based system for network-vulnerability analysis.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[21] The Millennium Project. 2009. *Futures Wheel, Futures Research Methodology Version 3.0* (3 ed.). The Millennium Project, Washington, DC.
[22] Noah Raford and Andrew (ed.) Curry. 2012. *The Future of Futures* (1 ed.). 1, Vol. 1. Association of Professional Futurists, Houston, TX, Chapter 8, From Design Fiction to Experiential Futures.
[23] Thomas Riechmann and Franz J. Hauck. 1998. Meta objects for access control: a formal model for role-based principals.. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.
[24] H. Sackman. 1974. *Delphi Assessment: Expert Opinion, Forecasting and Group Process*. Technical Report AD-786 878/R-1283-PR. RAND Corporation.
[25] Chris Salter, O. Sami Saydjari, Bruce Schneier, and Jim Wallner. 1998. Toward a secure system engineering methodology. In *Proceedings of the 1998 workshop on New security paradigms (NSPW '98)*. ACM, Charlottesville, VA, USA.