COVER

TITLE

May 2018

Introduction

THE ANNUAL CYBER SECURITY INTERNATIONAL CONFERENCE 2017

Upgrading the Internet

Menny Barzilay, CTO, Blavatnik ICRC, Tel Aviv University; CEO, FortyTwo

I wish to discuss here some of the current problems in Cyber Security. We invest so much money in cyber security, we have the best people and the best products, and we do everything right. But even so, it seems like we already know that in the years to come, more companies are going be hacked. These are going to be big companies, with a talented information security officer, and with a lot of money to invest in cyber security. It seems as if, to some extent, we have failed in providing certainty. When directors ask us, "how much money do we need to invest to make sure that we won't be hacked?" we don't have an answer to this question.

In order to understand why we can't answer this question, we have to go back in time to 1969. It was an amazing year; many things happened in 1969. Neil Armstrong walked on the surface of the moon, the Beatles published *Abbey Road*, the Internet was born. Back then we called it the ARPANET. Back then, all the users in the ARPANET knew each other and trusted each other in real life, and so they trusted each other in the cyberspace, too. All the devices were controlled and trusted, as well as all the applications.

In a very short time, the ARPANET became the Internet as we know it today. However, today we have billions of users who don't know each other, and definitely don't trust each other. We have no clue as for what is connected to the Internet, and we have so many applications, some of which are malicious by design. Trust has become a challenge. When the Internet was created, nobody expected it to be such a huge success, so in the beginning it was not created with security in mind, but only with connectivity as its main focus. To some extent, this is the secret of the Internet's success. The way it was designed allowed us to use it in a very simple manner, but most of the problems that we face in the world of cyber security are there simply because of the fact that the Internet was not designed with security in mind.

However, this is not the Internet's biggest problem. The biggest problem is that the Internet was designed in a way that doesn't allow us to upgrade it. We are stuck with this Internet, and this Internet allows so many people to do all the things that we see today. If we want to upgrade the Internet, we have to physically go to each and every one of those devices that create the Internet, and upgrade the systems there. The RFP for IPv6 was issued in 1998, and still, more than 95% of the Internet traffic is in IPv4. We know that we need IPv6; it provides more security, we need it for the Internet of Things, and yet, it is hard for us to deploy it on a global scale.

Speaking on the Internet of Things, we are about to connect billions of new devices to the Internet. We make things "smart" today; we take shoes and we make them smart shoes, we take tables and we make them smart tables, we take cars and we make them smart cars. Everything become smart. But we all know that if something becomes smart, it also becomes hackable – if it is smart, I can hijack it, which is problematic. The problem is that with the current limitation of technology, with the inherent problems of technology, it is hard for us to move forward.

There is a big, well-known secret – all of those amazing startups, all of those amazing innovations, all of those amazing products that we see today in the world of cyber security, to some extent (and I'm saying it very carefully) represent incremental innovation, and not disruption. They do improve things, but in the long term it seems that it becomes easier and easier to be an hacker than it is to be a security guy.

We all know the inherent asymmetry of security – if you are a hacker, you only have to succeed one time, while if you are a security guy, you have to succeed all the time. If you are a hacker, you can attack a single point, while if you are a security guy, you have to secure everything. Hacking has no rules, but security has so many rules; if you do something you are not allowed to do, you will get fired or someone will put you in jail. Hacking is very cheap, security is super-costly.

For this reason, and for many other different reasons, today security is like trying to guard a balloon with our bare hands while the hacker

has a pin. We need disruption; we need something that will completely change the rules of the game, that will make us believe that in the years to come we will be able to provide certainty. This is something that I hope will happen in the Interdisciplinary Cyber Research Center (ICRC), as well as in other places. We have amazing researchers on cyber security that could invent an entirely new way to think about technology, and so do other academic facilities. But this can also come from the industry and the amazing entrepreneurs that we see everywhere.

One thing that might happen is that we will have the an opportunity to redesign the Internet. Now, this is a strange concept. How can we redesign the Internet? Wireless connectivity networks, and wireless solutions to create networks, are actually evolving very rapidly. To some extent, this is happening in such a way that in the years to come we may have the ability to create worldwide global wireless networks. We already see Facebook doing something like that with Internet.org, and we already see Google trying to do something similar with Project Loon. The moment we will have the ability to create worldwide wireless networks, something very interesting could happen. First of all, we are trying to bring Internet to those places where they don't have Internet; there are many people in our world that don't have Internet access. To us, Internet access is a basic human right, and we believe everybody should have it. But what if one of those companies like Google or Facebook or anyone else will say: you know what, if I already have this global-wide wireless network, what happens if I choose not to use the traditional TCP/IP, but rather implement something else, and that something else will be upgradeable? That means that if I want to upgrade the core way that this wireless Internet works, I can do it.

Why is that such a game changer? Why will that allow us to do amazing things? Becayse such a move will create a paradigm shift in one of the core elements that we have today in the cyber security industry. Today we have a rule – new problems equal new security products. Every time someone figures out a new way to do something, a new attack vector, we all need to buy new security products. That means that for every problem we all spend money, and that is a lot of money to spend on every problem. But with an upgradable stack, an

upgradable Internet, this can change. Suppose we have a new Internet, which, for the sake of this argument, we can call AGN – Alternative Global Network. In this case, if someone finds a new way to carry out a DDoS attack, or if someone finds a new way to perform IP manipulations, spoofing or something like that, the AGN provider will say: Okay, I learned about this problem, and now I can just upgrade the Internet, preventing anyone else from exploiting it, and removing the need for everyone else to buy new security products. This will happen because the core Internet will be, mostly, secure. Do you need a new protocol to connect driverless cars? Don't worry, come back in twenty four hours, and we will have this new security protocol. Do you need a new way to connect drones? Don't worry, we can do that, too. An upgradable stack will allow us to do different and amazing things. This is the kind of disruption that I hope to see in the near future, coming from Israel or other places around the world.

The Future of Security Architecture

Gil Shwed, Founder & CEO, Check Point Software Technologies

I would like to talk about Check Point's vision, or what we see as the future of cyber security. I will try to share a little bit of what I see and what we are working on when we face the future of cyber security.

The major challenges of cyber security are quite known, but in 2017, we have seen different phenomena joining together to create something new. For one thing, in 2017 we have seen that everything is becoming connected. I have a newborn baby, and his baby crib is connected to the Internet. Everything is connected, in business and our personal life, and this leads to the fact that we are under constant attack, and these attacks are coming in many forms.

We have seen ransomware preventing hotel guests from getting into their rooms, and we have seen the more famous WannaCry type of attacks, which immediately took over big parts of the world. We are under attack, not just on the traditional attack vectors, like our personal computer or our traditional networks. Our mobile devices are really the back door to our lives and to our enterprises. Mobile attacks are not in the future, they are already here. Gooligan, for example, is an attack based on a mobile malware, which gave the attackers full access to the Google accounts of over a million people. We are seeing in Check Point that when we are going and doing security checkups for large companies, we find a mobile malware already in the network in 100% of the cases. This is not a theoretical risk.

The next part is the one that should really scare us all. Every cyber criminal, every member of a cyber force, and basically every person on the Internet has access to the most sophisticated tools, tools that were developed by cyber powers, by governments and so on. For example, in the case of WannaCry, we have seen how tools developed in the NSA were used by some anonymous cyber groups. I find that very scary, because if in the traditional warfare we know our enemies – we either deal with a big organization that has something to lose, or with small attackers with very basic tools that we can combat – in

cyberspace there are no boundaries. The single attacker has access to the superpower tools, which only a year ago were a secret of the NSA or some other superpower. When we put all these factors together – the collectivity, the increase in attacks, the risk in mobile, and the sophisticated tools that every hacker has – we can see that something new is happening. The future is not something we should be afraid of, it is already here and we should get ready for it.

What should we expect in the next couple of years, then? We should all understand that attacks will continue to grow, and that we are all targets; it doesn't matter if you are a small entity or a big one, if you are a government, a company – everyone is a target. The general hacker doesn't discriminate, they will try to get into the places where it is easy to break through.

Advanced threats will continue, and in the first half of 2017 we have seen that this is a reality, not a futuristic prediction. Attacks are becoming more sophisticated, the attackers use more tools and put them together, and they continue to target our networks. At the same time, we are moving more and more towards the cloud, extending our environment. The cloud is going to be one of the major targets for future attacks. The shift there is accelerating, and many companies and enterprises, and even many governments, believe that if we put a computer in the cloud it is secured, because we got it from some of the best companies in the world. The reality is that when we buy a computer on the cloud, it is just like buying a computer on our existing network. If we don't protect it, we leave our networks and data centers completely exposed.

I would also like to discuss the issue of mobility. These small devices that 100% of us are carrying in our pockets are the real back door to everything we do. Every move, every activity, every data piece that is moving around us can be hacked and can be accessed using our mobile devices. We need to keep in mind that when talking about a traditional computer, we know when it is on, we know when it is off, we know when it is connected and to what. We know that when a computer is connected to our internal network, it is protected by our perimeter security. Our mobile device is on 100% of the time, it is connected 100% of the time, and even when we are inside our

corporate headquarters, it is connected to the public network, and can broadcast what it sees.

With all of these factors under consideration, the next question we should ask ourselves is – are we taking the right approach? We go to conferences, we buy products from almost 1,600 cyber security companies around the world, but what is the right strategy? There are technologies that combat advanced threats, there are technologies that deal with mobile threats, and there are technologies that deal with cloud security. We surveyed major enterprises to see if they are using these technologies today, and found out that only 4%-7% of the companies use technologies for advanced threat prevention; only 5% of the companies use technologies to combat mobile threats; and only 1%-2% of the companies use cloud security. What we can clearly see, then, is that while we invest in and talk about future cyber threats, in reality we are not necessary taking the rights steps.

If that is the case, why are we all surprised that we are so vulnerable? As it turns out, there are several reasons for that. When we ask people and enterprises around the world what is the reason behind their investment strategy and lack of mobile/cloud protection, we hear many explanations and excuses, only some of them are good. People tell us that it is too complicated, which can sometimes be true. Others say they didn't realize it was such a problem. We also hear good reasons as well, with clients telling us that there are too many point products, and that there are not enough trained people.

What should we do about it? How should we look at the problems? We can look at our environment today, and compare it to the construction of a house. A house can be flooded, there can be leaks when it is rainy, and so we need to build a strong and resistant house. In cyberspace today our roof is very leaky, the rain keeps pouring on a daily basis, and every day in big enterprises the rain is getting through.

What are people doing about it? Some focus on remediation – bringing buckets and taking the water outside. Others focus on incident response – call the experts, and they will find the leaks and get the water out of our house. Others just pretend the leak doesn't happen, and ignore it until the inevitable high damage finally occurs. For example, in early 2017 the world's largest cyber insurance claim

was settled, and $115 Million were paid to people whose personal information was stolen during a cyber attack. Another option is to try to patch the leaky roof with many different point solutions, which never seem to be enough.

Let's try to see what this looks like in the technology and cyber world. In our enterprises, we have our main network, the people inside the buildings, the people that are mobile, and the people connecting from the outside – we need to protect all of them. Today the approach is to take many different solutions, a multi-vendor approach. We started with putting firewalls, and I'm proud to be the person who invented that technology, 24 years ago. However we need to cover more areas, and so we add more and more technology and products, and we hope that by putting in place this many point solutions, we will be protected against the rain. In reality, we continue to see that the rain is getting trough, and that there are inherent gaps in that architecture. We have a large investment in a lot of technology, we need many people to maintain that, but there are still many holes, the rain is coming through and our enterprises are flooded with attacks.

What are we doing about that? We focus on post-breach technologies that mitigate the attacks, and detect the breaches only after they had occurred. We see an immense investment in this field, and there are many companies in cyberspace today that deal in detecting attacks after they occur. We say that attacks are inevitable, and so we might as well mitigate the damage, rather than try prevent the attack itself. I don't think it is the right approach; I think we shouldn't be living in a leaky environment, just like if our house was leaking every day, we would fix the roof or build a better house using new technology.

Now the question is: can we find a better architecture to face all of these challenges? First of all, we need to first think about security as something that follows our business needs. Security should not be about where we have technology in place. Security should be present in all areas of our operation. Security should be done efficiently. We cannot use more and more resources and hire more and more people; we should be able, with a small enough group of people, to provide our enterprises and countries with excellent security. And mainly, we

shouldn't accept the fact that attacks are getting inside. We should keep the attacks outside and focus on prevention.

The first thing we should do is to secure every environment. If in the past we were focused on securing our networks and endpoints, today we should also secure the cloud, the datacenter and our mobile devices. As we do that, we also need to think about the future: we need to build security for the Internet of Things, for national infrastructure, for critical infrastructure, for industrial environments, and even for automotive and cars. All of these bases are currently under threat, and we need to have a security system that covers all them.

We need to take a different approach; a different architecture. We need to build multilayer solutions, with all the items we described earlier, that will represent a single vision and work together. We need a unified architecture with a 100% focus on prevention, with the outbound layers of next generation firewalls, advanced threat prevention, traditional threat prevention, and one single management console that addresses the cloud, mobile and the network. If we keep focusing only on patching what we currently have, and adding more and more point technology solutions, we will never get there. But if we start to layout this vision to our enterprises, I think we can get there, and I believe we have the solution. The solution I am suggesting is not something futuristic, but rather something that exists today.

Lately we conducted a survey with clients who adopted the multi-vendor architecture, and with clients who adopted the unified architecture that focuses on prevention. When we asked them if they were seeing attacks, and what were the costs, I was actually shocked to see the results. Clients from the first group presented an average of 40 days before they identify an attack, and the cost of the remediation was almost $700,000 per attack, which is very high and very consistent in the industry. However, clients that adopted the unified architecture, which is 100% focused on prevention, reported an average of only *two days* in order to identify an attack, and the cost of the remediation was only about $7,000. This is a 1:100 ratio between the effectiveness of these approaches. The right architecture should be much more efficient and much less costly in terms of the

investment required, and in terms of the amount of manpower we need to deliver security.

In summary, the future of cyber security its already here. We can build solutions for the future, and we need to think forward. We cannot continue to augment our existing infrastructure further and further; we need to adopt a new architecture when thinking about the ultimate security with no gaps. We need that security to focus on prevention. It is true that sometimes we will still need to do mitigation and other things on the inside, but most of the time we will need to prevent the attacks from getting in. That is the most effective way to keep ourselves clean, so don't settle on detection and post-breach mitigation. Also, last but not least, we need to cover all of our bases – the network, the cloud, and the mobile. It is not enough to deal with our traditional network, or to say we haven't seen huge damages coming from mobile, because these damages will come, and we need to be prepared for them. That is what we are building in Check Point, and that is what I think we should all think about when we build the architecture of the future.

A Retired General's Outlook on the Cyber World

(Ret.) Gen. Keith Alexander, Founder & CEO, IronNet Cybersecurity and Former Director, NSA

I started serving my government decades ago, when cyber was still called "sigint", and left in favor of the private sector several years ago. The difference between the two sectors, the government and the private, is amazing. There are two parts to this: first, working for the government and the military was great. The people and the missions were amazing. When you are on the other side, you don't realize what it is like to run a business, what you have to do. Suddenly you have to pay the people that work for you, every month. It is also surprising to see the problems that the cyber community faces. What surprised me the most is how difficult it is for CISOs to do all the things their job entails, and mostly to face all the problems related to the integration of all these different products to protect their network, when those products change almost every week. They have one of the most difficult jobs.

When it comes to the intersection between the government and enterprises, I believe that while working in the government sector I had a great opportunity to see what the threats are, from hackers and all the way up. Having lived in that world, I believe that I, as well as other people that used to work in the government and moved to the private sector, have unique insights on those topics. What I see now is what the CISOs and the companies need, and I am asking myself how I can help in giving that to them.

In my opinion, there should be some form of a public-private partnership, in the way of getting nations to help companies to stop threats that are increasingly sponsored by nation-states. I think those things are in the future, because cyber is clearly being used as an element of national power, and companies alone can't defend against it. We have to figure out how to get such partnerships to work, and I believe that this is where people like me, who came out of government, can help bridge that gap.

One of the great meetings I had recently was with President Trump, who was all straight forward. Without the press in the room, it was amazing to hear the questions and his focus on the topic of cyber. His question to the representative of the energy sector that was there was, "how do I help? What do you need? How do we protect the grid?" His cyber security team was also there, and their mission was to understand how the government can help defend the industry. The main topics of discussion included the public-private partnership and how to achieve it; and how to get people more ready to nation-sponsored threats by helping them understand what those threats are. We discussed the current threat matrix, and we also discussed the need for providing the necessary security clearances in order to achieve the above goals. My experience, despite what you read in the press, is that President Trump is focused on that problem, he understands the existential threat, and he wants to help protect the industries. In the commission, there were arguments over whether it is the industry's responsibility to defend itself, or the government's responsibility. We brought out the preamble to the Constitution and its "common defense"; it is the government's job to help, it can't just sit back and wait. I think that is visible for the government, it is understood and it is something they are going to take on.

When I first got to the NSA, I never expected that the first thing they were going to do was to tell me to go and brief the President, which is exactly what happened. What the NSA does every day is something that is viable to our nation, just like 8200 does for Israel. I will give you a great short story about leadership, from my meeting with President Bush. He came to Fort Meade, I said "welcome, Mr. President," and he replied, "General, get in the car, we got to talk," just like that. I got in the car, and he said, "there are two issues we got to talk about. First, I understand you have too many bosses." It was true, I had seven bosses; the President, the Vice President, the Secretary of Defense, the Director of National Intelligence, my wife and a couple of others. He said, "we are going to fix that right now!" I was wondering who we are going to throw under the bus, as they were all senior to me, so I told the President, "they are all good, we love them all." He replied, "if that is ever a problem, we will fix it

right away." The second thing he said was, "General, this issue with the terrorist surveillance program, it is going to get really bad. Here is the deal: you protect the country, I will take the heat." And that he did, every step of the way. It was the greatest act of leadership that I have ever seen, because what he did was to stand in front of the press and say that this program was started to protect the country, and he authorized it as President. He took the heat every step of the way, and he never squirmed. When the Snowden incident happened, I took the same approach, and that is what leadership is all about.

When it comes to current problems and dilemmas of future infrastructure building and security by design, there are definitely experts in this field, who have the ability to dive into the deep, technical parts of industrial control systems. However, there are only a few of those, and we have a problem mapping that knowledge for all the utility companies in our nations, and the problem of bringing those together. We have to move toward cognitive computational systems that integrate machine learning, AI, and expert systems. We have to have the ability to see events that hit each company, and share them across companies within the sector; to get sector-wide visibility and protection, and then share sector information with other sectors and with the government for a comprehensive program. That is technically doable, and that is where it will go. It will be driven by the Cloud and by mobility. We, the collective cyber community, need to come up with a way to have all this information fed to a machine, assist people in becoming much more proactive, and go after the threats by seeing the entirety of that layer. It will do something that we have done in the past: it will create partnerships among countries, working for the common good.

One of our most serious problems these days is small groups who, due to information leaks and breaches, put their hands on national-level capabilities, utilized in attacks such as WannaCry. To deal with such groups and cases, first of all we have to protect our assets and secrets better. That is clear. The second part is that while this area is changing so quickly, WannaCry would have been stopped if all the systems had been properly patched. Patches were pushed out for the newer systems, but not for the old ones that were no longer supported.

We have to acknowledge that there is a problem there. We are not patching, and humans are involved in the way we are doing it. If humans are involved in patching, identifying vulnerabilities and protecting against them, we are in the wrong place. Machines are going in milliseconds, and we are talking about patching things, a process that can take days, weeks, and even months. That is a wrong framework. We have to automate, and we have to get to an automated system that helps bring humans up to a different level. If this goes to a nation-state attacking a nation-state in cyberspace – which I believe it will – we have to have a system that can protect our nations. The protection of our nation is not just the government or the Defense Department, but also the critical infrastructure: finance, energy, healthcare, the communications; without that, there is no country.

Rethinking Secure Infrastructures

Nadav Zafrir, CEO, Team8 and Former Head, Unit 8200, Israel

I served in the military for decades, even before it was called "cyber", they called it "sigint" back then. Like everyone there, I tried to do our best to protect the nation. I retired from the military about four years ago, and joined the private sector.

I have to say that I miss that world, in which I don't have to compete for talent, and where everybody remains, on average, 20 or 21 years old, which is what I had in 8200. There was no Intellectual Property involved, so there was no money involved. The motivation of such a world is obvious, and everybody is highly passionate about what they are doing. One of the reasons that brought my associates and myself to decide on the Team8 model, was to at least try to bring some of that ambiance and passion into what we do today in cyber. If I was offered to be a CISO today, I would probably not take that job. It is overwhelming, it involves a lot of luck and chance, and you have to prepare for anything and everything.

There are many traditional things that we must do as we defend ourselves. We have to check all the boxes, and we have to get the right hygiene levels, but I believe we must also become more proactive. I think that one of the things that my unique experience as a former high ranking officer allows me, is to understand the mindset of what it means to be an attacker in the first place. Attackers are human, and will remain that way, at least for the foreseeable future. We can definitely say that the most sophisticated current attacks are carried out by humans. It is true that attackers now have access to capabilities that might be nation-sponsored, but they are still humans, and so they have their own abilities. We can take a more proactive approach, within our networks on the enterprise level, but there is only so much the enterprises can do before it has to go up to the country level, the government.

Speaking of governments, there is one comment I keep hearing, in terms of government responsibility. If there is a missile that comes

at a country, the government will do whatever it can to stop it, but if there is a "cyber missile", the government would only intervene if it has anything to do with national defense. That is an interesting phenomenon, and hopefully collaborations and conversations with governments will allow us to start working together to deal with such matters.

Today, with Team8, I am a part of the business world. From my perspective, there is one thing happening in the world right now, which should bring us to the next level of security. Right now we are in a situation where enterprises are at a bad dilemma, where they have to choose between productivity and security. In many cases, as an industry, this is where we must mature. I believe that there is a problem with our view of security as a binary situation – you are either infected or not, all black and white. What about a situation where we say that yes, we are infected, but we can live with it?

Another thing is the convergence between the physical and the virtual, which is happening in our national grids, meaning that everything is becoming more fragile. As things become more fragile and hyper-connected, I think we need a real paradigm shift in security, where we are not only building security as an afterthought, as another layer, but rather thinking about the infrastructure design to begin with. In many ways we need to be inventing and creating the infrastructure for the next Internet, which takes security as part of life, as part of the design, as part of the architecture. A bit like what Bitcoin has done to digitals currencies.

When we will look back, I believe that we are going to think of 2016 as an inflection point. For the first time, on a grand level, we started seeing attackers going after the real infrastructures, and trying to cause major disruptions, like with the US elections. They are going after our trust in our political systems, in our digital currencies, and so on. This means that in many ways, the very foundations of civilization are under attack. As the world becomes more hyper-connected, just attempting to come up with the same solutions, only slightly better, simply isn't going to work. One of the reasons is that nation-state capabilities from intelligence agencies and government organizations are being leaked. Unlike in the past, attackers with limited resources

can commoditize and weaponize these capabilities, extremely quickly, like what we saw in WannaCry.

Israel has been privileged to have a bilateral comradeship with other countries, such as the US, and I am privileged to have extended these relationships and collaborations with my colleagues in the business world, too. I think that this is essential. I do have a word of optimism for all the cyber experts, and even for the CISOs that have an almost impossible job: the hyper-connectivity brings prosperity, and we mustn't go back and disconnect ourselves. The benefits and the opportunities in a hyper-connected, big data, AI world, far surpasses the challenges that we face right now. I believe that if we join our hands together, different countries, industries, governments, academia – we will prevail.

Protecting a Country: Why Israel created the National Cyber Directorate

Eviatar Matania, director general of Israel's National Cyber Directorate in the Prime Minister's office

There are too many things to worry about in the field of cyber security today. From my experience, when we have so many things to worry about, there is always something, the most dangerous of things, that keeps you up at night. When I was asked about what keeps me up at night, I realized that I had so many reasons not to sleep well. This is true especially in the Middle East, where nobody can sleep well in our interesting neighborhood. However, after thinking about it for a while, I figured that what really frightens me is the things which I cannot anticipate. In Israel it's the Yom Kipur war, or the equivalent of the American 9/11. The things that we can anticipate, and that we know what we need to worry about – we have to worry about them, but we also know what to do about them, and we try, at least, to do something. But my worst nightmares are about the issues that I cannot anticipate; the idea that something is happening and I don't understand that it is happening. This was the atmosphere in our Yom Kipur war in 1973, much like the atmosphere in the US around 9/11, where everything was on the table, but you simply do not understand the strategic view of what is happening.

When it comes to how we deal with threat and how we develop our strategy, I can point towards three important elements in our strategy. Firstly, we understood that most of what we see worldwide is strategies of organizations on how to better protect themselves, and most of the nations that tried to go into the cyber issue, tried to duplicate what they knew about conventional threats. We understood that we needed to do something different, create a new discipline, and that is why we put in place a new concept of operations, of how to act, and this was the first important element of our strategy. We are doing a lot in order to build our robustness, and then our resilience. This means we are event-driven.

Secondly, we harness our security forces in order to better defend Israel against high-end mega attackers. This was the first notion for the understanding that we need a new discipline and strategy. The second element was the understanding that Israel is small, which is why we need only a single organization that has enough critical mass, and the right accountability, in order to lead Israel in the cyber domain. The government adopted our recommendation in February 2015, which led to the creation of the new cyber security agency, as the one and only organization in Israel responsible for the defense of the nation from cover to cover; from the doctrines, instructions, regulations, incident response and everything that is needed. This organization has been operating since 2016, and it has been developing, with already several hundreds of employees.

The third step, or the third element of our strategy, was to understand that without technology we will not be able to really mitigate the threat. We are working very hard to try to find what technologies should be adopted at the national level. If we look at the "Iron Dome", it does not protect a specific bank or energy company or some organization, but an entire nation. However, when you look at the cyber issue, most of the technologies and solutions aim towards defending a specific organization. We are trying to figure out – and we work very hard on that – to make the right R&D efforts, in order to find state solutions to better defend our nation. I think that going forward with the authority as the very important organization to lead the nation and find the right technologies is the best thing that we can do, and I believe and I hope that we are on the strategic course of defending our nation in the coming years.

As for the future, I have no doubt that the next step will be to better understand not just how to defend ourselves, but also how to react, how to retaliate, how to deter, and this is a policy that every state should adopt. That is a real challenge, because most of us simply do not understand yet the differences between the cyber arena and other arenas, where we know to talk about deterrence, reaction and retaliation. This is the next step for the future.

Deterrence Through Partnerships

Rob Joyce, cyber coordinator in the white house, and a special assistant to the president of the United States

I am quite new in the White House, although I have been in the field of cyber security for twenty seven years now. In my new position, the thing that worries me the most is how integrated computers have gotten into our whole lives, and the fabric of our societies. At the same time, what we are seeing are increasing threats from criminals, from nation-states and stateless actors. There are many opportunities for criminals to just go around and rattle those doorknobs, and see what they can do to make money, opportunistically steal things. That is the ransomware area, which is growing fast, and constitutes a huge problem across our businesses and societies. We also have the nation-states and the stateless actors, who will become even more persistent and focused, and they will come at our government networks and our critical infrastructures. They will come in and go after our key businesses with that focused exploitation, and they are starting to not just exploit, but also to attack, and they don't seem to have restraint and reasonable boundaries on some of the things they are doing. That is really worrisome.

More particularly, I am disturbed by the attacks that have been going on in Saudi Arabia. The Shamoon virus, a destructive malware, has been deleting data, and that's beyond the pale. There was a huge international cry about this topic. The real question is what are the responsible actions we can take and make sure that this malware or these attacks are not going to come at us – our businesses, our networks, our governments? We have to consider how to respond, and think about that in an international sense; I certainly don't want to see that spread into an area where its threatening us and our friends. Along these lines, I feel like we cannot protect against what we don't understand; if we don't understand our networks, what is important in them and what their components are, we don't have the chance to protect them.

Things like the Internet of Things trends are a big worry for me, as well. There is a huge amount of very inexpensive yet powerful computers going into our critical infrastructure, our sensitive networks and facilities, and all of those now become points of vulnerability if they aren't well secured.

In my home, I built a small device and connected it to my washing machine and dryer. It allows me to know when my clothes are dry, so I don't end up with wrinkled clothes. Is it really truly important that I have that in my home? No. But it's nice, it helps me. But now, all of a sudden, I have to think about patching my washing machine. And that's what I worry about with this Internet of Things trends; are we going to design them in a way that allows us to know the vulnerabilities that they present? When that surface area of exploitable devices and vulnerabilities, combined with the belligerent actions we are starting to see on the net, explodes, I think that this will be really dangerous, and I don't sleep much at night either because of that.

One of the challenges we have is figuring out the deterrence model in the strategies; how are we going to impose costs on other nations? How are we going to convince them that their attacks are either unable to achieve the outcomes they are seeking or that the cost of proceeding in those attacks is going to be too great for them to consider executing? In that space, I think that relationships, like the one between Israel and the US, is the thing that's going to convince others at times that there is a formidable opponent among us, between the capabilities we bring as well as the relationships and the norms we have set, that will prevent them from deciding that they need to come at us.

In the US we currently have a presidential policy directive that organizes us for big incidents, and it incorporates the lessons we have learned from significant cyber events in the past. In addition, it takes into consideration things we know from other domains, like disaster response and counter-terrorism; it gives us the clarity and the guidance to organize the government, but it also makes sure that the private industry knows what they could expect from the government in terms of an incident – what are the roles and responsibilities. For example, everybody knows that the FBI is responsible to go out and investigate, to find those responsible and to bring them to justice. The Department

of Homeland Security deals with the impact of the malicious activity on the targeted entity, and they help with mitigation. We have things, like our Cyber Threat Intelligence Integration Center (CTIIC), that synchronize all of the intel community and bring the intelligence the government has to bare. This often helps the private industry, so if it's a nation-state, the private industry can benefit from the national defense focus that we can bring. We put a lot of work into incident response schema, but we are still not mature enough. Evaluating how significant our response is at the beginning of the event has been challenging for us. Sometimes we significantly underestimate how bad something is, while at other times we get a little overexcited when we see that first indication and we don't know what it is. This is a work in progress, and we are trying to build some maturity there.

As the White House, our role is to try to stay out of the operations, but at the same time to make sure that everything is well-coordinated, and that all of those capabilities and resources are played in their proper position.

When I think about the near future, five years from now, I think about how much we have changed in this ten years' period. My iPhone didn't exist ten years ago, and smartphones have changed our civilization, the way we interact, the way we get knowledge. I am motivated to get us better organized with the strategies and concepts that will carry us forward in the next five years. I think that if we don't have the kind of bilateral partnerships that we are working on here with you in Israel, and if we don't have well-defined ways to impose costs and to do deterrence, we are going to fail. I think it's really important that we, as nations, have a vision on where we are and where we need to go. The good news is that I look at the talent that Israel has brought together, and I think about the partnerships we have across our governments, and I think about the incredible ideas that are coming up through industry, and those public-private partnerships, and I am pretty confident that if we get all that right, we are going to succeed.

Cyber Security: Perspectives from a Little Red Dot

Gabriel Lim, Permanent Secretary, Ministry of Communications and Information, Singapore

I would like to talk about Singapore's perspective on cyber security. I cannot claim to be an expert, I don't have a background of building a global giant like Check Point, but this is the view from a small nation state. Let me start with what some of you might consider a bold statement, which is that cyber security is existential for Singapore; it is a matter of life and death. Why do I say this? First, because Singapore is a very small country. We are only 3% the size of Israel, less than half the size of the Golan Heights. For small countries, security is number one. We have a saying back in Singapore: "when elephants fight, the ants gets trampled." There is a second part to that saying: "when elephant make love, ants get trampled too." The morale of that is that ants get trampled no matter what, and Singapore is an ant in a jungle of elephants in the international arena.

Therefore, I think it is imperative to do what is necessary to secure ourselves. We want to contribute to a peaceful international environment, to have domestic stability, race relations, social harmony, and political stability for a long term planning and execution. I think that most of all, security is a strategic tool. We have been working together with other countries, such as Israel. Our foreign services have looked together to enlarge our respective strategic space, and it has given us the room to prosper and to develop. That said, up until recently, our efforts have traditionally been applied to conventional threats in our region. As we all know, cyberspace is a new frontier, with no regard for geographical boundaries, and we are all under attack.

Cyber-attacks are powerful tools of intelligence and influence, and we have seen how this has impacted countries like the US and France. Frankly speaking, if big elephants, like the US, are vulnerable, more so small ants, like Singapore. In fact, from time to time we experience these attacks that are mounted by foreign actors, trying to

influence our societies and our government. This is why we have to be constantly on our guard, in terms of cyber security, because if we are not carful we might become a puppet of a foreign master, pulling our strings through cyberspace.

The second reason that makes cyber security existential to Singapore is that cyber security is the cornerstone of a digital economy. In Singapore, that is paramount to us, because economic stability and prosperity is essential to our survival. We are already very digitalized, and we have invested great resources in our security. We are ranked first in global indices such as the World Economic Forum's Network Readiness Index, which came out in 2016. We see the digital economy as the future. We aim to be a world leader, by maintaining a vibrant and successful digital economy. This year we had a ministerial committee on the future economy, chaired by our minister of finance. The recommendation that came out of that was very clear, and we are making huge investments in our digital economy infrastructure. We are developing a national trade platform, to make digital commerce easier and more efficient. We just launched an SME "go digital" program, to bring small and medium size companies into the digital age. Embracing the digital economy will create new opportunities, but at the same time, connecting ourselves digitally also opens up new risks. The theft of $80M from the Bangladesh bank was just an example of how vulnerable the digital economy is to cyber threats. Hence, the more we embrace the digital economy, the more it is also important to ensure that we are cyber secure. Otherwise, cyber attacks can easily undermine our economic prosperity and effect all our people.

The third and final reason why cyber security is existential for Singapore, is that it is critical to engendering trust; trust in a world where citizens are interacting much more with each other online. We are already in a rather unique position in Singapore. We are one of the world's most trusted nations worldwide, and the World Economic Forum ranks Singapore first in the world, in terms of public trust in politicians. As we deliver more public services over digital platforms, cyber security is becoming more and more important in preserving that trust. As an example, we are rolling out the national smart nation sensor platform, which is going to plug the entire island of Singapore into a

grid and network of sensors, including all 100,000 of our lampposts. It is going to collect data, analyze it, and the end product of all this is to create and deliver better public services for our citizens. I think this will create a better country for Singaporeans, but at the same time, the public will reasonably expect that their data will be taken care of, their transactions will be secure, and that the basis upon which they exchange information with the government is kept confidential and protected. If there are cyber breaches that compromise confidential personal information, or more importantly, corrupt the sources of data that form the basis of public policy, things like medical records, military data and so on and so forth, then that trust will be eroded. Trust is like a very expansive, very antique vase; once trust is broken, once the vase is broken, it cannot be put back together again.

I started off by saying why cyber security is essential for Singapore, even existential. That is true for three reasons: we are small; we are digitally connected; and we want to engender trust. These attributes are by no means unique to Singapore, but it is the important role that each of them plays in the Singaporean story, and more importantly, the interplay between all three factors, which makes cyber security much more critical to Singapore than to many other countries.

What are we doing about that? The first order of business was to setup a national cyber security agency, the CSA. This was done in 2015, under the prime minister's office. Previously, the work of securing our cyberspace fell on several agencies. Everybody worked hard, but coordination was a little difficult. I wouldn't go so far as saying it was a mess, but we knew that we could do better by putting it all together under a single agency. And so, after careful thought and many pieces of advice from experts around the world, we established the CSA. So far, I think we have made encouraging progress. The International Telecommunications Union (ITU) ranked Singapore first in their global cyber security index. It is a positive first step, but we know that we still have a long journey ahead of us. One of CSA's first tasks was to develop our national cyber security strategy. It is based on four pillars.

First, we are strengthening our critical information infrastructure, our CIIs. While CSA is the national authority, it cannot single-handedly

take care of all the CIIs. We have divided our critical infrastructures into sectors, and assigned a sector lead to each. These are companies owned by a government agency, with the mandate to take care of the CIIs under their charge. We have also put together a national cyber security response plan, and we conduct exercises to test that. Obviously, we also have incident response teams set up for times of crisis.

Second, we work to foster a safer cyberspace, by working with both the business sector and the public sector. We organize road shows to educate the public on cyber security, and we also work with our companies and business associations to level up our private enterprises, to make sure they have the minimum standards of defense in place. I thought our recent response to WannaCry was quite telling on just how important it is to take a national approach to cyber security. We opened hotlines to which people and business could call, and we helped them recover, when possible. We publish updated reports and help everybody take the necessary precautions to avoid the next attack.

Our third pillar is to deepen our cyber security capabilities, especially among our own people. We are launching our cyber security professional training, to raise the quality of cyber security professionals in the civil service and in the government. We have also been trying to convert professional ICT manpower to the public sector. One major move was to set up a cyber unit within our military. This is a completely new cyber command, which will oversee military cyber policy, training, network monitoring and defense. It will hand-pick relevant soldiers, and build up a force of about 2,600 people over the next few years. I wouldn't go as far as to say that it will be on the same level as 8200, but we believe that the defense cyber organization will serve an important role in recruiting young talents and preparing them for rewarding careers in cyber security. That said, the supply of cyber security professionals is only one side of the equation, we have to take care of the capabilities development side as well. This is why the cyber security agency is working very closely with companies to bring leading MNCs to Singapore, to grow our local enterprises, and to work with thinktanks, research institutions and universities, in order to enhance our capabilities. I call to all Israeli and international companies, to the academics, entrepreneurs and government officials,

if you are interested in coming to Singapore to see what we have to offer, Singapore is open for business and we welcome you here.

The fourth and final pillar is to strengthen our international partnerships. I think this goes without saying, as a small country we value international partnerships very much. We are doing it very extensively with our Southeast Asian neighbors. We have a ministerial conference on cyber security, and we also have a training program, where we have contributed money to build up cyber security capacities among our Southeast Asia neighbors. However, one area where we are trying to build upon is to work with other like-minded countries on protecting super-national CIIs. These are international platforms like payment gateways, port operations, and air traffic control, where the impact of a cyber security attack on one hub has transnational global repercussions. In areas of port management, Singapore is working closely with the Netherlands; we send our officials to each other, do joint cyber security training exercises, and we are doing more to exchange information about cyber threats and how to respond to that. We think this is a useful model for us, to work bilaterally or multilaterally with many other countries, and we look forward to doing our part to achieve this collective goal for the international community. Ultimately, though, international cooperation should not just be confined to formal partnerships like I described. Gatherings like Cyber Week are actually equally important, because they bring together a diverse range of stakeholders, to discuss these issues in an informal candid manner. In Singapore we are trying to do the same. We have the Singapore International Cyber Week, which started in 2016.

I hope my sharing has given you a useful glimpse into how a small country like Singapore sees cyber security. It is an existential issue to us and we have worked out a national strategy to address this. However, I think there is so much more we can learn from one another, and I look forward to hearing and collaborating with like-minded people.

Strengthening the Bilateral Cyber Partnership with Israel

Thomas Bossert, Assistant to the President for Homeland Security and Counterterrorism, The White House, USA

One of the main reasons I'm attending and talking at Cyberweek about cyber security, is because President Trump understands that the United States cannot lessen its engagement in this region of the world and cannot lessen our support for Israel. To do so would create a power vacuum for Iran, Isis, Hezbollah, Hamas, to fill, and that would create an even more dangerous place. And indeed, today, the Israeli people have a stronger, deeper relationship with the United States of America, due to the joint effort of our leaders.

The President's visit in May demonstrated our continued commitment to Israel, and our countries remain particularly close on security issues. America's security partnership with Israel is stronger than ever, and the Iron Dome missile defense program continues to keep the Israeli people safe from Hezbollah and Hamas. The David's Sling and Arrow weapon systems guard against long range missiles, and we hope, someday, to live in a world where children never rush towards shelters when sirens ring up. There is incredible technology in the Iron Dome system. In his speech in his recent trip here, President Trump said that Israel is a testament to the unbreakable spirit of its people, and aims to stand against violence, extremism, and unacceptable, intolerant behavior.

Cyberspace is a major arena of conflict between liberal and illiberal forces across the globe, making the inter-connected world of cyberspace one of the biggest strategic challenges since 9/11. Israel is a market-based, market-oriented, knowledge-based economy, with a strong technology sector. It has the highest research and development spending per GDP in the world, and it is one of the most talented technological forces in the world. It also has a system for developing that talent, one which we can all learn from and be envious of.

While physical borders can be extremely important, cyberspace knows no boundaries. Nations increasingly have the ability to steal

sensitive information, alter data, destroy systems and turning them in the wrong direction. Destructive attacks are being executed by belligerent nations. We know that North Korea attacked Sony, and that Iran attacked the Sands Casino and Saudi Aramco, and neither one of these countries have the sophistication and resources of China and Russia. We also cannot forget the challenges facing our small and mid-sized businesses, the backbone of our economy, which are facing threats from ransomware to theft of intellectual property, by complex, advanced intelligent services.

Cyber threats continue to grow, and the complexity of the challenge continues to elude us. The question is, therefore, what is standing in our adversaries' way. Part of the answer is the basic things: firewalls, antivirus, good network hygiene, etc. Better and faster information sharing also helps in suppressing malicious activity. These are all things which we, in the United States and the Trump administration, work on promoting and improving. And yet, this would have been the same answer fifteen years ago. While these are good and necessary things, the adversary doesn't encounter them until they have already compromised their target networks for the first time, at least.

Today we are introducing terms like artificial intelligence and machine learning. We have ways of sharing information and ways to orchestrate defenses in our networks faster than we could have ever done before. Again, all is good, and all is necessary. Things are better and faster, but not different, and we always react after the adversary is already in his target's network. The Israelis and others have adopted operational constructs between the various sectors, in order to discover what the adversary is doing, and how to thwart, impede, and otherwise inflict a defensive cost on the adversary; or, when necessary, to deter bad behavior with punitive measures.

We must recognize that while we have small differences, free and market-based nations must engage with the private sector in an operational way, to identify our cyber adversaries and increase our defenses considerably. We can also do it in a way that preserves our privacy and our security in our intelligence sources and methods. I believe we can operationalize our defenses in a higher level of Internet topography, and I challenge you all to think of ways to do it.

At the end of the day, cyber security is about risk management. Networks technology will never be completely secure, and we need to prioritize our work. We need to mitigate and manage that risk. This includes identifying key data and the functions that must be protected, and then deliberately planning for protection. We must centralize policies in government and industry, and decentralize their execution. We need standards and metrics to hold managers accountable. We must implement fundamental cyber security practices to include regular patching, multi-factor authentication, encrypting data in rest and in motion, and whitelisting applications. We must also secure our nations. This includes defending our critical infrastructure, and focusing on the energy sector, communications, financial services, and transportation, or as we call them – the lifeline sectors.

There is a clear role for government in this work, but although this priority is subject to countless discussions, it has not seen the progress it deserves. Across the globe, there are countries that can do this with greater success than others, and Israel is one of them. We cannot achieve the security we need without partnerships. Partnerships with industry, partnerships with academia, partnerships with the owners and operators of infrastructure, and partnerships with like-minded countries. Increased defense is critical, as is deterrence. We must seriously consider our deterrence strategy; the stakes are too high and the risks are too great not to do so. This requires a foundational understanding of what constitutes responsible behavior and what is considered unacceptable. Progress has been made in building consensus around responsible state behavior, and the Trump administration will work to expand that consensus.

We must move from taking steps and talking about norms to implementing them, but we must also hold those who violate these norms accountable. This may not be achievable through a UN effort. We have seen examples of the limit of the UN group of governmental experts, which have achieved some good results in the past but came out short. In some cases, they are even unable to agree on final reports. It is time to consider other approaches. We also work with smaller groups of like-minded partners to call out bad behavior and impose

cost on adversaries, and we will also pursue bilateral agreements and try to increase deterrence, which may require limiting bad actors.

We are not abandoning our multi-lateral efforts; the United States will move forward internationally in meaningful bilateral efforts, such the one we enjoy with Great Britain and now with Israel, while continuing to build a like-minded coalition of partners who can act together. The cyber strategies of the future must draw upon the clear experience of history. The only way to provide a safer and more secure future in a digitally connected world is to embrace the principals of individual property, the rules of law, and a non-wavering commitment to free market, and to exclude those who do not. We share these values with many nations around the world, including Israel. We know that nations that are economically and politically free will always be stronger than nations that are not, and there has been no greater engine for capitalism and growth than the Internet. If you consider the wealth and development that the cyberspace has enabled, the Internet reflecting our values is where we will find partners who share those values.

Nations that share these values also know that the role of government is to apply rules and to protect them. The free market succeeds because of basic rules that are observed between individuals, and also because of rules designed by governments to protect contracts and promises for transfer of goods and services. When this is threatened within a nation or internationally, it is appropriate for the government to respond. The system works, in part, because those who violate the rules suffer consequences, and those who act well, do well. If individuals or nations choose to manipulate cyberspace for financial gain or geopolitical advantage, we must act to protect our shared values.

The Internet is a great example of free market at work. No capitalist is surprised that the Internet was invented in a free society. The Internet was invented in America, by Americans, and at times forgotten one Brit, but with government help. And yet, it was private industry that turned the Internet into one of the world's greatest tools. The success of the Internet is vulnerable to fragmenting, and we need to push back. The next step must be gaining international cooperation to impose consequences on those who act contrary to the growing consensus. We should work to develop options for imposing consequences within a

coalition, if possible. Until then, the United States must seek partners bilaterally.

And so, it is with great pleasure that I can announce the commencement of a US-Israeli bilateral cyber working group. This group will be led by not only Mister Rob Joyce, who is working tirelessly in the White House, but also by Dr. Eviatar Matania from the Israeli side. They will be leading the group along with the Department of State, representatives from the Departments of Commerce and Defense, Homeland Security and the FBI. The US delegation will meet with senior leaders from Israel's National Cyber Bureau, Defense Force, the Shin Bet, and the Ministries of Foreign Affairs, Justice, and Defense. The meetings will focus on a range of cyber issues, critical infrastructure, advanced R&D, international cooperation, and work force development, among others. These high-level meetings represent the first step in the strengthening bilateral ties on cyber issues following president Trump's visit to Israel, and then make good on the promise he made to Prime Minister Netanyahu in their meeting in February 2017.

The bilateral working group of experts from across agencies will work with an eye towards developing a different operational construct, focused on finding and stopping cyber adversaries before they are in your networks, and before they reach critical infrastructure. They will also identify ways to hold bad actors accountable. We believe that the agility that Israel has in developing solutions will result in innovative cyber defenses, which we can test in Israel and take back to America.

The assembled group will develop ideas that advance cyber security and produce recommendations on best practices, implementation and execution concepts. Perfect security may not be achievable, but we have within our reach a safer, more secure Internet. I look forward to the progress we make together in this endeavor.

Israel: An International Cyber Security Powerhouse

Benjamin Netanyahu, Prime Minister of Israel

A few years ago I decided to establish Israel as one of the five leading cyber powers in the world, and I think that by all accounts we are there. However, the jury in cyber security is always out, and it is a constant challenge. Cyber security is a serious business, for two reasons: the first reason is that it is a serious and growing threat. It is also a growing threat everywhere, because every single thing has been digitized, and the distinction between high-tech and low-tech is rapidly disappearing. This is happening in one country after another, in one industry after another, in one critical infrastructure after another, and as we enter the world of the Internet of Things, the need for cyber security is growing exponentially. This is a problem for all governments, not only to protect their military or their security organizations or the critical infrastructure, but also their businesses and every single organization.

How you approach this is something that governments grapple with, and we decided to apply a rule that we use in the military when you have many forces scattered in the field and you cannot really organize them. You decide, "Okay, we will go that way," and everybody follows in line, and so we organize ourselves as we move. Our decision in this case was to create a national cyber defense authority and we are organizing them around the cyber net; everybody has access to the net, which allows secure information exchange between the government, the various organizations and the business entities.

We can communicate in a secure way and the parties inside the net can communicate with each other – not only to respond to attacks, but to prevent them by early warning as well as guidance; teaching them systemic doctrines, to the extent that you can be systemic in this business. We are experiencing dozens of cyber-attacks on the national level every month, and in every given moment, including right now,

there are probably 3-5 attacks on the national level that emanate from various sources; the usual suspects and a few others.

We have set up our system. It is constantly evolving, and a hundred or so big companies and organizations have already joined the cyber net, with more expected to follow. I predict that in the end every organization in Israel will join, because everybody needs it. We are ready to cooperate with other countries, and we are ready to cooperate with other governments. In general, with some reservations, we are better together.

The second reason why cyber security is a serious business, is because it is a great business. It is growing geometrically, and because there is never a permanent solution, it is an endless business. The battle for market share in a fast growing market is open to anyone, although there are certain requirements and certain barriers to entry, but this produces an endless crop of companies. In Israel alone we have 600-700 startups, with more constantly being added as some are weeded out.

This is a fast, geometrically-rising market, and you have to be chess players. In fact, you have to be speed-chess players. This is what Israel excels in, and this is why it is a great advantage today in this business to come from Israel. It used to be a disadvantage to say you are originally from Israel. Today, when you talk about cyber or about advance technology, it is and advantage to say: "I'm an Israeli company." It is, in fact, a great advantage. We are sought after all over the world. Most of the governments, and in fact all of the governments in the world that I have come into contact with, want cooperation with Israel on high-tech, and just about every one of them want cooperation with us on cyber technology, and cyber security technology.

This is an expression of the change in Israel's status. There used to be what was called "the Arab boycott", which has dissipated for many reasons, but the prominence of Israel in the technological field and in the cyber field has made Israeli companies very, very attractive. Because we have a lot of speed-chess players; because we have hundreds of startups; because we have demonstrable success in providing solutions, for a while, in this rapidly changing sphere, Israel is becoming an attractive target for cyber security investments.

According to estimates I have seen for 2016, we have about 20% of the global private cyber security investment. The success of the Cyberweek conference is a testament to this fact.

Cyber Security – Past and Future

Matthew Devost, Managing Director, Accenture Security, USA

In 1992, when I was a student getting ready to go to grad school, I wanted to write my thesis topic on what I called the "threat of information warfare", the threat to our national security posed by our increasing dependence on computer technology. It was denied. At that time, it was not viewed as a valid national security topic to be talking about. I think we can say now that the topic has been accepted into the mainstream. We have events around the world on the topic, we have great dialogues between different relevant parties, and it is becoming normalized. We are seeing the rise of cyber historians and cyber anthropologists, an entire field of study being incorporated into everything we do. It is permeating to every piece of our lives; critical infrastructure, our individual privacy, the way we do business – all has a cyber component today. I would like to discuss some of the issues and challenges we are facing, and some of my thoughts, having been in this space for 25 years, around what we need to be doing from a thinking perspective to address those.

Intellectual property theft is obviously a significant issue, and many over the years have been saying that it represents the greatest transfer of wealth in the history of civilization. But we have also seen that this has been diminishing over time. We have diplomatic recourse, we have ways in which we can have conversations with each other around intellectual property theft. We try and de-normalize some of that behavior. We also have an issue called "innovation in parity", where you steal a lot of intellectual property, to a point where you have enough basis to put smart people on developing your own intellectual property. We are actually seeing hubs of innovation emerge from areas that traditionally have relied on IP theft.

Cyber crime is another major issue, and specifically cyber crime has an impact on the gross domestic product of some of the nations that engage in those criminal activities. It is not only about breaking into bank accounts and stealing money, we have also been seeing

attacks where business processes are being targeted. In 2017 there were attacks against Facebook and Google, and the perpetrators stole $100M from each of them. These were not sophisticated attacks, in which the attackers hacked into the companies banking systems and stole the money. What happened was that they managed to convince someone to make a change in one database, which resulted in 200 million dollars being transferred from one place to another. We see this re-focus in the cybercrime domain, not just on what we consider to be traditional measures, but also on studying business processes, and figuring out how to exploit the ways in which companies do business. As we noted, everything is now in the cyber domain.

We have also seen what a resurgence of ransomware, primarily around the fact that we have technologies like Bitcoin available, which have allowed us to create a cybercrime ecosystem. We have crime organizations that install ransomware on people's computers, and then provide their victims with user friendly interfaces where they can put their credit card number, automating the process of buying Bitcoin and paying the ransom. We see an ecosystem that is sustainable, which in turn has caused a very large increase in that type of attack.

There is espionage. Our lives have moved into the cyber domain, so naturally our intelligence agencies are going to move there as well. We see every nation out there actively going and exploiting these networks, because they have to. It is a part of their mission statement, it is a part of their job. They have merged into that cyber domain as well. From an intelligence agency perspective, we have networks, including social media networks, which represent one of the greatest technological achievement ever. These are places where more than a billion people are putting their entirety of their lives into that digital domain.

We have the non-state actors, entities that are engaging in theft of content solely for the purpose of releasing it online. We see the WikiLeaks phenomena, we see tools that are being released online and those tools being exploited by bad actors, we see proprietary content being posted online. They might not be trying to steal your intellectual property and sell it in the black market, they might be stealing and releasing it on the Internet. If you search the web for the

name of one of my best friends, the number one result would be a WikiLeaks page that has a copy of his private email correspondence. His life is changed forever. This makes it very easy for attackers to create a dossier around us as individuals. In my own personal case, 2015 was a very bad year: my health records were stolen in a breach; there was a breach in a credit bureau, so somebody stole my financial information; and, of course, there was the breach at the OPM. For those of you who aren't familiar with the Office of Personnel and Management, this is where you fill out forms for folks with a security clearance in the United States. This is where you provide the entirety of your life, all of your direct family, every place you have ever lived in, references, everywhere you have ever worked in, your foreign contacts, your foreign trips etc. The targeting that you could do against me, just based on what was released on 2015, is very substantial. Health, financial, and personal information.

I spent a long part of my career thinking about this in the context of a critical infrastructure attack as well. One phenomenon that we always mention is this disconnect between capability and intent. If you think about a terrorist organization, they have the intent to target our critical infrastructure and cause some sustainable attack. However, today, they still don't have the capability to do so. Then you think about the nation-states, some of which have the capability to engage in a targeted attack against critical infrastructure, but not the intent to do so. We have global conscience and economic interdependence. We have many classic deterrents that prevent us from engaging in wide-scale critical infrastructure attacks against each other.

Deterrents play a significant role, but we need to recognize that we already do have some deterrence measures in place. What we need to worry about is where the lines start to converge. Where those attackers that have the intent develop a full capability, or where the geopolitics change, in which case an entity with the capability develops the intent to target critical infrastructure. When I get worried at night, I worry about the alignment of those two intersections, where we would have a sustainable attack on critical infrastructure, leading to some sort of a horrific event. A power grid going down would have an impact on public safety. We wouldn't be talking about Google losing $100M

and being able to recover it; this is a completely different context. A critical infrastructure attack is going to have a sustainable impact on public safety, it is going to have an economic impact. We don't need to worry about how those attackers are viewing the critical infrastructure space now, but how they will view it in the future.

We have the opportunity here to create a new, more secure version of the Internet. We have the opportunity, as it relates to critical infrastructure, as we develop new emerging technologies like the Internet of Things, to think about security up front. We have the opportunity to prevent what I call a strategic penetration for future exploitation, which is when an attacker would target critical infrastructure without a current intent. They don't currently want to take a power grid down or go after an oil refinery, but they can envision a world in which they might have that intent in the future. They might have a time-shifted intent. Those types of attack, I would argue, look much different from the types of attack that we are used to deal with. They are not going in to engage in destructive behavior, they are not going in to exfiltrate enormous amounts of data, they are going in to establish a capability that they can use in the future.

At this point, I would like to discuss how we can deal with some of these issues, and what are the frameworks that we need to be thinking about. We have a commercial that was famous in the 1980s in the United States for Dunkin' Donuts – the coffee and donut chain, where the guy woke up every day and said "it's time to make the donuts." I feel like in security we are somewhat faced with the same thing. It's time to make security. You wake up, you go to the conference, you buy the technologies, you implement the technologies, you are in a routine. I spent the better part of my career as a red-teamer, and this is the reality that I am often faced with: the best aspirations don't often lead to perfect implementation. I have seen environments that can be likened to a zip-lock trying to lock a metal chain. The security people accepted the very visible security deficiency, they had great intentions, but poor implementation. That is one of the things we need to address.

We also, for more than 25 years, have had a disparity between attackers and defenders. We like to mention the attacker's advantage. Neal Stephenson, in his book, talks about the ways in which the attack

can outpace the defense: "when you're wrestling for possession of the sword, the man with the handle always wins." In defense, we have always had our hands on the blade, and that is a disparity that we need to address. We need to be finding technologies that will allow us to gain some defender's advantage. Maybe we can do that by thinking about security a little bit differently. Thinking about what we determine to be a success, about the metrics that we use. I am a big fan of traditional business management, and Jack Welch, who is known as one of the greatest managers in the world, is quoted to say that "you get the behavior that you measure and reward." What things are we rewarding right now, from a cyber security perspective? What things are we measuring? Are they the right things?

I am guilty of it as well. I have sat in briefings with boards of directors and said: "you had 250 high level vulnerabilities, and now you have 245." That's better, right? We have decreased the numbers. However, maybe it is the wrong way to be thinking about that problem, maybe we should be thinking about time-to-detection. Maybe we should find the technologies that will allow us to reduce that time-to-detection from several months to days or even hours.

There is a famous psychological study, where participants were sat down in front of a TV screen showing a group of people playing with a ball, and were told to count the number of times that the players in the white shirt passed the ball amongst themselves. The video began to run, and halfway through a person in a gorilla suit walked across the room. At the end of the video, the viewers were asked if they saw the gorilla, and 50% of the participants said they did not. That, to me, is a great analogy for how we are engaging in cyber defense right now. We put a security operation in place, and we have our Tier-1, Tier-2 and Tier-3 levels of support, but what are we doing to actively hunt threats? What are we doing to break the blinders we put on ourselves and that are preventing us from seeing the gorilla that is walking into our network, that are preventing us from realizing that we are walking around in knee-deep water?

We need active hunting. Not just monitoring the network, but hunting on the network. How does an attacker see the network? We need to remember that at the end of the day, an attacker is not a set of

ones and zeros. Eventually, the attack can be traced back to a human being. This means we need to understand why attackers are targeting us, how they might target us, how they are targeting our peers, and recognize that there is a living breathing human adversary on the other end. The ones and the zeros are just the mechanism that they are using, which reinforces the notion that we need to integrate more Threat Intelligence into our operations.

We also need to recognize that some of our old concepts around cyber security need to be abandoned. Perimeter security is not sufficient anymore. The creator of the first firewall himself is saying that firewalls are not sufficient. What do you do if the attacker is already inside your network? How does that change your security posture? I had the opportunity a few years ago to help the Department of Defense of the United States create their first cyber strategy. One of the things that I was a huge proponent of, that I argued for the most, was presumptions of breach. How do you think differently about security when you assume the attacker is already in the network, and you have to keep conducting your business as normal. It gives you an entirely different model to be thinking about.

We engage in much planning around cyber security, but as military strategist Helmuth Von Moltke said, "no plan survives contact with the enemy." How are you exercising that plan? Mike Tyson said it a little bit better: "everyone has a plan until they get punched in the mouth." This highlights the importance of red-teaming. We need to be testing our network from the perspective of the attacker. Not with the gloves on; we need the "gloves off" approach. If you were training to fight Mike Tyson, would you train just hitting a static punching bag, or would you train with a sparring partner? You need a way that is going to allow you to exercise yourself, to see how a human attacker sees you, and to identify those weaknesses that you can correct. This gives us the advantage of learning from our failures. In cyber security, we have this bad habit of only learning from those failures that are a result of a breach, the result of something bad happening. However, if you engage in red-teaming and exercising, you can actually discover those critical vulnerabilities, and discover those issues in advance.

You can fail, maybe even fail at scale, without there being a real impact from it.

There are many reasons to be pessimistic about cyber security, such as the attackers having the advantage etc., but there are also many reasons to be excited. Technologies like behavioral analytics, machine learning, artificial intelligence, identity and access management, blockchain and obfuscated networks are going to change the domain from a cyber security perspective. We need to introduce all the things that we know the attackers are going to adopt in the way that they attack us, so that we can use them from a cyber defense perspective as well. This is something that I, personally, find very exciting.

The Major Requirements for Securing the IoT World

Esti Peshin, General Manager, Cyber Division, Israel Aerospace Industries

A well-executed IoT strategy creates a safe world, where devices are controlled by computers and are therefore safer, because they are more predictable. IoT provides us with better comfort, better efficiency, better decision making, and eventually creates more revenue. There are many advantages to IoT, but one of the major disadvantages of IoT is that a more connected world is also a more vulnerable world. Imagine a scenario where all the devices in a person's house are connected. Each of these devices can now be a part of an attack. Some of them can participate in a ransomware attack, and, for example, the little IP camera on the desk can participate in a DDoS attack. Therefore, the IoT cyber security challenge is that the world is becoming more and more connected.

Some studies indicate that by 2020, between 20 to 50 billion devices will be connected. The researchers at Morgan Stanly even predict 75 billion connected devices by 2020. I have even seen a prediction by Huawei of 100 billion connected devices by 2025. This is an enormous number; an ungraspable number of devices that need to be protected. Why do these devices need to be protected? Because they can cyber-attack us, and we have seen this at the end of 2016 and at the beginning of 2017. Initially, there was the Mirai attack, which was discovered in August 2016. This attack affected Linux-based devices, such as CCTVs, routers, and DVRs, and these devices actually became bots, and participated in botnet attacks. One of these attacks was the mega-attack on DynDNS in October 2016, which had a tremendous impact around the world, as various Cloud-based service providers went down. We couldn't work; it created a huge global impact. In May 2017 there was another attack, named Persirai, which impacted IP cameras. The current assessment is that over 1,000 models of IP cameras were impacted, and the most generous estimations put it at

120,000 devices, taking part in this attack. This is a very interesting attack, because after the malware is downloaded to the camera, the camera downloads some additional commands and other software, deletes the malware from itself, and starts propagating to other cameras.

IAI, the Israeli Aerospace Industries, is one of largest defense contractors in Israel. One of the main issues that we are looking into today is military IoT. IoT creates a great advantage for each and every one of us, for each and every consumer. However, it can also create a tremendous advantage for the military in terms of effectiveness and efficiency, starting from the personnel, i.e. the soldier, through situational awareness, autonomous systems, facility management, and so on and so forth. IoT is everywhere; it is the Internet of everything, and also the Internet of military things. With that in mind, we need to ask ourselves – what are the IoT cyber security challenges?

The initial challenge we need to address is the fact that it is about everything. Eventually, everything will be connected; my refrigerator will be connected, my IP cameras will be connected, my car will be connected, airplanes will be connected, military systems will be connected. In fact, many of those already are connected. It will create operability constraints, as there are system engineering limits. We need to secure various types of protocols and components, and I think one of the most important elements is that today hardware security is a must. Hardware is king. We need to secure the hardware level, because this is the most resilient level when we are talking about IoT. However, when we are talking about securing an enormous number of devices, physical security becomes an issue, because we can access these devices. If we can access these devices, we can hack into these devices, and from there we can hack into the entire IoT ecosystem. We can compromise critical missions and operators, for example the DynDNS attack mentioned earlier.

I have spoken many times about the fact that when we are referring to national-grade challenges, we require national-grade solutions. We require solutions that take into account technology, methodology, and constant innovation. Israel is a great provider of innovation to the global cyber eco-system.

We need collaboration both at the national level and the international level, and we need to constantly improve our skills capacity, build-up, and maintenance. However, when we are talking about IoT, we are talking about a global challenge, and when we are talking about a global challenge, we require a global approach, which also takes regulation into account. I would like to take this opportunity to clearly say that one of the most crucial elements for securing IoT is improved regulation. We need to regulate the protocols, we need to regulate the communication, we need to regulate how IoT is implemented, because otherwise we will be in a severely insecure global eco-system. We also need to take education into account. Today everyone is talking about cyber, but when we are speaking in conferences, we are preaching to the converted. Everyone reading this understands the importance of cyber, but does the normal Joe in the street understand it? I always give my mother as an example. She is an amazing lady; she knows practically nothing about cyber, but she is constantly using her computer; she is driving a car; she is utilizing connected devices; she needs to understand the potential cyber risks in all of these devices, and this can only be achieved by very early education.

When we are talking about IoT, we need to be able to combine various security layers. Some of these layers are traditional, and some of them are innovative, starting from the regulation and ending with cyber education, training and hygiene. This is true both for the end users, meaning consumers such as my mother, and the vendors. We also need solutions for protecting the hardware; we need IT and communication security best practices, and we need monitoring of the network layer, the data centers, the cores, the multi-service edge devices, the systems, and the sensors. We need anomaly detection in order to predict potential cyber-attacks and various mitigation capabilities. In other words, we need a multi-layer approach.

We also need to be able to conduct forensics at an IoT level. Today we are able to conduct forensics for a singular device – for an end-point, or potentially for a network. Can we conduct forensics on a global IoT network level? We need to be able to conduct ongoing vulnerability assessments and penetration testing of IoT networks, in order to be able to determine their potential risk, and harden them,

and secure them. As long as we approach the IoT cyber layers in a combination of traditional processes and innovation, we will be able to create a more secure world.

In my opinion, the most important message is that while IoT security is one of the most important and emerging challenges today, there is no single company that is able to solve this. In order to face these challenges, we need collaboration; we need collaboration within the private sector, we need collaboration between the private sector and the public sector, we need national-level collaboration, we need international collaboration, and this is the one topic almost everyone in this fields agrees on. In IAI, we practice what we preach, and we established, together with 9 other companies, the "ICCC – Israel Cyber Companies Consortium". It includes some of the best and finest of Israel's cyber companies, and advocates innovation. The idea is to be able to secure from the hardware level to the Cloud and beyond, and to help the eco-system to reduce the risk of incorporating technologies from startups when we are talking about national-level cyber security. Today we have 10 companies in the consortium, and we are constantly on the lookout for additional innovative and unique companies that would join us, so that together we can provide an end-to-end solution for our end users.

Protecting a Global Business

Nasrin Rezai, SVP, Global Chief Information & Product Cyber Security Officer, General Electric, USA

For those of you who may not be familiar with the GE business, we operate in eight verticals, from aviation to transportation to power, and we have a horizontal digital industrial business. The footprint that we support in our cyber organization defense is close to 350 thousand employees with over 400 thousand devices, and in 2013 we started an international strategy, to operate and exist in the Cloud. We started working in a hybrid mode, but our goal for the next few years is to primarily operate in the couple of Cloud providers that we work with, and ultimately our own platform. The journey of the practitioner, that I would like to discuss here, is the development, the readiness, and the architectural mindset that we need to have, to actually get to that existence.

When experts talk about cyber security today, it is not about throwing a whole bunch of technologies, and hoping that you get security as a result, but rather about having a systems and architecture approach, and looking at cyber security as a business of risk management. Many of the threat actors are working at the same speed in which GE is moving to the Cloud. They are operating through the adoption of Cloud, social engineering, and all aspects of phishing, or the old methods of DDoS, to attack us. Cyber security is a major business, so we need to tackle these threats with a comprehensive, systems-thinking approach. In the last few years, so many data breaches occurred, that most enterprises almost got numb when it came to attacks. However, the last few IoT major attacks woke us up, in the sense of thinking that what the cyber researchers have been telling us for some time has now become reality, and here it is. The argument that I make to my peers, especially the CISOs in the industrial space, is that we can no longer, with the accountability that we have to our boards, say that OT security is not our job. Many of us have been called in to have a discussion with our boards, and were told that OT security is our job as much as IT security, because it is all interconnected.

Cyber security is definitely a corporate risk in all and most enterprises. The other major corporate risk that we are dealing with, especially in the industrial space, is decline of productivity. That pattern has existed for some time now. We, in the industrial world, enjoyed productivity of 3%-5% for a long time, since the 1990s and up to 2005, consistently year after year. But starting in 2011 we have experienced a decline, and that productivity has gone down to about 0.3%. Therefore, many of the industrials have shifted their focus in leveraging technology, to create a digital transformation. With that came GE's introduction, three to five years ago, of a new horizontal we call "digital industrial". This change was the trigger for creating the first digital industrial platform, which we call Predix, with capabilities such as asset performance management, digital threads around automation, and modernization of manufacturing. Our aim is to enable "digital industrial" for our partners and our customers in all the verticals in which we operate, and to take an advantage of the new innovation and technology. We want to help them utilize the assets that they have in their infrastructure, through machine learning and analytics, to gain some of the productivity that they lost in the past few years.

A good example of this is a partnership that Predix and the digital business has with BP. They predict that through leveraging Predix and the industrial platform of GE in 30 of their oiling rigs, they can tackle the unplanned downtime that they deal with. This, in their own projection, is expected to bring about £200M in savings. We recognize that this is a tremendous opportunity for GE, as well as for "digital industrial", but at the same time we also recognize what this effectively offers as a solution. This is more about a connectedness of assets and devices, in an environment that for the past 40 years has been air-gapped and controlled under extreme regulation. In that regard, we take cyber security very seriously, as a key component of this end-to-end design, and under that light we created our own philosophy of what enterprise security and product security should look like.

At the highest level, we have three constructs. We think about how to build products, how to defend the enterprise and the commercial offering for our digital industrial customers. How do we ensure that

we build security into our products? How do we defend our enterprise, which our customers trust with their data, and find where the gap exists in the OT space? We are coming up with commercial offerings with partners and technology providers, many of them are in Israel, to take a holistic approach to our customers and to the market.

On the enterprise side, we think about this shift of the entire workload and platform to the Cloud in two parts. One – how do we adopt Cloud, with all the native capabilities that the maturity of new Cloud providers offer, without having to recreate them, and buy into the protection controls and mechanisms that they offer us? Two – how do we run our services natively in the Cloud, so that we can give the best security value back to our own employee base?

If you assume that we will get some of the capabilities from the Cloud providers, then we have to shift all of our investment, in terms of architecture and technology, to data protection and application level protection. This is what my peers and the other CISOs in the market are rooting for, because many cyber security practitioners still think in network level control terms, when it comes to protection. The difference between current detection and response capabilities and what we actually need, becomes something we cannot solve with the existing Cloud providers. Those are our biggest problems, based on the threat actors that we track, and the type of detection capabilities that we need in order to enable the existence of this big organization, which is GE, in the Cloud.

No matter how much regulatory control comes onto large enterprises, we still have a lot of money, and we can handle the process of building security. The weakest links are the third-part suppliers, both on the product side and on the enterprise side. They are a part of this ecosystem with us, but unfortunately they lack some of our security controls. Therefore, we put tremendous focus and attention on them. On the digital side, as we shift more and more into a digital organization, now we not only have hardware engineer, but software engineers that are generating millions of lines of code on a daily basis. The mindset of a security organization that assesses the code after it is already out is very old-school. We have to shift into secure software design, and development processes that are fully automated. It is a very big

mission, when you have to work that out for a 350 thousand person organization, but it is a big charter for us.

Our secure "Edge to Cloud" deployment model is a partnership between cyber security and all the engineering functions at GE, to ensure that we develop "Edge to Cloud" as a market differentiator for secure Cloud deployment. This means that if any of the industrials have a GE asset, e.g., a power turbine or a renewable engine, and they are sensor-enabling it, the data traverses out from our costumers' networks into our Cloud and Predix platform. We believe that it is not enough for each product to be secure, it is the horizontal that needs to be secure. We have instituted a program that aims to govern, protect, detect and respond, and applies that to all practices of engineering across GE. The "govern" phase or a component of our protection is based on the NIST framework, and every product has the responsibility to adhere to its principles, which are governed by us. When it comes to protection, we look through the lens of the costumer, asking ourselves how we can ensure that our products, our devices, Blue/Red Team assessments, threat models, have native trust assessment capabilities, which we can demonstrate to our customers. Detection is an area in which we are developing capabilities, but it is also a tremendous opportunity for technology providers to partner with us. Lastly, we provide an integrated CIRT/PSIRT response process.

When I think about defending the enterprise, an important piece of the puzzle is to establish a clear priority. To me, that is basic cyber hygiene. If you have $10 to spend on cyber hygiene, I would suggest you spend $6 on effective end-of-life, hardware, software and patching policies, and pay attention to the fact that the methods are different based on whether you have your workload on-premise or in the Cloud. Spend about $3 on identity and access and monitoring. Then spend the remaining $1 on your hardest problem in cyber hygiene. We, cyber professionals, tend to find ourselves attracted to solving small problems, and we put all of our technology effort in that, and that is not the right process. The next important element is the idea that offensive capabilities inform the defense. We need to use actual attacks to build defenses in IT and OT. Many of the products in services that we are developing are starting to integrate those elements, technically,

into our product offering. Third is automation, which is critical for cyber leaders. We no longer rely on "eyes on the glass", on analysts who delve into millions of lines of logs to really do the work that we need to do. We need automation to do defense at scale. That applies equally to how we think about risk and compliance, using automated continuous monitoring. On top of all of these, perhaps the centerpiece of the puzzle is exercising our resiliency muscle. This is not so much about us, as cyber security professionals, but rather about getting our leadership and our people ready. When the big incident arrives, which happened to some of us with WannaCry, you need people to cooperate with you. They will already know that there is a commander, they will know and understand their role, whether they are business people, your legal team, or your external council, and they can be part of the solution with you. Those are the five principles that we live by.

I want to take this opportunity and conclude by saying that I am very proud to be the first female global CISO for GE, and I am honored to have been awarded that. However, I am not mentioning this to boast, but to say that to the best of my knowledge, there aren't too many female technologists. Women constitute 55% of the world population, but in the US only 10% in engineering are female, and 25% in IT. This is something we need to fix, and I am happy that GE is taking steps towards that future.

Strategies for Securing IoT Devices

Matan Scarf, Strategic Advisor, Blavatnik ICRC, Tel Aviv University; CEO & Founder, Cycuro

I would like to discuss the equation that illustrates IoT cyber security issues. If we can understand what creates the problems, we can also remediate them. When we build an equation, we have to start with axioms, and my first axiom will be that we need to understand that IoT is a business. Not only is this a business, it is a factory that takes data and converts the data into revenue. This is a critical point, which we will refer to when discussing the strategy for resolving IoT cyber security. My second axiom is about the security property of the system: if something is connected, then you can hack it remotely. We know that everything is hackable, and what we need to understand is if you take something that is hackable, and connect it to a network, now anyone can hack it from anywhere in the world. My third and last axiom is that we need to understand the unique characteristics of the IoT domain, specifically when it comes to consumer apps of IoT, and mainly in the field that I work in, which is the automotive cyber security. When the complexity of a system is very high, building it, and therefore protecting it, is just as complex.

When the risk associated with the hack is possibly injury or death, for example in hospitals, critical systems, transportation, and so on, there can only be one conclusion, and this is my main argument: if there is money involved, if things are hackable, and if it is critical, you have to do security by design. This is the equation that I am about to describe, and this is the actual point that I will return to at the end of this article: how do we do security by design for IoT?

There is also something we need to understand about hacking. Hacking is not something that happens randomly. It is not something that is arbitrary. There is a force that drives hacking – or, to be more accurate, there are three different forces behind vulnerability research. One is what we all like to talk about all the time, and that is the state sponsored vulnerability research. This constitutes a small portion of the overall hacking we see happening in the world today. The

second component is the researchers/hackers, people like myself, who take it upon themselves to try to find vulnerabilities in systems for the purpose of protecting them, and for the purpose of creating awareness. The third, of course, are the cyber criminals. What we need to understand is that the first two groups that I mentioned have a very different ROI approach to vulnerability research than the latter. Cyber criminals are only doing things that have the potential to create revenue. For a state, if there is a specific individual or a specific system out there that they are interested in, the amount of time and resources they will invest in this campaign can be something that normally no one would invest if it was a commercial product. The same goes for academia researchers. A researcher, such as myself, can try to find a vulnerability in a specific vehicle system, which is something that hackers or cyber criminals would not do, because it would not scale as well. This, naturally, bring us to the point of Mirai.

The Mirai malware is a very interesting case, because it is unique in many ways. One of these is the fact that it was dedicated specifically to attack IoT devices. And this should create a question: why would anyone build malware that is specifically designed to attack IoT devices? Why not attack PCs? Why not attack mobile devices? And that goes back to the question of scaling. As an attacker, what you want is the ability to build, for example, a botnet, which will be undisrupted, which will be persistent, and which will scale in numbers, so that you could reach hundreds if not thousands if not millions of devices in a very short time. As a malware developer, your goal is to monetize it; you want to sell your botnet to cyber criminals, so that they can execute DDoS attacks, for example. This makes IoT devices the perfect target; the attacks surface is huge, these are devices that normally can't get software updates, or at least the mechanism of distributing them is very limited, and there is a significant gap in security products for them. There are no security products running on IP cameras, there is no antimalware, antivirus, or firewalls for most home routers, and if a certain one comes pre-installed with one, it is not an industry-grade product. For me, as a hacker, IoT is a really good way to build a botnet. However, I am mostly talking about general IoT, not just

consumer IoT; and ironically, it was a traditional malware that had the greatest effect on IoT, and that was WannaCry.

The WannaCry malware is a cryptographic ransomware, which targeted mainly Windows-based computers. That said, the effects of WannaCry reached beyond just workstations; it crippled hospitals, it crippled factories for automotive OEMs, it crippled delivery systems. How does a traditional malware affect the IoT environment so dramatically? This brings us once again to the fact that today everything is connected. We have production manufacturing machines that are connected to the backend network, and that backend is still mostly Windows based. You don't have to target the IoT directly, just find where it is connected and attack that spot. This means that my second axiom mentioned earlier is inaccurate or incomplete. Yes, if something is connected it is also hackable remotely, but the truth is that even if something is not connected to the Internet, it can still be hacked remotely. There are cyber security experts out there saying that some things should be disconnected, but I want to say out loud that we cannot assume that air-gapped systems are not going to be affected.

Recently we heard about the discovery of the CIA's Brutal Kangaroo anti air-gap device, a USB device that breaches aircraft systems, and I think the most famous example goes back to Stuxnet, the worm that affected the Iranian nuclear enrichment facility, which was supposedly disconnected. I can also tell you that in my field of automotive security, while there is only a fraction of vehicles out there today that is connected, meaning that there is a SIM card embedded in the vehicle or an after-market solution that creates connectivity. Despite that fact, we can still remotely attack non-connected vehicles. For example, one of the cases we demonstrated in our research in Cycuro is an attack utilizing the pairing of a mobile device with the infotainment system of a vehicle. That is something we need to consider when building a cyber security strategy. When we have an ecosystem that is populated by both connected and non-connected devices, the overall ecosystem is all connected.

If we want to build a strategy that is actually going to resolve the underlying issue of IoT security, we need to understand this is a financial issue. We need to devise a strategy that will make the cost

of the attack higher than the potential gain, or at least higher than the cost in other domains, so that we deflect the attackers somewhere else. To do that we have to have security by design, which could include things like regulation, components for a securer over-the-air update, an ecosystem where different security products work together, and we really need to work on building more sturdy platforms by using secure software development life-cycle practices. The potential end of this road is when money really becomes something that is not at the reach of the attackers.

Trends and Future Predictions in Cyber Security

Bob Kalka, Vice President, IBM Security

I've been in IBM for 28 years, and in 23 years out of that I've been helping to build the company's security business into one with more than 8,000 Cyber security professionals. One of the things that I've always enjoyed during that time, when looking at research that comes out in this area, is looking at the story behind the story.

In our field, we get many different kinds of data – numbers, percentages, costs, etc. And that is very interesting, but those of us that are in the security world are asking: "what do you want me to do with that data?" For example, I am going to refer to the Insights from the 2017 Cost of Breach study. I would like to show you how this study exposes some of the dirty little secrets that exist in cyber today. The reason why this is so important for us to do, is that when we take things that essentially exist but we choose not to look at, and bring them into the open, we actually have a chance of doing something about them. What we are going to see in this study ties into some of the most emergent issues that we see happening in cyber right now.

A good way to get into this discussion is to look at how the typical IT shops builds their security operations programs today. Across our 17,000 costumers we see many things that repeat themselves. We know what the right thing to do is, but the great challenge is often obtaining the ability to actually execute what needs to be executed, and get to where we know we need to get to.

If we look at the generic view of how most of our clients do security operations and response when we start working with them, it can generally be divided into four phases. Everybody starts with some kind of security analytics technology, and everybody starts the course with collecting, normalizing, correlating, reporting and monitoring logs; in other words, SIEM. 99.9% of our clients have a SIEM today. Roughly 80% of them are not really happy with how they are doing it, but at least they have it, so that they can mark a check on that box for the auditor.

What we realized is that these days, more than 50% of organizations have moved beyond just looking at logs, which is old stuff, and today they realize that looking at real-time data is important to catch use cases. For example, let's say you have contractors, and you know that they download on average three confidential documents per week, and right now you have a contractor downloading 300 confidential documents. If you are just collecting logs, you are probably not going to find it, and if you do, it will only happen later, and they've already left the company.

Today, the combination of real-time flow analysis plus logs has become the best practice. But we also point a lot of attention at user behavior analytics, because instead of 35 things that your analytics tools tell you, it is much more interesting to know that a single person has generated the traffic that indicated those 35 different things. Finally, there are many advanced things, such as incident forensics. That is what we do to identify what we should pay attention to, in order to figure out if we really have problem or not. Most people are actually still stuck down at the logs phase, but there is a lot of room to grow up from that.

Once you find suspicious stuff, then you can get into the trendy term of Threat Hunting, and we ask ourselves, how do you figure out if you really have a problem or not? We did a study of some of our clients last year, and we found that the most common practice of taking a suspicious indicator, and finding out whether you are actually having a problem or not, is an incredibly complex technology called the manual google search. According to the security analysts we interviewed, they find 5-13 sources that they trust, and in a sense they ask them if they have seen a configuration change of the mobile device that let this type of network traffic react in a certain way. They are looking for relevant documentation, or other means to confirm whether it is really a problem or just something benign.

Some clients go beyond that, into investigative analysis, by pulling feeds like social media, intelligent video analytics, as well as phone call logs and such. The core problem is that our researchers have discovered that only 20% of the available threat intelligence today is indexed and searchable by any search engine. 80% of it is on the

deep web. If you have your Level 1 security analyst running around doing this, it means that the most common practice today is to be manually searching through only 20% of the data.

After Threat Hunting comes Threat Intelligence, because you want to know what is going on with the rest of the world, and we found that most people do that by multiple feeds of Threat Intelligence. A study we conducted last year, asking clients what they were doing with those feeds in practice, resulted in the common answer: "we look at them." The inability to get actionable intelligence out of the Threat Intelligence tools is one of those dirty little secrets in our industry that only few people talk about, but it's very real. That said, we haven't found a programmatic way to do that either.

The next and last stage is Incident Response, which is actually dealing with a problem once you have it. This is probably the most naïve belief I've had in probably 20 years of cyber security. Two years ago I believed that at least 95% of the organizations had well-defined, up-to-date Incident Response plans, which listed what incidents are you worried about, and what's the detailed runbook you are going to execute when any of those incidents occurred. We've discovered that less than 2% of the organizations actually have such a plan. So in fact, nobody toady is doing proper Incident Response. If you don't have an up-to-date, detailed Incident Response plan instead of a 2-3 page guideline, it means that technically you are making it up as you are going. So to go from that to an actual automated response, and then pulling in cognitive stuff (i.e. Artificial/Augmented Intelligence), is clearly becoming a big issue. This was the biggest surprise to come out of our study.

Today the trendy term is Threat Hunting, but the next trendy term is going to be Orchestration, and that is simply is an extension of Incident Response. And so, if we come back to the 2017 Cost of Data Breach survey, it says that the number one factor that lowers your cost-per-record of a data breach is your Incident Response. You expect things like encryption, employee training, business continuity management, disaster recovery, threat sharing and security analysis. Incident Response, however, has turned from a dirty little secret to an obvious little secret that you just have to do. The story gets even

more interesting when you look at the part of the survey that describes what *increases* the cost-per-record in a breach: third parties, cloud and mobile. The fact of the matter is that none of these things are going to change. That is one of those fascinating things in the world of security.

In another recent study we conducted, 83% of the participating organizations stated that digitalization is a major priority for them, but only 15% of them felt like they have a mature plan of how they are going to accomplish that. There is a massive gap here, which calls out all the need for identity, data and application security.

If we look at the data according to industry, another interesting story appears. With Health, Financial and other general services at the top of the cost-per-record in a breach graph, there is no doubt that that IoT is the next big thing. And as security professionals, we need to understand that IoT devices are going to tax us in two major ways. The first way is obviously technology, and the question we have to ask ourselves is: do we even have the technology to protect the data, the transmission, and everything else as it relates to IoT? The second way in which IOT is going to challenge us, the cyber security professionals, is in management – how do we manage these threats within the context of the entire organization?

24 years ago, I got my MBA from Syracuse University in New York. I actually started with a marketing MBA program, but when I started working on cyber security, I immediately switched to a very unique program called Small Group Psychology. The reason I decided to get that degree is that you learn very quickly that in terms of group psychology, security is a shrink's paradise. In individual psychology, when you see a part of reality that doesn't match the reality you want, you usually have a dysfunctional response. In much the same way, cyber security represents the reality of what is happening in the IT systems. Not our hopes or wishes or thoughts of what's happening in the IT system, but what actually happens there. We represent reality, and the organizations around us don't necessarily want to see that reality.

IoT is going to bring the OT environment and IT environment together, and if you don't have an OT environment, you will now. That will force you to allow yourself to improve your management style as well, having to deal with some of the organizational psychology issues

that are related to those things. This is ground zero for psychology in cyber. The organization needs to work with us in collaboration, but there are many unsolved problems remaining, and there will be a lot of research coming on that.

The security division of IBM is a $2B business, with over 8,000 security professionals, so we have solutions to all of the current security needs. However, there are many dirty little secrets in this industry, and it's our ability to unearth those, and effectively deal with them, that is going to help us not only secure our businesses better, but it is actually going to lower the costs when a breach does occur.

Changing the Security Paradigm

Bharat Shah, Corporate Vice President, Microsoft Cloud & Enterprise Security Division, USA

Not a single day goes by without us hearing about somebody getting into some networks, stealing data, ransomware, etc. It does appear like the deck is pretty heavily loaded against the defenders. Typically, what we see happening is that there is much energy and efforts being putting into protecting the perimeter, and even analytics at the perimeter level. This, of course, assumes everybody is patching, keeping everything up to date and using multifactor. If they do all that, and you protect your perimeter, what you typically see happening is that the attackers, especially these days with automation, keep trying to find the chink in your armor. The defenders work relentlessly, trying to protect against these kind of attacks. Whenever there is scale, and humans involved, sooner or later something is going to go wrong, and it only takes one little mistake or one little gap, somewhere left unprotected for the attackers to get in. We need to rethink how we have been doing things, and what needs to be done.

We have to realize that the bad guys are probably already in, and maybe there are several of them. Even if you take care of the tech, as described above, you are going to have a tough time figuring out what have they done, how many backdoors they have created, where are they lurking, and so on. You can try to evict them, but if you don't know where they are, it will be tough. Worse, even if you do manage to evict them, you have to learn how they got in, run the forensics to figure out what was your mistake, where was the gap, to really go make things better. This is a very serious problem, which we really need to understand how to solve.

We have to change the way we are doing this. The way we in Microsoft have been thinking about it, and are acting on it, is that we expanded our perimeter defense further. Today, we have many sensors inside the network, in addition to all the perimeter sensors. I would assert that every resource you are trying to protect is actually a pretty good sensor, as long as you can collect their logs and telemetry

information. That said, we use even more sensors than that. We are big fans of identity-based sensors, and everything else that we can do with identity. For every attack that masquerades as an application or an end user or a machine, the more identity-based sensors you can leverage, the better it is. We use end-points, identities, resources, and we also found good correlations by looking at crashes that relate back to attacks. The more sensors you have, the wider the aperture, the faster you will be able to detect. Moreover, collecting large amounts of data, and actually keeping them for a while, will allow you to go back and do better forensics. It will allow you to understand where the attackers came in from, and what else they have done. As the attackers move along, somewhere along the way a sensor will fire, or maybe even more than one. In the classic model, the defender makes a mistake, and the attacker gets in. Now, if the attacker triggers even one of the sensors that we have, they will be detected.

The other thing to note about the sensors, is that we really don't think that you can look at all telemetry and the logs that you are collecting in isolation. Being able to connect the telemetry and the logs from your sensors, and have a broader context of what is going on, is the actual power you have. We call it the "intelligence security graph". According to one saying, attackers use graphs, while defenders use lists. We think that defenders need to look at graphs as well. If you are able to look at all the data and logs, and build a context around it, you will be in a much better position to not just detect faster, with a wide aperture of sensors, but also to have the forensics needed to clean up your network.

My next point is this: almost every customer I meet, every organization, is going through a digital transformation process. The Cloud is obviously one of the best ways for organizations to accelerate their digital transformations. Every customer that does a pretty good job with their on-premise resources is concerned about the Cloud, and that they will lose control over resources and lose visibility. The truth is that most Cloud vendors – and I can attest for Office 365 and Azure – do a very good job of looking at a wide range of sensors, and collecting large amounts of data, with almost no cost to the client; we are already doing it all the time. Your resources in the Cloud are just an

extension of how most well-managed organizations are running their infrastructure today. Another thing is what I call the "Cloud effect". From our point of view, your resources are just one tenant, but we have millions and millions of other tenants. Any "unique" attack on one tenant is actually not as unique when you look across millions and millions of tenants in the Cloud. The "Cloud effect" amplifies our ability to detect and do forensics and even go as far as attributions, discovering where the attacks are coming from.

I mentioned intelligence earlier, but the scale in which we operate gives us a unique value. We look at more than 300 billion authentications per month, billions of Windows machines and end-points, millions of servers, and billions of emails which we detonate and clean up and detect as malicious. Anything that we find, any malicious email or attack on a VM, we can apply within minutes the necessary rules to all the mail boxes and other VMs in our Cloud. This is a very unique power that we bring to bear, for all the assets in the Cloud.

The more you look at your network, the wider the aperture, the more sensors you have and the more you look at your resources, the better your ability to detect what is going on. This starts with your perimeter, data going inside your network, data going outside of your network, and across the Cloud. If you collect all of this data, it is going to be a lot of data; volumes and volumes of data that you will have to process. At some stage human limits will kick in. It is impossible for a human to look at all the alerts that are there. The answer to that is a much more fundamental adoption of machine learning, Cloud intelligence and artificial intelligence. We at Microsoft already applied a great deal of that in our Cloud, but I see a world where every organization, small or large, will take steps towards embracing machine learning and artificial intelligence in order to actually protect their resources. We are doing a pretty good job with these, but we are still in the infancy of where we potentially will be five years from now. I can guarantee you that five years from now, if you look back at what we are doing today and maybe even the next year, it will look like we are in the little leagues compared to professional baseball. The opportunity for what we can do, going ahead with Cloud harnessing power, analytics and machine learning, is tremendous.

To summarize, the more sensors you have, the faster your detections will be. The more data you collect and the more data you can process, the better near real-time detections you can achieve, and you will also have great forensics so you can see what is going on. Couple that with the Cloud, and you get some tremendous Cloud effects and benefits, just by being a part of an infrastructure with millions and millions of tenants, which allows us to find and defend against attacks in a very uniform way. Security organizations and Cloud vendors are making big bets on machine learning and artificial intelligence. These things together will fundamentally change the equation, and make defense a far more viable and easier option than what we see today.

Protecting Cloud Environments

Roy Adar, Senior Vice President, Product Management, CyberArk, Israel

One of the things that we believe is of huge importance to adopting Cloud technologies and DevOps technologies, is what we call privileged accounts. These are the keys to the Cloud kingdom, if you will. Privileged accounts are accounts that exist on every piece of equipment that you have, every software application and every system, which allow to manage, operate, control, and administer that system. By design, those privileged accounts allow the people who get a hold of them to do whatever they want, without restrictions or limitations, without filtering their privileges. In a good and normal usage, your IT teams or your developers are those who need to use them in order to do their job, and in order to operate or develop applications, or manage the IT infrastructure. When bad things happen, though – and there is a ton of research that was done to document and measure that – it seems that in almost all of the attacks that we heard of, one of the steps that the attackers went through is to compromise credentials, and especially privileged accounts. The reason for this is obvious, as they give the attacker the way in, which is fast, covers everything, difficult to detect, and difficult to prevent. Obviously, this applies to insiders as well. Insiders don't have to break into the perimeter to get into your organization. They are already on the inside, and they may even already have the right to use privileged accounts to do good things. However, on the occasion that they choose to do bad things, privileged accounts are a way to help them do that.

In CyberArk, one of the things that we do when we talk about protecting privileged accounts, is to come at it with the approach of the attacker. First of all, we hire many attackers who work for us on research, to help us understand how attackers go about compromising networks. Especially in Cloud, there are many new things that are more attractive to attackers when it comes to privileged accounts and controlling the keys to your Cloud kingdom. For example, in Cloud deployments – especially those that follow DevOps methodologies –

we found that, on average, more people have access on a daily basis to actual production data and production systems. Usually, the entire DevOps team has access, unlike in the on-premise world, where you had some kind of separation. There were developers, and there was a more controlled process for moving to production. With DevOps, it is all automatic. It is all happening all the time, so there are more people to go after, if you want to compromise the organizations.

There are also many more DevOps tools and technologies that are being used and adopted. I am not saying that there is anything wrong with them, but they are new, and because of that we can assume the normal maturity curve from a security point of view. The bottom line is that they have security holes in them. Over time, mature technologies fix their holes and their security issues. However, with so much new technology out there governing the DevOps processes, there are currently more holes and more places to poke holes in. One of the things we do for organizations is also to help them with red team exercises, and it also helps us understand scenarios that involve privileged accounts. One Cloud scenario, as an example, started with the attackers phishing users. Phishing, as you may know, eventually always works. Once a person from the DevOps team was compromised, the attackers found the root account to AWS on their desktop. With that key, they connected to AWS and cloned the servers. Now they had a copy of any server that they want from the organization, and they could take their time to pick and prod it. It was actually a very quick process. Some of you may think that this is too easy; that this is a made-up example. It is on the medium side of easy, but it is definitely not too easy, and it is definitely something that could happen in many of the organizations we have seen. The question is, what can you do about it? There are four major points to be taken care of.

The first point is to discover major risks. There are many solutions and tools that allow you to scan and find the privileged accounts and the access keys that exist in your environment, that allow control over Cloud assets, applications, and hardcoded credentials in applications connecting to a production database. The fact is that you will, in fact, get a very long list. The important thing, then, is to prioritize.

Which one of them is indeed important and puts you at risk? Start by understanding and mapping the solution.

Second thing we recommend is to look at your Cloud management consoles the same as you do for your own premise and similar to how you manage access to your physical data centers. Normally you would have stronger controls and measures before someone can get into your physical data center. Not everyone has a handprint or fingerprint type solution, but you would have some stronger controls. With Cloud management tools, sometimes it is just a file which is an access key, and anyone who has that key can immediately and directly access your Cloud console, and do whatever they want to your entire data center.

The third point is to properly design your security policy and controls. We definitely recommend looking at it from a heterogeneous point of view. If you have different security controls for your own premise from those you have for the Cloud, then you are just delaying some of the agility and flexibility down the road. Imagine having an on-premise application that you want to migrate to the Cloud. It is not just a simple adaption of the technology and the ops around it, but suddenly you have to manage credentials differently, and you have to manage other aspects that would delay that migration. Ideally, when you think of a broad strategy regarding your security and controls, you will also allow migration to the Cloud without the need to reinvent security and controls, every time, in every application that you are considering to migrate.

Finally, the fourth point and one of the more important recommendations that we can make, is to look at securing the credentials and the secrets these applications are using. Many applications, unfortunately, contain hardcoded usernames and passwords, which they need in order to connect to a database. This can come in multiple forms, such as configuration files; it could be hardcoded in source code; or it could be scripted. Either way, these credentials are just waiting for anyone who can see your source code to see, steal, go directly to the database bypassing the application, and own everything. The recommendation is to move to something more secure. Lock down the credentials, rotate them, and ensure that they are eliminated and are not hardcoded into source and applications when working in the Cloud.

We have seen that, on average, organizations are able to take the core of those four recommendations and implement them in a period of thirty to sixty days. This doesn't mean they will fix all their privileged accounts issues in sixty days, but we have seen many organizations look at the key aspects of securing their privileged accounts in thirty to sixty days. It is not as hard as some people may think.

Where Is My Security?

Iftach Ian Amit, Senior Manager of Security Engineering, Amazon, USA

I would like to discuss old school security. The Cloud is just a new version of our old servers that used to be in our basements, in our racks, in our closets, and in our data centers. In some sense, I get the feeling that we somewhat lost touch with what is the meaning of truly doing security. Just a quick disclaimer, this is my view of security; it has nothing to do with that of my employers.

In big Hollywood movies, we see scenes of hackers taking control over airplanes, enterprise systems, mobile phones and more. I have friends who have been practicing in real life each and every part of these hacking scenes, but I am not here to scare you; I just want to make a point about how connected we have become. Everything is now connected – our phones, our social lives our business lives, our home entertainment, our home automation, our cars, and other things. All the devices are really connected to each other, and we can add more connectivity and more services on top of those, to make our lives easier, to make our work more effective, and just to improve our daily lives and even our social lives. However, that comes with some kind of burden; the real question is: where is all that data? Can anyone really know where each and every piece of data and information that we are generating, processing, accumulating, and passing along is really saved? Where is that data?

In the old times we used to have floppy disks and hard drives. We knew that the data was over there, and it was a much more closed system. We had, at least, the feeling that we have control over where that data is. These days, with distributed systems that are available 24/7 in multiple areas of the world, we don't really know where our data resides. That causes a big problem, a big question, about how do we secure that data; how do we address it? How do we make sure that we grant the right access and deny the wrong access to the data that we generate? We also wonder what is really the data, because as anyone who has ever taken a picture on their phone knows, it is not just the

pixels that compose the visual imagery, there is a lot of metadata that comes with it, and is carried away with it through whatever service we upload that picture to. It is a big question. How do we treat that data? How do we really secure that data?

Let's say, for a minute, that you do know, or you think you know, where that data resides. What does that really mean? Do you know what are the legal implications of storing data in one country versus another? Do you know how to protect it? Do you know what happens if you lose data, or if someone steals that data? It really depends on where that data resigns, in terms of what your reaction should be and what your activities should be, and that should factor into your disasters recovery plan and incident response plans. I am not sure that everyone understands the full implications of it. When we are putting data out in the Cloud, when we are using third party services to process our data; we need to take it into account. We need to carry out the proper mapping of how do we treat remote data. Data, very much like animals flocking, and like people going places, does not want to stay in one place. Every time we consume data; every time we connect another service to enhance our lives, to enhance the connectivity of that big network of things connected to things; every time we do that, we open up an exponential complexity of where does that data resides. Can you provide me the assurances that a service that you think is just processing your data isn't keeping it somewhere? If it isn't, maybe it just rolls up and summarizes the data. That is also something that we need to take into account. How does that affect me? Do I need to account for it? Do I need to put legal bounds around it, or treat it differently? It is a big question.

All of these really boil down back to the fact that we, as security professionals, haven't been doing that great. You see many products out there, and you hear many people focusing on providing products to protect the digital element of our businesses, our organizations, our people and individuals. However, they often forget that as security professionals, we are not just tasked to deal with the digital element. It does not live by itself. It is not about the silos of having products protecting the digital realm. It is about making sure that we account for everything around it, including the physical and the social. As

any "red teamer" will tell you, red teaming is not about poking and prodding on a firewall, or stealing a password or some credentials. It is about simulating an adversary, a real adversary, which is not limited to the scope of a security control, or to the scope of a security product. That is what keeps creating the gap, what keeps providing attackers with the competitive advantage over defenders. That gap, in turn, is being filled, or supposedly filled by security products, and we have been have been indoctrinated just to click on them, to install them, and to say "oh yeah, that sounds right. That sounds like something that we should have." We add more products, while what we are actually doing is filling the digital controls, forgetting about the entire fabric of what an attacker sees.

Companies are used to think in terms of control. As anyone who travelled to the US and ran across the TSA knows, those kinds of controls are no more than security theater. It is a "feel good" button; it is a "feel good" mechanism; I have the data, I can see it. However, the new norm is moving to the Cloud, a facility that can secure your infrastructure better. Some say that this is more complicated and more complex, and I actually disagree with that. I think it is less complicated than setting up a new server; that used to take about 8-12 weeks, now you can do it in five seconds. It is really a matter of taking the skills that we learned when we were younger, when we just started in this business, and applying them again in this new environment. Nothing has really changed. The only thing that changed is the technology, and that is easy. At some point we forgot that we need to keep applying the basics of security to this new technology. With the Cloud, we have to remember that the basics still matter. For example, basic mapping of what we have; many companies miss out on that. We are setting up servers, setting up network, and forgetting that there are implications. Every server that you create, every service that you connect into that server, every third party company, every piece of code that you consume or use in your environment, needs to be mapped and fully understood, in terms of what is my fabric that I am now protecting, including the physical, the social, and the digital.

It really is going back to the basics, and if we look back this is what we used to do, this is nothing new. Even when we used to set

up data centers and rack servers, and connect networks, this is what we did. We took account of inventory; we took account of what is the perimeter, what is my attack surface; we locked those servers physically, we locked down that facility. We made sure that only certain people could have access. The same still applies, it is just a different technology, it is really not that complicated. I know I am supposed to scare you, and tell you that the Cloud is the new thing, and that you need to get certifications, and to learn everything about it, but really, security people still do the same basic things. We need to make sure that we are making life easier for our enterprises, not more complicated. You don't need the wizardry of launching a new server and securing it, like in the old days, when you had to rerack the production server versus the beta one. When you had to retest things. There is no wizardry. We are eliminating much of the hard work that we used to do, ten or fifteen or twenty years ago, as security professionals, because the new technology allows us to do it faster and more efficiently. We need to keep simplifying that. We need to keep inventing on top of the technologies that we have, on top of the Cloud, on top of whatever is going to come up next.

My message is to go back to basics. Be able to create a proper mapping of your assets, know where your data is at; know what services you are using, what third parties, and what components. Create a threat model. Really basic. When I run into a CISO or anyone who is dealing with security, my first questions are: what is your threat model? What are the assets the you are protecting? The businesses ones, not the technology ones. Who are your threats, and how do you account for them? What kind of controls are you putting in place? Not just the technology controls, your actual controls, which may include some tech products. These are the kind of questions that we, as security professionals, need to keep asking and to put in context, even when we are dealing with those newfangled technologies, even though they are not really new, they are just the new norm.

Self-Driving Information Security

Jim Reavis, Co-founder and CEO, Cloud Security Alliance, USA

The Cloud Security Alliance is a global non-profit organization. We develop vendor neutral security best practices for Cloud, and we also do a lot of work in IoT. We deal with anything that we feel is next generation technology. However, when we started building the Cloud Security Alliance, almost ten years ago, there were tremendous number of naysayers. Many organizations and enterprises said that they will never put any assets in the Cloud, beyond things that are completely low risk, and perhaps consider using a private Cloud. Over the years those same organizations have developed very aggressive Cloud adoption strategies, and many of them are already in there. We feel, to a degree, that we have really fulfilled that idea of having defensible security best practices for Cloud, which was our original mission. That said, we still need to look at how the world is changing, and we need to look at evolving our own mission. We feel that our mission has changed, and it is now about putting Cloud computing at the center, at the core, and at the foundation of information security, which we believe needs to be primary delivered as a service in the future. This is a change that is going to shift things quite a bit.

Self-driving information security, which is what I want to discuss here, is a little bit of a misnomer. It is this idea that technology is certainly disrupting many industries, and so I am using it more as a theme. I am using autonomous driving to portray how Cloud technology is going to disrupt information security. I think it is very important that we discuss this matter. I believe that we are going to automate information security. I am not saying that in 2020 half of the cyber security personnel will be robots. What I do think is that we are going to find that the magnitude and the scope of the challenges and the changes we need to run up against, meaning that we need to automate. We need to think about how to automate much of what we are doing today, because we have so many tasks ahead of us. Here is my perspective on how to break down these problems.

When I was born, there were probably about 40 thousand human beings for each computer in the world. In a few years from now, I think we are going to find it will be inversed, we are going to have 40 thousand computers for each person. It will actually be hundreds of thousands of computers, if you count everything; every chip out there for every human being on Earth. Many of the practices we have are rightfully not changing, but there are also going to be many things that we will need to change. We talk about the Internet of Things, and about many different types of technologies, but I don't think we've really put our minds into what that means for humanity or society. We are currently at a point where things are moving quickly, and this is as slow as we are ever going to move for the rest of our lives.

Today, when I talk to those organizations that were against Cloud when we were starting, it is clear that they gave up. Cloud computing is going to be the back-end. Organizations are decommissioning data centers, and even if they keep them around, it is, in fact, an orchestration. The point is kind of moot on whether the computer is here or the over at the Amazon or Microsoft data centers. The way that is orchestrated, the way computing is going to move around, that is really Cloud computing, and that is our back-end. The Internet of Things is the end-point. What I mean by all of this is that computing is literally everywhere. It is all around, surrounding us. We now have ambient technologies that dissolve in the background, and potentially always listening to us. It is everywhere, but you won't know where anything is. If you think you truly know where all your technologies and data are, you are not taking full advantage of any of this, because you are not letting it be free.

We have gone from software upgrades a few times per year to changing constantly every day. We have gone from this point of naming our severs, and keeping them up for a long time, to a point where everything is instantiated with virtual machines or micro-services in a very rapid way. What was here ten seconds ago is already gone now. To me, this really means that we need to think a lot about how do we cope with this type of scale, understanding that we are not going to get the same exponential growth in human beings on this earth.

Certainly that is true with professionals that are going to seek this particular discipline, even though we certainly want more of them.

I talk to a lot of the organizations out there, and I try to benchmark the market. I spend a lot of time on compliance, but I am also trying to spend more time on information security. I see a great investment in people working in security, but it is not enough. We are going to need more people, and we are going to grow this industry, but the world is going to depend on many other areas than just Cyber security. We won't have the manpower, and we are going to need to come up with a new cyber security framework. We believe this idea of software leading the world. Security will have to be delivered as a service, and to be orchestrated in a very rapid way. We have, for example, a project regarding software-defined perimeter, how to create very secure "black Clouds", meaning virtual private clouds that span multiple Cloud infrastructures.

We need to take things to the next level. There are two very important issues I must address. First is Blockchain. I think that this immutable log, this transparent ledger, is what allows us to have a digital currency. It is probably going to outlive Bitcoin itself, and I see that as something that is going to change information security fundamentally. This trusted way of communicating, of perhaps encapsulating SLA, contracts and all those sort of things, will make a huge difference. The second is Artificial intelligence. I think we are going to find that as being vital to how we assemble this sort of framework. Today, when we mange information security in a SOC, we take hundreds of millions of transactions, and then we filter that down to twenty that an analyst needs to look at. I have no confidence that we are getting the right twenty there. We are going to need to be able to apply AI and machine learning. We are going to need to have Blockchain. There is currently a great amount of work on autonomics, and how we automate crucial things, like DevOps.

This has been my call to action. I think we all need to work on this together. We need to think about automation, and we need to think about self-driving information security and how it is going to change the way we do our work.

Cyber Security for Autonomous Vehicles

Rick Echevarria, VP Software & Services Group, GM
Platform Security Division, Intel Corporation, USA

For most people in the US, Friday and Saturday nights are usually their favorite time. It is the time when everybody is either ready to relax are ready to party. For me, Friday and Saturday nights, around 11:15 PM, become a nightmare. I have four children, and most of them old enough to be out, enjoying themselves, but my wife insists that no matter how old they are, they should be home by midnight. So, at 11:15 PM, no matter how much I want to relax, my wife is nudging me and saying: "start texting them, they need to come home." As I am texting them to come home, the one thing I hope and pray for is that they are not reading that text while they are driving. Cars are something that we know so well, something that we have been dealing with for such a long time, and yet we see an increasing rate of accidents and loss of lives. Cars are more dangerous than ever, because in a data-driven world we want the information in real-time, and sometimes we are willing to risk ourselves just to get and consume that data. Therefore, today I want to about security in an autonomous driving world.

One important question to answer right off the bat is: why is Intel talking about autonomous driving? Autonomous driving is a great opportunity for the industry, and if you think about the concept of security and safety, it all starts with a hardware foundation. In that hardware foundation, we are committed to delivering an important number of usages; secure boot, which we don't think about in our car, but in an autonomous world we need and we want a assurances around that; secure storage; cryptography, and not only today's cryptography but future crypto-acceleration; trusted execution environments; and something that you are going to hear a lot about, and for which we are partnering with the industry, which is resilient, artificial intelligence algorithms.

The thing that is so fascinating to us, as we look at the autonomous car world, is that security in this context is about a lot of moving

parts. You have to protect electronic control units; you have to protect in-vehicle networks; you have to protect sensors from misbehaving; you have to protect the communications between the vehicle and the infrastructure; you also have to protect the data that is flowing all the way from the vehicle to the Cloud and back, and all the transactions that are going to happen in between. A lot of technologies equals a lot of data flows. You have to protect the user's privacy, and, as I already mentioned, you have to protect the algorithms that we are trusting in an autonomous world. That is a lot of technology to deal with, that is a lot of complexity to deal with, and that is a lot of an attack surface to deal with. As we all know, complexity and surface areas are the biggest enemies of defenders. That said, this is something that we have to deal with, and I am going to try and describe Intel's efforts in security for autonomous vehicles.

There are three areas that we are focusing on. The first one is protecting the vehicle and the foundation of the vehicle. It is going to require a lot of technologies, including a lot of hardware technologies, and not only we are committing to working with industries to secure the vehicle, we have to worry about the integrity of the communications of the vehicle. Which leads me to the second part of our strategy: the security of the communication and the connectivity from the vehicle to the infrastructure and everything else that the vehicle is going to be connecting with, like other vehicles, or other moving parts in the world that we live in. We have to ensure that any data that leaves that vehicle and goes back and forth is coming from a trusted source, and hasn't been tempered with. Furthermore, we have to do that at the lowest latency possible. This is a true real-time system. Finally, the third element of our strategy is the protection of three things: the data, the algorithms, and the privacy of the users. In an autonomous world, the privacy of the user is a very different context than from what we have known so far.

Securing the vehicle is quite a complex topic, with a lot of moving parts. A modern vehicle can have as many as a hundred electronic control units, two hundred microprocessors, and at least nine networks with a lot of different protocols. A vehicle like this also requires about a hundred million lines of code. However, in the world of autonomous

vehicles, the number of lines of code goes up by a factor of three. Three hundred millions lines of code. That is a tremendous challenge to secure. What are we going to do about this? First and foremost, we have to establish a basis of trust when it comes to operating the vehicle. Our plan is to collaborate with the industry to ensure that we are establishing the right chain of trust in vehicles. We have to attest to every system in the vehicle; and not only each individual system, but how the systems are going to interact. We also have to make sure that we enable critical code to be isolated within a trusted environment. We have to add and support a number of different technologies. Software isolation is going to be very important, same as whitelisting and being able to enforce control flows in the vehicle. There is also a significant need for key and certificate management, not only to manage their authenticity, but also their life cycle.

Another aspect of vehicle security is a research we are conducting about self-adapting mechanisms for the car. We want the car to evaluate and modify its behavior. Not only do we want to do this for efficiency purposes, but we want these car to have some level of self-protection, and we want these vehicles to have some level of self-correction. The self-adaptive and self-healing mechanisms are going to be absolutely necessary for autonomy.

The last area in vehicle security space we are dealing with, is working on the ways we can give the user a seamless experience, especially when it comes to dealing with their interaction with the vehicle. When we are talking about autonomous vehicles – and you are also talking about this concept of us transacting through the vehicles – we have to worry about a number of different areas that are typically reserved for computing: multi-factor authentications, and biometrics. We want to use multi-factor authentications to ensure that we have the right owner riding the right vehicle, accessing the right data and credentials. We also want to use biometric technologies in real-time interactions with the users. We can detect, for example, if somebody is falling asleep in the car, and we can then self-correct.

At this point I would like to discuss the way we secure the communication and the connectivity of the vehicle. This is important, because without real-time connectivity we are not going to be able to

reduce the number of accidents. We need connectivity and low latency, because these vehicles are going to be interacting with a real-life environment; not only with other vehicles, but also with pedestrians and pets, real-time objects. This is important because part of our motivation is efficiency. Traffic is becoming a very big deal, and we want to leverage infrastructure so that we can get to places faster and more efficient way. After all, time is our most valuable asset.

We are going to put a tremendous amount of focus on the connectivity being secure and with low latency, and that this is very challenging. Security and performance, security and low latency, this is usually not a natural interaction. It is going to require a lot of engineering, and this is where the promise of 5G is going to be realized. That is why we are so committed to our investment in 5G; the promise of high speed, low latency, and security. That is why, as an industry, we are all rallying behind 5G and 5G standards. In fact, we are putting our investment and our focus in participating and being very active in standards, including the 5G Automotive Association, of which we are a member. Standards are being written for 5G include identity, authentication and interaction between the vehicle and the infrastructure, and between the vehicle and other vehicles. We absolutely intend to support those standards.

The third part of our security strategy is the security of the data, the algorithms and the user privacy. You have probably heard Intel describe autonomous vehicles as data centers on wheels, and that is because our assumptions, our expectations, and what we have measured, is that these vehicles are going to generate four terabytes of data, per vehicle, per day. That is a lot of data to manage, to process, and to protect. We have to protect the data from the inside, going to the outside and coming back. There are many hops that the data is going to take through the infrastructure. In order to achieve this we will focus on the themes that I mentioned earlier: encryption, connectivity and the proper use of trusted execution environments along the path that the data traverses.

The algorithms are equally important. You have to think about the fact that these neuro-algorithms are becoming more and more popular, and are being used to process things like vision, speech recognition,

and image classification. We are going to have machines that are seeing for us, thinking for us, and making decisions for us, and we have to expect that they will do it better than humans, otherwise what is the point? The challenge is that these algorithms, just like anything digital, can be exploited. This is where our partnerships with the academia and the ecosystem come into play. We are very proud of the work we are doing in the Tel Aviv University, Georgia Tech and Princeton, specifically focusing on resiliency in this artificial intelligence algorithms.

The last point is around user privacy. When we think about user privacy, there is an aspect of vehicles that store our data, and they know something about us that we have to pay attention to. The life cycle of a car is usually five to ten years. This means that we take a car, we drive it for a certain amount of time, and then we sell that car. The last thing we want is for the new owner to have access to our credit card, our credentials, our contacts, or to anything related to the life that we have been living inside of the car. We have to make sure that we are building the right types of technologies and work flows that enable us to "wipe the slate clean", as we are moving in terms of ownership. There is actually one issue that I think really brought this problem to life, which is the fact that recently we had the movements of Mayor of a European city tracked through the smart parking infrastructure. That is why this concept of building privacy into the communications of the vehicles, and into the data of the vehicles, becomes very important. We have spent significant amount of time looking at this issue, and we have actually built technologies that are privacy-aware into our silicon, in order to address issues like this. We still want the data, we still want the benefits of autonomous driving, but we can't compromise on this level of privacy. It is a matter of safety.

To summarize, our focus, where it comes to autonomous driving and security autonomous driving, is in three areas: the first one is the security of the vehicle. There are lots of technologies in the vehicle for us to manage and secure. The second one is on the protection of the communications and the connectivity between the vehicle and the infrastructure, while maintaining low latency because these are real-time systems. The third element of our security strategy is protecting

the data, the algorithms, and the privacy of the user. We look forward to the road ahead, the collaborations with the ecosystem, and the great opportunities for innovation that autonomous driving, and specifically security for autonomous driving, present to all of us.

Catching the Next Snowden

Chris Inglis, Former Deputy Director, NSA; Managing Director, Paladin Capital; Distinguished Visiting Professor, US Naval Academy

I was previously the deputy director for the National Security Agency, and I served as their Chief Operating Officer for the better part of eight years. One of these years was the year that Mr. Snowden came out. In 2016, Snowden's story was turned into a big Hollywood movie. I think, and this is all I will ever say about Edward Snowden motives, that there are two actors associated with this movie – one of them just happens to be Edward Snowden himself. But having said that, what I would like to discuss here is what that story was about – not from a motivation perspective, or what the downstream damage was, but rather the system that allowed Snowden to do what he did in the first place. I believe in the "live and learn" modality, and I will look into the lessons we can take from it. At this point in time, this is not about how to catch Edward Snowden in terms of what he did. Rather, this is about the general principles you need to establish, both in general terms and with enough precision, in order to catch the next one. Unfortunately for all of us, the next one will not operate exactly like the last one.

There is a set of beliefs, or assumptions, that were established some 10, 20, or 30 years ago. Those of us who essentially started defending discreet computers, and then a discreet small set of computers, and then larger enterprises, perhaps still have this as a bias. We used to think that existential threats largely came from the outside, where most of the malicious actors live. That inherently causes us to think that the outsiders are the ones we should spend most of our time and attention on. We think that the insiders are inherently more trustworthy, because we went to great pains to find them, vet them and bring them in. We think that they deserve less of our time and attention.

The second assumption that comes from historical experience is that we know that there is a low probability of an insider threat in terms of quantity, the number of events that would be attributable to

an insider. And this assumption sometimes gives us a bias, because we will focus on those things that are more likely to happen, without regard to the potential consequences. A low probability event with high consequences probably should have more of our time and attention than a whole span of high probability events with low consequences.

The third thing that is implicit in the way we organize, and many organizations are still in this place, is that we perceive that security can be managed independently and sub-optimized in at least three different stovepipes. There is, of course, the information technology stovepipe, which includes the software, the hardware, the doctrine, and the procedures that are connected to that. However, there is also a human resources perspective, related to the trustworthiness or lack thereof of individuals, the permissions that you should or should not give to them, the way they exercise those, etc. There are many indicators that are related to general human behavior, rather than their network imprint. There is also physical security; locks, gates, guards, cameras, etc. If you still make the mistake of managing those independently, they will only cross over when one of them passes a certain threshold, something happens that is meritorious enough to get our attention. However, what we should be doing is thinking at all of them as one atomic unit from the bottom up, and you allow the tipping to occur before you get to a meritorious security event, such that combining two or three of these can show us something that only one of these could not.

Finally, the last bias is a strategy that doesn't work so well in the face of an insider threat today, let alone in 2013, and this is the strategy of "reacting well". According to this, we believe that we will be able to detect an insider threat in time, and catch it as it races to the virtual door. Of course we don't do that in the physical world anymore, despite the fact that cockpit doors are excessively secured, and have a very strong lock on them. Our strategy with respect to physical acts of terror is not: let's catch them on the airplane. Neither should our strategy with respect to insiders be "we will catch them in the act" or "we will catch them as they hit the door." This is very problematic to do, especially amongst a trusted population.

If we assume that these were the common beliefs back then, let me give you a very high level overview of the National Security Agency's strategy regarding the insider threat in the spring of 2013. This is what I would have told you if you asked me then how we defend the enterprise, which is comprised of the information technology, the people bound to that, and the procedures, against that threat. I would have said that we do a pretty robust job about the "protect" mission, traditional information insurance; the things you do with the software, the hardware, and the setup, the patching, and all the things you do in order to make sure that you have a defensible enterprise. We had a strong application of "defend" measures, believing that there was a real probability of breach. At the time, as mentioned above, we had a bias towards outsiders, literal outsiders, coming across the berm, and therefore we practiced defense-in-depth.

However, we were not mindless of insiders; we were certainly mindful of them. However, our traditions, our practice over the better part of 60 years, had been that if you go to great lengths to find people to whom you believe you can extend trust, you vet them in extraordinary ways. Some might even call that today extreme vetting. We vetted them to the point that the United States system gave them a polygraph before we placed them on the line. We had some reasons to believe that we could extend trust to them, and give them some ability to act without day to day, or perhaps moment by moment, oversight. We allowed them to live in an envelope of trust, and we essentially monitored the margins of that envelope. If they got to that margin, we assumed that at that point in time we would still have time to get them back to the straight and narrow, pick them up, and dust them off. If it was an inadvertent error, hug them and send them on their way. If it was an advertent or malicious error, the assumption was that the size of the box was small enough that we could contain the damage to a time and place where we can recover from that, and then take whatever HR actions are appropriate.

We, of course, also had forensics. In the face of something that we already knew had happened, we believed we could actually do some expert research and quickly determine what happened, and therefore

put ourselves in the place where we would make sure that it didn't happen again, or at least reduce that possibility.

All of those in combination were a reasonably robust strategy, but in the end it didn't work. The reality is that the insiders' privilege yields an increasingly existential vulnerability. It might not be an existential threat, meaning there might not be some person who is willing to exploit this, but it is in fact an existential vulnerability. We have fewer and fewer people who have more and more privileges, because we've taken people out of the system with automation, machinery, and all the things we've done to concentrate more leverage and ability in individuals. There are still going to be system administrators. It is an issue of whether you have too many, or whether you have not been careful and discreet about limiting their privilege. But no matter what, there will always be someone in your system who has enough privilege and leverage to do you harm. And this does, in fact, constitutes an existential threat, if you trust them but don't verify what they are doing. In the case of insider threat, if the previous point is true, then the combination of low probability high consequence puts you into a place where this is a serious issue, and probably you ought to change your point of view. It is not enough to ask yourself if you have an actual threat, but rather you need to ask if this is a vulnerability that can seriously hurt you, if and when that threat manifests. You have to deal with it when it is still "just" a vulnerability.

IT, behavior, physical security and HR, are all parts of the equation. Consider the possibility that a person who has a workplace incident, decides not to get mad but get even. They have their privileges, and they are attendant to their ability to get even. Then they elect to operate low and slow for period of months after that. You've got an HR clue that something perhaps is amiss. That clue might tip and cue the information technology security stovepipe, to say that this person with privilege is perhaps a heightened risk. This, in turn, might give you a sense that you ought to be watching the physical disposition of how this person operates on the campus; do they go into privilege spaces, do they have a work/life balance that perhaps is somewhat different, and if it is different on certain days, why is it different? You begin to gain an approximation of how you might combine those three looks

to better decide if you should be more worried about this person. If you don't do that, then what you actually do is to treat these as three independent stovepipes, and you will essentially revert back to what I would call a transaction mentality. If a given transaction in one of those stovepipes exceeds a threshold, then and only then you try to figure out if there is a correlation in the other stovepipes. The bottom line is that you fall prey to somebody who knows they are clever enough to beat you low and slow.

And then, finally, significant threats can move faster than you can react. A strategy that used to be "react well" simply can't work in the face of somebody that is clever enough, and perhaps devoted enough, to essentially seize the moment through his audacity or agility and race you to the door. They will win the race.

What, then, are the lessons learned? It is not enough to study the past, we need to say something about what do we do about that. There are some technologies that are clearly useful for this purpose. While not new or novel, they are not in a wide enough use. These technologies don't just look at transactions and compare them so some prior benchmark, but are looking at behavior. That's a fundamentally different look. You shouldn't look discreetly at your network through a soda straw, action after action after action. You stand back and try to understand the nature of this behavior, if it is anomalous or not in the context of itself. Some degree of automation can help in that regard, as does some degree of artificial intelligence, but it is a horizontal look at the system as opposed to a vertical look through the soda straw.

The second thing is that you need to do holistic security; combine those three stovepipes. That is not often done, and especially in corporations, where we have a traditional look at things like risks. We look at them through a very focused expert committee. What we have are reports, and those don't reasonably satisfy anyone except the COO or the CEO.

The third piece of the puzzle is culture, and for me it cuts in two ways. Culture is important because the people are the most variable, the most important, the most agile, the most creative, and the most unpredictable component of any system. Cyber, as a system, is much about the people as it is about the technology. You need to solve that

piece. Certainly there should be accountability. You need to make sure that people understand what their roles and obligations are, and hold them accountable to that. But more importantly, there needs to be the sense that you have bound that person to the larger enterprise, in all the positive ways that are appropriate, so that you don't suffer the fate of an organization of one – somebody who is in your organization, who physically looks like they are betwixt and amongst counterparts, that they have some kindred obligations, but in fact they are acting as an organization of one, exercising their own individual aspirations, wants and needs. In the case of Edward Snowden, he was a contractor in a federal organization within the United States. He was therefore, by law, treated as a commodity. He was not special, because, by law, he was told that he is simply supposed to exercise a function. He was not, therefore, given the same sort of feeling of any enduring relationship or proposition with the company, a feeling that he would have been given if he was a civilian employee. That in itself, in my view, does not justify in any way, shape or form everything Mr. Snowden did. However, if you miss this opportunity to have one more safety net, one more check on what perhaps motivates somebody to act as they do, then you missed bringing culture to work in your behalf.

Managing a National Level Cyber Attack

Ciaran Martin, CEO, National Cyber Security Center, UK

The UK's cyber security partnership with Israel is one of our most recent and fast developing one. We have recently agreed on a new strategic initiative between the UK and Israel, and we are having constant discussions about technological collaboration on operational detection, policy, communications, and how we manage responses to incidents that affect all of us. That, in my opinion, reflects that special and rapidly deepening partnership on this crucial issue. Why are we developing this partnership? Partly, it is because we see this issue in the same way. This is about safeguarding the freedom and prosperity that the Internet gives us. We want security that is good enough, so that our digital and economic experience can be fantastic, as it should be. We are moving beyond the era of what I would call "scare and sell", where we frighten people and companies about the threat, and then tell them to buy this or that solution. We are now, in the UK and Israel, working on hard problems; how do you block at scale? How do you stop government employees from going to bad websites and infecting crucial national systems? How do you stop spoofing? How do you stop that economic damage? We are going to work jointly with Israel on all of these.

In terms of living and learning, I just want to offer a few humble and tentative reflections on our recent experience of WannaCry. We were set up to improve UK national cyber defenses, that is our job, but realistically we have to expect some attacks, and we are the lead coordinator of the incident response. There is a rule in Britain, saying that crises always start on a Friday. And so, WannaCry started on a Friday, and while it didn't have the tragic human impact of a terrorist attack, make no mistake that the images of sick people, unable to get treatment in various parts of the country, was not a welcomed sight, and one that we had to deal with as a national priority. The way we handled it was all about the data and the crowdsourcing of the solutions. We, as part of GCHQ, the NSA equivalent in the UK, received data that only we could lawfully access. We mixed that with data from industry,

and we pushed it out as fast as we could. Our infrastructure helped, as we are an organization that works simultaneously with classified and unclassified data, and our new building in central London allows us to move between the classified and unclassified spaces within seconds. From Friday afternoon until Saturday afternoon, we worked around the clock, with global industry and other governments, just to share our findings, and collect what our partners outside found on their side. We shared information about what was and was not working, as well as recommendations and suggestions on what to do next. It worked. Within 24 hours we had a piece of specific detailed technical guidance out to the nation, on our website, providing guidance on how to contain this threat, and how to mitigate it. We were worried about Monday morning, so we worked on Sunday to promote that guidance through the media, and Monday turned out to be less damaging than we expected.

What did this teach us? What are the things that worked, and what are the things we need to do better? First, when we are dealing with a cyber incident, as a government – but I think this applies to companies and individuals as well – we need to predict the unpredictable. We need to be prepared for things that we cannot predict. If you would have asked me what will be the first national level cyber incident that will dominate the media in the age of the National Cyber Security Center, I wouldn't have predicted a seemingly random global attack, trying to blackmail bodies in the health service in the UK. By the way, by policy these bodies are not allowed to pay ransom, so the attackers had a strategic failure, even though this had a significant human impact. We have to ask ourselves, what will be the next attack? I think, in this maturing field of cyber security, that we are still in the era where we identify a major incident by using the good old criteria "you kind of know it when you see it." We need flexibility and agility, in order to be able to manage whatever comes.

The second thing is that we need the full range of covert/overt data and communication. We cannot have something in the classified space not put to use, and then have less effective attempts to tell people what to do. We need to fuse it as quickly and as safely as we can into

a product usable by other people. That is partly why our operating model has been set up the way it has.

The third thing that I feel really passionately about is that this is what partnership means. I think that we in the UK, and others, have been deeply frustrated, over a decade or more, with the constant calls for information sharing partnerships, private/public partnerships, government and industry collaboration, with nobody understanding what this means. What is the output? What is the purpose of this, beyond some subscriptions schemes, conferences, and very nice dinners? The answer is exactly the ability to call people in California, in Israel and all over the world, in the middle of the night, saying: "what have you got? Give it to us. Is this working?" and to get a specific piece of guidance that gets British hospitals back up and running, as quickly as possible. That is partnership. I will never forget, and always be grateful for, Dr. Eviatar Matania phoning me on a Sunday morning to tell me how the first day of the working week, in Israel, was working with the mitigations that have been put in the country. That is, again, partnership in action.

The fourth point is about public communications. These crises will sometimes have a classified dimension, but sometimes they won't. In either case, we have to communicate with the public. One of the reasons we were set up in the UK was a 2015 breach, which was not particularly serious. That said, it was a breach in one of our major telecom operators, and so it caused a public panic. At one point, four million citizens, in a country of 64 million, thought that they were at risk of having their bank accounts emptied, and there was no authorities or public voice to tell them what the real situation was. We have to change that. Our citizens need to know that there is an authority there, which has the situation under control, or is getting a grip and bringing it under control. However, there is another reason why communications are so important. We did all this great work, helped by partners across the world, but we needed people to get out there and do it. People needed to know it was there, so that relentless public messaging over the course of that weekend was absolutely central to bringing the situation under control.

There are things that we can learn from and do better in the future, but what all of these did was that on Monday morning, instead of dealing with a new spate of infections, and a new spate of cancelled operations and so forth, we were able to send our best technical people to the worst affected hospitals. That, in a nutshell, is how we want this model to work. We take away as much of the harm, from as many people, for as much of the time, and we deploy our very best experts on the toughest problems.

Two final reflections. One, there is no substitute for prevention, and most of the prevention we can do to raise our defenses is very simple. As the famous American political saying goes: "we hope we haven't wasted a good crisis." We hope that, in terms of ransomware, the British public is far more aware of the things they need to do to protect themselves better, and that they are all simple to implement and affordable. We hope that our message has gotten through. Second is the future. We may be in the era where we are detecting new types of attacks. We are used to look at attacks as a "one off" incident, but as you may be aware, in April 2017 we reported a major compromise of managed service providers, globally. This was not a "one off" attack, but rather a sustained campaign against us and our allies. How do we manage that? How do we tell the potentially hundreds and thousands of corporate victims about how to withstand that? Again, we do not need to fight the next war, we need to get our systems more resilient, our partnerships more effective, and to able to respond at speed and the put basic protections in place.

I know that we can count on friends in Israel, our friends in the US, and elsewhere. I think we should be clear about the moral purpose and value of what we are doing: the Internet is a tremendous opportunity for prosperity, so let us work together to secure it. It is doable.

Our Digital World: A Crisis of Quality

Trevor Rudolph, Chief Operating Officer, WhiteHawk, USA

In 2017 we found ourselves fully immersed in the digital age, an age that was made possible by some of the great technological advances of the 1960s and 1970s. Today, our modern devices possess a trillion fold increase in processing power, compared to their 1960s predecessors. According to *Experts Exchange*, the Cray-2 supercomputer from 1985 had approximately the same computing power as the iPhone 4 did in 2010. We are on the verge of our vehicles and medical devices being completely autonomous, not needing the very subjects they serve, in order to drive, pump and compute. But as we all know, these advances did not come without peril. In our interconnected world, the trade-off for convenience has often been an erosion of privacy, trust, and quality. My assessment is that in our heist to spur innovations, we have exacerbated a crisis of quality.

In 2016, Johnson & Johnson warned their customers that hackers could hijack insulin pump communications, and to deliver unauthorized injections to patients. In January 2017, the Food and Drug Administration (FDA) in the US issued an advisory that hackers could interfere with the interaction between Saint Jude's implantable cardiac devices and its corresponding transmitter, allowing hackers to rapidly drain batteries, and deliver deadly shots for heart rhythms. Some say that there is no inherent conflict between innovation and security; that when you consider the potential impact of a security event on a company's reputation, it should be intuitive for senior leadership to understand and prioritize security by design, when rushing a product to market. The same camp suggests that the solution is in a concept known as SecDevOps or DevSecOps, where collaborations between cyber security professionals and software developers will produce more security aware products from inception. They insist that this "marriage" will help ensure that security and privacy protection are baked into to the earliest stages of product design and development.

In my opinion, this view is naïve, and does not take inherent human behavior into account. SecDevOps can exist within established

empires like Google or GE, but most reasonable individuals and companies, which have finite resources at their disposal to build, secure and test products, will make trade-offs to produce the most marketable products as quickly as possible. Worries about security and quality upgrades come only after the revenue position allows it. This behavior is inherently rational, and will not change without intervention by outside forces.

I often compare today's technology industry to the automobile industry in the US of the 1950s and the 1960s. It was common in that era to have brand new vehicles without air bags, anti-lock brakes, or even seatbelts. It took a future presidential candidate to wake us up to the reality that we are driving in a world that was unsafe at any speed. Ralph Nader's book was, and is still, controversial, but its impact cannot be denied. Shortly after publication and many congressional hearings later, the National Highway Traffic Safety Administration (NITSA) in the US was created, with a mission to save lives and prevent injuries, through education, research, safety standards, and enforcement activity. Since that time, vehicle fatality has dramatically decreased, over 53% from 1965 to 2015.

I firmly believe that today we are riding our digital vehicles at one hundred miles per hour, with nothing to brace us for impact. You might say that this is an overly dramatic comparison; unlike the auto industries, the problems in the technology industry do not involve loss of lives, but don't be fooled. I firmly believe that loss of life is the next stage in our evolving crisis of quality in the digital world. The warning signs are mounting. As previously mentioned, products have been rushed to market, while any notions of security, privacy or trust have been lost.

It was not long ago that CNN and other media outlets reported that the CIA had pulled out operatives from Beijing, after the Office of Personnel Management (OPM) hacks, for fear about spies being held in captive or worse. As someone who helped lead the federal-wide response to this, I can tell you that the potential for loss of life was on everyone's minds. What was the culprit then? The federal government's aging IT infrastructure and devices, which by design allowed users to authenticate through insecure means. Looking ahead,

the most catastrophic outcome will likely not arrive from a single device or even a set of devices, but rather by the cumulative effect of many technology products, over time, being pushed into production before they are properly tested for security.

The news is not all grim, though. For starters, there is an increasing body of technological knowledge on how to build secure systems and products. For example, NIST's "Special Publication 800-160" on system security engineering addresses actions that are necessary to develop more defensible and survivable systems. There is also the EU Network and Information Security Agency (ENISA) that recently promoted baseline requirements for security and privacy in IoT devices. The foundational principles contain within these and other publications are nothing new. We have long had guidance for how to build more resilient systems. What we still lack, though, is a governance system to include enforcement, sentence and penalties, to ensure effective implementation of stronger security design practices.

President Obama's cyber security commission hinted at a solution in their final report, when they referred to a nutrition-label model, potentially being applied to technology products. As an example you can look at how the United States Department of Agriculture (USDA) grades beef, to get a clear sense of how that beef is assessed for quality before it ever reaches the kitchen table. Similarly, one can look at how the Israeli Ministry of Agriculture and Rural Development oversees the production and marketing of fresh food, to ensure quality and safety. Another example is the Consumer Products Safety Commission (CPSC) in the US, which is in charge of protecting consumers and families from products that pose fire, electrical, chemical or mechanical hazards. They execute this mission through a series of voluntarily and legislatively mandated, consumer products safety regulations, and processes of accreditation and certification, which use external laboratories to test products for compliance and standards. In early 2017, the FDA announced that it was working to create a similar third party certification program, which will allow low risk medical device products to be marketed without FDA premarket review.

Where is the analogous example for enforcement of the underlying information security of all technology products? Such an example

does not exist. There is no organization within the US government to ensure that all technology products are built to the latest information security design and practices. This is problematic, because market forces alone will not solve our crisis for quality. First, as discussed before, smaller, more innovative companies have little incentive to spend their precious resources on the underlying security of their products, before they go to the market. Second, the average consumer is far from being tech-savvy enough to question how a technology product makes our digital communications less or more secure.

In my assessment, to pull out of this crisis we must think through the proper framework to incentivize and enforce better security design practices. I propose that either an existing or a new federal agency will be put in charge of governing, incentivizing, and enforcing security design standards. The federal government could expand the authority of an existing agency, like the CPSC, or create a new agency to oversee the manufacturing of technology products. A similar model could be applied within the EU, or any other nation, to solve this problem. In order to build a sustainable and scalable model, I would leverage the CPSC's accreditations and certification process. The CPSC does not do this work alone. They are primarily responsible for coordinating standards, and accrediting third party assessors or laboratories, which are, in turn, responsible for certifying that products are built to relevant standards.

In my proposal, the existing or new agency would be responsible for the following: One, coordinating security design standards, and partnering with congress to mandate relevant standards. Two, building an accreditation and certification and a warranty program, leveraging the CPSC model were applicable. Three, enforcing quality through regular testing, third-party assessors, and conducting recalls when appropriate. It is important to note that what I propose is not simply more regulation, but potentially a boom to the industry. The accreditation and certification program could be used to encourage small and medium sized businesses to meet the demand. In order to make this proposal more palatable to the business community, I would encourage policy makers to consider a mix of tax subsidies and/or breaks, as well as labor market incentives, as sweeteners for early adoption and

participation. There is also, of course, the ultimate benefit, which is that if this program would work, we could save countless Dollars on security issues and compromises due to more sound security design principles. In our increasingly fractured political system, at least in the US, this may be an opportunity for various groups to come together to establish a common sense solution.

Earlier I described our cyber security challenges as matters of life and death, and I stand by that. Unfortunately, in the US we often wait until it is too late to take bold action. However, let us not forget the lack of sufficient building design standards prior to the San Francisco earthquake of 1906; the flood wall warnings in New Orleans prior to Katrina, and the laissez-faire approach to overseeing offshore drilling prior to Deepwater Horizon. There is still time for us to address our modern crisis of quality, but we are rapidly approaching the point of no return.

The IDF Digital Revolution

MG. Nadav Padan, Head of the IDF C4I and cyber Defense Directorate, Israel Defense Force (IDF), Israel

I am going to discuss the change we have undergone in the IDF over the past couple of years, and particularly in the Computer Service Directorate and the cyber aspect. This work is the result of two years of searching for a way to reach proper military preparedness, in which we tried, first and foremost, to describe what this cyber phenomenon is. Is it something that takes place in the familiar and traditional military operation dimensions, or is it a new dimension, for which we have to prepare ourselves in a different manner? This is not merely a phrasing of empowerment, this is a saying that implies within it organizational changes. If this is indeed a new dimension, then it requires command and control and organization and a life of its own.

The place I represent, within the system in which we were organized, is the place in charge of the operational freedom of action. This is my part in the cyber defense space, and in light of this change, my directorate has changed its name from the traditional "Directorate of Communication and Teleprocessing" to the "Directorate of Teleprocessing and Cyber Protection".

When we tried to mark for ourselves what was the challenge, or where is the challenge space in which we should be operating, we realized that we actually need to create something completely new, something that was not diagnosed in a sufficiently accurate manner in the beginning of the way. Anyone who has ever walked in the field, and has experience with navigations, knows that you only realize you've reached the top of the hill after you've started climbing down from it. While climbing down from this hill, we realized the power of the change we were going through. To describe it to our own people, we have drawn a graph that depicts the axis of time through the description of the great technological revolutions, and the quality axis through the changes and the transformations. We pointed at an inversion that has occurred during the information revolution and the digital revolution. In the past, man has waited for technology; he waited for the tractor

to replace the beast plowing the field. He waited, and every time he knew how to describe what he wanted the new technology to solve for him, in the human experience known to him. We are currently living in a reality in which technology presents to man much more than he requests for, and we are occupied with an enormous challenge, the purpose of which is to bring the available technology and human assimilation capability closer to one another. The challenge lies in the dimension of time. The challenge is in the ever-increasing gap in the capability of integration and pace of change that mankind is required to exhibit, and this gap also lies in the major changes in the essence of the existence in which we live.

These changes in existence are also related to the basis of our defense objects, things that are related to the national strength of the State of Israel. Today, a large part of the lives of the young people of this country are conducted in the digital space – not as a space of communication, of a superficial description of social networks, but their actual identity. The identity of young people is, in part, on the wall. We are in the midst of this challenge, and we are also occupied by educating our people, as well as creating technological accessibility, under the idea called by the Americans KISS, or "keep it stupid and simple".

In our view, there are three layers with which we deal in the digital revolution. The first layer is that same basic and traditional layer of teleprocessing. The IP, IT, and RF layer; the layer in which teleprocessing takes place. The second layer is a processive layer, which we call Information and Network Warfare, or in Hebrew, LOMAR. This is the area where the operational processes take place. The circle of fire, the circle of command and control, and the circles of the military work. The third layer, which we call cyber, is the layer of the malevolent, warlike occurrences, the purpose of which is to disrupt the other two layers In this sense, when we describe cyber, we don't only talk about bits, we also talk about our friction with the enemy, as well as the traditional one, in the worlds of electronic warfare. This is a language that enables us to describe and focus the operational work in the spaces that are central to our directorate.

Our purpose, as I mentioned, is to create operational freedom of action, and therefore we are working on three main areas. The first area is the network-related IDF, creating a network, IP and IT infrastructure that is protected, modern, fast, and in the Cloud. The second regards the defense worlds, the creation of defense tools, and moreover, the creation of patterns, perceptions, and dynamic defense processes. The third, which is actually a derivative of our definition of this cybernetic space as a dimension, is the space that allows us to control the occurrences with our enemies in cyber-time, in real-time, and this is a dynamic and active space in which there are many occurrences with enemies.

Our challenges have expanded in orders of magnitude. If in the past we knew how to kinetically draw the places that threaten us, the weapon ranges of the enemies that surround us, and the occurrences, using geographically based first, second and third circles. Today the space has grown in orders of magnitude. Not only has it grown in orders of magnitude, the rules of the game in it have changed, and they keep changing all the time.

It is not that there are no efforts of legislative regulations and writing of guides and legislative actions that are taken in various places in the world. Those exist, but anyone who understands political sciences and the legal world knows that in the international law, the strongest norm is custom. At this point in time, no person who can influence the international custom, or any wide agreement or a superpower with wide-reaching influence, has an interest to regulate this world. Without regulation, this world is wide open to attacks, and meets us in new places every week and every day anew. Within this dimension, there are actors from different worlds, ranging from superpowers, through international actors like Anonymous, international hackers, and even local ones of our own. It seems that adolescence rituals these days have changed completely. In the past, the rite-of-passage of adolescents in our area was the voyage to Petra, but today it is penetrating the FBI areas or the areas of the Israeli government. We meet very unsophisticated actors in the Hamas and Hezbollah sectors, who by manage to bother us quite a lot using very simple tools.

In the beginning of the way we described our top secret environments, our network centers, as the defense foci. We discover, time after time, that simple defacing or DDoS attacks, things that were seemingly out of our focus, actually bother us very much. This is especially true in the regular day-to-day, where they affect the national morale, but also in times of conflict, and definitely in places where we have dependencies on the online world.

When we look inwards, into our development in the cyber world, the basic and essential change the IDF is undergoing is a change from a language or survival in the IT and information security world, to a language or war in the cybernetic world. This is a monumental change. It affects the tools, the practices, the perceptions, and the spaces in which we act. This is a change from a language of viruses and computer network exploitations (CNE) to computer network interferences (CNI) and computer network attacks (CNA). This a change from a language of protection layers to a language of processive defense, starting from the IDF's and the State of Israel's secret operation centers and isolated networks, and ending with the weapon systems in the conflict frontlines, which are the IDF's representatives in the world of IoT.

Alongside these threats, in the past fifteen years we have been undergoing a transformation from a world of local LAN and WAN networks to a world of connected networks or a world of jointness. In the past five years we have been going towards shared spaces, and in the past year or so towards a modern, Cloud-based operational Internet, in which the operational processes will be shared by all of the command branches of the army. It is an ever-developing window, and it requires us to enter technological worlds in a much more profound and independent than we did in the past. In the midst of those challenges, we are moving from being enablers, that title that we carried for years – the commander's right hand for command and control – to operational shapers and leaders.

To that end, I would like to note our three main lines of work. One is in the world of Big Data and Data Science, and our need to manage, track and monitor an ever-growing scope of information. The second deals with the actions required to protect a modern, Cloud-based, network environment, as well as our ability to change in the paces

that this digital revolution requires of us, through going out to the worlds of open source and opening up to changes that occur around us, in the civilian world; and all the way to our need to take care, in a different manner, of all those weapon systems that represent, as I previously mentioned, the Internet of Things. The third line of work is the result of the convergence between the two previous ones. When these two worlds meet, we understand that we need a new type of process management, that of a digital campaign. This entails a basic change in the military state of mind. This means that we, the old directorate of communication and teleprocessing, are no longer passive technology providers, providers of tools of defense, providers of network capabilities for other entities; we are managing an active campaign, which involves changing from that lower circle of the network worlds through the worlds of operational processes and up to the cybernetic worlds. In that, there are two places in which we understand that we have to make a paradigmatic change.

The first is the world of creating the operational picture, the monitoring. We realize that we have to change from the old monitoring paradigms, which were mainly used to monitor our IT maintenance capabilities. In the past, we used to establish sensors that told us when the loads were high, when we had to route some of the processes from one data center to another, when all of our encrypting and isolating defense tools cannot handle the loads, and so on. Today we are doing to a new kind of monitoring, one that sees the defensive reality through the eyes of an operator in the cyber world, of a cyber defender. This is a new tool and a new perception.

The second place where we have to make a paradigmatic change, which I already mentioned as part of the Big Data worlds, is our need to analyze immense volumes of information. We have switched to monitoring methods coming from the world of systems, from a holistic, processive outlook on the military networks. This is inherently different from the processive or transactional lookout that we traditionally used to monitor in the past. We are developing tools for Big Data monitoring, which allow us to look at or detect anomalies in a way wasn't possible in the past. We don't only detect anomalies in this manner, or drill-down to them during later investigations, but also

in a manner in which, using a bird's view systematic analysis, we detect and attribute patterns and phenomena. This monitoring is based on two fundamental sources. One of them is intelligence, the same traditional intelligence we have been relying on for years, even in the classic kinetic worlds. The other one is the friction world. We realize that remote electronic surveillance, in the world of defense, is not enough, and that our need to operate in a proactive manner and remain in constant friction is a real existential condition.

Underneath the monitoring layer is our infrastructure, which is based on the fundamentals of security by design. As part of the project to move the IDF's infrastructures to the south of Israel, we are redesigning the IDF teleprocessing array. This project enables us to build new data centers and infrastructures, and we are building them with built-in defense tools and a new defensive perception. This perception, or this redesign, does not only create more sophisticated or more suitable tools, it mainly creates a base or a platform that allows us to take dynamic, proactive action to defend our space.

Transformations have always existed in this dimension and in this space, through the simple digital development and through change of defense tools, which mainly dealt with the network protection perimeter. We used to create active and dynamic gates to the IDF's isolated networks. Today, however, we understand that the transformation we have to represent or to create within our design is a major one, a transformation of an architecture that interacts with the processes themselves. We don't only change the defense tools in the isolated networks, we change the architecture and the operational processes within the networks themselves. This is an incredible technological challenge. We are only at its beginning – with initial successes – but this journey is one that we are trying to take into the design-related world of the IDF's transfer to the Negev. We understand that we have to play the entire field, in the sense that we have to create a defensive depth for ourselves. We cannot protect our networks from within the network space. We have to operate outside the space of IDF networks, even in the spaces of the civilian network, and know how to be the first to create friction with our enemies and create our own buffer zones in the defense worlds.

The last part is the integration story. This is a substantial challenge. I'm not only talking about integration of tools, patch management, as well as integration of perceptions; this requires the defenders in the various branches, in the Navy, in the Ground Forces, to act in a different manner. They are not providers of the defense tools as they used to be, they produce active attack and defense forces, the purpose of which is creating defense or achieving superiority that allows us that same operational freedom of which we speak.

In conclusion, our main mission is to achieve operational freedom of action. Unlike in the past, we understand that in order to obtain this we have to manage a campaign – this means constant friction with the enemy, actual warfare. In order to sustain this freedom of action, and in order to lead this campaign, we have to obtain technological and operational superiority. This can be obtained by changing the design of our networks. It cannot be obtained without collaborations and partnerships, both with our own civilian and national sector, and with our partners in the world. Additionally, it cannot be obtained if we don't know how to locate, build, and develop the proper human resources. People have always been the key to our success, and they are the ones producing that technological, design-related and operational superiority. We have a new and powerful challenge, an actual challenge to the free world. In my opinion, only through joint action, in which we understand our need to join hands in our country and outside it, will we be able to face this immense challenge.

Defending Israel's Cyber Borders

Nadav Argaman, Director of the Shin-Bet (Israeli Security Agency)

To my understanding, the Cyberweek conference creates a very important connection between all of us. A connection with a significant potential, which is crucial that we empower. The participation of a covert organization in a conference dealing with activities that take place in the communicative and overt range, marks, more than anything else, the characteristics of the world with which we are dealing, and its difference and distinction from the traditional world of intelligence. This difference is one of our challenges. Any term that includes, in proximity, both the word "borders" and the word "cyber" should be suspected as a paradox containing an internal contradiction, since we well understand that cyber is a borderless space.

As time passes, the cyberspace grows and expands, like the universe itself. The borders between the cybernetic space and the physical space will also become completely blurred in the future. In fact, most of the layers in our lives today already interface or connect to the cybernetic dimension. The Shin-Bet is responsible, by virtue of law, for safeguarding the country and its democratic institutions from threats of espionage, acts of terror and destruction. This responsibility comes to realization in the cyber arena as well, much like in any of our other arenas of operation. Nevertheless, the necessary response when facing cyber threats requires us to use unique strategies. In the real world we don't settle for passive defense, but rather strike the terrorists in their own territories, and the same goes for the cyber arena. We study the opponent's patterns of action, and know how to strike them and surprise them, using a variety of ways and methods. Hackers all over the world, who act in order to harm Israel, experience unexpected malfunctions from time to time.

In recent times, several cases have demonstrated the potential of the tactic and strategic threat posed by the global cyber arena. Among the rest, we have witnessed harm caused to critical infrastructure as part of signaling between countries; the use of cyber as an infrastructure for

blackmail or negotiation; takeovers of computers installed in vehicles; access gained to vital information, including technological IP and sensitive security information; use of network as tools of propaganda and incitement, which is directly connected to the individualistic terrorism. These threats turn the cyberspace into an especially challenging arena. They grow on the background of characteristics that are unique to this space, and here I will mention the major three characteristics.

First of all, this is a space where recovery from damage is usually quick and even immediate. Prevention of terrorism in the physical world will usually cause damage that will make it difficult for the terrorist infrastructure to recover from it. In cyber, however, once a certain Trojan horse was exposed, neutralized, and sealed, all that is required from the attacker is to change several code lines in order to relaunch it at its victim. Even once you've hurt and neutralized a threat, recovery is immediate.

The second characteristic is the extreme asymmetry between the resources required on the attacker's part to create real damage, and the resources required by the defending party in order to prevent the attack in advance.

The third characteristic is the weakness of the supply chains, which have become wide and branched. Those chains usually contain a large number of weak links, which the opponent can exploit for gaining access and penetrating the system. Over the past year, the Shin-Bet, along with its partners in the cyber defense and intelligence communities, has been battling many challenges in this area. From threats coming from nation-state powers through threats coming from terrorist organizations, all the way to individual hackers acting on their own behalf. To thwart these threats, we have executed dozens of sophisticated and successful activities and operations. And still, given the deceptive characteristics of cyber, it is appropriate that we should remain humble and cautious about our ability to get a hermetic picture of what is happening in this space. Therefore, I would like to focus here on our strategies for defending the security of Israel and its essential interests.

In the cyber arena, we are guided by three major perceptions. The first of these perceptions is a single, borderless space. Our strategy for

defense and prevention does not distinguish between defined borders, but rather strives to act in this world and all of its dimensions. We operate, and will continue to operate, everywhere in order to thwart the threats as far away from home as we possibly can.

The second perception is bi-directionality: the connection between the physical space and the cybernetic space is bi-directional. As our opponents aspire to influence the physical space using the cybernetic one, they should also expect surprises in the physical space, of the kind that will affect their freedom of action in the cybernetic space, too. These are communicating vessels in defense, much like in preventive attack.

The third perception is a strategy of collaborations and partnerships. As early as a decade ago, the Shin-Bet had already made Combined Prevention its main mission. This prevention utilizes integration and synergy between intelligence, operations, and technology. In cyber, jointness is the name of the game, and there is no other way to go. We operate against our opponents through a cyber coalition, which includes an extraordinary cooperation with the Military Intelligence Directorate, the Computer Service Directorate of the IDF, the Mossad, the National Cyber Security Authority, the Ministry of Defense, and friendly countries worldwide. The Shin-Bet has work relationships with colleagues from other intelligence organizations worldwide, and it is willing to assist, with the knowledge and experience we have gathered, with terror threats that Western countries face today.

However, collaborations in the arenas of the security and defense are not enough. The opponent is relentlessly looking for the weakest links of the chain, and focuses on the ever-growing connection of cyber to all the layers of our lives, and on the world of IoT. Therefore, the cyber industries, the academic institutions and the civilian societies, all constitute an inseparable part of our capability to create an effective strategy of defense and prevention. As a result of the Prime Minister's resolution in this matter, we are in the midst of a national effort in which we all bear the responsibility, and it requires the effective cooperation of those concerned.

In order to adequately address these challenges, we also needed to undergo a conceptual change that led to a structural change. The

Shin-Bet is currently undergoing an organizational revolution, the core of which involves the unification of all cyber technology fields into a single branch. The result is a powerful and concentrated technological fist. The power of this fist is derived from interdisciplinary combination: the cyber fields, in all of their diversity, along with the classical technology fields, which have developed with the organization since its inception. The Shin-Bet's cyber and technology array is an ever-growing start-up. By means of high-quality intelligence received in our cyber network, many terrorist attacks have already been thwarted.

Managing to locate a single terrorist is a huge challenge. Despite this complexity, the Shin-Bet, together with its partners, has succeeded, through technological, intelligence and operational adjustments, to locate more than 2,000 potential individual terrorists, in advance, since the beginning of 2016. The groundbreaking technological advances, along with familiarity with the field and operational work, greatly contributed to reducing the level of terrorism and the successful coping of the State of Israel with the threat of individual attackers. And just as the cyber threat is great, so is the opportunity. Naturally, as an entity that is responsible for thwarting and prevention, I have expanded a lot on the threats, but the opportunities that Cyber has opened for us in the operational and intelligence arenas are dramatic.

The Shin-Bet strives to realize the potential in cyberspace on the side of prevention of threats, but also on the side of opportunities, and we are able to accomplish all this only thanks to our human capital. Today more than a quarter of the organization's employees are technologists: hackers, programmers, cyber protection experts, electronic engineers, and more. We have set ourselves the goal of recruiting and developing employees who are capable of coping with the challenges and tasks that lay ahead before us. The training and career development tracks have been significantly improved, and are now competing in a worthy and respectful manner against the private sector. We strive to employ the best hackers, because part of the cyber challenges is facing all of our opponents, with an emphasis on countries and superpowers. On the side of the opponent are the best hackers that powerful organizations and the terrorist organizations manage to obtain. The asymmetry that plays in favor of the attacker

requires that the defending and preventing side will have creative, groundbreaking people who know how to think ahead.

Those who are in the world of education and academia and in the civilian arena have a key role in nurturing our future generation, in a way that will enable Israeli superiority in cyberspace. Recently we received an answer for the ever-present question of whether the generation is diminishing. It came in the form of five youths aged 15-16, who succeeded in solving the challenging cyber puzzle that we published. They have proven to us how important education for technology and excellence is, and what potential is embodied in the future generation. Those who join our organization are able to take part in groundbreaking operations, to face technological challenges that compete with the most advanced developments in the high-tech industry, and to be a spearhead in defending Israel's security. The Shin-Bet will continue to operate and renew, together with its partners in the intelligence community and the cyber authority, in order to prevent hostile activity in every arena and in every space, to serve as the unseen shield in the field and in front of the computer screen, for the security of the State of Israel and its residents.

Protecting Cyber Borders

Dr. Douglas Maughan, Director, Cyber Security Division, Science and Technology Directorate, Department of Homeland Security, USA

I am going to discuss how we are thinking about borders in the US. The Executive Order that was published from the White House in May 2017 focused on four key areas. One of the key areas, for us in the DHS, is focusing on our federal network borders. The goal of the Executive Order is to hold our government responsible for protecting our government systems; we had a slight problem with that in the past. The Executive Order talks a lot about risk management as well, but most of all, we need to change the game, and look at new architectures, and push towards more shared services. We have small government departments that are expected to have a CIO or an IT person, even though they actually do not know how to run IT. We want to take this to a model where large pieces of the government are protected, because we are utilizing shared services, like email, Cloud, mobility, and other types of things. That, we expect, will protect our federal government borders.

The second area is the private sector. The Department of Homeland Security has a responsibility to work with the private sector, to ensure that their infrastructure is also secure. We have risk management associated with the private sector, for what we call "Section 9" organizations, which are private sector companies that have national security significance, in sectors like the energy, banking, and finance. We are also looking at some of our federal policies, and how we actually work together with the private sector. We want them to be successful and make money, but at the same time to ensure that our infrastructure is secure. One of the biggest problems we see, of course, is botnets, and there is a stronger initiative coming out of the White House to look at defense against botnets. We are also looking at electricity disruption. Our biggest concern is the electric grid going down; almost everything is reliant on electricity. Last topic in this area is our defense department, another key piece of our federal

government infrastructure. Additionally, we talk about promoting an open Internet, and we are also looking at deterrence; how do we deter the "green hoodies"? If we can make it so much harder for them, then we will all be better off.

The third area of protecting the nation is our law enforcement mission and also the international strategy. We have to do a better job of working together to secure the borders. The fourth and last area is work force. There has been some discussion about that, but in the end, everybody has a shortage today. I think we can't start soon enough with creating the next generation of cyber defenders.

That was a brief overview of the Executive Order, but I am a science and technology guy. I fund research and development, and what we are really trying to do is to be part of the next generation of technologies, in order to secure the borders. The way we work is that we take requirements from both the government and the industry, and our job is to build new technologies. I fund startups, I fund academics, I fund labs, and through that we create new technologies in the very large list of technical areas that interest us. Our goal is to take those technologies and commercialize them. Our government does not create government-specific technologies anymore. Primarily, we are buying commercial solutions, so we want to make sure that the research that we fund produces commercial products. I have been at the DHS for thirteen years, and in that time have been successful in commercializing over 50 technologies, many of them are still in the marketplace.

We have produced a research and development of critical infrastructure security plan, which has already been published, as well as an implementation road map. We want to help people understand what is important for us, in research and development for critical infrastructure. There are some foundational understandings that we need to have about critical infrastructure systems and how they need to be secured. The electric grid is different from an oil refinery, which is different from a banking and finance system. There are some foundational understandings that we need to have in the community, and I think that is also the case for the researchers. It is hard for a researcher to actually get access to critical infrastructure systems,

which is what we are trying to do. The second area is risk assessment, and how do we deal with risk. We are looking to create technologies, so we are funding a research aimed at unified situational awareness, and most important is our collaboration across critical infrastructure.

We focus our research and development on five major areas: First is interdependency analysis; how does one critical infrastructure impact another? Second is the issue of position navigation in time; GPS or other types of positioning systems that are vulnerable. All of our critical infrastructure relies upon secure time protocols, and there are some vulnerabilities there. Third is water and waste water; a key sector for us. Fourth is transportation, and fifth is energy distribution.

One question to be asked is: what do we do with critical infrastructure from an R&D perspective? The companies own and operate the infrastructure, so what is the government's job? Our goal, from an R&D perspective, is to partner with the critical infrastructure, and we are doing that in several ways. First is our logic project, with the oil and gas sector. We put money on the table, they put money on the table, they provide the research projects, and we help them execute those. We currently have a number of large companies involved in the activities – not just US companies, because it is not just a US problem. The goal is to test new ideas, put them into their operational environment, and help the transition and commercialize. The key point is that they decide on the projects, we just support them. All of these projections are public, and you can look at the project reports online. One of the most recent of those is "safety instrumented systems". In an oil refinery, the biggest concern is that the security systems and safety systems are being connected together. If I am successful at compromising a security system, and I can get into the safety system, I can cause something bad to happen.

The next way is an academic center, which is operating out of the University of Illinoi, but includes about a dozen other partners. There are also international participants; academic entities in Israel can participate in this kind of a center. If you look at the topics, we are focused not only on the technical side, but also on the business side. One theme we are researching there is business cases; what is the business case, and how do you make the argument for the sector

and for some of these companies to invest in the cyber security of their systems. We also look at supply chains, as we have seen supply chain attacks, and we don't want to see those in critical sectors. Another thing we are looking at is interdependency analysis and communication.

The third way is a project that we initiated with the finance sector. We are actively, together with the Department of Treasury, funding new technologies to bring the finance sector into a commercial consortium structure. We have a contracting authority, called an "other transaction authority" (OTA), which allows us to produce quick contracts with commercial partners. The finance sector has identified their five key areas: software, dynamic defense, network characterization, malware detection, and insider threat. Over the course of the next five years, we expect to spend approximately $10M a year in the finance sector to bring solutions to them. This is open to anyone, including small business, startups, etc.

The last way I am going to mention here is international activities, in which I am a huge believer. Cyber security is a global sport, not a US only problem, and we have been working with fourteen countries, and have a strong engagement with Israel. We are conducting bilateral R&D; our first call for solutions was released in May 2017, with the Netherlands, and we are working on having similar type of bilateral research projects with Israel as well.

What have we learned over the last decade or so? First is that you have to build and keep trust. It is always difficult for the government to build trust, as a major part of the private sector doesn't trust the government. This means, which is the second lesson, that once you build that trust, you have to keep it. Everybody has to have skin in the game – not just the government should be doing this, the industry also has to bring money, people, resources to the table. We have to work together to secure our borders. The third lesson is that money and potential monetary gains are not important, we have to work on something that is challenging; we need hard problems for us to solve. If companies can do it quite easily, they should just go do it. We want to join from the government in order to solve some of the harder problems. Fourth is deliverables. There was a long time of what I would just call partnership "talkathons", where everybody

gets together to discuss the problem, but nobody actually delivers anything. A partnership has to produce some deliverables. Last but not least, there are many models for how this can work. We have discovered that having government as the center hub does not work. The government should just have a seat at the table, and as we do that, we will be able to secure our cyber borders.

Identity as the Great Enabler

Jeremy Grant, Managing Director,
Technology Business Strategy, Venable, USA

I have spent more than 20 years in the intersection of cyber security and identity management. Four and a half years of these, I spent at NIST, leading the cyber identity team, including implementation of a major US initiative – the National Strategy for Trusted Identities in Cyber Space. Today I work as the managing director at Venable, in the technology and business strategy unit. We are a law firm with the largest cyber security and privacy practice in the US. I wish to discuss here the theme of identity as the great enabler; not only as a fundamental building block of better cyber security, but also as something that is a foundation for more trusted high-value transactions online.

Over the past year I have received many invitations to conferences in the Fintech space, something which never happened before. One of the reasons I keep getting those invitations, as an identity guy, was summed up in a PayPal report that came out in 2017, called "Fintech from the Frontlines". In a nutshell, this report says that looking across everything that we are trying to do in Fintech these days – more innovative payments, better ways to manage your money, new ways to store or exchange your money – identities are is a major corenerstone, required to enable Fintech innovations. However, identity is continuing to languish in analog forms that are difficult to build upon for the provision of digital services. This makes it difficult to execute many of the more interesting Fintech use-cases that we have been looking for.

The problem today really dates back to a cartoon that is already 24 years old, drawn by Pete Steiner in the *New Yorker*. It is a great cartoon, where a dog is on the computer and he says to his friend the dog: "On the Internet, nobody knows you are a dog," and we are still dealing with this issue today. Naturally, there has been some evolution in the kinds of information we are able to figure out, certainly if you are a certain intelligence agency, with abilities to collect certain types of information. However, for those of us who are mere mortals, we are

still dealing with something that is a pretty simple problem, at least on paper. The dog is getting old; 24 in dog years. The dog should be up in heaven now, but here we are today.

The problem, at least when we start looking at Fintech, is in all the things that we want to do online; more high-value, high-trust services, things like "know your costumer" and anti-money laundering regulation. We have an identity/security problem. How do we deliver really easy-to-use secure login experiences, so you can't get phished and you can't get breached with stolen credentials? How do we deliver these solutions in a way that respects privacy and potentially even enhances it? How do we deliverer a better costumer experience; more streamlined, more bespoke personalized applications? How do we develop these things in a way they can actually reduce transaction cost? If we are adding on layers of security and other regulations and their requirements, and it makes things more expensive, that undermines a lot of the use cases for using this technology in the first place. A large part of this all really comes down, in the end of the day, to: how do we enable more trusted transactions online? Trust, it turns out, is something that is very hard to get right. But identity, when delivered right, can enable trust. Identity is the great enabler.

Digital identity, when you deliver it right, can provide a foundation for digital transactions and online experiences that are more secure, that are easier to use, and are more respectful of privacy than what we have in most online applications today. However, there is a big "but" here. It has to be secure, because identity is where the hackers are attacking the most; and it has to be easy to use, because at the end of the day consumers are not going to put up with any security that degrades their experience. In fact, they will just click out of the app and go to a competitor. One of my favorite reports, one that comes out every year, is Verizon's annual Data Breach Investigation report. It is not just Verizon who writes these; they work with other security companies, and they work with law enforcement agencies around the world to digest what happened over the previous year in breaches. The 2017 report said that 81% of last year's breaches leveraged either stolen or weak passwords. That is a stunning number. And this is not a new thing, every year for the past six or seven years, it has been

somewhere between 60% to 80%. This problem keeps happening time and time again.

Passwords are a recipe for disaster. There have been major identity breaches in the past few years: Target back in 2013; JP Morgan chase; Apple iCloud, also known as the JLE – the Jennifer Lawrence Event – where a number of different celebrities had their password phished, compromised, and suddenly some very exposed pictures and videos were exposed over the Internet; Home Depot; Sony Pictures; Anthem, which is the biggest health insurer in the US; the Office of Personnel Management; the "Bangladesh Bank Heist", revolving around stolen SWIFT credentials; the hack of the Democratic National Committee; it is a pretty common thing. And unfortunately, the problem is only getting worse.

The anti-phishing workgroup showed that from 2015 to 2016 there was a 65% increase in phishing attacks, and a 40% increase in breaches. The head of Microsoft's identity Services said that in 2016 Microsoft were seeing attempts to compromise 20 million accounts per day. Now it has been quintupled to 100 million accounts per day. In the 1930s, they asked the bank robber Billy Sutton why did he rob banks. He said: "Well, that's where the money is." Why are attackers going after passwords? Because that is the way to get to the data. When you have an increase of five times in a single year, there is a problem to be concerned about.

Unfortunately, the problem goes beyond passwords. We have been trying to add second factors on top of passwords for years. The security has been helpful, but the user experience – not so much. In 2011, Google did a great thing for the security community and consumers; they came up with a free app called Authenticator, which you can download to your smartphone and get a one-time password that will be good for 30 seconds as a second factor. That really helped the Google ecosystem for about four years, and then Eric Sax, who was in charge of all things related to identity at Google, got up on stage at the Google Cloud Identity summit, and said that the bad guys have caught up. That one-time password might only be good for 30 seconds, but these days an attacker can successfully phish for that one-time password just after they phish your original credentials. There is a

bigger message that comes out of this, which is that all shared secrets are phishable. It does not matter if it is a static password or one that keeps changing, these days the attackers have caught up.

So far I recounted the bad news, but here are some good ones: people are not just sitting by, looking at this, throwing up their hands and staring helplessly. The market has been responding. Strong authentication is suddenly getting much easier. What used to require a stand-alone smart card token and a reader that you would somehow hook up to your machine, or a dongle that you would carry with you for a one-time password, or a stand-alone biometric reader – all of the basic security functionalities, or in many cases the devices themselves, are now being shipped, out-of-the-box, in the smartphones and laptops that we all have readily available.

Mobile devices, it turns out, are pretty interesting, because there are some common elements that exist in all of them these days. The first is a secure hardware-based isolated execution environment. The Trusted Platform Module (TPM) chips in Windows laptops, the Trusted Execution Environment in Android phones, and the secure Enclave in apple iPhones, they all go by different names but they are all basically the same thing – something that is isolated from the rest of the device, and is capable of generating, securing and applying cryptographic keys. They also come with built-in biometric sensors. It is very hard to find a device these days that does not support face, fingers or voice authentication. Many of the cameras for face recognition can also be used for iris recognition. They also have other sensors and capabilities that can feed into identity and authentication decisions. Today we have devices that come with built-in security, strong, multi-factor authentication, and much of this is being enabled by a cross-industry effort which is called the FIDO alliance.

FIDO is an industry organization with more than 250 members, mostly companies but some governments, including the UK and the US. It has a single focus: to achieve better authentication standards, based on public key cryptography that is unphishable, to solve the password problem. Today in the marketplace there are more than 340 FIDO-certified solutions, and generally speaking, FIDO is now available to protect more than 3.5 billion user accounts worldwide,

in part thanks to firms like Google and Facebook that have deployed them across their users. In order to explain why FIDO is different, I will briefly review how old authentication works. "Shared secrets" is a pretty simple concept – I know something, the server I am logging into knows something, and I pass over to the server a proof that I know the secret. Whether that secret is good for 30 seconds or forever and ever, it is still phishable. FIDO introduces a different concept into the equation, called an authenticator, that sits in the middle. In the first step, a user authenticates locally to their device, leveraging those secured protected isolated execution environments, generally with a biometric pattern. Then, behind the scenes, the authenticator leverages public key cryptography to log the user into the server itself.

FIDO is interesting in terms of technology and standards, but it is important to talk about who is driving it, because it really makes the point of what is actually happening in the market place. For starters, Google and Microsoft are generally like cats and dogs fighting, but here they are working together side by side. Major banks, such as Bank of America, ING, USAA; all three of the major payment cards networks; other players in payment like PayPal and Alibaba; Athena; NTT Docomo; hardware manufacturers like Lenovo and Samsung; chip makers like Intel, Qualcomm, ARM, Infineon and NXP, which make the components for all of our mobile devices; and many other security vendors. And that is just the board, 34 members. The other 220 or so are sponsor members, associates and liaisons.

Today it is hard to buy a phone that is not FIDO certified, in part because Qualcomm has built it into their chipset, and Google is building it into the Android OS. Apple, which usually don't join standards' organizations, have conveniently architected their devices to support FIDO through third party software, and Bank of America, NTT Docomo and eBay are all utilizing that. Microsoft's Windows 10's password-less login screen is enabled by the FIDO standards. You can use your finger or your face to login, but it is not just a biometric login, it is biometric plus the public key cryptography behind the scenes. W3C's web authentication standard includes, or is supposed to include, an easy way to embed FIDO in all major browsers. Governments take notice and approve. In the UK, the *National Cyber Security Strategy*

puts a great focus on making consumer grade products secure by design, and it specifically named FIDO as one standard to invest in. The strategy clearly states that we need to focus on things that don't rely on passwords for user authentication, but instead use the machine and other devices in the user's possession to authenticate. And finally, NIST, a part of the US government, released its new Digital Identity Guidelines, where they specifically declared FIDO as the highest authenticator assurance level recognized by the agency.

In summary, we are on the cusp of a no-password world. Nearly every device we use will soon be shipped with unphishable authentication. It is easier to use than passwords and other types of multi-factor authentication, basically leveraging biometrically enabled single gesture strong authentication, without a password, backed by both strong standards and a security certification program, to measure the security of the implementations. We are finally on the verge of delivering identity as the great enabler.

Preparing for Cyber Influences on National Elections

Yair Lapid, Founder and Chairperson, Yesh Atid, Israel

Sometime in the next year and a half there will be elections in the State of Israel. Like all other election campaigns in Israel, these elections will be emotional, aggressive, offensive and tense, but they are also going to be different. This is because there is a great chance that these elections will also entail a profound cyber involvement, which would influence the results; perhaps even determine them. A large part of the actors we already know are going to be involved: superpowers, nation states, terror organizations, ideological groups, and so on. How can I determine this with such certainty? Because this is exactly what had already happened in the elections in the US; in the elections in France; and this is what is expected to happen in the elections in Germany, too. This is the new world in which we live. An election campaign in an inherently technological democracy such as Israel, which is also located at the heart of the most explosive conflict zone in the world, is really an invitation to attack. Think about it for a moment in terms of a state struggle. What a great advantage it would be for a country who could determine who will be the next Prime Minister or the President of its rival, or even of its ally.

I was a soldier during the first Lebanon war, in 1982. Back then, Israel had sent its army into an enemy country, in severe conditions, in order to try to bring Bachir Gemayel, leader of the Christian Phalanx in Lebanon, to power. Hundreds of soldiers and members of the Israeli defense forces were killed then. It cost billions. We got stuck in Lebanon until the year 2000, and we have failed. Today, we could have executed the same operation, but much more successfully, using twenty guys wearing round glasses, sitting in an air-conditioned room. The political cyber, much like the military cyber, specializes in hiding its tracks. The superpowers impersonate as other nation states, nation states impersonate as terror organizations, and the terror organizations impersonate as other terror organizations.

Under these circumstances, one has to be a saint in order to refrain from taking an action. I know the international political space fairly well, and there are very few saints there. Besides, even if you get caught, the political risk is not that great. Totalitarian countries are much more daring than democratic countries. Assume we catch Russian hackers, or the cyber division of the Iranian Revolutionary Guard Corps, interfering with the elections in Israel. What exactly are we going to do to them? They will deny that they have ever done anything and continue doing the same thing.

In the next elections there will be cyber attacks on Israel coming from two threat axes. The first is the direct axis, which will include cyber attacks and breaches on the Central Elections Committee, on the computer systems of the various parties, on information databases. In the last elections in the US, there have been reports on breaches into election committees in no less than 39 states. We can also expect cyber attacks on critical infrastructure, in order to influence voting patterns, and to undermine the voters' feeling of security. Think, for example, what would happen if all the traffic lights in the center of Israel stop working on the election day.

The second axis is the broader one, because it will revolve around the arena of the public mind. It will try to influence the voter in a wide variety of ways. There will be bots, lurking in social media and spreading rumors; there will be fake news, of course; there will be fake photos and fake testimonies. There will be the thing called "influence operations", and negative information about candidates will suddenly appear at the last minute, right before people will be on their way to vote. In the last elections in Israel, for example, the Prime Minister sent, on the election day, a text message to about a million voters, in which he said that the Arabs were surging to the voting places in buses funded by the left wing associations. That text message, sent on the last possible day, gave him 3-4 additional seats in the Knesset. What is the problem, really, to do something similar, but from another country?

What happens, for example, if a week before the elections there will be a breach in the Prime Minister's website, and someone will plant there a speech containing racist quotes against certain ethnicities

in Israel, or a derogatory statement about bereaved families? If you find it unreasonable, let me remind you that this is exactly what the Qataris claim that was done to their Emir, Sheikh Tamim bin Hamad Al Thani. One day, a speech attacking Saudi Arabia and supporting Iran appeared on his website, and while the Qatari claimed that it was a cyber breach, the Saudis didn't really care. This was exactly what they needed to escalate a conflict that existed anyway, and as a result of this publication, Egypt, Saudi Arabia, the United Arab Emirates, Bahrain, and Yemen cut their diplomatic connections with Qatar, and the Saudis closed their continental border. All that happened because of a breach into a single website, that was not sufficiently secured. This is the new war. This is the thing with which we should be dealing.

The great potential damage, the great risk is no longer the penetrations to computers in order to obtain intelligence, nor cyber attacks that were meant to harm critical systems. Today cyber is around something completely different – it is about sovereignty. It revolves around something for the purpose of which we used to require an all-out war: the ability to frontally attack a foreign country and replace its leadership. This constitutes a dramatic change not only for democracies, but also to the world of cyber itself, because it started its way as a part of the covert war. Cyber people preferred to remain in the shadows. However, in the new world cyber operates in the most public arena there is, in the arena of politics and media. It does not rely on very high technological capabilities – these are relatively simple breaches – but rather on psychological capabilities. The thing that makes political cyber attacks so effective is the fact that people are not really interested in obtaining objective information, which will allow them to make rational decisions. There is a long series of studies that prove this. Out loud, people say that there are rational, but in fact they keep looking for information that will strengthen the opinions they brought with them from home. Cyber can shape our political space because the citizens are already used to consume information in a biased manner.

Even today, each one of us can type the same search word on Google, and we will receive completely different results. The algorithm makes sure we get the result we like. The Internet is not objective;

it is programmed to capture our hearts, to create a focused message for a specific person. In the same way in which the large advertising companies and the world syndicates have shaped commercial fashions and trends on the Internet, the cyber can shape our political space. It has the tools, it has the infrastructure, and it has the power.

Last year, a major study was published by two researchers, Roland Bénabou from Princeton and Jean Tirole from Toulouse School of Economics, who have proven that most people refer to beliefs and opinions and emotions like consumer products. Similar to products, they consume their emotions from the Internet to feel better, to improve their self-image. People wear their political stand like a pretty shirt or like new sports shoes. The fact that this is an emotional consumer product makes people reject any information that contradicts their emotions – attribute it to fake news, attribute it to conspiracy – and at the same time to adopt any information that supports their feelings, even if it is feeble and partial, and even if it is false.

Technologically speaking, a political cyber attack isn't more complicated, but rather much less complicated. It is much easier to hack into the computers of a newspaper desk or television broadcasting stations or news websites than it is to hack into the computers of the Pentagon or the Mossad. This is not about writing 15,000 lines of code, but about writing a short, mean story that is easily believable.

To complicate things a bit more, dealing with political cyber becomes even more complex due to the characteristics of the political arena. Even without attacks from the outside, the political arena already contains defamations and lies and fake news, and it lives off partial rumors anyway. It might not be very aesthetic, I wish our democracy wouldn't look like that, but this is already our own internal business. The problem is that those things make it harder for us to distinguish between an external cyber attack and the legitimate internal democratic struggle. Really, how can we know what comes from Russia or from Iran, and what is an integral part of the campaign of the Likud or the Labor parties? How can we defend the democracy without causing damage to the democracy?

This is a tough question, but that doesn't mean that we are exempt from finding an answer. On the contrary, it means that we are very late

in providing that answer. Much like in other areas of cyber, Israel has to lead the world in the battle for protecting democracy, too. We have the capabilities, we have the required level of sophistication. The State of Israel should have, long ago, established a special task force, which would include 8200, the Shin-Bet, and the Mossad – three entities that are the best of their kinds in the world. This special task force should be given resources and it should be given an instruction, to start preparing to face cyber attacks on our election campaigns. Due to the complexity and the sensitivity of the matter, this special task force should report to an objective entity, as well as receive its instructions from that entity. I hereby suggest that the person to which this team will report shall be the President of Israel. He is the only person who both belongs to the political system in Israel and is also above it.

In the next few years, there will be more and more cases of democratic countries whose elections will be determined by external cyber attacks. Israel cannot afford to be one of those countries. It has to start preparing for this, it should have started preparing for this yesterday morning.

Introducing the New Cyber Technology Unit in the Israel National Cyber Directorate

Yigal Unna, Chief Executive Director, Cyber Technologies Unit, Israel National Cyber Directorate, Prime Minister's Office

As the director of the Cyber Technology Unit, and previously the head of cyber and Sigint division of the Shin Bet and the ISA, I would like to discuss the new Cyber Technology Unit in the Israel National Cyber Directorate, who we are, and what we do. But, in order to talk about cyber technology, we must agree on what is cyber technology.

The way I see it, cyber technology is comprised of three major elements. First, the knowledge, which we achieve through our cyber research centers in all the major universities in Israel. Then, we have the physical infrastructures, such as our labs. And last but not least, perhaps the most important element is, of course, the human capital, without it nothing will work. Like in everything else that is cyber related, we need integration, the mixture of all the elements, in order to work. The Cyber Technology Unit's main goal is to strengthen Israel as a cyber technology key player. We already have a great industry, a great academia, and a human resource pool that is growing constantly and continuously, but they should all be emphasized, strengthened, and getting even better results.

We deal with a national capacity build-up strategy, we empower the technology ecosystem, we generate and develop new concepts and even practical working solutions for the state-level problems and challenges we meet. We are the main technology unit of the National Cyber Security Authority, and, of course, we are the national knowledge center of cyber technology in the government. The unit is comprised of three parts: the CTO, the R&D, and the projects, but these elements won't work unless we have an overall strategy, which we title National Empowerment.

The main activities of the unit revolve around the government backbone, technology platforms, state-level projects, state-level

activities and the national cyber robustness, but they all fall under the holistic concept of implementation; implementing all the layers required for these kinds of solutions and working very closely with our colleagues from the UK and other countries with which we share similar challenges.

We have research capacities in all the universities in Israel, and we are financing and funding researchers in large scale, and with large sums of money. We operate the Cybernet, which is the national network of information sharing, and is already active and providing results. We also operate the Meteor, the fast lane for cyber security platform of governmental procurement. This is not a technology, but it is still a crucial to the success of the all the startups, small businesses, and even larger ones, who face a lot of difficulties with bureaucratic bids and other ways of conducting business with the government. We launched a dedicated website, where any government agency facing a cyber security challenge can publish information, and using the website, any company can view the challenges and offer solutions, therefore getting a foot in the governmental door. We have a secure information sharing technique, working with one of the universities and some vendors, to assure safe sharing of sensitive information between government agencies and other agencies, using multipoint and other sophisticated computer science techniques.

We also deal with another state-level activity that we call "rating", which is a working model for cyber security measurement. One of the problems we all meet, in the government and in the private sector, is how we can measure if what we put in is worth what we actually get out of it; if the outcomes are what we were expecting. We developed a model that assists in evaluating the readiness level of any agency and any company that wishes to use it.

Lastly, we have a project titled "Enabling Trusted Transaction". We have a framework that is agnostic to technology and agnostic to telecom, dealing with the endpoints, and protecting against Man in the Middle threats and attacks. Our plan is to spread it through the right vendors, all over Israel and even internationally.

As for the human capital in Israel, we have special programs, beginning with middle school, through high school and going all

the way to the academia, aiming to meet the gap, the shortage of competent workforce in R&D and security.

Finally, I would like to present some of the main challenges we are tackling. We counter the influence cyber-attacks, those that aim to influence election results or other high-stake events. We deal with the influx in both criminal attacks and national security level attacks. We look into phishing and identity thefts, and also with big data and IoT, taking us towards a world where everything is online, and we must deal with it properly. Lastly, we deal with "recycling", the fact that national cyber weapons are leaking from governments to terror groups and from there to individuals. The accelerating cyber arms race is the common denominator of all these problems and challenges, and that is our main mission.

Security Transformation

Rohit Ghai, President, RSA, USA

I think that cyber security people, the defenders, are superheroes. They fight the good fight, and it must be noted that they fight it against overwhelming odds. Exactly how overwhelming? Here is an example of a mission that faces such odds, just for illustration and inspiration. In 2006 Nasa, after some delays, launched a mission called New Horizons, whose objective was to explore the outer edges of our planetary system – the Kuiper Belt and Pluto. A piano sized spacecraft, weighing roughly a thousand pounds, had to take a lot of scientific equipment, and travel 3.6 billion miles to accomplish that mission. Timing was of essence, because there was an opportunity to catch Pluto when it was closest to the sun, which happens once every 283 years. It was literally a once in a lifetime opportunity. For this mission to work, the stars didn't have to align, but the planets certainly did, because we were going to use Jupiter and its gravitational pull to slingshot the aircraft to the destination to get it there on time. The power budget for the aircraft's onboard power system was 245 watts, which is about the power of three light bulbs, no more. There were many setbacks along the way, one of them was heart-wrenching: nine years and a few months into the journey, at the footsteps of the planet Pluto, the spacecraft went radio silent. But through deft risk mitigation, the mission was accomplished with a margin error of exactly 72 seconds behind the originally planned schedule. How on Earth could they beat these odds? In this article I will try to analyze that situation and learn something from it, in the context of cyber security.

Before I do that, I would like to introduce RSA. RSA is the company known for throwing the big RSA Conference, or for our cryptography background, or for the tokens. The new RSA is now part of Dell technology, the world's largest privately controlled company. We are also currently working on a very simple insight, an idea that might help us balance the force, and perhaps give us defenders a shot, a fighting chance, a sliver of a shot. We call this idea Business Driven Security. I want to share this one big idea with you, as well as a few

small ones, that may be useful as you craft the cyber security strategy for your companies and organizations.

We have resorted to fear mongering and motivation driven through fear. Let us ponder the ethos of the defender: is it the manifest destiny of the defender to be defined by the attacker? Is it the defender's destiny to always be doomed to be one step behind, to spin up a firewall when the network is under siege? In the heat of the fight to thwart the attacker, have we actually lost sense of why we were fighting in the first place? Have we lost sense of our core purpose? Our mission? Our business? And I use the word business broadly. If you are part of the cyber security team for a government, your business is to serve the citizens, and to uphold the values of your country, like democracy. Throughout centuries, military commanders and thinkers have pondered the notion of what factors influence motivation and morale in high stakes combat situations, which is the kind of situations cyber defenders face every day of their lives. A Greek philosopher in 400 BC termed this "the force of the soul". It is what drives you to accomplish and tame seemingly unsurmountable odds. I ask you, what motivates you? Are you in love with the mission and the business that you are part of? Does it stir your soul? Love, soul, two words that you probably have not heard in any cyber security presentation ever. We need to stop being driven in response to what the attacker does, but being business driven in our approach to fuel the force of the soul. We have to motivate ourselves differently, we have to take a business driven approach, and that is not just a great motivator, but also a phenomenal tool to make us operationally more successful.

The business guys are thinking about changing the world, and changing the industry through digital transformation. They are imagining the fourth industrial revolution, with new applications that unleash unprecedented productivity. They are imagining quantum computing, and precision medicine, and autonomous vehicles, and crypto-currencies. Certainly, a worthwhile and inspiring mission. If you think about what that drives, this transformation is driving a fundamental re-architecture from a technology and IT perspective – IoT, Mobile, and Cloud. The perimeter has all but disappeared. It is also driving a fundamental change in the workplace. People want

access to information anywhere, anytime. Every facet of our lives is now digital, and therefore vulnerable. And finally, the business thinkers, the business teams, obsess about managing risk. Managing risk to tide through this spirit of change and of disruption.

The regulators have also decided to help, creating policies that protect the consumer, the civic and legal considerations, as this change washes over us. In the US alone, there were more than 25,000 regulations since 2008, and it takes 11% of our GDP spend to take care of these regulations. If that 11% was another country's whole GDP, it would have been the 10th largest economy in the world, just to put matters to perspective. The specter of GDPR looms large in the UK and in the EU, and boards of directors everywhere are worried about what that means, in terms of cyber security considerations.

But what are the defenders thinking about? First of all, they are just worried about keeping up with the attacker. We always talk about that asymmetrical advantage that the attacker has; they have to be right only once, and their ROI is instantaneous and continuous when the attack transpires. The defenders, on the other hand, have to prove the negative, that they actually prevented an attack, which is between hard and impossible. In addition, the fight against the attacker is against a sophisticated adversary that has the same technology as the defenders. Moreover, the attackers are great at collaborating and sharing the technology. Just keeping up is a phenomenal challenge for the defender.

When you think about it, when the odds are tremendous, what do you have to do? When you live in an industry that has the highest negative unemployment rate? When you don't have enough good guys to fight the good fight on your behalf? What do you do? Well, when you don't have enough human intelligence on your payroll, you have to recruit machines – AI and automation. Another options is to recruit from the ecosystem; you have to outsource.

Second, the defenders worry about the new architectures that are coming up. We have to re-think security for these architectures, and we should demand that the infrastructure providers design built-in security. For example, if you set up a software-defined network, that network better have a micro-segmentation capability, to create security

enclaves for the applications that get provisioned. This software-defined network also needs to have the intelligence to reconfigure policy as applications change contexts, either on premises or in the Cloud. Security needs to be instrumented in the modern new infrastructure. You have to build a home alarm in, as you construct your new home. You certainly have to meet the new expectations of the modern user, who does not want any friction. Therefore, your security posture has to either step up or step down, in response to the risk of the situation. Every posture around security needs to be informed with the notion of risk management.

And finally, it is all about enabling the business. Cyber security is squarely a business problem. It is the domain and the consideration, and a hot topic in boards of directors. At the RSA Conference we saw 400% increase in attendees that are members of boards of directors, wanting to learn more about security considerations and what it means for the business. There is motivation on the part of the business teams, as well as security operations teams, to get on the same page. This is harder than it seems, because fundamentally, they worry about different things, they focus on different things, and they use different languages. The business teams want to know when a vulnerability occurs, and when a security incident occurs. They want to know what happened, how bad it is, what it means for their business, in terms of reputation, costumers, P&L, and intellectual property. They do not care if the attack was perpetrated because of a vulnerability due to cross-side scripting or SQL injection or any other technology buzzword of your choice. These two teams, despite best intentions, find it hard to communicate, and to be on the same page.

Going back to our example from the beginning, the New Horizon mission – how were they able to tame the odds of a journey over 3.6 billion miles, which was fraught with all kinds of risks? What they did was to make hard business decisions on how to use that power budget of 245 watts to either navigate the spacecraft, do processing, do image capture, or do transmissions. Those business decisions were made based on advice from the technical teams, which were able to translate incident risk to business risk, so that the business teams could make the right decision. Then, the business teams orchestrated the

technology teams to focus on the right problems at hand, every mile along the way. It is all about robust risk management. It is about taming what at RSA we call the "gap of grief", the gap between the business teams and the security operations teams. There is magic if you tame the "gap of grief", and there is potential for disaster if you do not.

Our security industry has gone through various phases. We started with centralized computing, a small perimeter, and fewer humans. We took the motes and castles model of building protection whereby you put defenses around the perimeter, initially signature-based and, over time, rules-based and anomalous behaviors-based. Then we realized that the bad guys were already inside. The perimeter model was not cutting it, so we shifted to the paradigm of detection and response. The bad guys are already inside, let's find them, and get them before they get to the good stuff. Then we realized, as the architecture changed, as more and more users got connected, as things got connected, that it has become overwhelming. You didn't have enough good guys to detect and respond. Then, the mindset shifted to automation and AI. However, as I said before, all of these things are not adventitious to the good guys, because the bad guys have all of these advantages as well.

There is but one asymmetric advantage, that we, the good guys, can wield. The one asymmetric advantage, which can give us that one edge, is the knowledge and understanding of our business context; nobody knows our business better than we do. The timing of our business events, how our business is organized, our higher priority today, compared to yesterday. We know that, and the adversary does not. This notion of business-driven security is the idea and the simple insight of bringing to bear the one singular asymmetric advantage that we can apply to the context of securing our world. It is the confluence and the merging of security technology with robust risk management.

I would like to address the key takeaways of some the other ideas that I communicated. First, as you think about building this new, glorious digital world, with transformed IT and transformed technology, demand that security be designed in as a consideration at the right abstraction layers in the technology stack. Second, think about recruiting machines to your team, but remember that the bad guys have also been recruiting machines for more than a decade. It is not the advantage that tips the

battle over. Furthermore, using supervised machine learning, train the machines just like you train the humans, as you recruit them, because machines learn through patterns, and the adversary is great at disrupting patterns. Business-driven security is the one advantage you can bring to bear.

Finally, I implore you to not just worry about the technology of cyber security, but also about the phycology of cyber security. Think about what drives your motivation; think of security as a risk management problem. Our mission, our definition of success is not to create the un-hackable world, it is to create a safer world. It is about harm reduction. You have to redefine success, because success is a motivator. Nobody wants to work on a mission where you constantly fear that you are failing. You have to redefine success, and you need to think of security as risk domain and apply robust risk management technologies. And last but not least, remember why we fight in the first place. Fall in love with the mission, with the business of your organization, and may the force of the soul be with you.

Will We Ever Get Cyber Security Strategy Right?

Benny Czarny, Founder and CEO, OPSWAT, USA

I would like to share with you some of OPSWAT's philosophy, and hopefully help you build a better cyber security strategy. OPSWAT is a global diversified cyber security company based in San Francisco, with over 150 employees. We have five global offices, and we are in the midst of opening an R&D and BD center in Israel. We play across several Gartner categories of cyber security, and we are a diversified cyber security company. We have over 1,000 customers across the traditional cyber security verticals including government, defense, energy, and financials, and have several key technology partners, such as RSA and Check point. We have some key and strategic customers in Israel, such as the IAA, and most of the nuclear facilities in the US are protected by OPSWAT technology. We play three different roles and we have three different product lines, each one of them is extremely successful. Our major platform is protecting against advanced threats in the organization. We have a holistic framework, a multichannel protection against advanced threats, whether the threats are coming through the USB channel, file downloads, web portals, or any other channel. We are a known leader in end-point security technologies for traditional network control, and we are currently innovating in Cloud access control. As for the third platform, which is a Threat Intelligence solution, we are very open to the API-driven world, that is a part of our philosophy. Many organizations are integrating and looking to integrate platforms via APIs, so many of our technologies are accessible via APIs.

In order to start talking about our cyber security strategy philosophy, we need to look at where we stand in the industry. The first question we look at is: how much malware is out there? To try and answer this question, we look at two different sources, AV comparatives and McAfee, both very reputable vendors. Clearly, there are many malware, but there is also an inconsistency. Different vendors see different threats on any given time. When we look specifically at advanced threats, we see a significant increase of over 900% in advanced threats related to

JavaScript since 2016. Some of that is related to Ransomware, and a major part is related to targeted attacks based on email and documents. When we look at CVEs in terms of total numbers year after year, we see a giant increase there as well. It also shows us something that has become quite constant; as an industry, we are failing to come up with patches, to a point where the rate of new CVEs is greater than the number of CVEs we can close. In short, we see more and more vulnerabilities, over 6,500 new ones every year, and a big jump in advanced threats based on documents, but we don't really know how many, because there are no consistent numbers. The result of that is, obviously, an increase in actual breaches. According to the latest research, 75% of these breeches are related to malware. The biggest problem here is that the spending on cyber security jumped at approximately the same rate as the threats; we are spending much more on cyber security, but still we have more breaches, and there are more malware. So if we go back to the question of where do we stand as an industry, I can say that we are failing.

Given that information, the second question we need to ask ourselves is: what is the capability of our solutions? There are currently over 1,800 cyber security vendors worldwide, and different reports and consulting firms classify them into all kinds of different technology groups. It is very hard for CISOs to understand what they need, let alone going ahead and actually choosing a solution. One thing that we do at OPSWAT is putting together a market-share report for antimalware on BYOD Windows devices. It is a very popular report, and fairly known in this specific space. If you go to our website you can find it and other reports, all free to download, we collect data from millions of machines and share it. From the data we collected, it is clear that the market is quite segregated; there isn't a dominant vendor that controls everything, and this is the same in many other areas, not just antimalware. How can we go ahead and measure the quality of those?

There are six very known companies – Westcoast Labs, AV Comparatives, ICSA Labs, AV Test, NSS Labs, and Virus Bulletin – with a mission to rate and rank cyber security products. Each of these companies measures the quality of total protection, based on

different metrics, such as response time, product quality, operating system compatibility, how fast you can detect a new outbreak, what is the global stability of the product, and more. We closely follow the reports of these companies, and when you compare them, quarter by quarter, you can see that the quality of the cyber security rises very rapidly, especially over the last couple of quarters. In fact, it was so rapid that it got us to question how these test are actually being done. Seeing a double digit percent jump in a quarter makes us question the quality of the test, and therefore the quality of the cyber security products themselves. This becomes even more problematic when you compare the reports of the different labs to one another. We found significant inconsistencies between the results of multiple labs testing the same product, be it an antimalware, a web proxy or any other type of solution. Even the vendors themselves are contesting these solutions, and when results are not favorable to any specific vendor, they tend to sue.

Will we ever get cyber security privacy right? With all these unknowns and contradicting information, how can a CISO still build a cyber strategy that is going to work? I want to share with you some of our philosophies, things that we very much believe in, that can help you to go and build a better cyber security strategy. The number one thing that we promote is that we don't trust files. From our perspective, any file is bad; executables are bad, documents are bad, JPEGs are bad, videos are bad, AutoCADs are bad, XMLs are bad, every file is bad. The one technology that we very much believe in is called Content Disarm and Reconstruction (CDR), sometime known as "data sanitization". The idea is to assume everything is bad. When you go camping, before you drink the water, you boil them; you don't trust anything you are going to consume. In much the same way, we have technologies today to reconstruct files before we actually use them. In our company and research we have proven again and again how we can build malicious files that most antimalware vendors will miss, but with data sanitization you can clean these files before they touch your network.

The second philosophy is that we believe in multi-layered, multi-vendor defense. Because you don't know how many threats are out

there, and you don't know what is the true capability of any given solution, try to use as many as you can. Obviously try to base your decision on reports, and find the solutions that you trust the most over others, but try to use as many as you can. We looked at the top 10,000 threats in our costumers' networks, and we scanned them using over 30 cyber security solutions, including Next Generation AVs, traditional AVs, and others. We could clearly see that the more solutions you integrate, the better your chances of stopping a threat. In specific integrations, we were able to reach a 99.99% detection rate, and reduced the average time-to-detection from twelve hours to fifteen minutes.

The third philosophy I would like to share is block installation of vulnerable software. The idea is that you can block vulnerabilities of applications before you install them on a network. You don't have to wait to install them before you scan them. This way you can block many malicious binaries, rogue installers, and rogue firmware that should never be installed.

The fourth philosophy is integration by design. When we say "integrating the technologies together", it is not only about installing them at the same time, but making them work as a unit. It is an art. You need to know what is the capacity and throughput of each one of your engines, and then go and integrate them together.

I sincerely hope that the information given here, and the main points of our philosophy, will help you in building a robust, future-proof, cyber security strategy.

Uncover the Unknown

Mark Gazit, CEO, ThetaRay, Israel

I would like to share with you why we, at ThetaRay, believe that AI is the next stage of cyber security protection. If you think about it, today cyber security protection, protecting your operational data, protecting your organization, is almost like finding a needle in a haystack. We have a huge amount of hackers that produce huge amounts of data. Normal operational behavior produces huge amounts of data as well. How do you find this one thread? But I guess before that, we just need to realize that the world of threats has changed. Today it is not only about viruses and malware, and defacing websites and stealing information, today bad guys are after the crown jewels of the organizations. If you are a financial institution, hackers will use their means to steal real money; if you are an aviation company, hackers will try to cause damage; if you are an operational company, maybe they would like to shut down your facilities. Also, the border between the inside and the outside of the organization has changed. We have people that bring their mobile devices into the organization, and actually bring the world into your protected domain. If you use the Cloud, then your organization is in fact outside of your organization, so perimeter-based solutions don't work anymore.

We all understand this, but we still see cyber security means that let people to steal money. The world of cyber crime has changed. Today nobody would really try break into a bank to steal money from the vault, like it happened in *Ocean's 11*. Today it is much easier to steal $81M through SWIFT, and allegedly sneak the money to North Korea. Today they use automatic machines to steal just $1 from a bank account, label the transaction something like "iTtunes purchase", do it for 20 million accounts, and nobody notices. This is especially true if they steal money from people that have more than about 40 transactions a month. They run their machines for a month, steal $20M, and disassemble their operation. No risk, and a large reward. To stop these type of attacks, we understand that we need to analyze operational data, but it isn't as easy as it sounds.

We work with organizations like General Electric, which are also investors in our company, and they deal with airplane engines. One engine of a 787 airplane produces 20TB of data in one single flight. That is 400 times the size of the Library of Congress. Human beings cannot analyze this data anymore, and if you look at almost any type of organization today, we collect more and data, and the amounts are growing exponentially.

Traditionally, security was carried out using rule-based solutions, but they don't work anymore. Rules are written by people, and when we build rules to find some suspicious activity, we can only look for things we know about, or at best, think about. Suddenly, when there is much more data coming in, we need to create more and more rules. Cyber people will create rules for the firewalls, and then financial people will create rules to deal with money laundering, operational people will might create rules that try to catch failures in engines, but it is very difficult to combine them all together, and when you have so many rules, you also create many false alarms. Eventually, you still miss those attacks that you do not expect, and the real danger is coming not from the things we know, but from the things that we don't know, and we don't even know that we don't know.

We can use AI to solve these problems, and our solution is just one example. Machines digest all of the data, and they don't take or give any meaning to the data. Human knowledge and domain expertise is just one input, and we believe that machines can analyze much more. Machines can look at all the data in real-time, simultaneously, with very high speed and very high precision, identify issues, and notify us about them. The term AI sounds very sophisticated, and the truth is that it is. Our company was born in Tel Aviv University by Prof. Amir Averbuch. In order to explain it, I will give a real example about a financial institution, how they use operational data, actual business data, to identify cyber security attacks. The client in question is a commercial bank that gives loans to consumers. Originally, everything was good, and margins were great. Then, the bank took the loan business and made it online, so a client doesn't need to come to the branch of the bank anymore, they can use their mobile phone to take a loan. Simple. Suddenly, the bank noticed that

that the entire business stopped working; suddenly they started to lose money in this very profitable business. They tried to understand what was going on, brought the best financial people, and could not find any problem because everything looked good. And yet, they were still losing money. This bank installed an AI system, in this case our system, which digests historical data. It took the system two hours to digest two months of data, and within seconds the system identified a group of fraudulent transactions, with a total value of more than $10M.

The system analyzed and found a relationship between several parameters. Each one of them, by itself, was okay, so human beings could not identify them, but the relationship between them was suspicious. The first parameter was the age of the consumer. All the fraudulent transactions were made by consumers between the ages 16 and 19. For the bank investigators, this is not an interesting parameter, because anything above 16 may be young, but legal. They disregarded it. For a computer, though, the number by itself has no meaning, and it is taken into consideration along with the rest. The second parameter was transaction size, and in this case all the loans were for less than the average amount. From the bank's perspective, this is somewhere between not interesting and low risk. The third parameter was transaction type, meaning what type of loan did people take. In this case, all fraudulent transactions were of a product called "small mortgage", something you would take to buy a car or a small apartment. This is also a very low risk type of transaction. However, if you look at the relationship between the three of these parameters, there is a problem. The bank would have never given mortgages to minors; that is illegal. As soon as the bank people saw these transactions, they immediately knew that they were either a mistake or fraudulent.

In the physical world, governed by people, this could never have happened. They had a physical customer, talking with a physical clerk, showing a physical ID, and signing physical papers. In the virtual world, there was a cyber security attack. The people in the bank are not stupid; they developed a system that checked the age of the costumer. They even had a connection to a government database, getting information about the costumers and comparing the numbers. However, hackers discovered that if they disabled the connection

to the government server, then there was a bug in the software that approved transactions that shouldn't have been approved. Those hackers used real kids, found real assets, took loans under the name of those kids, and then didn't repay the loan. When the bank came to seize the asset, the parents obviously got angry that the bank approved the loans in the first place, and the bank's legal or public relations people preferred to cut a deal and take a loss. This is just one example of finding security issues by looking at operational data. It has nothing to do with viruses, network data, etc. Machines are able to identify cyber security attacks with a very high precision. Moreover, this was something that was truly an unknown unknown; the people in the bank couldn't even imagine this thing happening, they made mistakes starting even from the loan system's user interface.

Machine learning, from our experience, finds between three to six times more operational issues than existing solutions. However, this is not the most important part. I believe that the most important part is the level of false alarms. Today, the level of false alarms is extremely high. In the case of the Target attack, where more than 40 million credit cards were stolen, they had 25,000 alerts on a daily basis. The day of the breach was no different, except that out of the 25,000 alerts, there were three low-level alerts that showed something wrong was going on. Nobody paid attention to them. Machines don't suffer from this problem. When we talk about a very deep level of machine learning, every transaction, every amount of data will train the machine better and better. This is very different from the world that we are used to, where the more sensors you have, the higher the level of false alarms you get. In machine learning, it is exactly the opposite; it breaks this paradigm. The more data there is, the more findings there are, the lower the level of false alarms. In our case, we usually show that we can reduce the number by 100 times, which is huge. Of course, because machines use a very deep level of machine learning, you don't need to update them, they update themselves, and the deployment time is very fast, so it is very convenient.

The benefit for enterprise defenders is that it allows them to increase the immunity of systems. Many people compare biological viruses to network viruses. Real viruses, which are very dangerous to us, cannot

be dealt with using antibiotics. We actually now know that sometimes it can even make them stronger, and we know that the best ways to deal with future viruses is to increase the immune systems. The same, we believe, works for large organizations. When you use these types of solutions, not only do you increase your strength and your immunity, but you also reduce operational cost because of the very low level of false alarms. You can also mitigate damages, because just as your body's immune system can identify a virus and kill it very quickly, machines can catch problems extremely quickly before they become more serious problems.

I mentioned earlier the huge problem of finding a needle in a haystack, but if you ask our customers, they will tell you that the problem is more difficult; it is not finding a needle in a haystack, it is actually looking for a needle in a needlestack, because all the problems look the same. We also believe that this world of Big Data and cyber security and machine learning doesn't only create infinite threats, but also infinite opportunities, and that is certainly one of the reasons why our company has won so many awards and received so much appreciation from our customers.

To summarize, we believe that there is a real need, which is becoming more and more important, as cyber security can affect real life more and more, maybe even win elections. We believe that in order to deal with the unknown unknowns, very innovative way to attack systems, you need very innovative technologies. We believe that, in a way, the world is governed by machines, and hackers do use machines. You need machines to help human beings to make the world a safer place, and to do it in a way that will allow us, as human beings, to understand machine input. I am a strong believer that machines will not replace human beings, but as history shows, machines help human beings become safer. I am often asked what is the future, with hackers advancing so fast, and I believe that this movie has a happy ending, because maybe the bad guys are faster, smarter, and don't have restrictions, but there are many more good guys. I believe that eventually, with the help of machines, we will win.

Enterprise Protection: Create a Strong Cyber Strategy

David Grout, Pre-Sales Director, Southern Europe, FireEye, USA

I am not American, and I am not English. I am just European, and I think that we, in Europe, are seeing things a bit differently. The maturity level is slightly different, if you are looking at European companies, and even the regulations are somewhat different. What I want to convey here is my European vision on where we need to go, and how we need to build up the strategy to be strong on cyber. I want to show what is the plan every organization needs to put in place, or what is the journey they need to take to build up a strong cyber security strategy.

I will start with a description of the landscape; if you want to build up a strategy, you need to understand what you are looking for, and that is clearly innovation. When you look at the topics discussed in major security conferences such as Cyber Week, you can see innovation everywhere; IoT, drones, connected cars, all those kinds of innovations are clearly the way we need to go if we want to improve our business, our profitability, and our margin. This is all about business. Security needs to be here to enable the business, and to do this you need to take innovations into account.

The second main item in the landscape is the complexification of malware, and the rest of the hacking tools. We need to keep in mind that malware is only used for gaining a foothold, during an APT attack for example. Once the foothold is achieved, everything is made by using standard tools that are already deployed in your own organization. If you are only thinking about the malware, and how to tackle them, you are only thinking about ten to twelve percent of what you need to focus on. The reality is that the attackers, in general, will use Windows services and tools, Linux, different rights on Active Directories you already have in your network, etc. If you think only of the malware, you will lose.

Looking at the cyber security strategy we have been putting in place for the last decade, it seems that we missed something. It is not easy for marketing people to hear that we need to rebuild the global strategy of all the world, but the reality is that ten years ago we didn't talk about the same things. Back then we dealt with worms, e.g., Conficker and Sasser, and the approach was completely different. We thought about something named "reactive security strategy". The idea was that we only needed to collect logs, build some correlation rules, and wait for alerts. Today, however, we need to think about another way to approach security. We need to be more proactive, and change the way we are tackling the security. This will not happen in a day. We need to continue to manage what we did in the past, and clean up the noise coming from all the solutions we already deployed. But at the same time, we need to build up a new process and a new strategy. We need to think about the future, how we want to be protected for the future and what are the next steps. My vision to this topic is something I call the "intelligence-led security approach".

The first step is to make sure you know what you are looking for. You need to be sure that you have a plan, because everything starts with a plan. If you don't have any plan, if you don't know what you are looking for, you cannot do it. In order to build a good cyber strategy plan, you need to put in place metrics. To say that you blocked one of the thousand malware attacks you get throughout the entire day is not a metric; nobody knows if it is good or bad. If you come to me with only this information, there is nothing I can do with it. I don't know your business, I don't know how you are exposed, and I don't know what you operation looks like. You need to be sure that you build a plan where you define good metrics: what are you looking for, how you want to be protected, and what is the level of protection you want to achieve. You need to assess, design and redesign in cycles, the plan you are putting in place, because you will introduce new usages, new applications, new tools, and maybe even new organizations within your organization. You need to be sure that you understand how each of these will impact your plan.

When building a cyber security strategy plan, you need to take three pillars into consideration. The first one is intelligence-led

security. Up until now, most companies, regardless of whether they are startups, institutions or governments, are working with a one-size-fits-all approach. When companies realize that everyone is buying something new, like a new AV, firewall, or machine learning, then they decide to buy it as well, because they want to be protected. This needs to change, and you need start thinking about the attackers, not about yourselves. The army has been working like this for decades. You need to understand who is behind the attacks, what they are looking for, and what are their techniques, tools, and procedures, or TTP for short. Based on that, you will be able to start building up your plans. You will be able to start building up the threat modeling, threat profiling, and attackers profiling, and to define what you want to do, and what you want to be protected against. This will help you look in your network, and understand if the attackers are already there. Intelligence should be the center of your new strategy.

The second pillar is gaining visibility into your network, in the end-points and all the components, including applications. This is something that we messed up. Everybody is going to the Cloud, as well as several other new options, but at the end of the day, if you don't get visibility, you are dying. If you are looking at the APTs, the advanced attacks, most of them are using WMI and PowerShell. If you are unable to log information coming from those components, you are blind.

The third and final part is responsiveness. You need to build up capabilities to collect information with exactly one idea in the back of your mind – speed. Time is the only thing you cannot buy. You can buy technologies, you can buy people, you can outsource, but you cannot buy time. When you are building your plan, you need to think about which tools, processes and types of people you need to put in place, in order to be able to reduce the time between when the attackers are trying to enter your network, and when you are kicking them out. Within FireEye we name it the "dual time", which is where we calculate the time between when the attackers are able to enter a network, and when a consumer or an enterprise is able to understand that they are there.

The good news is that in general, the time is decreasing. Five years ago it was more than 400 days, at present it is about 99 days. The only issue is that when we are running any red-team exercises, it takes us between seven minutes and three days before our consultant becomes a domain admin, exfiltrate information, or is ready to pass a SWIFT order for $1M. If you need 99 days to understand that we are there, it is too late. I am not trying to scare you, regarding the time aspects. My only point is that when you are building your plan, you need to put in place a metric, making sure that you are reducing the dual time. There are many ways to achieve this, be it tools, technologies, people or, of course, processes, which will be linked up to the intelligence.

When we are working with governments, we are talking a lot about the "playbook". When the President wants to talk to you, he has already read the documentation, he doesn't need you for that. All the playbooks already exist. There are many people who are working in the back, thinking about all the scenarios and all the possibilities. In cyber, you need to behave the same way. You need to build up the plan on the intelligence, understanding the TTPs, to be able to build up the playbook on how to react and how to be fast when you are reacting. This is something we are not doing right now. What we are doing today in the cyber world, building some correlation rules, taking the alerts and trying to manage these alerts. You need to change your way of thinking about this. You need to know what you are looking for, what is your playbook, and how you will tackle this situation.

In term of takeaways, you need to think about the fact that intelligence is key. Being equipped with the right tools is not about how much we protect ourselves, and it is not about the percentage of the malware detection. It is about how the tools will fit with your processes and people, to give you the right metrics in your plan. This is a drastic change from where we are today. I am not saying that you need to buy a product with a 99.9999% of something, but rather that you need to understand what you want to do, which kind of processes and metrics you are going to put in place for protection, and you need to define the tools that will fit with this – not the opposite. The bottom line is: build a plan, assess it, try it, and repeat until you get it right.

Disrupting Adversaries by Changing the Game

Michael Daniel, President, Cyber Threat Alliance;
Former Special Assistant to President Obama, USA

I want to discuss security at a level that is somewhat higher than usual, about the challenges in cyberspace. What actually are the challenges that we face and why is this problem so difficult? And how can we actually play the game differently? From my perspective of both being at the White House for four and a half years, and for now being in the private sector with Cyber Threat Alliance, it is like we are in Las Vegas and we are playing against the house. We all know what happens; the house always wins. That is not a game I want to play anymore; I want to play a different game, and I will return to this later.

As we all know, the strategic context is not good. The threat will continue to get worse. The Internet of Things is growing. The numbers of actors in cyber space is increasing. All of those things are making it worse, and states and criminal groups are going to continue to expand their activities in cyberspace. These are the very obvious reasons. There is, however, a third reason, that we don't give as much credence to, and that is because the analogies and models that we use to talk about cyberspace all come from the physical world, and they are all wrong. The physics and math of cyberspace, as those of you who work on the technical side know, are completely different. A Lightspeed Nodal Network operates very differently from the rules of the physical world, and we have not fully internalized it into how we are using the mental models with which we approach cyberspace. When I was in the White House, I would often get asked if we can't just treat cyberspace like we do border security, and have the federal government provide border security for the United States. I would answer that it may be possible, but to do it we would need to add NSA boxes in every American network, and build a great firewall surrounding all of the connections either outgoing of or incoming to the United States. Usually this was the end of the conversation.

The thing is, we don't know what we really need, because the border security model is wrong. The borders and boundaries in cyberspace

do exist, but they don't follow the international boundaries that we have. They exist where networks touch each other, at peering points, firewalls, routers. It is as if everyone in the entire US, for example, lived on the Rio Grande. In that kind of world, how would you assign border security to one particular group?

Another analogy we can use is missile defense. In four and a half years in the White House, I never once had anyone come running to me and say: "if you don't do something in the next thirty minutes, the bad malware is going to hit. The flaming ball of cyber death is coming right at us." Never happened. I either got: "we discovered that we were had, eighteen months ago, and have been being had ever since," or "we have intelligence that somebody somewhere might do something in a few months." The idea that we are going to be able to stop malware as it is coming like a missile is just not going to work. The physics and math of cyberspace mean that we need to come up with some new mental models, new constructs for how we are thinking about cyberspace. If we don't, we will continue to get the problem wrong.

It is not all bad. The fact is that all the malicious actors face some constraints. First of all, Hollywood does not equal real life. All of us have seen the shows where they flip open a laptop, and with three keystrokes they hack into the whatever. The truth is that doing something to somebody, somewhere, randomly on the internet, is easy. Having the specific impact that you want, at the time and place of your choosing, and only the impact that you want, is still really hard. It requires true effort on the part of the bad guy. Second, we actually don't really know what collateral damage means in cyberspace. From a broad perspective, we don't know when we run something what the ultimate effect is going to be. From the nation-state perspective, this puts some constraints on them, but even for the bad guy perspective, they have to be careful about what they do if they don't want to get discovered. All of the bad guys also face capacity limits. There are only so many people. Yes, a large part of it is automated, but they face the capacity constraints that there are only so many smart people to actually do the development of new malware. Finally, there is also what I call "restricted action sequence", by which I mean that if you have a

goal of stealing money or stealing information to turn it into money, there are certain things you have to do in order to achieve that goal.

Nation-states face some additional constraints, one of which is the "intelligence dilemma". Most nation-states rely on their intelligence apparatuses in order to actually carry out their cyber operations. This means that they always face a dilemma: taking an action in the cyber world might expose sources and methods, and compromise their ability to collect intelligence or do other things. There is also the "third country conundrum". If you are going to take action against another nation-state or its cyber operations, you are often talking about doing it in a neutral third country. This is because any given country's cyber operations are spread around the planet, and if you trying to interrupt them or disrupt them, you are carrying out activities on a box or machine that is in some third country. This, of course, poses some interesting policies and diplomatic concerns.

And lastly, in order to actually be really effective at using cyber as a tool of statecraft, you actually need to integrate it with all your other tools of statecraft. In order to actually make it effective, you have to carry out that integration process, and that is not easy. Most countries even find that very difficult to do, and that lack of integration often hinders the ability of nations to be truly effective.

Given all of the above, how do we actually change the game? How do we not play the same old game? The first thing is that we need to change how competition occurs in the cyber security industry. For a very long time, the cyber security industry has competed on the basis of "I know something you don't know," or what I call horded information. But the truth is that nobody's set of information is ever going to be big enough. Instead, what we want to do is encourage a change in the competitive landscape. We definitely want our cyber security vendors to compete, but we want them to compete on the basis of "I do better things with the data than you do. I'm faster. I understand your business model better. I integrate with your stack better," and so on. We want them to compete on that basis, on a broader set of *shared* information. If we do that at speed and at scale, that will enable us to get inside the bad guy's decision loop, so we can actually begin getting ahead of what the bad guys are doing.

The second thing we need is to think in terms of undermining the bad guys' business model, and especially for the criminal groups. We need to stop focusing on their malware or command and control nodes, taking them out in one-off operations. We need to be going after their entire business process. We want to undermine their business model, so that we are focused on disrupting them for the maximum effect, for the maximum amount of time possible. In essence, we want to make them have to do a lot of business process reengineering, over and over again. We want to raise the level of costs on them, to the point where the low-end guys are driven out of the market, and we can impose significant higher friction costs on the more sophisticated adversaries.

Lastly, these days many people say that governments need to get more aggressive about imposing costs and disrupting what the adversaries are doing. While I agree with that, we need to be able to coordinate the actions between government and the private sector in a much more effective manner. I am not talking about authorizing hack-back, letters of marque and reprisal, because I think that is a terrible idea. I do, however, think that we can do a better job of using the comparative advantage that both governments and the private sector have, to enable a coordinated response to cyber threats. I think that the allocation of responsibility between governments and the private sector is one of the key public policy questions that, particularly in the west, we need to answer. I think that we need to look at some different models. We could, perhaps, treat it like disaster response, in the sense that it starts locally; if a hurricane is coming, you are still responsible for ensuring that you've done the right things to protect your house. However, if the storm is bad enough, you are going to get progressively higher and higher levels of government intervention. We need to find similar settings. The division is never going to be sharp and precise, it is always going to be a sliding continuum, but it will enable us to do a much better job, overtime, of disrupting those adversaries for a longer term basis.

To summarize, I would say that we all know that the situation is getting worse, but we have been trying a lot of the same things for the last fifteen or twenty years. We are now at the point where we really do need to play a different game. We need to try to change the way

we compete in the industry, undermining and disrupting the criminals, and coordinating disruptive actions between the government and the private sector. These will help us go after the problem in a completely different way.

When Things Start Killing People, Governments Get Involved

Bruce Schneier, Fellow, Harvard Kennedy School of Government, USA

I want to talk about the IoT, which I see it as something that is going to change our industry. The Internet of Things is the Internet of the physical world. We are creating an Internet that affects the world in a direct and physical manner. This means both the sensors that collect data about us and our environment, and the actuators that do things. This is coming at all levels. No longer do we have a refrigerator, we have a computer that keeps things cold; a micro-wave oven is a computer which makes things hot; and we have a computer that makes phone calls. We now know that an airplane is in fact a several thousand big computer network that flies. This happens through all aspects of our lives. For us, this means that Internet security becomes everything security. Everything we know about how the Internet works, how security works, starts affecting everything, everywhere. There is, however, one very important difference; the threats are greater. Cyber physical systems have real world consequences. It is the difference between a spreadsheet and a heart monitor, or an automobile. This difference is going to affect us greatly.

Normally, we think of security in the CIA triad: confidentiality, integrity, availability. Most of the time, we deal with confidentiality; how do we protect our data? How do we keep it from the wrong hands? How do you make sure that it is safe? When you start dealing with the Internet of the physical world, the integrity and availability threats become much greater. Obviously, I don't want someone to eavesdrop to my conversations in my car, but I really don't want them to change the database of where the roads are. The WannaCry attack was an attack against availability; the data was locked up. The Dyn attack of 2016 was an attack against availability; websites were taken down. When you think about the threats we are facing, that is where they are: attacks against elections, attacks against integrity; there is

someone manipulating the data. These things are going to become true and very real, very physical. There is a world of difference between "your spreadsheet crashes and you lose your data" and "your car crashes and you lose your life." It might be the same CPU, the same operating system, the same software, the same vulnerability, yet very different consequences. To me, this means that the government will get involved, because that is what happens when people start dying.

There are five lessons from the world of IT security, which I think we need to promulgate to the world everywhere. First, most software is poorly written and insecure. We know this; software tends to be lousy, because we don't want to pay for quality software. The joke of "good, fast, cheap – pick any two" is largely true, and in software, we have picked fast and cheap over quality. Poor software is full of bugs, and some bugs are security vulnerabilities. Therefore, poor software is vulnerable software, simply because the market doesn't want to pay for better software.

Second, the extensibility of computerized systems means that they can be used against us in ways conventional systems can't. This is important. The functionality of a computerized device cannot be constrained. No matter how much I try, I cannot make a bottle of water do anything but be a bottle of water, because that is all it is. A phone, however, can do whatever I want it to do, because it is actually a computer. I saw a news report that someone took an Internet-connected thermostat, and made it play the computer game Doom. You might laugh, but you are not surprised, because it is a computer. A water bottle could never play Doom. It means that a phone is really hard to secure, because it is continuously evolving; every month it gets a new app, it does something different, there are new updates. It is a different creature. Moreover, some of the upgrades on these devices are carried out by other people without my knowledge. Effectively, malware is just a software upgrade, just one I didn't consent to, and does things against my interest. It is going to be the same with everything, like my refrigerator or my car.

Third, the complexity of these systems results in new insecurities. Complexity is the worst enemy of security for many reasons, but the natural complexity of computer systems makes them much harder

to secure, gives them a larger attack surface, and makes them more vulnerable. This is why attack is easier than defense, this is why testing is hard. Many things fail due to complexity. This also leads to lesson number four, there are new vulnerabilities in the interconnections of things. When you have complex systems, and when you compose them together, you can get new vulnerabilities. You can take two systems that are secure, put them together, and you get insecurities. You get an insecurity in this system, and cascading insecurity in that system, affecting something else. That is the core of the Dyn attack, exploiting vulnerabilities in Internet routers, Internet recorders, DVRs and webcams, which were used to commit DDoS attacks against a name server, which dropped a bunch of websites. There was a cascade of vulnerabilities leading to a catastrophic result.

Five, computers are vulnerable in a different ways than traditional systems. Because of software, because of natural empowerment of software, vulnerabilities look different. Attacks can scale to a degree that was not possible otherwise. A bank robbery of stealing a single Dollar from 20 million different accounts – that is a scale attack that computers can do, but burglars can't. It is the same with the notion of "class breaks". The automobile manufacturers know how cars fail. Cars have parts, and parts have average times between failures. You calculate them, you do the math, and you know how often a car will fail. Computers fail differently. All the cars can work perfectly, until one day none of them do. That is the notion of a "class break", that one morning you find out that every instance of Windows 7 is vulnerable, or every instance of this iPhone is vulnerable. This is why we have patching mechanisms that are different from those of automobiles, because the threats are different. This becomes even more critical as systems become more critical.

Additionally, I am concerned about hackers in a different way than I do about "regular" attackers. I live in Minneapolis, I have a home there, and there are burglars there. There is a kind of a bell curve of burglar quality. I am concerned about the average burglar in Minneapolis. If there is a really great burglar in Tel Aviv, I don't care, he is not getting on an airplane, he is not going to fly to my house, and he is not going to brake in. The Internet doesn't work that way. I

am concerned about the most skilled attacker. The hacker who wrote the Dyn botnet released the code into the wild. For a time, we saw 3-4 new variance a day. The first one had skill, and the rest use software.

These are all differences that we know are moving in the Internet of Things, and are moving into the real world. So far, we have largely been okay with leaving computer security to the market, because it didn't matter that much. It worked, more or less. We had many problems, many externalities, and many interdependencies, but we have muddled along, because the effects were not that great. That is what's changing. The economics of the IoT are also different in some very important respects. Our computers and our phones and all of these things are secure as they are, for a few very important reasons. The first is that there are teams of engineers at Apple and Microsoft and Google, who are designing them correctly, or as correctly as they can the first time. Those same teams of engineers are on call when there is a vulnerability, to rapidly write and disseminate a security patch. That is what we are used to, that is how these things are secured. That doesn't work for low-cost embedded systems like DVRs and home routers. They are built on a much lower profit margin, they are often built by off-shore third parties, and they just don't have security teams associated with them.

For many of the devices vulnerable to the Dyn attack, there was nobody around to write the patch, and even worse, a lot of the devices have no way to be patched. Right now, if you want to patch your DVR, the way to do it is to throw it away and buy a new one. That is just not going to work. Moreover, we also get our security from the fact that devices like our phones have a very short life span. We replace our phone every 2-3 years, and our computer every 3-4 years. IoT doesn't work that way. If you imagine someone writing software for an automobile today, it is going to be in a car in 2020, driven for ten years, and then put on a boat and sent down to Nairobi, where it is going to be bought and driven for another twenty years. We, as industry, don't know how to maintain 35 years old software. We can't do that. If you pull a 35 year old computer off a shelf and plug it in, you are not going to be able to upgrade it. It is orphaned, it is gone, and that is the way the IoT works. My DVR is going to last

5-10 years, my refrigerator is going to last 25 years. Two years ago, I bought a smart thermostat which is connected to the Internet, and I expect to replace that – never.

We have very interesting market failures. Neither the buyer nor the seller cares here. They both want cheap devices that work, there is not a great demand in the market for security, and much of it is an externality. My prediction is that as the Internet gets more visceral, more deadly, more critical, we are going to see government involvement at the level that we now see for things like aircraft avionics, or health care or food safety. When things start killing people, governments get involved.

Our Changing Threat Landscape

David G. DeWalt, Former Chairman & CEO,
FireEye and McAfee; Chairman, Claroty, USA

I am going to discuss four big threats, and then I'm going to discuss four big trends and opportunities in cyberspace. I have been pretty fortunate in the last two decades to be a part of a lot of interesting companies. I have been a part of McAfee, RSA, Mandiant, FireEye, and more. I was lucky to drive more than $25B of market value for my companies, and hire more than 10,000 people into this environment. This gives you a front line seat for what is going on, and before I go into the future, I want to spend a moment learning from the past. After responding to thousands and thousands of incidents, and being on the front line between a lot of nation state activity and criminal activity, you get a little bird's eye view on what is happening.

I think that whenever mankind discovers a new domain, we conflict over that domain. This has been true during the discovery of lands, discovery of oceans, discovery of air, space, and now cyberspace. Over the centuries, we have always had major conflicts whenever a new domain has been discovered, and cyberspace is really no different than what we have been seeing for centuries on this planet. It is interesting to watch, but there are some fundamental differences in cyberspace. The biggest difference is that we don't know who our enemies are. It is very difficult to understand attribution. Anonymity on the Internet has changed everything. Obfuscations of the attackers has changed the way we do things. When we had battles and land and armies and armadas in the oceans and air wars, we always knew who our enemy was. We have a very different situation in cyberspace.

Another major difference is the asymmetrical theater. We live in a very different world, where the offense and the defense are very different from one another. In my twenty years in cyber-security, it has been interesting to watch how powerful the offense is, how weak the defense has been, and how the trends and technologies get exploited. Add to that all the innovation that is occurring, and you get a perfect storm in cyberspace. We saw an escalation of attacks and danger levels

over the last two decades, from hacktivism to crime to espionage to terrorism. We used to have quantity-based viruses, like "I love you", "Melissa", and "code red", hitting millions and millions of computers simultaneously, but that gave way to many high quality and highly targeted attacks. I bore witness to a lot of those with FireEye and Mandiant. We ended up responding to more than 5,500 American breaches, related to the Chinese military. This was an IP war between two superpowers, which ended being a pretty big epiphany for me as I watched what was occurring on there. It used to be about stealing Intellectual Property, gaining advantages and innovations, but we moved to the information warfare that we are now witnessing between two other superpowers. We have very interesting chemistry here, a petri dish of dangers and scenarios that are unfolding.

The biggest threats we are seeing are oldies but goodies. The one thing we see in every attack is the exploitation of human trust. If you go back to the 5,000 and more APT campaigns that we responded to, almost every single one of them started with spearphishing, stealing credentials, logging in and taking over, and establishing footholds. There was a whole series of commands and controls that could be put in place, but they all started with human trust exploitation. There are many thoughts and trends in trying to fix that, and about authentication intelligence. I invested in a company called Callsign, which is trying to create analytical behavioral intelligence, to authenticate one peer to another peer. If we can stop valid credentials being used maliciously, we can, maybe, slow down spearphishing and credential type attacks, which are still the largest ones we see. Other threats in the top list include cyber extortion, meaning Ransomware, and attacks on ICS systems and IoT types of attacks, where offense and defense have a massively wide gap. These are the things nation states worry about, and these are the things vendors respond to. However, the biggest nemesis of security vendors is privacy. We have a pendulum of privacy vs. security. As the pendulum swings more towards privacy, security becomes more difficult. As it swings more towards security, privacy becomes more difficult. As that pendulum goes back and forth, we have many trends, many threats, and many opportunities in this marketplace.

I see four big trends in the cyber security market today: Industrial IoT; Critical Infrastructure; Social Networks; and Satellite Networks. The rate and speed of the Internet of Things is amazing. Many people have a watch connecting to a satellite, or at least to a mobile device, and we see more and more of these consumer electronics connecting, which are becoming pervasive as part of our critical infrastructure. We need to ask ourselves, what kind of defensive tools do we have in place? What kind of controls of safety and standards do we have for consumer electronics that are plugging in to our critical infrastructure? The answer is zero. Think about what is at your home, or what is at your office or at your corporation. Where are the weakest links in your architecture? Is it that thermostat you just plugged in, that is connected to your network? Is it that HVAC piece of equipment in your building? Is it that Bose speaker that is connected up to a satellite, pushing marketing intelligence? It could be any one of those. More and more of these connected environments has created tremendous vulnerabilities, and someone is going to help solve that problem.

This problem is even bigger when you think about how it relates to the industrial markets. I'm working with a company called Clarity, which is putting behavioral analysis into the operational technology protocols. If you look at the 2,200 vendors in cybersecurity, most of them are focused on IP protocols. What about the OT protocols? How many people know what Modbus is? DNP3? There are hundreds of protocols that are a part of our critical infrastructure, which run with no authentication, no encryption, and they are not air-gapped from our IT Network. The lack of segmentation between networks offers the attackers a major advantage, and of course the operational technology is mission critical to all of our industries, like our water supply, energy supply or transportation. I sit on the Delta Airlines board, where I chair safety and security, so I know that there are 20,000 SCADA sensors on every single aircraft. That said, there are only 14 data gateways, including air-to-ground and air-to-satellite. This transformation into connected systems have created so many vulnerabilities. It is very interesting to watch what we are going to do, from a security trend point of view and a threat point of view, as we industrialize our entire infrastructure, in the pace and rate of which it is moving.

Social domains are the most new domains being hit. The entire notion of Social Networks is incredible, and not many people saw that one coming. I sat in a Delta board meeting where we saw a viral video go out. We had a hundred million views before we could even finish the conversation of what to do in a crisis. But social media is a very difficult environment right now. It is used to recruit terrorists, it is used as part of the kill chain, it is being used to falsify information, and more. The offense has a tremendous advantage in this domain, and we don't really have serious security tools, cutting across all the social networks, monitoring the integrity of the social systems. This domain is presenting an unbelievable problem in the information warfare that we are about to see, or are already seeing. This is particularly true in our national election processes, and if you start to understand the flow of information, there are many more attack vectors coming because of this. The biggest victims here are the social media companies, suffering an extensive amount of breaches over the last 18 months. Thinking about what is coming downstream, I look back at 2008, when I responded to a breach related to Operation Aurora, orchestrated by China. We discovered 153 breaches, and discovered they were after source code and bug databases. They were trying to do vulnerability research, trying to discover zero-days, so they could perpetrate more attacks downstream. We are seeing exactly the same things happening now in the social media system.

Last but not least, an amazing phenomenon. Not many of us pay enough attention to what is happening in the sky – the satellites. We have created a mesh. Our new Internet is now being presented in geostationary orbiting satellites. Most of the transponders are connecting up. Most of the devices are going from the ground to the sky, as opposed to a fiber link on the ocean floor. There are so many problems and issues that need to be dealt with when thinking about ground-to-satellite communications. All of our current problems, such as data exfiltration, and command and control servers, are enhanced when discussing ground-to-satellite communications. We have to think about air defense systems, of ways to conduct and stop DDoS attacks on a satellite, of an entirely new set of protocols, of proprietary authentication and identity controls, and many more. This

is also, in my opinion, where governments need to collaborate with the private sector. We need to avoid all the mistakes we made in the past with broadband and fiber, and integrate satellite approaches in our nationwide defenses.

When you think about the offense and defense across these four vectors, you see what is coming in cyber space, and that is what I have been paying attention to. The talent shortage is a major issue. It causes a great demand on almost every security operations team in every Fortune big company. We don't have a great talent pool, so we are putting a real strain on being able to manage the alerts, being able to understand the fidelity of those alerts, and the context of the alerts. This has led to many artificial intelligence tools that came into the market, tools for Big Data and machine learning, to ensure the fidelity of the alerts, so we can really lift the strain. The transition to the Cloud is also a great opportunity for improvement. I often talk about the Maginot Line, which the French built between World War I and World War II. We learned from that almost the same thing we are learning in today's cyber security: the defense-in-depth architecture at the perimeter is easy to evade. As we move to the Cloud, we need to improve that security, and there are many improvements we can do. Last but not least, it is fun to watch artificial intelligence integrate with everything that we are a part of, solving some of the problems we are seeing in the domains we discussed.

I hope this helps you a little bit, and hope that you understand, at least from my perspective, what I'm seeing in the marketplace.

Bits and Bytes Meet Flesh and Blood

Joshua Corman, Founder, I am the Cavalry; Director, Cyber Statecraft Initiative for the Atlantic Council, USA

I have grown deeply concerned with the relationship between technology and the human condition. About four years ago, I grew very worried about the fact that if you look at cyber security, our failure rate is about a hundred percent. Almost all of the Fortune 100 companies have lost Intellectual Property and trade secrets over the past several years. Nearly every single one of us has had a compromised credit card. We are spending about $80B per year on cyber security. We essentially gave up on the idea that we can really prevent things, and we just rely on detection and response. It is one thing when it is a credit card, but what about the less replaceable items, like public safety, human lives, trust in key markets, failures that affect GDP or national security?

Four years ago, I tried to ring the bell and call out that if the cavalry isn't coming, if our trust and dependence in technologies is growing faster than our ability to secure it, in areas affecting public safety and human life, where bits and bytes meet flesh and blood, if no one is going to come and save us, then we have to help ourselves. At Defcon, the largest hacker conference in the world, we essentially said that if the cavalry isn't coming, it falls to all of us to be a voice of reason, a translator, an educator and an ambassador, to work with the public sector, the public and safety critical industries to try and drive much more attention on high consequence failures.

What is a high consequence failure? It is not even a breach. It is not a loss of intellectual property. No one has gone out of business due to a high consequence failure. The Ashley Madison breach did, indeed, lead to some suicides, and we had the OPM and Pentagon attacks and nation states attacks, but nothing yet has sufficiently triggered a corrective action, or the will to do something different. The industry is doing exactly what we incentivize them to do. The problem with safety-critical IoT is that these aren't simply markets, they are market-enabling markets. This isn't even a metaphor anymore. Software is

now foundational to modern society. We don't sit in auditoriums in perpetual fear that the building would collapse on us and kill us. No, we depend upon it, because it is dependable. We are becoming as dependent on software as we are on steel and concrete, and it is infinitely more vulnerable.

We keep saying that we expect things to stay as bad as they are until we have a catalyzing event, which is going to trigger the political will to do something different. The problem is that this is just like the Cuyahoga river in Ohio. We actually had a burning river. How do you put out a burning river? It turned out that having a river on fire because of pollution wasn't enough to trigger a corrective action. This was the 22nd time that river caught on fire across a 70 year period. It is not as if we are going to have one big failure that will tell us we should get better, we are actually having a series of failures. Here are a few, to remind us all what we saw in 2016, which seems to be the year when the tide started turning.

While we have known, for many years, that one could possibly hack critical infrastructures and power grids, we now see the Russians actually do it with BlackEnergy. Similarly, we know our water facilities, oil and gas, our critical infrastructure, can be exposed on things like Shodan, where you can easily find hardcoded passwords and internet connected devices. Someone with no hacking skills whatsoever can now manipulate things, and we are seeing it done, at least in a few public cases. But the one that really got me, of all the safety-critical industries, was to see how Hollywood Presbyterian hospital in southern California had a single flaw in a single device in February of 2016, without a deliberate adversary, and this shut down patient care for one week. They had to cancel surgeries and transfer patients to other facilities in cases where seconds mattered. Moreover, it is one thing if it is just one hospital, but it is entirely different when it is an entire system of hospitals.

Many years ago, I researched the rise of chaotic actors like Anonymous, and people were not very impressed with them. They weren't very talented, there weren't that many hackers, and perhaps people thought it was a fad. What I saw in it, however, wasn't what they looked like when they were born, but the beginning of a trend.

We call that research "building a better Anonymous". We researched that for about a year and a half, because what we saw in Anonymous was a blueprint, which would be adopted and perfected by others. At the time we didn't have extremist groups like the ones we were depicting, but a concern of ours was that people with a clear or more extreme ideology would perfect the innovation of high-intent, low-capability adversaries, imposing their will on others.

There is a lot of talk out there about nation states, but I'm actually a lot more concerned about extreme groups with very little talent, but with a great willpower to use it. It is impossible to look at the Cyber Caliphate and say that they have not picked up the blueprint of Anonymous. This is not merely a metaphor of the very few hackers in the global movement. One of the hackers from Team Poison in the UK was arrested and went to jail for hacking Tony Blair's website. At some point after leaving prison he radicalized, left the UK, moved to Syria and was a founder of the Cyber Caliphate, before making his way to number three on the ISIS list, ending with drone missile strike which took his life a few years ago. What we had in TriCk, also known as Junaid Hussain, is someone with the means, motive and opportunity to take life in cyberspace. He was not a very talented hacker, but talented enough to attack Windows XP in a hospital. When February 2016 came around, I said to myself "thank god TriCk is dead", because if an accident can take out a hospital, what could someone do on purpose?

We, in the international policy community, are so focused on norms and treaties and deterrents and offense versus defense between rational nations states. We have almost completely ignored the intersection of high-intent, low-capability adversaries with the willpower to use it. When these attacks come, it won't be a quick fix, because there isn't the infrastructure, the staff or the experience to deal with them quickly. This will be more like the BP oil spill in the gulf, where it is just gushing on the news for days and weeks.

I spent the last year on a congressional task force for healthcare cyber security. One of the things we asked ourselves there is if there are any technical barriers to a sustained denial of patient care, in any or all US hospitals, or globally, and the answer was no. As a founding

premise, we had to work much more aggressively at how to make these systems more resilient. We tend to focus on nation states because they are very talented, and our defense is so low that even if you double or triple it, we are still not going to prevent ourselves from being attacked. However, we ignore the high-intent, low-capabilities adversaries. When we don't take enough care for our cyber hygiene for SMBs on the Internet, for patching a Windows vulnerability that had been around for months, we expose ourselves to a whole group of accidents and adversaries, in entirely avoidable ways. Any one of these could trigger the crisis of confidence in the public, causing them not to trust connected medicine. While we tend to use norms and treaties and deterrents and economic sanctions for nation state adversaries, we could also focus on raising hygiene; to make more defensible IoT; to incentivize and encourage basic patching. While we will never be able to stop a determined advanced persistent threat, we can't typically even stop a "script-kiddie" or someone using free attack tools.

We have had wave after wave of tsunamis of technical security incidents. The Mirai botnet, for example, was one of the largest attacks, using about one terabyte of information per second, which is a fraction of its true attack capacity. If we ignore this long tail of low-cost, low-hygiene devices, we will get to a point of unstoppable attacks. You can't look at IT or IT security as a market; it is a market-enabling market. For the longest time we said that if we regulate IoT we might stifle innovation or hurt the economy, but what we are finding globally is policy makers realizing that if we fail to regulate IoT, we may stifle innovation or hurt the economy. We are in some very uncomfortable areas. Also, it is one thing when it is a Ransomware that is locking up your data, but it is another thing entirely when it is doing a denial of service attack, with a destructive malware, which is just a different payload that could destroy the device. In fact, many people wanted to respond to Mirai by "bricking" or destroying these $110 cameras, which is okay when it is a disposable device, but it is another case when the Internet of Things device is a actually a multi-million Dollar MRI machine in a hospital, which became a part of this botnet.

The three distinguished characteristics of Mirai-affected devices were fixed passwords, Internet connection and unpatchable devices. I just described almost every single medical device. The next Mirai botnet will not be composed of disposable IoT, but of life-critical IoT hooked up to patients. In fact, one of our worst fears was just realized recently. In May 2017, WannaCry, without deliberately targeting healthcare, took out 65 hospitals in a single day in the UK. That was 20% of their entire capacity. Losing one hospital means diverting ambulances up the street or a few blocks away. Losing every hospital means degraded and delayed patient care, and people will die and have died. Luckily, we had spent a year working on these really hard problems, and with great candor we put out our healthcare cyber security task force. While we are a US-based task force, we have a global supply chain and a global attack surface for most of these things.

I want to punctuate five of the uncomfortable truths we uncovered. Essentially, we see the situation as quite dire. Number one, in the US, about 85% of our hospitals lack a single qualified security person on staff. This doesn't mean security products or security monitoring – there is no one there to even apply a patch or take a warning from law enforcement. Number two, hospitals tend to be maintain really old devices. Windows XP is often the best-case scenario. These are obsolete technologies, which are unsupported and much harder to defend. Number three, the push towards electronic health records caused us to have unsegmented networks, with devices that were never designed to be connected to anything but are now connected to everything. It is often very common to see devices naked on the Internet and exposed to accidents and to adversaries like WannaCry. Number four, as I mentioned earlier, a single flaw in a single device was enough to take out patient care for a week at Hollywood Presbyterian, and neutralize 20% of the capacity in the UK. And the really bad news, number five, is that the average medical device has over a thousand known vulnerabilities.

The combination of these things means we have much work to do, and many of these fixes will take a decade or more if we chose to do so today. However, with the next Mirai right around the corner,

we don't have time for that. With the next WannaCry knocking at our doorsteps, we are out of time. Normally, we would cry about this, saying that there is no money, there are no resources, we cannot possibly respond to this; but if we can't afford to protect it, why do we think we can afford to connect it? I'm not arguing to go back to the dark ages, but a crisis in confidence will cause Chief Medical Officers and physicians to retreat from otherwise superior and life-saving technologies, and that is the real harm. It is not the number of deaths, it is the number of potential lives we could have saved, if we had more confidence and trust.

We recently conducted the first-ever clinical simulations of hacking, with real physicians in a real ER. We hacked three different devices, with unwitting suspects, using real actors and real dummies, which had organs and blood and responded the same way that humans do. All three patients died, and they died gloriously. One died three times before they finally gave up. While these three simulations look great on TV, the really disturbing things came up while we did tabletop simulations on the second day. We discovered that if you take out a single hospital, the region can more or less cope with it, but if you take out two, the whole system falls apart very quickly. The big takeaway from this is that physicians are not trained for these type of events. They can, and will adapt, once they understand they have to be skeptical of some of the technologies they depend upon. We are working on this in the government, both in the US and abroad, but we are only starting now, and this process will take five to ten years. Instead of looking at the fact that we are in a bad shape, the question becomes how much political will do we have right now, to deal with the next wave and the wave that comes after that. While we got lucky with WannaCry, that luck will run out. The bell has been rung. Adversaries have seen this. There will be other waves, and we need to start acting now.

My closing thought is that through our overdependence on undependable things, we have created the conditions such that the actions of any single outlier can have a profound impact on our economic and national security, but mostly for bits and bytes on flesh

and blood. We had fires before, how many fires will it take before we have the will to do something?

Developing Cyber Defense as a Means to a Better World

Rudolph W. Giuliani, Chair, Greenberg Traurig Cybersecurity, Privacy and Crisis Management Practice, Former Mayor of New York City, USA

The whole cyber world developed in a very strange way, a way in which cyber defense is far behind cyber offense. The development of our ability to gather, categorize, analyze, and disseminate information is so powerful, and there is such a tremendous enthusiasm for doing that. It is such an unbelievably profitable area for companies and private individuals, that enthusiasm has often led us to develop very complex information systems and not think about defending them. This is because defending them is more boring, and it is not profitable. Yet. It actually is becoming profitable for bad reasons, as people are being invaded and they are losing money and they are being sued and losing reputation. But at least at the beginning of the cyber world, which I remember way back in the 1990s, the whole emphasis was on gathering information, using information, disseminating information, and turning it into services like Uber and thousands of other applications.

However, the whole idea of defending the information you were getting is not something really emerged until about 2002, meaning the offense had a lot of time to advance. It was hard, back in the early 2000s, to get major companies to invest the money that was necessary in defense, because it was seen as a basic loss to the bottom line, just a large expense. If you spend $5M on a new IT program and it made you twenty percent more efficient, you can write that down on the bottom line, and it turns itself into profit. But if you had to spend $1M on cyber defense, that was a million Dollars right down at the bottom line and it became part of your loss. It wasn't until the tremendous number of hacks that took place in the later part of the first decade of the 21th century, that people became focused on cyber defense and now we are catching up.

Israel is one of the leaders in helping us catch up on cyber defense, because you do a lot of the experimental work and a lot of the innovative work that is needed to accomplish something that doesn't have a hundred percent solution, which is very frustrating. I analogize cyber security, or cyber defense which is what I prefer to call it, to trying to cure cancer. We have been trying to cure cancer for a rather long time. There are actually some forms of cancer that we can now cure. In most cases, however, what we can do about cancer is arrest it, slow it down, and when we can cure it, we can cure it because we catch it early. If we have developed a test that can show you have cancer in an early stage, you have a much better chance of curing it than in a later stage. The same thing is true with cyber. There is no perfect defense for a cyber system. The perfect defense for a cyber system is to put no information in a cyber world. I do security for companies and I do security for nuclear power plants, and there is one nuclear power plant that is perfectly secure from cyber attacks, only one that I do work for. The reason is that it so old, it has no cyber. It is completely manual, and therefore it is immune to cyber attacks. The minute information gets out there, somebody can get it. If a client tells me they have encapsulated the information, I can find the person who will figure out how to un-encapsulate it.

Between 2003 and 2006, I was involved in an attack and penetration business with Ernst & Young. The way we sold ourselves to clients was by telling them: "let us hack you for one month. We will sign an NDA, we will give you all the information back, but if we don't get in, don't hire us. If we do get in, hire us." We won forty eight out of fifty. Today we probably would have won fifty out of fifty. The reality is that the only way you can defend yourself is with a combination of disciplines, and that becomes expensive. The way to think about it is that each one of the disciplines is reducing the percentage of risk. The head of one of the major companies that does perimeter security, outside security, told me that they have a forty eight percent success rate in stopping penetrations. That means fifty two percent of the time he fails, which is pretty bad. It is not their fault, and there aren't many other perimeters security companies that do much better than that. I can think about a few that manage to prevent about sixty or sixty

five percent of the penetrations. What that tells us is that you need a backup to perimeter security. You have to have some form of internal security of which there are a thousand different varieties.

The job of these solutions is to pick up the intruder as quickly as possible. If they can pick up the intruder within a day or two or three, in most cases, the intruder can do very little damage. Forexample, Yahoo lost hundreds of millions of identities, because the intruder was in Yahoo for two years, maybe longer than that. The United States Office of Personnel Management was invaded, and lost two hundred thousand identities of people who had undergone the FBI background checks, which obviously contain very sensitive information. That intruder was inside the OPM for nine months. Had that person been there for two, three or four days, they may have gotten a few identities, a few social security numbers, but they probably would have gotten no files. Actually getting to the files took a long time, because the files were more protected than just the identities.

As is often the case, you have gradations of protection, depending on how important or how sensitive the information is. This is similar to airplanes, where the cockpit is protected with a lot more cyber security than what the airline puts in order to protect the in-plane Wi-Fi network. They separate the two systems, so that all those people who are using the Wi-Fi aren't allowed to have access to the cockpit and take the airplane down. There are different levels and different gradations and many different ways in which you can protect the inside of your infrastructure, but the most important is to have a system that is good at finding the intrusion quickly. The quicker you find it, the less damaging it is. In fact, if you find it quick enough, there is no damage at all. The longer the person is in the system, it starts to become devastating, and you find out that information on 30 million identities were taken from you.

In addition to that, if you really want to follow most of the suggestions and most of the regulations, which eventually will become laws, you probably have to attack yourself from the outside on a regular basis. You need to employ an independent company that attacks you, and even though you have put perimeter security and internal security, the people that you hired to attack you are probably still going to find

ways to get through. You will take their findings, and you will fix them. However, even though you are paying a great deal of money for outside and inside security, you still can't get to a hundred percent security. By now, if you do these three things, and do them right, you are closer to about eighty percent. The fourth, which is the one that is ignored the most, is identity protection. If a hacker wants to target me, and be successful in doing it, meaning getting a great deal of information about my law firm or my security farm, they need to target my secretary. Don't hack me, I have all kinds of protection. I have so much protection, I can't even use my own system. But do I protect my secretary as much? No. When companies implement protection, they tend to forget about the people in the copy room and the secretaries and the people on the building department. They don't protect their identities, preventing them from getting hacked and penetrated, and therefore give up all their information. They have a great deal of the information about a particular company or a law firm or a government agency. They are the people that are most often ignored, and according to what I see, that is where the attacks happen. They happen at the lower levels, because the upper level is very well protected.

Finally, if you've done all that, you also have to have an investigatory company. You need someone who can take all this information coming from all of these sources; your external security, your internal security, the independent penetration testers you hired, and your identity protection systems, and monitor at it. The job of that investigatory company is, first of all, to find the person, persons or group that is attacking you. Hopefully, they can catch them and stop them, or at least start categorizing them into the types of companies, countries, agencies, or the people that would like to come after you. Once you do that, you can narrow your cyber defense, and put your defense in the right place, rather than the wrong place. Each of those disciplines is important and necessary. It's not a matter of how many different solutions you buy, eventually they all fall under these disciplines.

If you implement each one of those disciplines correctly, you probably have gotten yourself close to ninety percent security. You are still going to get penetrated, and you are still going to get victimized,

because we haven't come up with a perfect solution. That said, we are much better than where we were ten years ago, when we were not doing any of this. We are also going to be better off ten years from now, as we develop better and better technologies for defending ourselves, without interfering with our ability to use the Internet, and to use cyber for all the wonderful, great and terrific things it does, including lifesaving things. Defense is a necessary part of the development of the Internet.

Let me give you one example, that maybe that makes my point. One of the things that doctors in the healthcare industry in America would like, and I think would be terrific, is if every American had a little card with their entire medical history. I just went for a physical two weeks ago, and my doctor sent me all the results. He would have loved to put those on a little card, showing everything about me, so that I could carry it around with me. If, god forbid, I have a heart attack and get to a hospital, the doctor there could take that card, put it in the computer and see all the information about me. If I am unconscious, I can't tell them what I'm allergic to, what illnesses I have or don't have, what medicines I take, but it will be right there.

Most doctors agree that such a card would probably reduce the number of deaths due to accidents, heart attacks, and things like that by ten to twenty percent. In other words, it will make them ten to twenty percent more accurate in how they care for you. The reason most people are reluctant to adopt cards like that, is that they are afraid that other people can steal that information if it is on the Internet. That is why defense is so critical. If people felt perfectly comfortable and confident that their information couldn't be taken, they would want a card like that. The only reason people would hesitate having a card like that, is because most people want their medical history to be personal. In the US, it is such an important thing that we have a law called HIPAA, that makes the disclosure of medical information into a crime. We treat it very seriously. There are numerous applications like that, and that is why when people are developing better methods cyber defense, they enable more advances that can make us safer, wealthier, happier, and better entertained.

An Overview on the UN's 2017 GGE Sessions

Karsten Geier, Head of the Cyber Policy Coordination Staff, Federal Foreign Office, Germany

In June 2017, the GGE, the United Nations' Group of Governmental Experts on communication technology in the context of international security, including 25 experts in the field – of which I am the chairman – did not reach a consensus on a report. That said, we, they still had some very important discussions and agreed on a number of elements worth retaining, and I would like to discuss them here.

First of all, I will provide a short review of who we are: the GGE's mandate was to study existing and potential threats in the sphere of information security, examine possible cooperative measures to address them, and look at how international laws apply to the use of information and communication technology by states. We also looked at norms, rules and principles of responsible state behavior, confidence-building measures, and capacity building. Following our convention, we were supposed to present a report to the general assembly.

Despite some very substantial difficulties, the members of the GGE managed to identify a large number of valuable points, related to most elements of our mandate. Looking at existing and potential threats, the experts shared an understanding that ICTs (Information Communication Technologies) continue to provide immense opportunities. At the same time, these technologies, their development and rapid adoption open new risks, including the malicious use of ICTs by non-state actors and actors acting as proxies. These trends are continuing. The experts expressed growing concern that malicious activity can significantly disrupt or impair the general functionality of globalized ICT systems. Another new element, a new threat identified, was actions to undermine the security of data. The experts also voiced concerns about the use of ICTs for terrorists and criminal purposes.

Looking at capacity building, the experts commented on capacity building in the context of international peace and security, and emphasized the importance of engaging in such capacity building activities, and contributing to open, more secure, stable, accessible, and

peaceful ICT environment. They have also offered concrete guidance to this end, identifying the number of concrete measures that states can undertake, including raising awareness, promoting educational and professional training activities, information sharing, and so on.

An important suggestion was that ICT security capacity building efforts should be integrated into national development policies. Experts also made very good progress on confidence-building. They offered concrete measures for implementations at the national, regional, and international levels. At the national level, one of the offers was to include raising awareness among single decision makers, identification of appropriate points of context. The experts also suggested standard incident severity schemas and encouraged related activities, including exercises. At the international and regional levels, the ideas that have been developed included templates and procedures to facilitate information requests between states about serious ICT incidents, activities on preventing, responding and managing serious ICT incidents, dialogue with all relevant actors, and an exchange of best practices.

Equally, GGE discussions saw a number of very helpful observations on binding norms, rules, and principles. The experts have offered recommendations to support implementation to the voluntary, non-binding norms of responsible state behavior, that were represented in the 2015 report of the group of government experts. These concrete recommendations include establishing national structures, policies, processes and coordinating mechanisms, necessary to facilitate careful considerations of serious cyber security ICT incidents, and to determine appropriate responses, as well as ICT assessment or severity templates, to evaluate and asses ICT incidents.

The experts suggested a number of other points, such as procedures for official notification protocols, from one state to another, and provided clarifications on how to respond to incidents affecting critical infrastructures, warning that states considered the potentially harmful effects of their ICT activities on the general functionality of global ICT systems. They offered valuable recommendations concerning the proliferation of malicious ICT tools and techniques, including on preventing non-state actors from conducting malicious ICT activities.

The only parts of the GGE's mandate where consensus and progress were really missing, concerns certain aspects of how international law applies to the use of ICT by states. Another aspect that was a point of contention, even though it was not a part of the formal mandates of the GGE, was the discussion on how to take this process further, in the United Nations. Experts shared the understandings, and reaffirmed the findings of the 2013 and 2015 reports on how international law applies to the use of ICTs by states. In an intense exchange of views, a consideration was made that as a response to an internationally wrongful act involving the use of ICTs, states may consider using counter-measures or even using their inherent right to self-defense. The experts felt that on these issues and on the issue of how to apply international humanitarian law, further studies are required, which is why we ultimately did not agree on a consensus report.

The question now is, how to proceed from here? At the end of the deliberations, the overwhelming majority of experts in the GGE confirmed that they could work on the basis of a draft report put forward by me, in my hat as the GGE chairman, and that this text seemed ready for consensus with minor changes. The deliberations of the group have not been formally concluded, and this opens a variety of avenues, which I, as the chairman of the group, am still exploring, including with the help and the request of the UN's High Representative for Disarmament.

As for what we are going to do next, it is a bit of the issue of squaring the circle. On one hand, I believe that we are faced with a global issue, because ICTs and the Internet are global, so you need security in the context of a global system. On the other hand, it is very difficult to build a consensus involving 193 countries. I think we need bilateral work, as well as regional and like-minded initiatives, but also some sort of a global process. I think it is possible to do, and that it is not a fundamentally difficult issue, but it does require a little bit of fantasy. I will be very clear on this: I do not believe that you can simply set a global process like the UN aside and say, "we will go with selective partners and allies only." We are very happy to work in regional and like-minded context, but it just cannot replace a global debate.

As for norm violations or rule violations and their consequences, there should be a clear understanding for all parties involves: A. what the rules are, that should not be violated? And B. what are the consequences? This leads us into the issue of deterrence. The problem with deterrence by retaliation, which is a traditional, narrow definition of deterrence, is that it doesn't necessarily work well in cyber. The attacker always hopes that they can hide behind the attribution problem. That is why we have come to a wider concept of deterrence, where you also include issues, and not only of deterrence by denial, but also of entanglement and even of taboo. But for this kind of deterrence to work, you really need to have a clear understanding with the other side of what the rules are, and you cannot reach that by only speaking with like-minded partners. You really need that global discourse, with both friends and foes.

When I am asked which is more important, punishing bad actors in cyberspace or keeping this sanctity of the Internet as an open, free, and global space, I say that it depends a little on how you define the bad actor. Who is the bad actor? Is it a state? Is it proxies? Or maybe it is individuals? If you want to really decouple an entire state from the Internet or slow down Internet services for an entire state, that is, of course, a fairly big tool to use. I thought a much more nimble response was actually what happened after and in response to the DNC hack. The indicators of compromise were made public, which hurt the actual bad actors, but didn't affect the integrity of the system. This is just one of the topics that require in-depth conversations and discussions, perhaps for future GGE sessions. When I think of the future and of what we can do next, and looking at the GGE process, I keep recalling the phrase "it ain't over until the fat lady sings", and she has not sung yet.

In conclusion, I would like to make two remarks. The first is on attribution. Legally, I agree, every country has the right to do its own attribution, but politically, you also have to be able to convince others, and we have a problem there. The second is on tools and responses. The European Union recently took an important decision: the European Council agreed to continue working on what they call a "cyber toolbox", which outlines a variety of possible responses to

malicious cyber actions. I think it is quite useful, a set of instruments that other might also find instructive, and I would encourage everybody to have a look at that.

Winged Ninja Cyber-Monkeys: Harm Reduction as a National Cyber Security Strategy

Dr. Ian Levy, Technical Director,
National Cyber Security Centre, UK

When you look around in life, every window you see is square, apart from windows on airplanes. The reason for that is something that happened in 1950, when three "De Havilland Comet-1" airplanes exploded in mid-air, which is considered a bad thing for an airplane to do. Eventually, they found some of the wreckage in Italy, and found that there were diagonal stress fractures on the corners of the windows. They figured out that the windows undergo too much stress and they pop out. What do you do in this situation? You now have a system with a fielded vulnerability. The aircraft industry took a "De Havilland Comet-1", put it in a big tank, filled it with water, and drained it 1,500 times to simulate the cycle count, and closely observed that indeed – cracks appeared. What do you do now that you have proven the vulnerability? In the aircraft industry they decided to change; they stopped building aircrafts with square windows, which was a good start. They also added extra monitoring in place, put buttresses on the diagonals, and in general made the system safer by reducing the potential damage/harm.

What would you do in cyber, if it was the same thing? You would give it a really cool code name, like "Certain Death", and run around panicking, saying that everybody is going to die the moment they get on a plane. We have a language problem. The aircraft industry is doing what they refer to as Harm Reduction. According to the FIA Harm Reduction data – or as I call it, "the chance of dying when you get on a plane" – if you got on a plane in 1970 you were relatively likely to be in an accident, and you were relatively unlikely to survive. By 2005, thanks to a Harm Reduction strategy, if you got on a plane you were relatively unlikely to be in an accident, and if you were, you were relatively likely to survive. That is harm reduction, not vulnerability reduction, and I think that is what we should be doing in cyber. We

should be working about reducing the harm of attacks, not panicking about vulnerabilities. The first public discovery of a buffer overflow, to the best of my knowledge, is from 1972. Heart Bleed, which came out 42 years later, in April 2014, is the same type of software defect. If we can't fix buffer overflows in 42 years, I think we should stop trying and do something else; that strategy is not working.

Everybody know two things about hackers: they wear hoodies and they are surrounded by strange, green java code. I think this is the real problem in cyber; we call things Advanced Persistent Threats. Also, cyber is the only part of public policy where there is no independent data. It is the only part of public policy where the public view of what is going on is driven by a massively incentivized group of people – the security industry people. And it is incentivized to make it sound like these attacks are perpetrated by winged ninja cyber-monkeys in China or Russia, who can compromise a machine just by thinking about it. That isn't true, of course. Back in Medieval England, if you got a strange rash on your leg, you would go to the wise woman in the village and buy a magic amulet from her. One of two things would happen: the rash would go away and you would get better, in which case the magic amulet was awesome, or you would die, In which case, you didn't buy a big enough magic amulet. That is where we are in cyber security today. If we allow people to talk like that, it drives a fear response that roughly translates into: "I don't really understand what is going on, I don't really understand what the problem is, but this guy over here says if I pay him some money, he can fix it for me."

When comic strips like XKCD make jokes about how cyber is deployed, I think we have a problem. The proper name for most attacks should not be called Advanced Persistent Threats, but rather Adequate Pernicious Toerags. They are adequate because the hackers only do the minimum necessary, and we don't make it very hard for them. Do not confuse technical sophistication with level of impact, these are totally different things. If we are honest about how most of these attacks work, if we are honest about the technical capabilities that most of the attackers use, we can do something about it. We can change the conversation and make it much more about risk management, and much less about hype and emotion.

The National Cyber Security Strategy 2016-2021, published by the UK government, is actually pretty good, for a governmental document. It is 80 pages long, and it says what we are going to do. I am only going to talk about one topic of what's in there, but in general it is a much more interventionist strategy. Instead of sitting for ten years, which is what we have done so far, saying: "you should all do better, and you should all share information," which probably does not work, we are going to actually go and do something. This strategy covers everything; it covers skills, it covers making the economy better, it covers offensive cyber, as well as active defense.

As security experts, we often give very bad advice. My favorite piece of bad advice is "Don't open attachments or click links unless you trust the source of the email." I find this really dumb, because how do you trust an email? It is true that some people know how to read Internet headers; I reckon I can do it about 90% of the time. For a simple message such as "thanks", you have rows upon rows of headers, and we expect everyday people to see if they can work out if a message is trustworthy or not. That piece of advice is monumentally stupid, because the average person cannot do it, and it just reinforces the response of fear.

Maybe we should do something else, then. DMARC, for example, is a protocol that allows a domain owner to take control of who can send email from that domain. We, at the National Cyber Security Centre, put 2 DMARC records for the @gov.uk domain. The first one was basically an instruction that says, "if anybody other than us tries to send an email from @gov.uk, meaning they are spoofed, don't deliver it to the end user, deliver it to us instead." On the first day after we turned that on, we got 58,000 spoofed emails, not delivered to the end users, but delivered to us. All of them were from the address taxrefund@gov.uk. On the second day, we got about 58,000. On the third day, we got four, all different. On the fourth day, about 58,000, and from then on, zero. They have gone somewhere else, because those spoofed emails are not being delivered to their victims.

The second DMARC record regards Her Majesty's Revenue and Customs, which is our tax authority, and is the most phished brand

in the UK. 300 million spoofed emails were not delivered because of that one DNS entry, which is pretty awesome.

Coming back to the topic of Harm Reduction, we have been experimenting with, and more precisely, we have been doing three things: if someone spots a phishing scam physically hosted anywhere in the UK, we will go take it down – we will email the hosting provider, asking them to stop the phishing. If it is a web inject in a piece of UK infrastructure, or anything government branded, anywhere in the world, we will go after that too. Before we started, a phishing site physically hosted in the UK would last about a day, and now it lasts about 45 minutes. This is harm reduction: the chance of somebody clicking the link doesn't change, but the chance of them being harmed when they do – does. Simple things can have a big effect.

We have built a public sector scale DNS, so that all of the public sector in the UK will use our DNS service. If someone tries to resolve something that is bad, they will not get to go there. It also generates huge amounts of data about what the government is doing, what government systems are doing, what they are trying to contact and why. We can start to build data at a government scale, to say – this is what the world looks like. This is all part of a program called "the active cyber defense program", and all the information is on our website, https://ncse.gov.uk.

This is about protecting the government, and it is about protecting the UK. First we secure the government, and then we scale it out. For example, we had conversations with UK Internet service providers, in which we told them: "it is probably not okay that your customers can hurt themselves without knowing, how about you protect them by default? And if you need data, you can have ours for free." Those are the sort of things that you can do on the national scale, if you start to look at the problem in a different way and you are honest about how some of these attacks work.

Our scaling strategy is really simple, we refer to it as "Point and Laugh". Suppose I'm able to get DMARC on every government system by the end of the year. I can then point out all the other sectors – the retail sector, the banking sector, and so on – and tell them: "hey, look, we are stupid, we are government and we could do it. How about

you do it too?" Because all this data is public, I can generate it, I can analyze it, and I can make evidence out of it, in public. I can help people make better investment decisions, to focus on things that work. The target of the strategy is moving the conversation in the UK from fear, which is what it is today, to published evidence and analysis, in public, transparently. If you want to change how cyber security is done, globally, change how you talk about it.

Conducting Global Discussions on Cyber Norms, While Still Taking Action

Christopher Painter, Coordinator for Cyber Issues,
Office of the Secretary of State,
United States Department of State

One of the things the United Nations' GGE session, which took place in June 2017, didn't reach consensus on, is the important issue of how international law applies in cyberspace. I think that this is an issue which we have been talking about for over a decade now. This means that it is not that the conversation is immature, but that some countries really don't want to reach the conclusion of how international law applies. It is very unfortunate, because I think that those countries want complete freedom of action. They don't want the structures, which have helped keep the peace in the last hundred years, to go forward, and to really apply in cyberspace in a way they can be a stabilizing factor. That is troubling, and the fact that a conclusion was not reached is not insignificant. There is more work to be done here, and I think we do need to think about those issues, and we shouldn't abandon them.

On the other hand, I think that over the last several sessions, the 2013 and 2015 reports in particular, and further work in this session, the GGE reached some very significant conclusions, and great progress was made in terms of the rules of the road, the voluntary norms of behavior below the threshold of armed conflict, the kinds of confidence-building and capacity-building measures that help better transparency and dial down the risk of escalation and misperception. There have been many good efforts, particularly, in terms of the norms that I mentioned, which I think are very important.

How can we move forward with this? We have already been talking, even beyond the GGE; the GGE is an important venue, but it is not the only venue. For instance, this year in the G7, the foreign minister statement had a more fulsome explanation of how international law applies. In terms of norms, we have been trying to get other countries,

even beyond the 25 that participated in this GGE, to really embrace this stability framework, including international law, norms, and confidence-building measures. We do this because this really isn't just for the developed world or a small group of countries, it is for all countries, for the entire world. The idea of these norms, of the "rules of the road", is that they really set expectations for states in cyber space.

In my view, the next step is that we continue to work with like-minded countries in the world – which could be a very large group – to pursue affirmations of those norms, and then use them as a foundation for deterrence. We need to band together collectively, as countries who embrace those concepts, against bad actors. People ask, what is the use of norms if people violate them anyway? Well, people violate norms in the physical world, too. I think the Russian invasion to Ukraine is a good example for that. It doesn't mean that we abandon the norms just because they are violated in the physical world, and it doesn't mean that we are going to abandon the norms just because they are violated in the cyber world. However, if we can get countries to agree, and then work together to sanction trespassers, using a whole range of different tools – our trade tools, our law enforcement tools, our diplomatic tools, and creating new tools, thinking of new consequences – we can work together, a "flexible" group of countries. It doesn't mean we don't need to continue to work in multi-lateral bodies, but we also have to think of how we can join together these like-minded countries.

When we talk about defining international norms and regulations, I think this is a process that requires a combination of partners and agreements – of course you are going to have a global debate, but it is going to be very hard to reach consensus on every issue, especially when countries have vastly different views of how they treat cyberspace. They have different views of sovereignty and control of information, and there are going to be limits. But that doesn't mean that while that process is going along, or while we are thinking of how that process should be, we cannot do other things.

I should make it clear that the global process should be a consensus-driven process, led by experts with a focused mandate, not some generally open-ended gathering that is majority-ruled. However, we

have wo work with the regional organizations, we have to work with like-minded states, and this is particularly important when it comes to the deterrence aspect and the consequences aspect. We have some good agreements in the set of norms that came out of the GGE, but the last thing I would like to do is create a "norms factory", where we are just inventing more and more new norms all the time. Yes, there may be things that are not covered because our imagination can't think of everything, but I think we have a good set of norms now, and that has been agreed upon. I think countries are comfortable with these norms, and that we need to expand the consensus.

However, when we start thinking about consequences, we also need to think creatively about how we can do things that are temporary and reversible, which will shape the adversary's behavior. Such consequences will allow us to communicate to them and say: we are going to hold you accountable, and if you don't change your actions, just like in the physical world, we are going to keep doing whatever we are doing. It is better if we do that collectively than alone. Think of an analogy to the proliferation security initiative, where voluntarily countries came together to stop the transport of fissile materials, and to act against transgressors. This is a flexible initiative, the parties don't have to agree on every case, but they bring different capabilities to the table, and then start thinking creatively about what the consequences are. What are the things beyond the toolset we talked about, that we can do to hold countries accountable? I think there is a lot of productive work to be done there in a large like-minded group, a small like-minded group, as well as in bilateral settings like the one we have with Israel and with other countries around the world.

There have been some calls for binding global conventions in cyberspace. Some countries, Russia and China among them, have been advocating for these type of conventions for more than fifteen years. They do it for many reasons, but one of their reasons has to do with something that democratic societies don't do, which is trying to control information. For them, the information itself is the threat, not the attacks and the intrusions that we are worried about. In that sense, it might be very dangerous to go down that road.

You can also build consensus from the ground up. For instance, we reached an agreement with China to not steal our Intellectual Property for the benefit of the commercial sector. After that, the UK reached an agreement, Germany reached an agreement, Australia just reached an agreement, Canada reached an agreement, and the G20 endorsed this. The useful work of the GGE was getting some agreement among the P5 and other countries on things like not attacking critical infrastructure of another country, which provides services to the public, outside of war time; war time has a different set of rules. We can take that forward in regional and bilateral agreements.

Regarding deterrence and violation of norms and rules, I think we see many bad things that happen in cyberspace, that we haven't necessarily thought of or considered before. Things have happened that we didn't expect, and we got surprised by them, but that was clearly a bad act by a bad actor. In these circumstances, I think you *can* act against those bad actors or violators of norms. That said, you need to be transparent in acting against them, and you need to communicate why you are acting, and why you are doing what you chose to do. I am not saying we that shouldn't have global discussions on such matters, but I think we shouldn't wait to have a discussion in close like-minded circles, and then in larger group of like-minded parties, and then continue enlarging the group until we create a complete consensus.

Some say that as part of acting against bad actors, we should disrupt or slow their access to the Internet; others say that it is more important to keep the sanctity of the Internet as an open, global, and free space. Personally, I think you can have both. You want to safeguard your core values, but at the same time, you want to have consequences that will make a difference, and will hold people down. It is not retaliation, it is a constant position to try to get a bad actor to change their behavior, and that is different under the international law as well. I think that in order to do this we have to be creative, not just in the governments, but also to talk to people in the private sector, people who do international trade, people who deal with some of the technical issues. We need to find out if there are things we can do to restrict the access of people or countries to the benefits the Internet

brings, or even those that the global financial system brings. You don't want to disrupt this system though, so you have to have a balance. You have to talk to the people about what the positive parts of this are, as well as the negative parts. These are creative discussions we have not had yet, to expand those tools beyond the fairly narrow set of tools we have.

This brings me to attribution. People raise attribution all the time, but in my opinion, at least, attribution is not as huge a problem as people say it is. We were able to make attribution in number of cases, including the Sony Pictures attack by North Korea, and we recently released a set of indicators about North Korean activity. This is something that can be done, especially with prolonged contact, and attribution doesn't have to be a hundred percent. It is a political decision; if it walks like a duck and talks like a duck, it is a duck, and you can take action. I think people who look at the space often use that as something to let them say they can't do anything, and that is a mistake. If you don't do anything then that signals to the adversary that everything is Okay.

Looking at the future, there are multiple tracks we can take. Conversations with the like-minded, and building a flexible coalition to take action, is an important thing that really needs to start now. Those are conversations I look forward to having, while at the same time trying to advance the ball, in terms of how international law applies, getting more acceptances and norms. None of these are exclusive, but I think we really need to move to that stage of how we work to go against transgressors, to make a more stable and peaceful cyberspace for everyone.

Defining International Norms in Cyberspace

Paul Ash, Director, National Cyber Policy Office,
Department of the Prime Minister and Cabinet,
New Zealand

States are the primary bad actors in cyberspace, as no non-state actor has the capability to launch the most damaging cyber-attack. We are in a transitional moment in international negotiation on cyber security, a unique situation where we have people who are our opponents, and they don't agree with us on the rules and norms that have been put in place a long time ago. All this was evident at the GGE session of June 2017. I think we are in somewhat of an inflection point, in terms of how we think about state behavior, and how we operationalize the really good work the two previous GGE reports delivered.

First, we need applicability of existing international law to the online behavior. Second, a relatively tightly focused group of norms. The GGE has always struggled with trying to deliver consensus around norms. We get into difficulties when we start having a norm for every letter of the alphabet. We even have a difficulty on who's alphabet to use. However, in the absence of normative behavior, norms become quite challenging when you start to get to a panoply of different ideas. What that leaves us with is two reports that really set out the basis amongst a small group of experts. I think it is important to remember that we have gone from 15 to 20 to 25 states in the GGE, meaning we don't have a whole range of states that participate actively in that debate yet. From here, the question is how do we take the good work that is been done, at the GGE and at the global conference of cyberspace series, as well as a number of other discussions, and try actually operationalize those, see what they look like in practice. In many respects, this will be what determines what normative behavior looks like, because we actually start to put some boundaries on behavior that is not normative, behavior which is beyond the bounds of acceptability. Over time, that is going to require much work amongst like-minded states. We need to broaden our thinking beyond the organizations or the states we currently work with, while trying to work hard to understand where

that core of common agreement is, and what kind of responses we have in order to push back behavior that is beyond those bounds.

One of the most fertile discussions of the moment is: what does deterrence look like in cyberspace, and what do the consequences look like in cyberspace? And there is a couple of points I would make there. First, we are at a fairly early stage of that discussion. We know what bad behavior looks like, because we are starting to see it regularly, but we are really just starting to push up the bounds of what the responses will look like. The first assumption we need to challenge is quite a common one today, which is that bad behavior using cyber measures necessarily means a response using cyber means. I think that would lead us down a rabbit hole. It is quite heartening to see a number of states starting to broaden the types of tools they have – diplomatic tools, economic sanctions, law enforcement tools, and so on.

I think that all the states involved will probably reserve the right to use force, if they need to use it, or some of their cyber means. That said, we really are still at a very formative stage in this process. It is important to look at recent events and at some of the hiccups we have seen in our efforts for joint discussions, acknowledge that there were some bumps in the road, but focus on the fact that we actually start to form that kind of customary international law, wrapped around on our own framework that will deliver outcomes.

I think the challenge in trying to create a global discussion process is to avoid leaping to hasty conclusions about how we should take this work forward in practice. This is not a binary choice between a multilateral approach or to each his own. If we look back at the lessons of history, whether in trade policy or in security policy, it has taken us quite a while, using a range of tools, whether they are multilateral, regional, bilateral, or whether they were just collective security tools with the "coalition of the willing", to take forward the debate where we have had new types of threats to grapple with, or in the trade arena, where we had to work our way bit by bit through some of the complications of trade policy. I think that along the same lines, we should be looking to build as many tools in to our toolkit as possible, as we take that forward.

There are some caveats I would place around the entire notion of retaliation and global norms. From my government's perspective, I certainly don't think we are in a position where we should be stepping forward into a negotiation of some formal treaty, as we are still trying to understand what the main actors in this area are doing, and we are still building a picture of what the capabilities of cyber armed activities are. If we are not careful, we could find ourselves locked into fairly poor outcomes, with unintended consequences, if we push directly to that multilateral treaty space at the moment. The discussion needs to continue, but formalizing it at the moment would not be a good place to go to. That leaves us with a range of other discussions.

There has been useful confidence-building work going on in several regional organizations. We do that in the AESEN Regional Forum, with a range of partners, and I know the OSCE has done the same. That is very illustrative of the challenge we have. The early discussion in those forums tried to draw on confidence-building measures from an earlier era and from earlier threats, and we quickly found that they didn't work. We need to be looking for new solutions, and I suspect that will be best done by groupings of like-minded organizations and states coming together, to really test the boundaries of the tools we have. The ones I have named are quite traditional, but there may be others in the technical realm or elsewhere, that we can start form together if we got that core of countries working together.

I think we somewhat disregard our need to reach out beyond like-minded nations, and start having conversations around the points where we differ from other states as well. Keeping those communications channels open is going to be quite important, as we try to understand how we can get to a better place, of relative stability in cyber space.

As I mentioned earlier, when we talk about consequences, usually people think about cyber vs. cyber, but there are other tools we can use, and this is exactly the place where we need to be taking our discussion; what is the full range of actions we can take? The first place all of us will start is deterrence by denial, making it very hard for threat actors to carry out bad behavior. I would prefer punishment or deterrence by consequences, rather than retaliation, because I think retaliation makes it seem like you are looking at a "tit for tat"

response, and we need to be far more creative with our toolkit. That can be something in a whole range of diplomatic activities, and we have seen it in public. We have seen people being expelled or pushed back, and we have seen statements of concern published by attacked parties. Additionally, not all diplomatic activities have to overt. In private, conversations teams can be much more direct, and I think that is a useful thing.

We are starting to think around what economic sanctions as means of deterrence to cyber attacks can look like, and how you start to implement them. There is a whole range of those tools that we can use as well. Another important thing here is the emerging conversation around what collective security really looks like in cyber. We have already had 70 years of that conversation when it comes to conventional threats, but we are still at the front page of what that would look like for cyber, and we saw that in some of the NATO discussions after the Estonians issues, back in the early 2000s. We have yet to work this out, and in a sense that is really one of points of contention in these discussions around whether or not existing international law applies online. We have to understand what thresholds we are going to apply for some of the existing pieces of international law around collective security.

What I do know and what I am sure of, is that not responding to a threat – rather than responding to it – is going to get highly problematic over time. In customary international law, such a response (or the lack thereof) has the consequence of making the activity appear to be a tolerated activity. Over time, that will be something we are going to need to push back on. In terms of response, I think some of the principles that already apply in international law, in terms of how you respond to an attack, would apply in a range of other places as well. It should be proportional, it should be precisely targeted, and it should not involve non-combatants or others in the response. That, I think, starts to provide some useful principles right away through the spectrum of bad activity, for how you might respond to it. However, there are governments who don't want to take options off the table, simply because they might have an impact on the free functioning of bits of Internet over time. We certainly see that from threat actors.

We know of DNS amplification attacks that have taken down large parts of the Internet in small states. We absolutely need to balance the ability to prevent that from happening with things related to precision and care.

The key thing I would say is that the threshold for that, the accepted practices for that, will only emerge through a really robust set of discussions, both with those with whom we are very like-minded, and actually with those with whom we are not. If you think about deterrence, it is very difficult to deter unless there is, at least, some overlap of what acceptable behavior looks like. We should work with like-minded on the full circle discussions, and with others on where the Venn diagram overlaps.

Thinking of the future, we need to work with what we have already achieved around the application of existing international law around norms and safety, and seek to grow those things out of the core of organizations and countries that take similar approaches. We have been through that before with security issues, where we have to work through things, see what works, see what doesn't, maybe even fail fast at some of them. But in the end, with some joint effort – we will succeed.

The Cyber Effect: Introduction to Cyber Psychology

Dr. Mary Aiken, Forensic Cyber Psychologist and Member of Strategic Advisory Group, Paladin Capital, Ireland

Cyber psychology is the study of the impact of technology on human behavior. It is an advanced discipline within applied psychology, and it is only about fifteen years old, as a discipline. My area is forensic cyber psychology, and I specialize in deviant, criminal and abnormal behavior online. There is a symbiotic relationship between the so-called real world and the cyber world. What happens on the cyber world impacts on the real world, and vice versa. We can go back to the work of Licklider, and the paper he wrote in 1960 about man-computer symbiosis. Ultimately, this was pre-technology, as we know it now, pre-Internet, but it is a wonderful example of how, instead of looking at AI, artificial intelligence, one can focus on intelligence amplifications, IA, and place the human at the center of technology. My job as a cyber psychologist is to work at the intersection between humans and technology, or as law enforcement people say: "Mary works where humans and technology collide."

I am the academic advisor to the Europol Cyber Crime Center, and I have been involved in a dozen different research areas. One thing I have observed is that whenever technology interfaces with base human behavior, the result is amplification and escalation. From a scientific point of view, this points to the mutation of behavior in cyber context. We have established state and trait behaviors within cyber psychology. For example, the online disinhibition effect, which is a work of Suler in 2004, dictates that people do things in a cyber context that they would not do in the real world. Whether you are a toddler with an iPad, a teenager, somebody in your twenties, or a criminal, human behavior can mutate and change in cyber context. I wrote a book about it, called *The Cyber Effect.* I think this effect is the "$E=mc^2$" of the century. If we can figure out this amplification, then we can also figure out how to deescalate it. Whether you are a cyber

juvenile delinquent, a lone cyber criminal, a hacker, a state-sponsored threat actor, or a member of an organized cyber crime gang, the one thing that you all have in common is that you are human, and the one thing that is unique to human behavior is motive. In cyber psychology we talk about primary motive, sustaining motive, overlapping motives, and primary and secondary gains.

I want to talk about the Sony hack. I didn't actually work on this, so I am going to talk about this theoretically and hypothetically. Effectively, what was the motive for hacking Sony? Conventional wisdom says that it was intended to stop the release of the movie. But in the end, the movie was indeed released by a company called Jedward, on a URL that they bought for $4.99 from GoDaddy. If the initial motive, the primary motive, was to stop the movie being released, then it was not a sustaining motive. That URL was not attacked, the account wasn't compromised, so therefore, was it about the movie? If we look at it in terms of typology, and this typology leads to the creation of a deductive cyber criminal profile, if the movie was released after all, then maybe it wasn't about the movie. I am not saying that North Korea didn't jump in once the door was open, but the question is who opened the door. Effectively, if you look at this in another way, you would say that Amy Pascal was the victim, because she actually lost her job as a result of what happened. If you are looking at creating a group of suspects, in terms of motive, I would look at anybody that she has sent an aggressive email to or had had confrontation with, within the industry. Allow the data to speak to you, and don't simply jump to side with conventional wisdom, because there may be more to it than it seems. Sometimes the truth is there, hidden in plain sight.

I am going to shortly discuss anonymity online, or should I say perceived anonymity online. I think that one of the big truisms about human behavior online is that anonymity is a superpower, the power of invisibility, but it comes with great responsibility. As humans, are we capable of dealing with that level of power? I think there might come a time where we have to question anonymous protocols online, because what happens in a cyber context, in cyber society, will impact our real world society. I don't think we will ever be able to regulate

at the speed in which the technology evolves, but it certainly doesn't mean that we cannot have governess or some form of good practice.

Another topic is online syndication. This is a theoretical construct that comes from my work in cyber psychology, and it is really the mathematics of criminal or abnormal behavior online. Think about it this way: to date, the incidence of the general population of this type of negative behavior online, has been bound or capped by the laws of probability and domain. What does that mean? If I am a sex offender living in the north of the country, and you are a sex offender living in the south of the country, then what is the probability of us coming across each other, and then normalizing and socializing our behavior? It used to be limited by the geographical factor, but now, under the cover of anonymity and fueled by online disinhibition, negative actors and threat actors can actually syndicate to find each other online, and then normalize and socialize that behavior. As humans, we are reasonably simple creatures in some ways. If it looks like everybody is doing it, then sometimes we will convince ourselves that it is okay to do it. My prediction is, and I hope I am wrong, but I am probably not, that this will actually drive the incidence of these abnormal and criminal behaviors in terms of the general population, and that is a problem. A friend of mine from Interpol has said that we are facing a tsunami of criminality, coming at us down the line, online. My work is really to try to get people to pay attention to the space.

In 2016, NATO declared cyberspace as a domain of operations, which effectively means a domain of war. People like me, cyber psychologists, have been talking for years about cyberspace as an environment, somewhere you can go, a domain, chatrooms, forums, etc. Have you ever sat down, just to finish an email before you go out for dinner, and all of a sudden an hour has gone by? There is a time distortion effect, as you get into this space. At the moment, don't forget, we are talking about transactional interactions, using our phones or our computers. The point in which we get into virtual reality and head-mounted displays, is a point where we are not going to be only psychologically immersed in cyber space, we are going to be physically in the space as well, and this is powerful. If we go back to psychology, and look at the environmental psychology, environment

actually impacts on behavior. Again, I believe that to a certain extent, the behavioral scientists have been blindsided by rapid evolutions in technology. When I studied psychology back in the day, everything changed when I first came across artificial intelligence in the late 1990s. What occurred to me was that nothing in my studies, up to that point, equipped me to understand the profound and pervasive impact of technology on the human, as an individual, and on society as a group.

There is a concept called "routine activity theory", or RAT. This is a theoretical construct that comes from real world criminology, in terms of geographical profiling. Coming back to the scientific empirical aspect, the one question that we have to ask ourselves as scientists, as academics, as researchers, as students, is if our theories, that have been conceptualized, hypothesized, tested and validated in real world constructs, are still applicable online? If we look at theories of criminal behavior profiling, we have personality theories, we have labeling theories, and we have geographical profiling theories. In this theory, we take a typography, like the map of London, and we might have a cluster of crimes. Effectively, we would look at triangulations. We would look at a hypothetical positioning of the perpetrator; where they live, where they work, where they play. Then we would rotate that triangulation with the crime clusters along the connecting buffer zones, because the point is there is going to be a buffer around these areas, and crimes will always take place on those pathways. We can take that concept of routine activity theory, and transpose it to cyberspace, which is a domain. Effectively, what we get is cyber routine activity theory, surface web to deep web.

I am going to conclude with a philosophical existential question. In terms of causation and correlation, does technology cause bad behavior? I don't think so. In terms of the connectedness of the Internet, unprecedented connectedness in fact, it may be that technology and the Internet per se have enabled us to shine a very bright light into the darkest reaches of the human psyche, what Jung called the two million year old man. Maybe we are all just Game of Thrones underneath it all.

Hacking the Human

Chris Roberts, Chief Security Architect,
Acalvio Technologies, USA

Usually, when we see discussions and articles titled "Hacking Humans", we hear about the psychology side of it, or about hacking the human from a social engineering stand point. I want to actually hack humans. It is much easier. I'm done dealing with humanity, let's hack the humans themselves. Specifically, I want to discuss building better humans. Why are we doing this? Because if we look at how we are advancing ourselves, the logic is really, we want to build better humans. We tried endoskeletons, we tried exoskeletons, but they don't work as well as we expected. Now we start thinking about what is the actual construct of humans, and realize that we are 18% carbon, give or take a bit. This leads us to wonder whether or not we can use carbon nanotubes, and the answer is probably yes. This, naturally, leads to another question, the eternal question of mother nature: how hard can it be?

Let's take a look at how nature hacks humans; we are hacked every day by the viruses that we get. We are hacked on a regular basis through RNA coding. If you take a human network and you take an actual network, you can have a lot of overlay. You take a look at the viruses that get into our systems, they access the cells and use our bloodstream as the transport methodology, to get around us. They communicate by using an infected host to say to other viruses the host is already infected, and that they should move on. They have a contagious exfiltration – I sneeze and you get infected. We want to know how mother nature does this, and so we look the core genetic material of viruses. It is important to note, though, that we are not taking nice civilized stuff, we are taking Ebola, and we are smiling, thinking we can make it civilized and inject it back into the body with all sorts of carbon nanotubes, and it will be perfectly okay. We have managed to this before. The idea is that if we can't beat mother nature, we can copy it.

How do we do this? We have broken it down fairly well. We need to define the nanotechnologies we will be using, we need to build them, encode them, and deploy them. Also, we need a panic button.

We begin by defining them. What is nanotechnology, really? This is the fun stuff. This is where we take something that is a couple of atoms across, we take a look at something that is one of the strongest and hardest elements, one that has some fantastically, freakish thermal properties. We realize that when you get to the molecular level, the standard laws of physics that we have today don't necessarily apply. We start walking around in the quantum side of the world. At this point, we have these fantastic materials, that we can do all sorts of interesting things with, and they are very small.

We can build with these things. We have been able to build with them since the mid-1990s, and probably even before that. We have three different methodologies of building with nanotechnologies. We have the ability to use nuclear acid robots; we have bacteria-based technologies; and we have nanoscopic assemblers, which are being built today. We are creating the ability for a nanotech architecture to build itself. We are creating self-replicating machines, at a molecular level. This is what it comes down to, the whole concept of manipulating atoms. Basically, we are doing here the same things we do with computation technologies, the way we hack a computer system. We are thinking about the building materials, the programing, the code, the payload, the infiltration, and the exfiltration, in much the same way. We also do the same things we do with an assembly line, but on a slightly smaller scale.

How far along are we? In 2016, the EPFL labs released a study showing that they can manipulate and teach nanotube carbon architectures to mimic letters in the alphabet, and do simple mathematics. The next step will be to teach them how to recognize chemicals in the body. Manufacturing is also a useful indicator as to where everything is and where everything used to be. In 2010, it cost us $1,500 to produce one gram of nanotubes. In 2016, the cost is only $2 per gram. We are getting better, faster, and smarter at building this next generation of technology.

Moving on to encoding, or: if I have my material, how do I actually do something with it? We started by doing DNA computing, the ability to move something through a DNA life cycle, and encode it into nanotechnology. Nowadays, we have come up with a programing languages. We can start using code, similar to the regular simple coding languages we use today. In essence, what we have is hardware languages, like Verilog, and then languages like Cello will help you design DNA sequences. Eventually, you will end up saying: "I want to build a sequence, I want to have the sequence, I want to design the sequence," and you will be able to do it in a language that most of us can encode and program in. For example, we can use Cello to design a program, using a simple design language available today. We take that design coding and encode it into a constraint file, which we then use to build gate assignments. This is simple mechanical engineering; this is electrical engineering 101. Eventually, we take our circuit, and turn it into a signal. This is no different than what we do on a computer network, except in this case the signal is transmitted through either a laser or near-field wireless style communications.

After we managed to build our technology, we have encoded it, and now we want to deploy it. We are hacking humans, so we need to get it into them. In order to do that, we grow our carbon nanotubes and give them little tails. We take a classic propulsion system, apply a heat source like laser, put them close to the skin, and the thing starts streaming around the bloodstream. We just built a motor at a micronic level.

Lastly, we can talk about the panic button. Yes, we have a panic button, but we are not going to hit it, because there are some myths that need to be dispelled first. The whole concept of getting these nanotechnologies into the body is currently possible only by injecting it into a test subject, human or otherwise. At present, it is impossible to drink tainted water or breathe tainted air supply, and "honey, I shrunk the kids" doesn't work, despite all of our best efforts. We, as the entire scientific community who is working on this, still have a long way to go. The mechanisms are in place. We have the nanotubes that can communicate from the internal to the external; we can make them move around the bloodstream; we can have them actually attach to cells; we

have the internal communication system. What we are working on, and what is being researched by many teams and laboratories across the globe, is what we can do in the next ten to fifteen years. What we, as a security industry, need to come together and do, is to pause for a second, put a picture of the big scary hairy six foot three guy with the green beard on the screen, and say: "don't let him hack things." We don't really want to be hacked.

But if we did, this is how it would be done. We have actually created a hexametric ring with receptive binding capabilities. In other words, we created a virus that gets into the system and says: "hi, I would like to make friends with something. I have receptors that fit with some of your blood cells. Let's be friends." And so it does. It has the ability of transport, Bird Flu in this case, nice simple things that can get into the body; it has bypass tools; it has reporting tools that can tell us what is it doing in the body, and how does it achieve it; And finally, it has the payload. In simple English, we took Bird Flu, and then put carbon nanotubes on it. We convinced the body this was something it should accept. We put a tracking device into it, and we put a deployment technology into it. In a good world – in other words, in a laboratory – at the moment, we can deploy this close to cancer cells and target them. That is the good side of this technology. The bad side of it is that we can also target red blood cells, for example. One might think it takes an entire computer room full of technologies to communicate with it, but the truth is that we can do it with a hundred Dollars' worth of Arduino communication systems. This is game over. We can hack the human with an Arduino and a couple of instruments of equipment. This would be a good time to panic. If that wasn't bad enough, they decided to drop it into the agriculture environment as well. The agriculture industry decided that this technology was cool enough to put in the food supply. We are going put it in the food processing chain, and in the supplements we are taking.

Best case scenarios are awesome. We can be targeting a cure for cancer, breaking blood clots, delivering drugs and medicine to where they are needed in the body, breaking up kidney stones, removing parasites, and doing many other good things. The downside is the possibility of targeted attacks. In a worst case future, we will become

nothing more than a walking USB stick. We will weaponize it, lose control of it, technology will gain the upper hand, and this will not be a really good situation in general.

What do we do about it? How do we solve this issue? We need to communicate and collaborate better. We obviously need to start taking a look at the deceptive technologies, the technologies that are outside the normal stack of equipment that we have that sits there with blinking lights and doesn't protect us. We need to take a look at addressing this with organizations, and we need to educate them and get them to listen. That said, we have the power to do so much good. I call everyone reading this to do some good; gather, communicate collaborate, because this is the future that we are working with. Let's guide it in a slightly better way that we have done in the past, please.

Methodical Social Engineering

Chris Nickerson, CEO, Lares Consulting, USA

I need to give a disclaimer that I'm a hacker. I developed my consultancy by coming from the hacking perspective. I have been an adversary my whole life, as my mother would tell you. I don't have the same understanding of the human mind that scholars may have, or that mentalists use to manipulate and shock people. I do it professionally, because I saw that there is a huge landscape out there of things to attack. There are people focused on the landscape of IoT, computers, Internet, hyper- connectivity or any other buzz word that we talk about these days. However, there are many things that are very easy to attack, which are not powered by batteries. I want to examine that a little bit deeper. I tried to do some of these things professionally, working with law firms, large telecommunications companies, and some larger distributors. What I found through this is that I was just looking at it from a hacker's perspective. I wanted to see exactly how far I could get, with a very limited amount of information.

I have been hearing many answer to the question what is Social Engineering. Unlike what some people believe, it is not phishing nor lying to people. It isn't even deception, really, but rather it is trickery. Think of the photographs of Liu Bolin, where we are looking at one thing and completely missing some of the hidden details. Even our eyes can be tricked. No matter how much you pay attention at first glance, you might completely miss what's going on, and it doesn't matter how trained you really are. We can look at those pictures over and over again, and as we start to see the face and the shoulder of the man in the middle, as we start to see the center line of the dress and attire, we start to realize that he is actually painted into the picture. But when you first saw it, was he there? This is the idea of the tree falling in the woods. If I was to walk into your data center as the electrician, and you thought I was the electrician, was I ever really there? It is complicated to think about it, because no level of IDS is ever going to detect me, because that human barrier is gone. I have actually erased it from your memory because it is so common,

and it is so normal to see these types of things, that you won't even remember that it happened.

I'm selling chutzpa as a service. Can I just walk into somewhere, and say: "give me all the things?" Maybe. But is that a repeatable process? Is that Social Engineering? According to the dictionary definition, being social is having the ability to intercommunicate with people, something having to do with society. Engineering is a repeatable process. Social Engineering is not just the basic lie. This is not just me going and willing my way into it, and using my skills to try and manipulate someone. This is me, understanding exactly what I'm going to do and how I'm going to do it, much like an attack structure. There are basic structures of attacks, that are repeatable, and are fundamentals of engineering. This is myself, living as an attacker, running a team of attackers all over the world for the last twenty years or so of my life, writing standards on how to attack humans in a repeatable way. We need to gather intelligence about them, we need a threat model, we need to understand where the vulnerabilities exist, and then we need exploitation. We need to be able to find that tipping point between where the vulnerabilities exist, and how I am going to move to the next stage. Once I moved to the next stage, I need to be able to get the access and the information I need, pull back, and make sure that I can regain my access.

It is an interesting thing to think about, which goes well beyond the lie. Sometimes I will get to a point where I know the building security is so advanced that I am not going to be able to do it on the first try. To overcome that, I fail intentionally in the beginning. I walk into the building late at night, expecting somebody to ask for my credentials. I tell them that I forgot them, they explain that I need to have credentials to get in the building, I tell them that my boss is going to tear me up because there is a server down or some other excuse, and they insist that I have to have credentials, and that they can't let me in. Meanwhile, all I'm doing is poking for holes. It is like a sports game; I'm trying to find where the vulnerabilities are. Even though I know I'm destined to fail in this exercise, I have started to identify where the vulnerabilities exist. Instead of coming back that night, knowing the guards' schedule, knowing what's going on, I

come back the next night. This night I have a fake badge, because I was able to go to a copy center and print a copy of the badge with my face on it. I present the badge to the guards and tell them I'm back. Just one more night of work. By this time, they already recognize me, so they let me in. They don't bother to check my credentials, because I've already anchored myself in their memory. They already know that I have had a tough time getting in, and the thing that they said to do was my exploit.

This is a repeatable process. You try to identify what's going on in their mind during those situations, in order to get the most effectiveness for every piece of speech that you have. Interestingly enough, speech is almost irrelevant. When you look at the human mind, and what we do and how we do it, speech is only about 7% of all of our communication. This is a very small part of how people communicate. Most of the ways we communicate are non-verbal. The reason for that is because unlike IoT devices, which are supposed to be smart, and can make several calculations per second, our brain makes trillions of calculations a second. The last time IBM did the scientific study to identify what the processing power of the brain was, they said it was somewhere around 38 petaflops. If somebody wants to tell me that their advanced computer system is so much harder to hack, I refer them to the stats of human mind. I refer them to the fact that our eyes have been calculated somewhere in the neighborhood of about 127 million megapixels. That doesn't sound like a very slow and non-advanced computer to me, but wow, what an opportunity there is to hack there. If it runs that fast, it must be able to do all sort of things, meaning we can put an infinite number of things into that channel in order to exploit.

How do we do it? Port scanning, just like you would in a network. How things taste, how they look, what they sound like, how they smell, even what they feel like. If we go through all of these different topics, and we start breaking down what the journey of this looks like, we get a very complicated chart. However, from my experience it is easier to look at these things from basic views. Let's take the physical aspect for example. Everybody breathes, so I went to India and I studied breathing. I needed to know how I can breathe with someone, because

this was a non-verbal that created trust. I could learn how to breathe out of trust with someone. I could learn to take someone's aggression, and just by breathing move them into a different state, and by that I'm beginning to run my exploit. By using things like family therapy, by using neurolinguistic programming, by looking at different Gestalt therapies and transformational grammar and reading loads of Chomsky books, you start to learn these wonderful tricks of the mind. There's a great video of Darren Brown, the magician, showing him using sound anchoring techniques to prove that he could win on a losing ticket at a horse track. He does it over and over and over again. Granted, I don't normally get paid that way, but for fun you can do it.

The next topic is exploitation. How do I take all of these different techniques of breathing with someone and looking like someone and push that into a point where I'm able to get access? I managed to get access to a data center just by wearing the right shirt. I knew, from my initial survey and external testing, that they were having problems with one of the power systems, and that it was making an insufferable noise. I tried to present myself as an employee of a power company, but the guard the data center insisted that there was no possible way I was getting in. What I did, eventually, was to send an email, impersonating someone from the power company, to the person in charge of the data center. In the email, "the person" basically told the data center manager that they know there are some people from the power company coming over to work on a different matter, but maybe we can help them as well. Within exactly one minute I saw someone running from the data center towards the gate, and said to me: "oh, thank god you are here." They rushed me into the data center, and told me that this noise was driving them crazy. What did I do? I simply walked over to the audio button and I pushed it, held it, and finally let go. I had no idea if that would work, but fortunately it did. From that point I had full authorization from the data center manager to do whatever I wanted. The funniest part about that story is that I went a little bit further, to make sure I can maintain my persistence. I picked up the phone, called the actual company that they were trying to get for service, and I started screaming at them; "I can't believe you are going to let this data center go down and not send technicians

out here. What is your problem?" and all that time, the guy in the corner is yelling and cheering me on. I now created access forever. All it took was a simple t-shirt.

When I look at it and deconstruct it, I am able to pick all of these different things that I was doing, throughout the entire time. It wasn't a simple lie, it was planned. By using that plan and executing on the plan, I am able to not only get access, but to maintain access forever.

ACADEMIC PERSPECTIVES ON CYBERSECURITY CHALLENGES

Trends in Cyber Research

Dr. Yaniv Harel, Head of Research Strategy,
Blavatnik ICRC, Tel Aviv University and General Manager,
Cyber Solutions Group, Dell EMC, Israel

The Cyberweek conference is one of the unique conferences that take the interdisciplinary definition of cyber and make one big conference with various sessions, each focusing on a different aspect of the cyber: the technical side, privacy, law behavioral and others. We always aim to get a wide view of this topic, and we try to present the most recent ideas and researches in this conference, as well as discuss cyber challenges, which today constitute an integral part of our society challenges.

In recent years we see more and more budget efforts and initiatives within the cyber field, in the industrial sector, in the academia, and in other sectors as well. As we follow the cyber topics of the last few years, we find that originally, the cyber research was mostly focused on the technical part of cyber, expanding and investigating our technical knowledge, while at the same time promoting collaborations between the various sectors. Later on, other aspects have also joined this discussion.

In 2016, we saw many studies and papers revolving around security of mobile platforms, and specifically platforms that are relevant to Android. We also saw a lot of work done on cloud security. Finally, we saw that a lot of research is being done on the human side of attacks, namely social engineering. We are seeing work being done on new ways to view and think about cyber security, such as node network dismantling, reducing cascading failures, and looking at our networks as social structures. In 2016 we started seeing initial work that tries to move the conversation from potential threats to actual threats, or lack thereof, and also about attribution, which is a major and important question in the cyber world today..

All of these topics and more are of interest to us here in the ICRC, the Israeli Cyber Research Center in Tel Aviv University. This center was established in 2014, and one of its purposes is to take the wide,

interdisciplinary definition of cyber, and run many initiatives and research efforts in almost any of its aspects. When we try to track, after more than two years of research, how our researches are divided between the soft and the core technical cyber, we mostly find a balance, which is what we were aiming for, in the present and in the future.

Classifying Android Malware into Families and Determining New Families

Prof. V.S. Subrahmanian, Professor of Computer Science, University of Maryland; Head of the Center for Digital International Government

First, I would like to mention that the work presented in this article is a direct result of my coming to Tel Aviv University several years ago, even though the relevant published paper has no authors from the Tel Aviv University. I met one of my co-authors, Mr. Fabio Piratzy, after I met his advisor at Cyberweek a couple of years back. Following our meeting, the advisor sent Mr. Piratzy to my lab for a year, which resulted in this paper and others like it. Therefore, I would like to thank Tel Aviv University for indirectly, and perhaps unintentionally, facilitating this collaboration. I would also like to say that even though I will be discussing Android samples, the techniques presented here are much more general.

Android is the dominant mobile system today, and that is a known fact. A lesser known fact is that the number of new Android malware samples is also increasing in a very high rate. In 2016 alone almost 3.5 million new strains of Android malware emerged, which is almost 9,000 new samples per day. I started thinking about the problem of classifying samples into specific families many years ago, when I was working on some data from Symantec. They were interested in the following problem: given that a binary file is known to be malware, can we say which family it belongs to? This is interesting because often, when a malware is released into the wild, other hackers take that malware and adapt it so that they can defeat methods that detect it. And so, the same piece of malware evolves into many different samples, based on who makes the modifications, what those modifications are, etc.

For antivirus companies it is very important to be able to understand that the new sample that they are seeing today belongs to a family that they have already dealt with in the past. If they know that, they also know that they can adapt the signatures they developed for the

past samples, i.e. the family to which it belongs. They know that they can then adapt patches from what they already have, which reduces the amount of required work. At the same time, it is very important for them to detect new families that are emerging, new samples that do not seem to correspond with any of the families that they have encountered in the past. What we try to do in our research is to capture both types.

In this research, we worked on three data sets. Two of these data sets, Drebin and Koodous, are publicly available, and you can get them online after filling out a request form. The nice thing about Drebin is that about 5,000 of its samples have labels, and these labels corresponded to 156 families, to one of which each of those samples belong. Therefore, we can see that, on average, the size of a family is about 30 samples. For machine learning people this is not a very big number. Koodous has much more data, about 100,000 samples, split about 50/50 between benign software and malicious software. However, only about 7,000 samples are labeled with families, and moreover, the family distinction in this case is much cruder, having only about 12 very large families. Finally, we built our own data set by leveraging VirusTotal. In this article I will only discuss the first data set, Drebin, but the results we obtained apply to the others as well.

We took these samples, and much like epidemiologists study biological samples, what we are trying to do is to study them in a controlled environment. Just like biologists use a biocontainment facility to study pathogens of different types, we are studying our samples in a sandbox environment. Obviously, there are problems with using any kind of artificial environment. The behavior of a tiger in a cage, which is a controlled environment, may be different from its behavior in the wild. In the same way, the behavior of the malware in a control environment may, in fact, be different than it is in the wild. This is a potential flaw, not just of our study, but of all studies out there that use controlled environments. Obviously, my university and many others will be very unhappy if I will try to run all these malware in the wild, so I can't actually do that without great personal risk to myself. And so, we have a sandbox environment which we have built ourselves, consisting of network simulators, emulators

and some commercial sandbox solutions. We also make sure to log all kind of things, including file system activity, network activity, cryptographic operations, and so on.

As mentioned before, the Derbin data is classified into 156 families, and in the world of machine learning this is referred to as a 156 Way Classification problem. In other words, it is not a simple binary classifier, and you have to classify things into very fine-grained boxes. The second challenge is that 47 of these families have only a single sample, and any machine learning person knows that it is very hard to do anything with that. We also know that of the 156 families, 112 have a sample size of less than or equal to nine, which is also considered a small number. This means that only 44 families have more than 10 samples. The reason I am stressing this point is because this makes it very hard to generalize, and build good classifiers.

We use the sandbox environment to inspect both the binaries and the source code, if it is available. We also perform both static and dynamic analysis. Our static analysis was conducted on three features, dealing with the author, the structure, and the type of permissions requested. The dynamic analysis ran the code and identified sequences of operations that were executed. We sometimes like to imagine sequences of operations executed as a piece of text, and then identify within it bigrams, unigrams, trigrams and so on. We then take our findings, and derive features from that. In other words, the kinds of dynamic features we are looking at are sequences of operations that are executed and we are counting various factors that associated with them.

We started by looking at families that contain at least 10 samples, and asked ourselves – what if we just take standard classifiers and run them? We took the results and looked at three types of segmentations: static features alone, dynamic features alone, and a combination of both static and dynamic features. We then looked at results that received a score of 0.9 and above in terms of micro F-score, macro F-score, micro area under a ROC curve, and macro area under a ROC curve, all of these are all well-known statistical measures. ROC curves are generic solutions, and so we had to use both macro and micro statistics, since there were significant skews in the sizes of the samples. Using

random forests, we got, on aggregate, scores across all four measures that exceeded 0.9, and in many cases 0.95 and even 0.97.

Another thing that we did was to define a very simple score called "the composite performance of an algorithm", with a range between 0 and 1. The score of an algorithm is set to 1 if it is the best algorithm, according to a given measure. If an algorithm is best in all four measures, its composite performance score is 4. That is how we came up with random forests, which gave us the best results in all four metrics. The composite score of other algorithms are their accuracy according to a given measure, divided by the best accuracy according to any of the other algorithms in that measure. We found out that decision trees are doing pretty well with a score of 3.92, but other algorithms did not do as well.

An important question to ask is: why haven't we stopped there, when we had 0.9+ results on four measures with a random forest alone? The answer is the constraint I mentioned earlier, the fact that we were looking only at families containing 10 and more samples. I should also say that most of the competing papers and prior work in this field only looked at families that were bigger, usually with a sample size of 100 or more. We are looking at very small families, and most of our data is about very small families. What we discovered is that according to micro F-score, macro F-score, micro AUC and macro AUC, the big families are getting very good results when trying to detect them by classifiers. However, as the families get smaller, we get less and less accuracy. Singleton families, of course, cannot be detected by classifiers at all, because you cannot train for them. In order to address this problem, we took a different approach.

We know that in other fields, unsupervised clustering works pretty well in order to identify small families, so we tried that as well. The results, it turned out, were quite good. Looking at measures such as K-Means and DB scans, we got consistent positive results for small families. At this point, we used a technique we used before during the DARPA Twitter Bot Challenge of 2015, which led to our team winning the competition. Back then, because we only had 28 days, and we were not limited by the constraints of writing an academic paper, and because we wanted to win, we had to improvise a solution to identify

bots in a situation where we have no training data. This is, of course, analogous to the situation we have here with the malware families. What we did back then was to try an ad-hoc solution of mixing both classification and clustering, and now we wanted to see if we can repeat the experiment in a more formalized and clear way. We have already proven that we can get good results with classification for large families, and good results with clustering for small families, and now we just needed to find a way to combine these two techniques together. Eventually, we came up with a very simple algorithm we call the Ensemble Clustering & Classification, or EC2 for short. Explaining the full algorithm is more complicated than the constraints of this article, and so I will demonstrate it using a few examples.

Imagine we have a bunch of six objects, or types of malware, and two classes. We start out by using versions of classifiers which don't classify per se, but rather give you a probability of membership in a particular class. And so, in our example, this tells us that Object 1 has a 0% chance of being in Class 1, and a 20% chance of being in Class 2. Object 2 has a 60% of being in Class 1, and 40% for Class 2, and so on. After getting the probabilities for all objects being in all classes, we use a cutoff, let's say 0.7, which we mark as $\delta 1$. If the classifiers tell us that a certain object has a probability greater than the cutoff threshold to be in a certain class, then we allocate it to that class. At the end of this process we get some objects allocated to certain classes, but we still remain with questions marks, as we don't know what to do with objects that didn't pass the threshold for any class.

At the same time we do clustering. What this tells us, in our example, is that objects 1, 2 and 3 all belong to the same cluster (marked as Group 1), objects 5 and 6 belong to the same cluster (Group 2), and object 4 belongs to a cluster of its own (Group 3). Now we can leverage the clustering in order to refine the classification results, and use another cutoff, in this case 0.6. If 60% of objects in the same cluster actually belong in a known class, let's say Class 1, then in fact we can safely assume that all other objects in the cluster also belong in Class 1. What happens is that we now have a highly accurate grouping, where Objects 1, 2 and 3 belong in Class 1, Objects 5 and 6 belong in Class

2, and Object 4 does not belong in either, which are the results we were hoping for.

We compared our algorithm with several others in the academic literature, commonly used for classifying malwares into families. When used with small data sets, they all get relatively weak accuracy, around 0.5 or 0.6, while we consistently get 0.7 with EC2, using random forests and DB scans. Obviously, the process is more complicated then what I described above, as there is much hyper-parameter optimization involved, setting the values of $\delta 1$ and $\delta 2$ right, deciding what kind of combinations of algorithms work and more. That said, according to all measures, we do significantly better than the past work on classifying Android malware into families.

More importantly, EC2 is able to detect singleton families. Because of our judicious use of a combination of clustering and classification, we are able to detect singleton families with a precision of 43% and a recall of 31%. These numbers may seem pretty bad, compared to random precision and recall statistics, but what it actually says is that in 43% of the cases, when we say that a certain malware belongs to a new malware family, previously unknown, we are right. It also says that we are able to get almost 1 out of 3 new families that are emerging, and that is a huge deal, and an incredible improvement for cyber security firms.

To summarize, I will try to show how our algorithm works on a real-world malware, for example the Geinimi malware. This specific malware is an exfiltration malware, stealing personal data from the device and sending it to a Command & Control server somewhere else. We are able to say that because a certain malware has a crypto feature that behaves in a certain way, there is a good chance that it belongs to the Geinimi family. However, this is not good enough on its own, and so we also look for the existence and behavior of other features of the malware, such as the fact that it takes control over the Google keyboard. We also take into consideration the absence of features that we would have expected to find, were this malware a part of the Geinimi. All of these together, run through both the classification and the clustering algorithms as described above, give us a fairly accurate knowledge on whether or not a specific malware

belongs to a known family, or is it an emerging new one. This has important applications in the real world, but more research is still required before it can be commercialized.

Secure Distributed Computation

Prof. Jonathan Katz, Professor of Computer Science, University of Maryland; Director, Maryland Cybersecurity Center, USA

I am a cryptographer, but this discussion is not going to be a technical one. Rather, what I want to do is to discuss a particular class of cryptographic protocols that my group and I been have been working on for the past several years. Without going into any low-level details, I will present an idea of what this technology can do. I also want to share some recent results regarding the improvement in efficiency and scalability we were able to achieve over the past three to five years. Specifically, I would like to discuss techniques for secure distributed computation, and to start off I want to show you how it applies to learning from Big Data. We are collecting more and more data, and we are just beginning to explore how much we can learn from it. However, there is a fundamental tension that is rather obvious, but yet most of the time people dismiss it when they are focusing on the learning itself; there is a tension between the need to collect data from multiple sources, and maintaining privacy.

Suppose we have two data collections, reordered on different servers or under different agreements, and perhaps even from different users who have different expectations of privacy or trust relationships with these entities. If we want to maintain maximum privacy, what we would do is to put a firewall around each of these computers, don't let any of the data out, and don't let anyone external utilize the data. On the other hand, if you want to get the maximum utility from this data, learn from the union of the two data sets to glean as much as possible from this collection of data sources, then it is the opposite; what we want to do is to completely share this data, and run algorithms on the total collection of data to see what we can obtain.

In general, resolving this tension would be much easier if there was somebody we could trust, a central authority or a trusted third party, that we can all agree upon and trust to hold our data. In that case, the solution would be easy. What we would do is to take this

collection of data sources, as many as we like, and everybody could send the data to the central authority. Then, the central authority could run the different machine learning algorithms on the data, compute the desired results, send them back to each of the individual parties, and then securely delete the data. That would give us both the utility that we hope to get by learning from this collection of data, while at the same time it would guarantee the privacy of the individual data sources.

We can imagine several applications for this technology, and the one that people often talk about is homeland security. Suppose we have a terrorist no-fly list, maintained by the Department of Homeland Security, and also a list of passengers manifests for all passengers who are boarding airlines in a particular day. What we could do is to use the central authority to check if any of the passengers of the second list also appears on the no-fly list. The to use a trusted party is that the Department of Homeland Security may not want to publicly share their no-fly list with the airlines. This could contain, or be derived from, classified information. At the same time, we would also like to maintain, as much as possible, the privacy of the passengers flying on the airlines. You can also add an additional party. For example, the Electronic Frontier Foundation (EFF) can make sure that the computation carried out by the trusted party is following some policy constrains that have been agreed upon in advance, like making sure that nobody stays on the no-fly list unless there is valid evidence for putting them there in in the first place. Naturally, there are other applications we can imagine for this technique.

The fundamental problem is, obviously, that there is no entity we all agree upon to act as our trusted party. When I gave a talk about this topic a year and half ago, the only example I could come up was, maybe, the Supreme Court, but even that is only relevant to the US, and nowadays I'm not even sure about that. Perhaps many people trust Google, because if you are using any of their apps you're essentially sharing a great deal of your information with them. Nevertheless, we may not be willing to trust Google with other information, such as medical information. Also, there may be certain classes of information, and certainly governmental classified information, which cannot be

shared with Google. This means that Google can't really serve as a trusted party, even though they, de facto, do serve as a semi-trusted party for many purposes.

This problem even intensifies when we talk about interactions among people in different countries. Even if we can imagine there may be an institution or some government authority that people within a country might trust, it is very difficult to imagine an authority that people from different countries will be willing to agree upon and trust. I also argue that it is unlikely that there will ever be a trusted party that can serve in this role for general computations. There are many reasons for this, and one that is particularly interesting to think about is that it is not really an economical business model. It doesn't really make sense to imagine a company serving as a trusted party on behalf of individuals, because I think that the amount of money people would be willing to pay for that service is limited, the cost to the company in order to ensure the security of the data they are obtaining is very high, and the liability is high as well. Lastly, a central authority would also introduce a central point of attack, which would be particularly inviting for an adversary to try and get the data from. Even if you could trust them to do their best to maintain the security of the data, it will still be very difficult for them to withstand a targeted attack.

In theory, it would be nice if there was someone we could trust with all of our data, but in fact it will be even better if we could avoid the need for trust in the first place. In that case, what we could do is to simply replace this model of a trusted party, where all the other entities send their data to, with a distributed protocol that the parties themselves could run, and that ends up having the same effect. This is exactly what a class of protocols called Secure Computation Protocols guarantee. A distributed protocol run by several parties is defined as a secure computation protocol, if an execution of this protocol provides an exact emulation of what you would obtain if you indeed had that central trusted party that everybody could rely on. A secure computation protocol guarantees all the properties that you would like to have in such a distrusted computation, and in particular guarantee that the privacy of each party is maintained, up to the point of revealing the final result of the computation.

This system also guarantees integrity, that the correct result is being computed from the data, and that no party manipulated the outcome. It ensures availability, meaning that no party can abort the computation early and prevent other entities from receiving the output of the computation. In this sense, it also protects you against a certain class of Denial of Service attacks. It can also guarantee properties that you wouldn't necessarily even think about in advance. One example is that secure computation protocols also guarantee input independence, which means that no party can manipulate their data to depend upon other parties' inputs. Just like in the model with the central trusted party, where we have all the entities sending their data independently to that party, party A can't make their inputs depend on party B's inputs.

However, there are some things we have to be aware of when using these protocols. The first is that there are assumptions being made about the behavior of corrupted parties. We need to be concerned about the fact that we have a collection of parties running this protocol, but some of the entities may either be corrupted themselves or become corrupted by an external attacker during the course of the computation. The other question that we need to ask ourselves is what is the total numbers of the corrupted parties. For example, if we have a protocol running among 20 parties, you might have a protocol which is secure, under the assumption that, at most, two of the parties are corrupted and the other 18 are honest. Or, if you are extremely pessimistic, you may want to take a worst case assumption, and assume that all the other 19 parties are corrupted and you are the only one that wasn't hacked. In that case, you would want to protect yourself against malicious behavior from the others parties. Different protocols have different trade-offs between the efficiency of the protocol and the number of corruptions that can be tolerated.

I would like to mention briefly two main threat models that have been considered regarding this problem. The first is called "semi-honest", where the assumption is that the other people running the protocol are indeed following the protocol as described, but the concern is that they might learn something sensitive from the execution of protocol, based on the messages that are being exchanged. This assumption is

realistic in settings where you actually have trusted parties running the protocol. For example, if you have different government agencies running the protocol together, then it may be reasonable to assume that all of them are trusted. However, there still may be legal or policy constraints about what data can be shared between these different entities, so you may want to run a secure computation protocol anyway. There are also scenarios where you can use other techniques, such as software attestation or auditing after the fact, to ensure that people are running the correct software. This means that you can be fairly certain that all parties are running the correct software, but there still might be dangerous consequences.

The strongest threat model is the one called the malicious model, where the protocol protects against arbitrary behavior of all other parties. There are no assumptions made on what software they are running, and they can arbitrary change the software and the code of the protocol, but as long as *you* are running the correct copy of the protocol, you are guaranteed to be secure against arbitrary behavior of the others parties. Once again, there are trade-offs here between what kind of adversary behavior you want to adjust for, and the efficiency of the model.

It sounds quite amazing that such a secure computation protocol can exist, but in fact it has been known for about 30 years now that such things are possible. Actually, anything that you might hope to achieve in this model is doable, at least from a feasibility point of view: a secure computation of any program or any function that you like, which protects you against arbitrary malicious behavior of some or even all the others parties, is possible. The main question, however, has been the efficiency. For a long time, the general perception was that using generic secure computation will be inefficient, and that it had no chance of ever making its way into practice. There are several reasons for that, but maybe the most obvious one is that in order to run these protocols, we have to take the program we are interested in executing and map it to a boolean circuit, computing the same function. The problem is that even a very small and simple program can turn into a very large boolean circuit, and the efficiency of the

protocol often depends on the size of that circuit. If you just look at that, you immediately think that there is no hope.

However, it turns out that there has been steady progress over the past fifteen years. The first implementation of a secure, two-party, semi-honest computation protocol was the Fair Play protocol, published in 2004, but the initial results were somewhat disappointing. The number of boolean gates that could be processed in a second was very low. In 2011, however, in a joint work with collaborators of mine from the University of Indiana, we managed to present techniques that can dramatically improve both the efficiency and the scalability of secure computation. We managed to process a billion gate circuits at a rate of 10 microseconds per gate, which was a dramatic improvement over prior work. This was the first time we were able to show that semi-honest secure computation can, in fact, be efficient, and can be implemented and usable, at least for average sized circuits.

As I mentioned before, the semi-honest setting is a very basic model, where you assume that most parties are following the protocol. Because of that, since 2011 other researchers have turned their focus towards the malicious setting, which is much harder to deal with. Here, too, we have seen dramatic improvements in the efficiency of protocols over time. For example, we measured the time to compute an AES cypher, where one party has the data to be encrypted and the other party has the key. We have seen exponential improvement over time, from 1,000,000ms in 2011 to just under 100ms in 2016, which was the result of the work done by several students of mine. Naturally, we also want to solve these problem for multi-party settings, and in that area we have developed an exciting new protocol for the malicious model, which handles an unlimited number of corruptions and any number of participating parties.

What we were able to show, for the first time, is a global scale protocol for a secure computation of AES, where we had a distributed computation run among 128 parties across five continents, with an online time of only two seconds. Of course, two seconds is much longer than the time it takes to compute AES locally, but nevertheless, this is a good indication of where secure computation is heading, and what it can potentially do.

This is not only an academic interest, there are several companies that are trying to commercialize and use this technology. Today Google, for example, is using two-party computation techniques in some of the things they are doing. Other companies have products that focus on secure computation, such as Sharemind, which is an Estonian company, Dyadic, which has one founder in Israel, and Caltopia, another European company. We know that the US government is very interested in this, and that there are a couple of DARPA and IARPA programs focused on applications of secure computation for processing classified data.

To conclude, there have been tremendous advances in this field in recent years, which bring the technology of secure computation much closer to reality. Also, we see there is recent interest in this field in the industry, getting this technology into the real world. I hope and expect that we will see greater deployment of these techniques in the near future. The code that we have developed is available for anyone to download, and it allows you to express, in Java, the program that you want to securely compute. You don't need to know anything about the underlying cryptology, and you can generate a secure protocol to compute your program of interest. We also published the papers explaining the technical details of how everything works, and you can find them online as well.

A Layered Approach to Secure Mobile Computing

Robert Deng, AXA Chair Professor of Cybersecurity and Director of the Secure Mobile Center, Singapore Management University (SMU), Singapore

The Secure Mobile Center, which I head, was set up two years ago, supported by the Singapore National Research Foundation. In this article, I would like to give a brief overview on the four major research projects we currently have, working on the different layers of the mobile computing security. Two of these projects are in the field of mobile platform security, one in mobile application security, and the last in mobile internet service security.

The first project aims to fortify mobile platforms with a user-centric trust anchor. If you look at mobile operating systems, they are huge, and because of that we know they will always have vulnerabilities. The purpose of the project is to design and implement a trust anchor, and we use that for security purposes. It is a lightweight solution, similar to a hypervisor, but not for any function other than security. This trust anchor will leverage the hardware features, such as ARM's Trust Zone and virtualization extension, and the TPM in the PC environment. We are essentially building an extension of the root of trust. Within the root of trust, we know we have security we can trust, via the hardware component. However, once you extend beyond the root of trust, we want to have additional extensions, and this is exactly what the hypervisor is. It is designed to support three security features. One is a fully isolated execution environment, meaning it is isolated from other processes and applications, where one can have security critical operations carried out. The second aspect is presence attestation, so that users can make sure that their hardware is functioning as expected and not corrupted. The third is a secure user interface. So far we have finished the design of the fully isolated execution environment, and we are still working on the other two.

Many previous researches have been done in the academia in regards to an isolated environment, using hardware-based memory isolation. However, most of those existing techniques are not adequate, not secure enough, when the platform is multi-core. They work well only in a single-core environment, but they fail in multi-core, because of the incomplete isolation boundary. In a paper we published in 2017 during the European Symposium on Security and Privacy, we showed a number of attacks on the existing techniques. The other approach of isolation is the so-called "domain isolation", which takes place, more or less, in the Cloud. The idea is that you isolate the virtual machine, including the operating system. You assume the operating system is trusted, and because of that, this solution has a very large trusted computing base (TCB), which is not ideal. Our approach is different. We only isolate a certain memory region, a CPU core, and sometimes an I/O device. The design is simple, because the code is very small. It is not reliant on specific hardware, and it is not reliant on the operating system itself. It is secure because it gives us a minimal attack surface; the physical resources which are within the isolated region, including memory, core, cache, and paging structures, all of those are completely isolated from the guest operating systems. We provide a small TCB, and in addition, our solution supports I/O devices.

The second project is about secure and usable authentication systems in mobile platforms. So far, we have two designs for different kinds of platforms. The first one, FaceLive, is a solution that you can use in conjunction with your face-based authentication system. It makes sure that there is a real face in front of the mobile device, and not a photo. This is a known weakness for some of the current face-based authentication solutions, and we have solved it. The second solution is a leakage-resilient password entry mechanism for Google Glass. We designed it just recently, and presented it in Asia CCS 2017. We created a way to authenticate yourself to Google Glass, without the help of a mobile device or a laptop computer.

I am going to focus on the first solution, FaceLive, and present the high level ideas of the design. There are already certain similar solutions in the market, which are designed to detect liveliness of the face. For example, when you try to authenticate, some systems

will detect the blinking of the eye, or a rotation of the head. These solutions are indeed resistant to photo-based spoofing attacks, but if you put a video in front of the device, they fail to detect the difference. We observed that during the authentication process, you move your mobile device in front of your face to a certain distance, maybe 30-40cm. During that process, in our solution, the camera will capture different views of the face, and the built-in sensors in the mobile device recapture the movement information for the mobile device. Our solution takes these two sets of information, and makes sure that they match, in order for the user to be authenticated. If we just show a video, it will not be accepted because the mobile device is not moving. We conducted a comprehensive user study of the solution to assess its performance, 73 participants. The equal error rate (ERR) is quite good, 4.7%, and the solution is quite robust.

Our third project is about three aspects of the mobile malware. The first aspect is to study what are the factors that affect mobile malware propagation among users. We conducted a user study, where we designed an app to let users customize their message, to select a media of communication, and then we monitored the recipients' behavior. The second aspect is to detect mobile malware. Our approach is to build behavior models for the apps. So far we have built a behavioral model for 10,000 apps in the Android Play store, and we put them into our database. During the execution of the app, if its behavior ever deviates from the model, that is an indication of malicious behavior. The third aspect is to contain malware damage, because not every malware can be detected. We automatically split running apps into two processes. One is the process in the native environment, say C++, which is vulnerable. The other process runs in the JAVA environment, which is much more secure, and most of the security critical applications are executed in the JAVA environment. The process splitting is automatic, and, of course, for the app to function, the two processes must be able to communicate with each other. We need to build a manager or firewall in between the processes, to regulate their communications, which is still a work in progress.

The fourth project is a system for scalable access control for encrypted data and trusted servers. We assume that the service provider

is honest but curious. Honest in the sense that it will keep the data and provide you with the service. You upload data to the Cloud, expecting you can access the data later. Curious, in the sense that we do not trust the service provider to keep our data confidential. The second part is that we do not even trust the service provider to enforce access control correctly. We believe that access control must be enforced by the users themselves, not by the service provider. Given these assumptions, the objective of our solution is to encrypt the data before it is uploaded, so we can have end-to-end data security and privacy, and at the same time the user must be able to access the data, to search over the data, and, maybe in the future, to conduct data analytics and to perform secure computations on the data.

The basic solution is based on Cyphertext-Policy Attribute-Based Encryption (CPABE). This is a technique invented by Shahai and Waters, some time ago. In CPABE, each user have a set of attributes. For example, you have your identification number, your position, and your affiliation. There is a key generation center, which generates decryption keys for each user based on what kind of attributes they have. The data owner will encrypt the data, after which they will attach an access policy to the data; for example, only allow access and decryption to the user holding a certain identification number, or if the user is working for a certain hospital, or is a cardiologist. This solution has a number of advantages. First of all, it has a one-to-many public key encryption. If you have one message, you just need to encrypt it one time, and many people can access it. As long as the recipient's attributes satisfy the access policy, they can access the data. The control policy can be expressed in a boolean algebraic expression, and more importantly, the access control is built into the message, it is not provided by the service provider.

This, as I explained, already exists. We only made the basic solution practical. One of the issues we solved is that while the data in ABE solutions is encrypted, the access policy itself may leak a lot of sensitive information. In the hospital example mentioned above, anyone holding the encrypted message can now know that a certain patient, allowed to access the data, is suffering from a heart condition, because except for them, only cardiologists are allowed to access it. We improved

the basic design, and came up with an ABE solution equipped with a privacy-preserving access control policy. The second contribution is to also allow users to perform keyword search on encrypted data. Users can now go to the Cloud or the storage space, search and retrieve the cyphertext they are looking for, while all the data is still encrypted and while preserving the access control. The third contribution is a verifiable outsourced decryption of ABE. If the access policy is complicated, the decryption process can be complicated as well. This is especially true for mobile devices, as it will consume a lot of energy. We outsource the decryption process to the Cloud, but without disclosing any information to the Cloud, and the user is able to verify that the decryption is correct. Our last contribution is to solve the very challenging problem of user revocation. In our design, we have an efficient user and attribute revocation technique, delegated to the Cloud, and eliminating many of the issues. Once you are issued a key to perform decryption, you don't have to have any communication with the key generation center again. All revocation operations are handled by the Cloud system, so if a certain user's access is revoked, they cannot access the data. This is done without the public server disclosing any sensitive information.

Based on all of these, we designed a practical system, which is an attribute-based secure messaging system, in the Cloud. This allows, for example, secure group chat, and supports all the features I mentioned previously. It has been published, and is available for usage to all.

Attacking the 3D Printing Process for Critical Components

Prof. Yuval Elovici, Director of Deutsche Telekom Laboratories; Head of the BGU Cyber Security Research Center; Research Director of the iTrust of SUTD

My topic of discussion in this article is somewhat interdisciplinary, and relatively very disruptive. Three years ago, we started to work in a new field called "how to secure additive manufacturing". First of all, I would like to demonstrate that the challenge indeed exists, and then I will present several mitigation technologies that we are currently developing.

First of all, we need to establish what "additive manufacturing" means. When we say additive manufacturing, we actually talk about 3D printers, which take existing materials, and by adding the materials to themselves, can generate new objects. The actual process of generating such objects is very complex, but how does this relate to cyber security? Over the last years, it was approved to 3D-print critical components that are being used, for example, by the aviation industry. If cyber attackers will be able to intervene in the process of creating 3D objects, they will be able to insert a void that eventually may create a critical infrastructure to collapse, for example, an airplane.

What is the process that involves additive manufacturing? Usually, some mechanical engineer takes a CAD tool, and designs a new type of object. Then, the object is converted to G-Code, containing instructions to a specific 3D printer, which is supposed to take the instructions and convert them into a printed object. When we analyze the additive manufacture workflows, there are many opportunities for attacks. For example, an attacker may attack the designer's computer, or even the 3D printer itself. There are, in fact, multiple opportunities to damage the process, and to create a situation in which the printed component include such a big defect that it may influence the operation of the system that is going to use it in the future.

We did a very detailed systematic analysis of how attackers may interfere in an additive manufacturing process. We started by identifying the attack vectors: the software, the hardware, all of the compromised elements that may be attacked by the attacker. It could be the controller PC, it could be the network, or the printer itself. There are also the influenced elements: the files that I want to print or the code that I am going to send to the printer, and so on. We tried to examine what are the potential influences on the process itself, and what is the attack target.

We realized very quickly that the understanding of the impact was not clear in our community, so we decided to do a PoC, and to demonstrate what attackers may do if they are going to harm the additive manufacturing process. We decided, from the point of view of the attack vector, to use a very simple phishing attack. We targeted the controller PC, and we wanted to modify the target element, influence elements in the file that the user wants to print in order to create a critical component. We did this by inserting gaps, which I will explain further on. We picked a DJI drone and chose a design of a drone blade as a potential attack target. We could not find the design file of this specific propeller, so we 3D printed it, to obtain a fake propeller that actually works. Then, we added a defect – we added a gap of 0.1mm in a specific joint, which is very difficult to identify. In order for the propeller not to break immediately, we added two pins that connect (although very weakly) the blade to the main part of the propeller.

The full process was fairly simple. We emulated a phishing attack against the users that are supposed to print the blade, in order to be used on a drone. In this case, we used a zero-day vulnerability in WinZip, to demonstrate the attack. We sent the victim an email, and the only goal was to gain access to the computer of the user who is supposed to print the object itself. The payload connected to a C&C server, which sent the proper commands to give us access to the controller PC of the user. Once we got in, we looked for the design file of the propeller, and when we found it, the attacker – meaning us – downloaded the design file, and "improved" this design by adding some kind of a void. Then we uploaded the modified version of the blade back to the user's PC. At this point, the victim came to work in the morning and wanted to print the new object. In their eyes and

mind, nothing had happened, and they don't have anything to fear. The user prepared the blade for the printing, as usual, and printed the blade. At no point did the victim see anything strange, not even in the final result. When the user tested the drone, exactly one minute and forty seconds into the test, the blade broke, and the drone crashed. What we did was, in fact, to design a fatigue attack. Our target was not the blade, but the drone it was attached to.

What we did in this first part of our research was to demonstrate the process of attacking an additive manufacturing process. What we are currently working on is to try and automate the process in a way that doesn't require us to download the file in order to compromise it. We want to write a malware that is going to look for STL design files, look for weak spots, and add voids completely automatically. We believe it can be done, and we are currently working on it. But what about mitigation?

Usually, when cyber attackers attack a specific network, we assume that they can compromise any part of this network. Therefore, we looked for an out-of-band solution that will verify that whatever we are printing is exactly as the designer designed it. We developed a new method which we called "audio fingerprint". This method consists of taking a mobile phone, placing it close to the printer, and listening to the printing process. The 3D printer that we use has three motors that are generating a lot of noise. We are able to listen to that noise, which is the printing process, and generate an audio signature of the specific printed object, under the assumption that there is no current cyber-attack. This is based on the audio fingerprint reference model. Later on, during the detection phase, we place a mobile phone by the printer once again and listen to the printing of the object. If there is even a slight deviation, we immediately instruct the user to stop the printing, because somebody managed somehow to execute a cyber-attack. We don't know what type of attack this might be; it could be on the design file, it could be on the firmware of the printer, but the important point is that whatever result you are going to get is not exactly what the original designer planned. This signature is very complex. We took a time window and conducted a Fast Fourier Transform (FFT), then

used the PCA to find the most important frequencies, and finally we created a signature for a moving window along the printed object.

The idea behind this method is that once you download an object design from the web, and there is a risk that somebody modified it, you are going to also download the audio fingerprint, so that you can compare it to the actual printing process. This will allow you to verify that nobody managed to execute a cyber-attack in order to make you print something that eventually is going to fail once it is embedded inside a real system. We did some evaluations for this approach and it works very well. We can detect any deviation of less than one second on the printed object. We believe that 3D printing will be a critical target in the near future, and we hope that our research will inspire and aid in finding the ways to mitigate this threat.

The Good, the Bad and the Ugly

Prof. Tal Zarsky, Vice Dean of the Faculty of Law,
University of Haifa, Israel

I am going to address an aspect pertaining to privacy, data protection, and autonomy in the age of cyber, and that is automated decisions that impact humans, based on personal data. What are the elements of these dynamics, and how are they relevant to cyber? I will discuss those, as well as the current law, the roles of law and policy, the theoretical issues arising from it, as well as solutions and prospects. I will provide some examples, which will illustrate that we are confronting such automated situations and dynamics throughout our lives, and that they exist in almost all realms. We will see more and more of them as time goes by, and we need to deeply understand the problems involved, both real and pretended problems, and only then we may move on to solutions.

In my first set of examples, algorithms crunch personal data and provide recommendations. One popular example pertains to the justice system, the COMPAS system in the US, which makes recommendations as to the sentence which should be rendered in a given case. The system can give recommendations regarding the extent of the sentence, meaning perhaps if a sentence should be limited, and even gives recommendations as to the actual judgment in criminal cases. Beyond the criminal world, we see decisions made automatically in terms of credit allocation. We see many applications and new companies online, trying to circumvent the existing credit models. Unilever, for instance, is now recruiting more and more workers through a mostly automated process, which is driven by various games and algorithms, analyzing and crunching the data of applicants. We also see this discussion in a very different realm, the realm of military operations, e.g., the use of drones, as well as other unarmed vehicles and weapons. These devices are making automated decisions in various aspects: navigating, moving towards targets, identifying a target, and moving in to destroy a target. We see this in

the security realm and in the intelligence realm, so we can, of course, see how this relates to this notion of cyber.

Given the speed and complexity of cyber issues, sometimes we use automated processes in defense and in offense. If we see this topic of cyber broadly, like something that involves analyzing data pertaining to personal information and making meaningful decisions regarding individuals, we can clearly see the link. We will see expansions of such uses everywhere. There is a big dilemma in Israel regarding healthcare and hospitals, and I believe that we will be seeing many more recommendation systems in hospitals, assisting doctors and other medical staff. Other elements we will see in banking, in finance, and in education.

How does this pertain to law? We see it both in the regulation and in the courts. Regarding the COMPAS system, there was a famous case, Loomis vs. Wisconsin, which went all the way up to the Wisconsin Supreme Court. In this case, there was an argument that this system does not provide due process, due to its limited transparency. We also see it in "title 7" cases in the US. There are questions that arise as to the outcomes of this system, leading to disparate impact. Are they discriminating? Are they discriminating in employment? Are they discriminating in housing? These are issues that could be addressed in the US, and in Europe there are other directives and regulations which even provide for additional legal arguments in courts.

As for regulation, many of you probably heard about GDPR, the General Data Protection Regulation, which will come into force on May, 2018. Article 22 of the GDPR pertains precisely to this issue, to automated decisions made entirely by a machine, have a substantial impact on an individual, and rely on personal information. Article 22 provides various rights to individuals, not to be subjected to such process, and even if exceptions do apply, they will have the right to receive the "logic" that this process entails. As you can see, there are many legal issues, going through various systems, and we will have many discussions on this issue in the years to come.

When addressing this issue, first we need to try and see what this matter encompasses. We are talking about instances where there is a meaningful impact on an individual, and we need to ask ourselves:

what might that be? Obviously, if it is a drone, trying to decide if someone is a target or merely an innocent bystander, or a system that is going to decide upon the extent of credit to be afforded, or if it is about employment, or even things related to medical treatment or insurance, it is probably meaningful. But what about an automated ad on Google, offering to check if you have a criminal record based solely on your name, but it is targeted to people with African-American sounding names? Is that meaningful? Maybe, maybe not. We have to decide upon this issue.

The second issue is the extent of the automation. Article 22 of the GDPR is very easy to sidestep, because it only applies to decisions that are entirely automated. If there is a human involved in the loop, the problem can be sidestepped. In the Loomis case, the Wisconsin Supreme Court was discussing a system that merely made a recommendation to an individual. That said, we all know how strong recommendations from computers are. We all have been to the airport and told the agent we made a booking, only to have the agent say that they are sorry, but it is not in the system. We have all heard somebody telling us: "it's not me, it's the computer." To err is human, but there is a strong notion that computers can't really be wrong. The recommendation of a computer has a very powerful suggestion on an individual, and the Wisconsin court understood this. They said that perhaps, if there is a very substantial impact to the automated process, it should be overseen and reviewed by a human.

The last element here is the use of personal information. That is actually a legal hook, that allowed the GDPR to come into action. However, in many instances, an automated decision could have a substantial impact on a person, even if no personal data is involved. We have seen cases where an attacker manages to disrupt the automated manufacturing process of a drone, which could then randomly fall on someone's head. There might be a problem here, it does not involve personal data, but still this could have substantial effects. Shouldn't we have laws on this issue? This is something that the European committee is currently studying.

When we think about this issue on a theoretical level, there is the good, the bad and the ugly. Why is this good? This is very simple,

and some might say naïve. A process that is fully automated is highly efficient. It allows us to grasp all the analytics. It allows us to go beyond the boundaries of the legal human mind, which has a problem correctly consulting all the needed sources and then bring them all together, and come out with a quick response. Also, it could potentially limit the ills of human decisions. The individual that wrote the COMPAS software which was used by the Wisconsin judicial system, was a criminologist. He was a professor that wanted to use evidence-based systems; he wanted to only look at evidence when deciding on rulings. Individuals have problems, they have biases. In many cases these are unintentional, but a computer system only does what it is told, always. Therefore, many of these human ills, which especially impact minorities, could potentially be limited and even sidestepped.

What is bad about the system? Loomis, arguing before the Wisconsin Supreme Court, said that this was unacceptable because there is no due process. We know data is coming in, we know data is coming out, but inside it is a "black box", we don't know what is happening inside. Loomis wants to know what is happening, and because he does not know, there is a problem with due process. There is a demand for additional transparency. Actually, it is not really true that we don't know what is happening, because based on what is coming in and what is coming out, we could potentially understand what is happening inside. Also, we must remember, we have the same problem with the judge; we know what is coming in, we know what is coming out, but we don't know what is happening in his brain, and it could be affected by many things, such as what he had for breakfast, if he fought with his spouse that day, what is happening in internal politics, and so on. We have to think about these differences, and about the alternatives.

Other problems we might have are errors in the system. All systems have errors, but we might have systematic errors, which we need to be concerned with. For example, systematic errors that lead to biases towards the weaker parts of society. We might have errors resulting from the data that is being fed into the system. The data itself might have been biased from years of biases in the way data is collected. There might be errors which result from the computation process, of structuring the algorithm, and once again, these might be intentional or

unintentional. In many instances, the people who write the algorithms are not skilled in understanding and interoperating legal rules. There are projects today about teaching engineers to understand the concepts of law, exactly for this reason. But there are other sets of even deeper problems. There might be problems in a machine learning context, where the system learns the biases from the individual. There was an article in *Science Magazine* a few months ago, that demonstrated exactly how decision trees, which were learned from individuals, started mimicking the same racial understanding that individuals have, associating specific words with others.

Another aspect, which is a cyber aspect, is that these systems could also be attacked. An Artificial Intelligence system could be attacked, it can be fed various inputs and in that way the outputs can be influenced. Of course, we all know that humans can be compromised as well, we just don't call it attack. Humans can be manipulated just the same as a machine, and so we need to acknowledge that machines, as well as humans, have their problems and their errors. In the Loomis case in Wisconsin, indeed the justices decided that the system could be used, but a long list of disclosures should be presented to the judges, which should make them aware of the fact that machines make mistakes, that it is merely a statistical finding, that this system and its outcome might be biased towards minorities, and so on.

We mentioned the good and we mentioned the bad, now it is time to talk about the ugly. The public generally fears automation. It fears that its jobs might be taken over by robots, and this is something we see often in popular culture. If you think that lawyers or even law professors are protected, you should know that symphonies are currently being written by machine learning. There is probably no real reason why law review articles should not be written by machines, and they will probably be even better than many of the things that they are reading now. Coming back to the public, it fears automated machines, and so it starts imagining problems.

A very popular imagined problem is that decisions made in an automated process lack morality, dignity and mercy, and I would like to present a few responses to that. First of all, a response might be that we could teach the machine to engage in these so-called human

traits. There are currently applications in which you feed the system your moral standings, could be religious standings as well, and then, when faced in various situations, the system makes recommendations based on the moral standing that you provided. In short, this doesn't really seem to be a problem. Regarding dignity and mercy, indeed the system might not be merciful, but perhaps we need to think, what does mercy actually mean? We might look at it as a codeword for having an ability in the system to treat those that are perhaps in a higher status, those that are closer to the decision makers, in a more lenient way? In other words, doesn't mercy actually means: "he deserves mercy because he is more like us, unlike the others?" The computerized automated system would not do that, but perhaps that is actually a good thing, something we should aspire to.

To summarize, we generally see three strategies to dealing with current issues with automated decision making machines. One strategy is to try and limit the reach of such automated systems in various instances, and assure that we have humans maintaining their position. Another strategy would be breaking the black box, helping us laymen to understand what the process is, and enabling us to exercise due process. We have to be skeptical about this strategy, because we have so much information already. We are seeing a stream of studies in health law, contract law, and other realms, where you see that all these additional disclosures do is take space on our screen or destroying the rain forests. We, humans, don't really do anything with all these words floating around. Another way to deal with it, as we saw, is to provide more information to the decision makers, letting them know that these automated systems have problems. Current studies are showing that this approach doesn't really help as well.

How can we move forward? First of all, we need to strive to distinguish between the good, the bad, and the ugly, and especially between the bad and the ugly. We need to understand what is a real concern, what is a pretended one, and on a case-by-case basis, try and limit the problems related to various errors. This will, possibly, allow us to have our cake and eat it too; to benefit from automation, while assuring that we don't have the problems that we discussed.

Managing Law Enforcement Access to Data Across Borders

Prof. Jennifer Daskal, Associate Professor of Law, American University's Washington College of Law, USA

My article is part of a broader project, looking at territoriality, data, and the ways in which data in an inherently unterritorial medium (i.e. the Internet) is being regulated by territory based states. We are looking at the many challenges this poses, and the implications for privacy, security, speech rights, sovereignty, and democracy. I first got interested in this issue several years ago, when I learned about what is known in the United States as the "Microsoft Ireland" case.

This case began back in December of 2013, and it may very soon make it is way up to the United States' Supreme Court. In this case, the United States government sought to compel Microsoft, via a warrant issued by a judge, based on a finding of probable cause, to turn over emails associated with a particular account in a law enforcement investigation. Microsoft refused, saying that the data was stored in Ireland, and that therefore it was an impermissible, extraterritorial exercise of it is warrant authority; the warrant authority only extended to data within the United States. The government insisted that because Microsoft is based in Redmond, Washington, they can access the information from there. The government's claim was that what matters is where a company is based, not where the data is stored, and therefore this is clearly a territorial search, not an extraterritorial one. The magistrate judge and the lowest courts all sided with the government, but when it made its way up to the appellate level, to the Second Circuit, the decision was reversed. The result in the United States, at least in the Second Circuit, is that the United States warrant authority only reaches communication data that happens to be stored within the territorial boundaries of the United States.

This is causing many problems for law enforcement, investigating a local crime with a local victim and a local perpetrator, but where the data or some parts of the data happen to be stored abroad. In these

cases, the United States government can't access it. It has negative consequence, in my view, for privacy. In this case, the government proceeded based on a warrant issued by a judge, based on a finding of probable cause, and it importantly reflects a total mismatch between the rule that was adopted in this case, and the sovereign interest, which the presumption against extraterritoriality that was applied in this case was meant to protect.

From what I can tell, this stands as a blight transmission of rules, governing other forms of tangible and intangible property, to the regulation of data. There have been five magistrate judges, who are the judges who look at the search warrants when they are first requested by the government, who have ruled the other way precisely because some of these concerns. I use this case as an example to highlight the ways in which data is requiring us to reexamine, rethink, and reevaluate what is territorial, what is extraterritorial, and the underpinnings of those assumptions. To that end, I want to make four broad claims. First, data is different in key ways from other forms of tangible and intangible property. These differences matter to our understanding of what is territorial and what is extraterritorial, and thus matter to the enforcement of prescriptive and adjudicated jurisdictional rules, that do and should apply. Simply looking at the jurisdictional rules that have been applied to other forms of property, ultimately doesn't work, and they undermine some of the key normative values that this rules are meant to protect.

Second, this is not the same as saying that territory-based rules do not matter, they do. There was a long standing debate in the 1990s amongst academics. On the one side were unterritorial people, scholars like David Post, who argued that data, because of its unique properties, will defy national regulation and lead, organically, to the creation of "super-national", global structures to govern the new Internet. That, obviously, has not come to pass. On the other side were the territorial scholars, such as Jack Goldsmith and Tim Wu, who insisted that territory-based sovereigns will find a way to assert territory-based controls. They, obviously, were right.

That leads to my third point, the interconnectedness of data. This means that there is, as a result of the above, a new form of international

rule-making via territory-based regulation. To the extent that the Internet remains interconnected, which I hope it will, rather than retreating into our silos, territory-based regulations can and do have broad extraterritorial effects, and those extraterritorial effects need to be mediated and moderated.

Fourth and final, these cross-bordering effects are becoming increasingly regulated, controlled, and mediated, not by governments, but by the private multi-national big corporations that manage our data across international borders. They get to determine which rules apply, how they are interpreted, and how conflicts across borders are being mediated.

The issue of law enforcement access to data across borders is an increasingly important topic. It is a topic of conversation not just with respect to the one United States case that I mentioned. It is actively being debated and discussed in the Council of Europe, it is a topic of a recent year-long study that was released by the European Commission, and it is obviously a topic on the agenda of many other places as well. Briefly, the problem arises when law enforcement investigating a crime in state A, involving local victims, local perpetrators, and local witnesses, discover that either the data or the provider that controls the data, or both, happen to be across an international border. This is an increasingly common phenomenon, and it is also relatively new, deriving from the growth and development of the Cloud.

The "background rule", the rule that was previously applied, and the one that was relied on by the Second Circuit in the Microsoft case, is that under international law, state A cannot unilaterally send its law enforcement officials into state B, without state's B's consent. FBI agents cannot unilaterally cross into Israel and seize data out of the home of somebody in Tel Aviv, doing so would be a clear violation of Israel's sovereignty and a violation of the international law. Instead, the United States would have to submit a request to the Israeli authorities, and the Israeli authorities would, to the extent that they agreed, conduct a search according to Israeli law, and then ultimately send the information back to the United States.

There are, however, good reasons to why the rules governing transmission of data across borders ought to be different then the

rules governing traditional searches and seizures of physical property. First, there are no agents crossing borders in the example I am talking about. Data, as we know, is highly mobile, highly devisable, and it is a highly unstable basis for delimiting jurisdiction. It is often held in places that are completely disconnected from the data user or any other relevant factor, other than the decision of a third party provider to place a server in a particular location, for reasons like efficiency, tax rates, and energy costs. There is often no normative link between where the data happens to be held, and the other legitimate sovereign interest in either accessing data or controlling its access.

The divisibility also means that when either a provider copies data and provides it to law enforcement agents, or if the government does so itself, the requesting agency is not in any way depriving the host state, the state where the data is located, from being able to access it. This is obviously different than the example where FBI agents walk into a home and take out the only diary, the only photobook, the only piece of property, in which case they deprive the host state of the ability to access it. Data can be copied without altering it.

That raises important questions about what are the appropriate jurisdictional hooks that are ought to apply, and I would say that the location of data is probably the worst one. It has substantial negative security costs; it is bad for efficiency; it encourages mandatory data localization rules as a way of controlling and ensuring access to data, in ways that have big efficiency costs, which small start-ups will not be able to manage. In the United States context, I would argue that it is bad for privacy, in the sense that instead of the United States government getting a warrant based on probable cause, the rule effectively requires them to go to foreign governments and ask for their assistance. In many cases, the privacy rules of the foreign government are less protective than the US requirement of a warrant based on probable cause.

On the other hand, the idea that states can simply access anything, anywhere, is also concerning. This was an approach that was adopted by the courts in Belgium, in two cases where the courts said that as long as data is received in Belgium, it is a territorial search, even if both the provider and the data are outside of Belgium. This basically

makes every single request for data territorial, because it will ultimately be received in the jurisdiction that requests it. This, too, is concerning. It obviously serves the law enforcement interest of the Belgium government, but it creates a concerning precedent via which states could unilaterally reach across borders any time they so choose, in order to access data, without regard to anything like baseline privacy rights, or baseline concerns about due process. It also fails to respect what is a legitimate interest of sovereign countries, which is controlling access to the data concerning people within in their country. This is not about controlling the data in their country, but about having control over the individuals within their country. It is about setting limits on when and how foreign governments can access the data of those individuals, rather than the zero's and one's themselves.

As a practical matter, this kind of rule also runs headlong into conflicting, blocking provisions, which are imposed by other states that prohibit disclosure to foreign governments, and therefore puts big companies in the middle of this two competing obligations. I argue that we need a new solution that mediates between these two extremes, which takes into account the various interests at stake. These include, at least, the following: the sovereign interest in security, which translates into the ability to access data relevant to an investigation of serious crime, regardless of the location of the data; the sovereign interest in controlling and setting the rules for data of one's own citizens and residents, and potentially setting limits on the ability of foreign actors to access that data; the interest in privacy protections and promotion of due process, human rights and the rule of law, both as a normative value and also from a more self-interested sovereign perspective. The fact is that all of our data is so inter-mingled, that the rules matter even if one's only interest is to protect one's own citizens and residents, or creating mechanisms for accountability and some sort of transparency.

Coming back to my original point, and to conclude this discussion, the companies, the major multi-corporations that manage our data across borders, make critical decisions about a whole range of rights and security. They make decisions regarding speech rights, privacy rights, and what information governments are able to access. They

do this by making business decisions, ranging from minute decisions about where to place servers, to bigger decisions about how they are organized, where they place their law enforcement personnel, when to comply with government demand for information, and when to resist it. This means that one of the most important players in determining the scope of individual privacy rights, and determining the scope of speech rights, is no longer just governments, but increasingly these private companies. They are acting as intermediaries between the citizenry and governments, and not just one's own government, but effectively governments across the globe. I think it is a power that is likely to grow over time, as our lives become even more digitalized. This is going to require some hard thinking about accountability, responsibility, and the rules that ought to apply.

Privacy Complications in Cyber Physical Systems

Laura DeNardis, Professor of Internet Architecture and Governance, Faculty Director of the Internet, Governance Lab, American University, USA

As a backdrop to this article, I want to mention the conceptual framework of all of my research; people describe me as an Internet governance scholar. The basic thesis of my work is that arrangements of technical architecture are also arrangements of power. I study the politics of things that you cannot see when using the Internet, e.g., systems of interconnection, routing, protocols, addressing, and all the things that keep the Internet operational behind the scenes, and the enactment of substantive policy around that. I study this from three different perspectives. One is a background, both professional and educational, in information engineering, but also a doctorate in science and technology studies. I spent five years at Yale law school at the information society project. I really am an interdisciplinary scholar that looks at infrastructure.

The main point that I would like to make is from my new book that I am working on, and it is about rethinking Internet freedom and governance in the age of cyber control. The primary thesis is that the Internet is no longer a communication network, but rather a control network in which more things and people are connected, and in which infrastructure control is now a proxy for political power. Therefore, the cyberspace is not just a discursive public sphere, connecting people and content; it is a pervasive background infrastructure that connects cars, wearable technologies, home appliances, drones, biological systems, currency, and every conceivable industry sector. There is a diffusion of boundaries between offline and online, and between virtual and physical spaces, which is very relevant to public policy and how we look at Internet governance. The real world, connected by digital systems, also includes currency, cryptocurrencies, manufacturing processes, and even the body is now definitely a part of this space.

If you think about cardiac implants that are connected, biological systems related to medical monitoring, and biometric identification devices, these things create a new set of technology policy concerns, particularly around privacy, and the connection between privacy and human safety. An outage of the Internet today is no longer about disrupting my ability to email someone, or access knowledge, or really nothing about content, but about losing day-to-day functioning in the real world. This raises the stakes over how cyber infrastructure is designed and governed, including, of course, privacy.

A number of recent incidents help us understand this. In fall 2016, popular websites, including Amazon and Reddit, went down because of a massive DDoS attack. The outage, at least in the US, received a lot of attention from both the public and the media, because of the high-profile nature of the sites that went down. But from the standpoint of what I study, the more interesting thing ,to me, were some alarming characteristics. First, and like other attacks, these attacks did not actually target those sites; they disrupted the infrastructure of the Internet, and in this case the domain system, and even more specifically, a company that manages domain name system services. This is actually the subject of my last book, the turn to infrastructure and Internet governance, that the infrastructure of the Internet is the place where all of the control is happening.

The second characteristic of this attack is its massive scope. As the chief strategy officer of the affected infrastructure company described it, tens of millions of IP addresses across multiple attack vectors and Internet locations were involved. However, much more consequential for the future of the Internet, they were carried out by hijacking home IoT devices, like digital video recorders and security cameras. These were trivial to infect, because they had known security vulnerabilities, or in some cases weak or no passwords at all. The IoT is not only a potential target, but also a potential attack vector, where security incidents can arise and where privacy can be massively complicated. These trends help shed light on emerging global geopolitical cyber policy issues, in this diffusion of the physical and the virtual world. Cyber security, of course, is the most obvious area, but also the enormous scale and reach of IoT deployments. It creates many challenges around critical

Internet resources, not just spectrum, but also IP addresses. There are resource constraints that arise in this area, and human safety is a cyber issue now. Attacks can cause a disruption of a medical device or a malfunction of a car, disrupt additive manufacturing, and there are other examples of how this bleeds into the real world.

A major policy issue is the question of accountability, and who has responsibility for outages, upgrades and the insurance of basic human safety in this environment. I would also mention that cyber conflict, already a concern, is much more of a concern when you think about how this bleeds into real world devices. We have had examples of this for years, but what counts as the threat matrix, what counts as the types of foreign surveillance that are possible, and the disruptions that can occur, expands widely.

The tech policy communities haven't really caught up with this privacy challenge, in the material diffusion of the IoT, where data is not just collected about our communications, but also about what we do in all areas of our life. Corporate data surveillance, which takes place in this sphere, is much more massive, much more pervasive, and much more invasive, because we are not even aware of the ambient data collection that can occur through things like televisions, and different types of control devices we have in our homes and in our cars. This data gathering in the corporate sphere, of course, is the reason that much more government surveillance can be conducted. What data is gathered and how is it used? How is it retained? I am looking at the privacy policies of some of these companies that provide the cyber physical systems, and it is very unclear, there is not much transparency on the question to what extent our information is personalized and anonymized. Is data encrypted, both in transit and at rest? We don't know. These are very complicated issues that are evolving. Consent is very complicated in this area, and opt-out is complicated. If your IoT device doesn't even have a screen, how can you consent to its terms of service? People who don't even buy the products can be caught with the ambient data gathering. Privacy is very complicated, and there are not necessarily sufficient market incentives for companies to build in certain kinds of protection for privacy. The rush to come

to market is a challenge, we want to incentivize innovation, but how do you do this and also have security and privacy?

Considering the broader use of infrastructure as a proxy for power, geopolitically speaking, and in light of how the IoT systems are both a target and an attack vector, cyber physical systems are the emerging sphere in which geopolitical cyber conflict will play out. What raises the stakes considerably is how this relates to human safety. The rationale for my new book is to reconceptualize human rights in this environment, and to reconceptualize tech policy around this diffusion and enmeshment of the Internet into the material world, and in light of the geopolitical context, in which infrastructure is increasingly politicized. We have to think about all firms as tech firms now. Technology policy has to expand notions of what counts as a technology company. Distinguishing between tech companies and non-tech companies no longer makes any sense, and doing so is actually detrimental to our understanding of human rights and tech policy. The same types of questions about civil liberties, that arise when we look at large content intermediaries like Google, Facebook, and other kinds of intermediaries, have to apply to the personal data collection, and even more so in this context, blending the virtual and the material. Many of the firms, which are now suddenly also digital technology companies, don't have as much experience with cyber security in the public facing component, or dealing with privacy. Conversely, those who address policy concerns sometimes continue to focus on content intermediaries, rather than these real world digital intermediaries.

I would also suggest that cyber security, considering some of these trends, is the great human rights issue of our time. It isn't just about freedom of expression, democracy, and privacy of content anymore, but also about basic human safety. Some researchers easily demonstrated the ability to wirelessly connect to a car and disrupt the braking and acceleration systems, by hacking into the car's emergency communications system. I constantly monitor the safety warnings published by the various government agencies, and one that caught my eye recently was from the US Food and Drug Administration. They issued a safety warning about cyber security vulnerabilities in

the radio-frequency enabled implantable cardiac monitors. We also know that democracy now depends on cyber security; for example, the Russian hacking of personal email accounts, belonging to US political institutions and individuals, during this last presidential election.

Privacy and security are related to democracy. I am an Internet governance scholar, and I describe the mantra of Internet governance, just like everyone else in my field, as multi-stakeholder. It is distributed in a decentralized way. One of the goals of Internet governance is to preserve a universal Internet. We have to call into question all of these standard Internet governance philosophical positions, in light of the human safety concerns that arise in this sphere.

What is an intermediary in the IoT environment? We have many approaches to immunity from liability for information intermediaries, but what does the liability landscape need to look like for the intermediaries that are not sending content between humans, but between manufacturing devices or medical systems, etc.? Under what conditions should they be liable for the data that passes through them, when it affects not only freedom of expression, but also someone's life? There are enormous risks around liability in this space, and we ask ourselves, how can we have liability and also have innovation? That is something that needs to be paid more attention to outside of the content intermediaries, and into the physical intermediaries that are now completely digitized.

The role of Artificial Intelligence in risk determinations complicates this issue even further. At some point, there is going to be a decision in a driverless car: do you hit the stroller or do you kill the passenger in the car? I am slightly exaggerating, but those are decisions that will be made in a split second, and they will be made by AI. Who is accountable for these decisions? Another question that relates to Internet governance is, how does multi-stakeholder governance apply? It may be time to rethink that description of how policy works in the cyber sphere. There are certain unclear areas here . Should traffic management systems be multi-stakeholder, where civil society has input, as well as governments, as well as the private sector? I think that, more than ever before, they will have to really think about the layered approach to what the different services are, in this environment, and where is it

appropriate for governments to have authority, where is it appropriate for the private sector to have an authority, and where are the spaces that really should be more of a shared governance agreements.

I have one final challenge to Internet governance. I wrote a book on open standards, I have spent my whole life working on interoperability, and trying to have a universal Internet, but in this particular space, fragmentation has some desirable affects. Privacy in this space is similar to cyber security, in that sometimes the only firewall to having human safety protected, is having some lack of interoperability between systems. We have to rethink where fragmentation can be applied as a solution to some of the privacy an cyber security issues that arise. We cannot just always say we need an open and universal and free Internet; in some cases we need to have fragmentation that helps to preserve human liberty, and those cases include security.

In my book, which will be published in 2018 with Yale University Press, I am thinking about the Internet as a control network for the physical world, rather than as a communication network. I am using this as a provocation to both see the infrastructure and make it visible, and to reimagine the politics that are embedded in this infrastructure, so that we can have Internet freedom that includes basic human stability and security.

Protecting Freedom of Speech in the Age of Surveillance

Prof. Jacob Rowbottom, Associate Professor of Law at the Faculty of Law, University College, Oxford, UK

I am going to discuss here the relationship between freedom of expression and privacy, and the impacts that digital surveillance has on the relationship between those two rights. I come from a background as a lawyer specializing in freedom of expression and media freedom, and typically we have tended to see privacy and freedom of expression as being in opposition to one another. We see privacy as a restriction on freedom of expression, and that largely comes about by the way these issues have arisen in the United Kingdom. Normally in the UK, when we look at privacy and free speech, it concerns a newspaper that wants to publish details about a celebrity's private life, and the celebrity then tries to restrict that publication; then we see privacy as a restriction on free speech or media freedom.

But that is just one particular aspect of the relationship between privacy and freedom of expression. In many contexts we can argue that privacy and freedom of expression are mutually supportive of one another. By that I mean that before we actually speak, before we actually write anything, or say things to other people, we sometimes need privacy to develop our thoughts, maybe to have private discussions, to gather information, and so on. This connection between the two is nothing new. Over 250 years ago, there was a famous case in the UK called *Entick v. Carrington*, which concerns the power of the government to raid a journalist's home, cease their papers and go through their belongings. In this landmark judgment, Lord Camden said: "a person's papers are among their dearest property," and by that he was getting at the idea that your personal papers are special; they are not like other property, because they somehow reflect your freedom of mind; that is what you do before you speak, and what you do to develop your ideas.

You need your privacy protected, and it is not something new, but when I think of an era of digital surveillance, that aspect of the relationship becomes increasingly more important. What I argue is that when we look at free speech, we don't just look at the actual moment of communication, but at other stages in the communication process as well. We don't just look at the moment of publication, we also look at the stages of gathering information and the reading of information. This issue of digital surveillance has become incredibly topical in the United Kingdom. In 2016, the UK government enacted a law called the Investigatory Powers Act, and Edward Snowden described this as the most extensive form of surveillance in any western democracy; it even went beyond what was available in some authoritarian regimes. You can debate that, but it is nonetheless a significant piece of legislation. It is also quite a long and complex piece of legislation, so I am not going to go through it all, but some of the provisions are that it gives power to the government to intercept communication, and actually look at the content of communications, things like emails, telephone messages and so on.

The Investigatory Powers Act also gives powers to the government to do something called "equipment interference", which is a polite way of saying hacking into people's devices. These are things like going into someone phone or laptop and use that as a device to eavesdrop on people's conversations. That power is probably going to become increasingly important as we witness the growth of the Internet of Things. As people have smart TVs and things like Amazon Alexa, there are more devices that can be interfered with, and be used to eavesdrop and to listen to what you are saying.

The third aspect of the act is that it gives the power to the government to require telecommunication companies to retain communications data and allow certain public authorities to have access to that data. Specifically, communications data can be retained for up to 12 months. By communication data I don't mean the content of communications, but rather their details and circumstances. We don't know what people are saying, but it is about ways of finding out who spoke to who, at what time they spoke to that person, and through what medium. If we get all of that communications data, knowing the context of who

is talking to who and when, we can get quite a big picture about a person's life and we can make certain assumptions about them.

Obviously, all of these powers can only be exercised for certain purposes, and the main ones are the obvious candidates, like issues of national security, prevention and detection of crime, as well as the country's economic wellbeing. That said, these are very broad powers of surveillance that have been granted, yet the enactment of this legislation hasn't caused a massive public outcry. Some people wonder as to why this has not been such a big source of controversy; it has been with certain liberty campaigns, but it hasn't really been in mainstream politics. There are certain thoughts as to the answer. First of all, the various digital companies have access to much of our personal data, and there is an argument that says, if they can access this stuff to decide what adverts they show to us, should we not let the government access to some of this data to actually prevent terrorism or prevent crime? The other argument I have heard for this legislation is that in the digital era, surveillance is simply a fact of life. Edward Snowden showed that this happens, regardless of whether or not there is a legislative framework in place, and this argument is asking if we shouldn't have this put in legislation and at least, if it is done in the open, it can be more transparent and subject to significant safeguards. That is indeed one of the things that they put in to the Investigatory Powers Act. In order to exercise these powers you need a warrant from a minister, and there needs to be judicial oversight of that warrant as well. Some people say these are not sufficient; the warrants can be cast in very broad terms, and some people question if judicial oversight is sufficiently rigorous.

What is the criticism of this system of surveillance? There are three main points we can make. The first relates to the utility of the system of surveillance. Some people say that the government granted itself broader powers than are necessary. All of this level of surveillance just isn't needed to prevent crime or terrorism, they could simply have drafted the legislation in narrower terms. Closely related to that as well is the question on what the utility of this level of surveillance is, given that they can harvest masses of data, but that is no good unless

you have the resources and people in place to actually process that data, and have some interpretation of it.

The second criticism is an obvious one, relating to people's privacy. The idea is to say that this level of surveillance is a gross interference of privacy, and that has really been an issue post Edward Snowden. There have been some concerns about this, and as a result, in the Investigatory Powers Act there is a clause that states that privacy has to be protected, and when they are issuing warrants and there is a judicial oversight of warrants, privacy rights have to be taken into account. Once again, people wonder just how effective that type of safeguard might be.

The main point I want to make is that it is not just about privacy, there are certainly some major free speech issues at stake, when we look at these powers of surveillance. On one area, which was actually a topic of discussion when this statute was being enacted, is that concerning the protection of journalist sources. As you all know, before journalists print information, they need, in some cases, to rely on confidential sources, and they don't want to disclose the identity of such sources. If they do so, it creates a chilling effect, and people might be afraid to come forward and inform the journalist. Traditionally, if the government wants to find out who has been informing journalists, they have to go to court and try to get an order for the journalists to disclose the identity of their source. However, if you have access to communications data, you don't need to find out, you don't even need to ask the journalist who their source was, you could just go to their telephone records, go to their Internet history, and you can pretty much find out who they have been talking to; join the dots, and it quickly becomes apparent who is informing them. The older powers of surveillance in the UK had been used to do that on a number of occasions, and that was taken to court, and it was found to be a violation of the right to freedom of expression. As a consequence of this, the Investigatory Powers Act says that if the purpose for accessing communications data, or internet records, is to discover who the identity of a journalist's source is, then there has to be some judicial approval of that access. It provides a safeguard, but we will have to see how well that works once the act is fully implemented.

These all go to show that surveillance can impact not just the moment of publication, but the prior stages as well. Surveillance can be used to regulate the gathering of information that takes place prior to publication. But more broadly than that, because that is very specific to media and journalistic freedom, you've also got the impact of surveillance on the rights of an audience. If people know what websites you are looking at, that can have an inhibiting effect. Some people might say that only people that are up to no good would want to keep things private. However, there are many good reasons why people might be wary of visiting certain websites, having certain conversations, or speaking to certain people, if they think that this information will become known to the authorities. They might also fear that there will be certain advert consequences if you are visiting certain websites; you might find that you have become a target for investigation, or that you came to the attention of the authorities in some way. It can also have a chilling effect on people's reading habits, if they think they are going to be held accountable for what they read or what they look at.

I am in the process of finishing a book on media law, where I have been examining various controls on freedom of expression. One of the interesting things that I found is that we have a range of publication offenses in the UK. There are criminal offenses and other types of offenses, which say you cannot publish certain types of content. Several decades ago they might have been the primary control, but what I have seen is that in many areas there is a declining reliance on the classic offenses related to publication. Obscenity laws, certain hate speech laws and others aren't as widely enforced as you might imagine. One conclusion that we might have from that is that countries are becoming a little more liberal, and more tolerant of free speech, but I'm not sure that is the case. I think it is more the nature of the controls on freedom of expression that is starting to change. I am referring to things like the role of Internet intermediaries; there is an increasing push in the UK for people to say the government should have some monitoring roles, or that it should be proactive in blocking certain sites and taking them down. The public is not imposing any liability on the publisher, it relies on the intermediary. In other cases,

governments or other speakers might try to counter speech, not by censoring it, but by engaging in communications of their own. Rather than invoking the laws of defamation, which protect your reputation, you engage a reputation management company that will try to put out positive information with you. I also think that another piece of this jigsaw, one of the new controls, is that of government surveillance, watching what people are looking at and so on, and that can create an environment that inhibits the free speech, and inhibits people to choose what to look at.

My central point is that in an era of wide spread surveillance, free speech needs to emphasize not just the right to publish, not just the right to speak, but also, the rights of the audience and the right to read and receive information as well.

Defending Critical Infrastructures Against Cyber Attacks

Prof. Sandeep K. Shukla, Poonam & Prabhu Goel Chair Professor, Department of Computers, Science and Engineering, Indian Institute of Technology, Kanpur, India

Let me tell you a little bit of history. I was working at Virginia Tech in the US for fourteen years, and in 2015 I arrived back home to India, to work in the Indian Institute of Technology. I started looking at the underground situation in India, in terms of defending our critical infrastructure, power system, manufacturing system, and so on. We were being sold nuclear reactors by various companies in the US, Russia and France, and we were very worried about their cyber security as well.

While doing that, we realized that there was no testbed for Supervisory Control and Data Acquisition (SCADA) systems in India, which would allow us to do realistic experiments, looking at the cyber security of various equipment and software that go into such systems and are deployed all over the grid. We started with a very small scale testbed at our lab, but now we have been funded by the government to put together a national testbed that is similar to the one in Idaho, where companies can come and do hardware-in-the-loop and software-in-the-loop testing for vulnerabilities, test mitigations techniques, and so on.

The number of cyber crimes in India is skyrocketing, with an exponential growth rate over the last couple of years. This has been a phenomenon seen all over the world, and India is not immune to that. Government websites are getting hacked, as are websites of various companies. We had cases where government websites have been defaced, but this could very soon escalate. There are unconfirmed rumors that the Israeli power grid was hacked in 2016, although there were no substantial damages. In 2015, the Ukrainian power company got hit by the Black Energy malware and the country suffered a massive blackout, and then got hit again in 2016, although this time they were able to recover more quickly. Television networks get hacked. The

German Steel company was hacked in 2014 and suffered equipment damaged. Turkish banks suffered cyber attacks last year, with a massive DDoS attack that impacted the entire Turkish economy. An Austrian hotel got hacked by an IoT ransomware, and people got locked out of their rooms. The hotel had to pay a ransom to get guests in or out of the rooms. And then, of course, in 2016 there was the famous Dyn attack, which was an IoT-based DDoS attack, and we are seeing an exponential rise on mobile malware.

This is a global situation, which exists in India as well. We believe that a hundred percent of the companies were hacked last year, but only the enlightened ones actually know that they have indeed been hacked. According to a survey conducted by Assocham Mahindra, 72% of the financial companies in India were hacked in 2015, as well as 37% of oil, gas and utilities companies. In that light, our contention is that we will be hacked, and we are probably already hacked. We have put honeypots in our institute, and we see about 300,000 attacks per day, even though we are not critical infrastructure. Therefore, the question is: how do you design cyber resilient physical systems, which are able to detect attacks as soon as possible, contain them, and isolate areas that have been attacked or infected? That is the goal that we are progressing towards, and in order to do this, we need a realistic system.

In the past we have used virtual SCADA testbeds, where everything was virtual except for the front-end. We took an actual front-end from a SCADA vendor, and we simulated our power systems using a standard simulator. Unfortunately, that didn't allow us to do various kinds of attacks, like we can with our current testing environment. Our approach is multidisciplinary: we have had people in various universities, working in the fields of machine learning, program analytics, computer architecture, memory systems, formal methods, robotics, cryptography and more. We wanted to bring all of them together, so we proposed the Center for Cybersecurity and Cyber Defence of Critical Infrastructure at IIT Kanpur, a plan that came true.

At our center, we look at critical infrastructures as an "extra-large" cyber-physical system, so what we did was to break it into different layers, and map attacks accordingly. There are the physical dynamics

of the system, which have to be constantly checked for anomalies. Many times an attack will reflect something in the physical dynamics without any sign elsewhere, because it could be an insider attack. There are electromechanical components, like relays and breakers; we have the industrial automation and control, which is basically the SCADA system, with its various communication and computation components; we have to inspect the firmware for issues like secure boot; and naturally, we have to keep an eye on all the traditional components like software, network, perimeter defense, and lately also the Cloud.

We then reviewed the various methods we wanted to explore, and the required disciplinary knowledge and expertise that will be used. For example, we have electrical engineering and power system people looking into a PMU-based power, and conducting wide-area monitoring and control. On the other side of the scale, we have number theorists and cryptographers looking at lightweight cryptography and also breaking various types of crypto. We have system simulations and code simulations that often give you a cheaper means to do proof-of-concept ideas, and we are currently building a full scale emulation testbed.

What we did first was to build a lab-scaled testbed. One of the most prominent challenges today in most of these manufacturing systems, or any other critical infrastructure, is the notion of IP convergence, meaning that your business network is somehow connected to your field network, where your measurement devices and other various equipment are connected. These networks are usually isolated by firewalls, but once a malware gets into the business network, like in the Sony, RSA and other cases, it is very possible that sensitive passwords will be stolen, the firewall will be breached, and eventually the hacker can get access to the critical infrastructure, and that is a problem. We needed to replicate this kind of environment in order to recreate this and other attack scenarios. We started by building a test environment very quickly, with the idea that we wanted to first demonstrate that a lab scale testbed is possible, and that we can successfully point to weaknesses in various industry standards, equipment and devices. In the first year, we built a model of a power system and distribution

panel, where we have multiple types of transformers: three phases to three phases, three phases to single phase, single phase to single phase, and so on. We also had relays, power meters, and protocol switches. Of course, all data coming in or out of the system was collected, whether through dedicated protocols like 61850, 104, or just plain TCP, which gave us the capability to attack from various angles.

The academic literature is referring to something called the "industrial automation pyramid". The bottom level includes the instruments and the devices. The second layer includes the program logic controllers, the remote terminal units, and various others IEDs. The third layer is the SCADA layer, the super-visionary layer; and the top layer pertains to the business parts – reporting and so on. What we managed to create today, our biggest progress, was a lab environment with parts pertaining to all of these layers, containing PLCs, various meters, GPS time sync units, relays, and switches. On the other side we have a LAN environment with engineering workstations, meaning we can emulate a real working SCADA system.

With this setup, based on equipment we got from Schneider Electric, we started looking for vulnerabilities. Our first finding was in a Moxa switch, which had a privilege escalation vulnerability. This meant that any user in the system was able to become an administrator of that switch, inject data and perform various actions. We were able to exploit this vulnerability, and we caused the SCADA to make wrong decisions. We also found that all the standard protocols that are being used today, including 61850, are actually not encrypted. We were able to successfully conduct Man in the Middle attacks, false data injections, and DDoS attacks. We managed to carry out ARP poisoning, making the SCADA use the wrong configuration. All of our findings are reported to the Indian CERT.

As I mentioned earlier, we recently received the funding to extend our lab environment into a real-scale system, and it will probably take us another year to build this. Up until now we could only emulate transmission of data, but this new system will have generation, transmission and distribution of data. It will also have manufacturing automation, both process-based and discrete. It will be a two-zone system, where we will have solar and diesel generators on each zone,

as well as a power transmission. We are also looking to emulate a home automation environment, along with the necessary residential infrastructure. We have just started planning that new building where these things will be put in. Each zone will have its own control center, with multiple substations connecting the transmission lines, and the process automation and discrete automation testbeds will serve as load centers for both zones

Our students are having a lot of fun designing and using our environments, but in fact it is quite scary to discover all of these problems with the devices powering up the critical infrastructures in the country. We also keep discovering that these problem exists in other infrastructures as well. After the 2016 demonetization process in India began, the government started pushing for digital financial transactions, and the urgency of this situation really escalated. We received calls from banks to give talks on mobile banking security, even though we knew anything about it back then. We found out that they invested so much money in creating the necessary infrastructure, without thinking about how to secure it. Eventually we got connected to the finance ministry, and we are currently expanding our research into that world as well, applying the same interdisciplinary approach as before. We plan on building a similar testbed, adjusted for the Fintech world. It has not been funded yet, but we will have various labs working together to research Fintech security. We are also looking at other major problems, such as IoT authentication with Blockchain, honeypots, various methods for code replacement attacks, malware analysis, and more. We have multiple collaborations, including one that is still under work with the ICRC, as well as several industry collaborators, including the Bombay Stock Exchange.

To conclude, we have an ambitious project with many goals and activities, we have an industrial testbed that is unparalleled in India, and we have been partnering with top places around the country and the world. We truly believe that can make a difference in the cyber security posture of our country.

Secure and Privacy-Preserving Data Analytics and Machine Learning

Prof. Dawn Song, University of California, Berkeley, USA

I would like to discuss one of our projects, on secure and privacy-preserving data analytics and machine learning. This work is a joint work with my students, Noah Johnson, Joe Near, and Joe Hellerstein. One of today's biggest problems in Big Data and machine learning, is that we have a huge amount of data, but much this data is sensitive. Because of this, much of the value is actually being locked in, because people don't have secure and privacy-preserving ways to analyze this data. As a consequence, the value of this Big Data is not being utilized. The question, then, is how can we utilize it better? On the one hand, we could potentially get many business intelligence insights from the data. On the other hand, if we don't do this well, then in terms of security and privacy, privacy breaches and data exfiltration can become serious issues.

As we all know, security and privacy risks are real. There are many severe consequences for data breaches, and there are many real world incidents. What kind of solutions can we offer for these issues? One solution people have proposed in the past is anonymization. Unfortunately, as we know, anonymization is ineffective. For example, Netflix released some of their anonymized movie review data, NYC taxi released some anonymized taxi trip data, and in both cases researches have shown that you can combine the released anonymized data with other public data, to deanonymize users.

Another solution people typically use is access control, but this is an "all or nothing" solution. Either someone gets no access to the data at all, and they cannot use it, or you give access to some employees or some analysts, but then that person has all the access to the data, including raw data. In this case, you can have an insider attack, or the person's account may be compromised. This does not provide sufficient protection to the sensitive data. In other words, solutions

like these are not only insufficient, they also degrade the value of the data, and often times destroy data utility.

In our project, called Allegro, we try to solve this problem in a better way. We propose a 3-in-1 solution for secure and privacy preserving data analytics. In particular, we have three components in Allegro: privacy preserving in general; data analytics, where we support general analytics but at the same time provide privacy protection; and privacy preserving machine learning, where we support machine learning with protection of privacy. We also enable fine-grained access control and auditing capabilities. Some of the work is still in progress, and so I will mainly focus on the second aspect, privacy preserving and data analytics.

We find it important to be able to deploy the system in a very easy way. We don't want the clients to change the existing workflow, and we also want them to keep the existing backend data store. In particular, what we want to achieve is that the analysts will keep doing their typical workflow, where they write their analytics programs to pose their queries, while we enforce certain security and privacy policies.

Generally speaking, Allegro transforms the original analytics program into a secure version of itself, which is guaranteed to satisfy the security and privacy policies. Using this approach, the secure analytics program can then be run on the existing data store, either in the backend or out or Big Data computing framework, meaning that the client doesn't need to change the analyst workflow or the backend system. In other words, the backend system is executing the secure analytics program version, computes the necessary information, automatically generates secure and privacy preserving program results, and sends them back to the analyst.

The major advantage of Allegro is that it enables the analysts to do their job, meaning they can still analyze the data and extract value from the data while enforcing the security and privacy policies. Additionally, Allegro is designed to support a variety of different types of security and privacy policies. We have some built-in security and privacy policies, including differential privacy, SGX security analytics and others. These are all customizable using a rich policy language. Allegro is compatible with almost any database and types of datastores,

and supports general data analytics and machine learning. Its work is transparent to the analyst, so they don't even have to know how to do secure and privacy preserving analytics and machine learning, the system will do it for them.

How do we enforce differential privacy as a security and privacy policy? Differential privacy, essentially, is a very powerful notion for providing privacy for data analytics. In a sense, it provides a formal guarantee of a indistinguishability. For example, let's say that we want to compute a differential private query 'F', meaning it satisfies two properties: the first is that adding or removing one row does not significantly affect the query's output; the second is that the results are similar with or without information on an individual. We use two data sets. The first is the original data set, and the second is called a "neighboring" data set, which only differs from the original by one data point, meaning it was either removed, added or changed. The computation of 'F' satisfies the differential privacy requirements if the computation results of 'F', using each of these two data sets, are almost the same, meaning you almost cannot distinguish between the two results. Using these two properties, we can enforce the anonymized version and protect the user's data, without hurting the results.

Differential privacy is a very strong notion, it gives you a theoretical guarantee about preserving users' privacy. Various privacy mechanisms have been proposed in order to achieve differential privacy. One approach is to add random noise to the query results, scaled to the function's sensitivity. The main challenge is to enable differential privacy in practice, in the real world. There has been growing interest from the industry on differential privacy. For example, both Google and Apple have recently deployed differential privacy in certain parts of their products. However, these are very rare examples from the real world, and in both cases, it has only been deployed in a very limited way. Most of the current work done on differential privacy is theoretical in nature, and have only been evaluated on a very small scale, using synthetic data sets. We want to take it to the next step, enabling differential privacy in practice, supporting real world requirements, including performance, scalability, compatibility with existing data systems, and flexible access for the analysts.

There are different mechanisms to implement differential privacy, each with their strengths and weaknesses, and may be particularly suited for different types of data analytics queries. We also proposed a new mechanism, called "elastic sensitivity", which is more powerful than other existing mechanisms. However, most of these are not compatible with existing databases. Also, if you wanted to run different queries on the same data, you would often be required to have copies of the data, stored in different ways. This is, of course, not practical. What we propose with Allegro is a completely new way of work, called "intrinsically private queries", which is multi-purpose, supports existing data stores, and is very easy to deploy and use. It was made possible only due to the query and program rewriting approach, which I mentioned earlier.

An intrinsically private query is a query that already has a privacy mechanism built-in, and hence the query itself automatically enforces the differential privacy guarantees on its results. The main advantage of this is that it doesn't require a custom run-time or database modifications. Additionally, because it is a very flexible mechanism, it can support all existing differential privacy mechanisms. We built a system that can take an existing query, and rewrite it, in real-time, into a version of a certain differential privacy mechanism, while keeping its original data store. We support different privacy aspects, including different sampling mechanisms and adding the noise perturbation. Compared to previous approaches, where in order to deploy these privacy mechanisms you would have to change the backend databases and replicate the data to support different mechanisms, our approach is very easy to deploy, and the analysts only need to write the original queries as before. The analyst doesn't even need to know about differential privacy at all. The Allegro system would automatically rewrite the query into an intrinsically private query, and it will get executed on the original database, which doesn't need to be changed at all. This execution will then provide differentially private results.

Inside the Allegro system, in order to support this automatic query writing and to enable intrinsically private queries, we have to develop several different components. We need to be able to automatically analyze the original query, using our analysis engine. This lets us

analyze the sensitivity of the original query, and figure out how much noise perturbation to add to the final results. We also have a system for privacy mechanism selection, which uses a machine learning model. Different queries may require different privacy mechanisms that are best suited for them, and so this machine learning model takes the features of the original query, and figures out which privacy mechanism is the best. Last, we need a transformation engine, which will transform the original query into the intrinsically private query, using the particular privacy mechanism that was selected.

We have actually built the system, and have conducted experiments. We have a collaboration with Uber, which enables us to analyze the meanings of real-world queries, and analyze the utility of our mechanisms on a very high number of queries. Initial results show that with our approach, for a majority of the queries we can provide very good utility, and at the same time provide privacy protection. We have also shown that our mechanism selection process does in fact select the best privacy mechanism for each query. So far I have only mentioned the differentially private analytics queries, but we also use a similar approach for differentially private machine learning, although this is still a work in progress. Additionally, we use a similar rewriting approach to enable fine-grained access control and auditing.

We hope that by using this system we can enable secure and privacy-preserving data analytics and machine learning, allowing to unlock the value of Big Data, while diminishing the risk of data and privacy breaches.

Beauty and the Burst: Remote Identification of Encrypted Video Streams

Roei Schuster, PhD. Candidate, University of Tel Aviv, Israel; Research Assistant, Cornell Tech, USA

I am going to discuss a research we called "beauty and the burst", about remote identification of encrypted video streams. This is important, because many people stream video contents in the privacy of their own homes, and they don't expect any third party to know what they are streaming. However, there are many parties who are interested in what people are streaming, for targeted marketing purposes, market characterizations, and other purposes. Luckily, Netflix and YouTube and most other mainstream streaming services encrypt their traffic, so the traffic between the service and the client is encrypted. On the face of it, an attacker that is positioned on the wire, or has some information about the packets that go from the streaming service to the client, cannot learn anything about the content that is being streamed. A possible attack vector would be for the attacker to understand that they don't know the contents of the packets, but they can know their timing, their size, and maybe some other metadata; and perhaps this metadata can actually identify the video being streamed.

Videos are especially susceptible to this, because of their variable bitrate (VBR), which means that different seconds of the video take different amount of bytes to encode. This is well demonstrated by the "Iguana vs. Snake" video, which is a famous video of an iguana being chased by a snake. In the beginning of the movie we have a lot of scenery, the tension is rising, and there is much movement. This means it takes many bytes to encode all that. Then, when the tension is peaking, everything is still, and we get a headshot of the snake, a headshot of the iguana. This still scene, of course, takes very few bytes to encode. Then we have another chase, everything is moving rapidly and the bytes per second peak. Finally, we get the iguana resting in its safe haven, and everything is good again, so the bitrate

drops. We can, technically, create a signature of a video, based on its variable bitrate patterns.

We know two other things about stream videos. One is that, on the server, they are saved in chunks. Each chunk corresponds with about five seconds of video, and because we have a variable bitrate, and we know that it takes a different number of bytes to encode different segments of video, then these chunks are variably sized. The second thing we know about video streaming traffic is that it comes in bursts. Every few seconds we have burst, and apparently this is a result of the way that the streaming client and server interact. This means that every few seconds the client fetches one segment that is stored on the server. We discovered that these bursts are visible to almost anyone, and correspond with the sizes of segments on the server. To recap so far, the content affects the variable bitrate pattern, the variable bitrate pattern affects the segments as their stored in the server, and the segments affect the bursts size. The question we asked ourselves was whether or not these bursts can be used to identify videos.

The first scenario we are going to consider is when somebody is actually sitting on the path from the streaming server to the streaming client. This scenario also assumes that this somebody is quite a powerful adversary, and that they have a certain degree of control over your Wi-Fi access point. It could also be a malicious ISP or a network moderator. These adversaries work by using machine learning and traffic analysis. These type of attacks have an offline phase, where the attacker simulates the attack in their own network. This enables the attacker to learn what the bytes of a specific video look like on the wire, while they are encrypted, and also what the video's metadata looks like. Using this information, they are able to build a detector that receives an encrypted stream as an input, and computes which video it corresponds to. After that, there is the online phase, where the attacker observes the actual communication between the streaming service and the victim network, and uses his detector to identify the video being streamed.

In order to simulate this attack, we used deep neural networks, which are good for many tasks. They were especially good for us because they can handle noise well, which is important for a side-channel

hack. They are also generic, so we can attack multiple services with them, and we don't have to fit them for a specific protocol or specific features that we derive from the communication trace. We defined a few classification tasks for some mainstream video services. For instance, in Netflix we set the algorithm to differentiate one video from 100 classes, where one video equals one class. We did similar things with YouTube, Amazon and Vimeo, but at a smaller scale. These classification tasks demonstrate the power of an adversary to actually lean inside to what you are watching, even though its encrypted, and we got good results across the streaming services with 98.5% in Netflix and 99.5% in YouTube. It was actually very scary. Another thing we did was to adjust our classifiers for precision rather than accuracy, because when you use them in the real world, you care about the very low base rate in this sort of attack, and they have very few false positives.

We also simulated an off-path attack, which is a side-channel attack. In this case, let us assume that our attacker is not that strong, they are not an ISP, or they don't have a hold on your access point. In fact, the attacker just uses a webpage that you visit, or they are someone with a laptop connected to your home network, or an app on your phone. These adversaries are very weak but they are ubiquitous, so they could be very dangerous if this leak is exploited by them. There have been several attempts to model such attackers in traffic analysis works, but we haven't seen the network congestion side-channel being utilized as a vector of attack, or brought up. This attack is best demonstrated together with a cross-site attack, in which we have a server, and we assume there is a browser open to that server and its streaming content. We also assume the browser is also open on a different website, which is a malicious.

In our scenario, the malicious website can send the client its own JavaScript to the open browser instance of the victim, and it can start sending messages to this client. These messages cause congestion on the presumably limited bandwidth link between the victim and the internet, and when a traffic burst arrive they cause the congestion to grow – which, in turn, causes delays that the attacker can pick up. We discovered that although these delays are a little noisy, meaning

they don't represent exactly the sizes of the traffic burst, they still represent them fairly well. They are very close, and the correlation is very high. We were able to use the delays we caused in order to get an approximately 96% detection rate of the streamed videos, which is just slightly less than the first scenario. It is important to note that while this was a cross-site attack, we could have just as easily implemented it as a cross-device attack, in a scenario where both a regular computer and a malicious computer have access to the same Smart TV used for streaming videos.

It seems that there are two causes for the leak. First, the videos are being recorded in variable bitrate, which I don't think anybody has enough motivation to change, because it is very important for the efficiency of the encoding and the presentation. The second thing is that content is fetched at segment granularity. Apart from being widely deployed, it also seems like this is designed to maximize some notions of streaming efficiency, such as quality of experience, server load and so on. Any attempt to change one of these elements will have to trade-off some notion of streaming efficiency. Therefore, it is unlikely that streaming services will be inclined to change this in the future, and so we don't have a probable mitigation in the near future.

Proportionate Response to Cyber Attacks

Dr. Jarno Limnéll, Professor Cybersecurity,
Aalto University; Adjunct Professor Cybersecurity,
Tampere University of Technology, Finland

I want to talk about cyber issues from the political point of view, and specifically about the topics Finland has been examining in this world, which are hybrid threats and hybrid warfare. I think it was no coincidence that the establishment of the European Center of Excellence for Conquering Hybrid Threats was signed in Helsinki, our capital, a few months ago. I think that this center really represents a concrete step in building resilience to hybrid threats, and is also a real boost in EU-NATO cooperation in this field. Since cyber and hybrid issues are international, the response should be international too. What is really needed, from the political point of view, is a comprehensive multinational, multidisciplinary approach to these issues.

Many of the concepts and dichotomies that we, security experts, usually use, such as war/peace, our side/their side, military/civilian, win/lose, are blurring. That is the reality in which we have to create security, and that is the reality in which we have to be stronger with cyber issues. Our opponents, state and non-state actors, are utilizing this reality in a very active way. In Finland, we have experienced our fair share of different kind of hybrid influencing, and we have understood that cyber issues should never be separated from the political context. We are looking at the publications on this topic, as well as the various conferences, and we see and hear that there is a lack in political understanding, but there is also a political decisiveness to try and solve them. I argue that cyberspace is primarily a political domain, and issues related to the cyberspace have evolved into the highest realm of the high politics. Only a few years ago, these issues were a little bit like low politics, almost background noise, but now it is part of the high politics.

That said, there are plenty of different kinds of high cyber hostilities. Just read the recent news. Moreover, new ones are appearing constantly, and according to the US Cyber Command outline of the cyber operations

spectrum, based on past events, we can see that both state and non-state actors are testing the boundaries of the so-called cyber battlefield. They are testing the red lines, what they can do, and in what ways their political objectives can be achieved via different digital activities. From the political point of view, this raises very important observations, and I will mention three. First of all, political response to different kind of cyber hostilities is still an exceedingly untested phenomenon. Second, the cyber domain is a relatively new area or arena of conflict, especially for policy makers. The third one is that much more research is needed. Multinational, multidisciplinary research is needed in this field, because there seems to be a situation in which, when we face different kinds of cyber hostilities, our policy makers start thinking that they have to do something, but they don't know what. All of the actual responses, at least those we observed, were done ad hoc.

Let me give you a few examples. Elections security is a very important issue in Europe at the moment. The United States, in several different documents and statements, has officially said that Russia interfered in their elections. From the researchers' point of view, these kind of statements, like the joint statements from the Department of the Homeland Security and the intelligence community of the United States last year, are very strong messages. The US is saying clearly that they know who was behind this interference. At the same time, of course, the opponent is denying it. Russia is claiming that perhaps some patriotic hackers have done something against the US elections, but the state had nothing to do with it. I think that a very interesting question, in this context, is how the United States, which in my eyes is the only real superpower in the whole world, is responding to these kind of activities. As we know, former president Obama said many times that there will be a "proportionate response" to these issues. Eventually, the United States expelled 35 diplomats, closed two compounds, and put in place some sanctions against the Russians. This was not a very strong message.

The same situation happened in Europe, in the presidential campaign of Emmanuel Macron. His emails and some documents were hacked and leaked to the public. Immediately after those leaks and hacks, the then president François Hollande promised that France will respond

to these activities. I wanted to know how exactly France will respond, but so far we haven't seen any response, at least not publicly. Another example comes from the European Commission on Cyber. The European Union has agreed on a cyber diplomacy toolbox, and the council of the EU launched an initiative in hopes it will strengthen the blocks' or unions' ability to deter and respond to cyber threats. In its report, the council said that a cyber attack on any member state would be met by a joint response, including sanctions. The word "proportionate" was also used in this case, but what does that mean in reality, if the EU or even a single country faces that kind of situation? It is unclear what sort of tools that toolbox will contain, and what "proportionate" means in concrete actions.

The message I want to convey is that we have to be able, or at least encourage our political decision makers and others, to create a framework of proportionate ways to politically respond to different kind of cyber hostilities. Of course, the proportionate response is complicated, situation-dependent, and it is always a very political question. When we determine appropriate responses to cyber hostilities, I think that political decision makers need to consider the following six variables, a list of questions which must be answered before responding.

The first one is of course the most difficult one, the degree of the attribution. When considering proportionate response, policy makers should understand the level of confidence they have in attributing the attack. For instance, if the level of that attribution is low, decision makers will be limited in their choice of response, even if the severity of the attack is high. Of course, the government needs to calculate the costs that it would incur, if they wrongly attribute attackers.

The second issue that the policy makers have to consider is what was the impact of the attack. They need to understand what the real impacts are, and the type and the level of the response is determined, more or less, by the extent of its impact: how harmful the attack has been to the national security, the society, or its digital services, and what kind of services are being affected. As we very well know, it can take weeks, if not months or years, for computer forensic experts

to find out the extent of the damage done to the target organizations networks.

When the policy makers are considering proportionate response to a cyber attack, the decision is always about options, what the state is able to do. It is said that every nation that can respond using at least four instruments: diplomatic, informational, military, and economic. The political decision makers need to consider the full range of responses at their disposal, from a quiet, diplomatic rebuke to a military strike.

The fourth point is the policy guidelines. Political decision makers need to take into consideration the current national security and cyber security strategies that they have outlined. Additionally, if the state is a member of the international allies, meaning they are a part of the NATO, their politic guidelines must also be taken into consideration when thinking about the proportionate response.

The fifth point is time. How urgent is the response? How fast does this kind of decision have to be made? The time is a relevant issue to take into consideration when discussing politics.

Finally, the political decision makers have to think about the political risk. If we do that, what will that mean in the next phase? What is the risk of the escalation, if we respond in this determinant way? I think that whatever the response is, it will have an impact on the state's diplomatic relations, reputation, power, military and intelligence operations.

There are many ways in which a nation state can respond, using these kind of different elements. I would really like to encourage the policy decision makers to think about these different variables and these different ways, when considering how to respond to different cyber hostilities. In politics – and in cyber politics, which is an increasing research branch at the moment in Europe – there will always be flexibility, depending on both the current decision makers who have the power, and the ambiguity of the situation. Naturally, this kind of framework is different from one nation to another.

My last point is that when we are operating in today's unpredictable hybrid security environment, we really need more political expertise and preparation regarding these issues, and especially in cyber issues. Policy makers, at least from Finland's and Europe's point of view,

are probably forced to reconceptualize cyber war or cyber conflict definitions as a form of the larger concept of the hybrid warfare.

What Does the Brain Tell Us about Usable Security?

Cyber Security in the New US Administration: A Work in Progress

Dr. Herbert Lin, Senior Research Scholar for Cyber Policy and Security, Stanford University, USA

The information and thoughts I will present here are not based on detailed discussions with anybody in the new administration. I offer it to you as somewhat informed speculations, based on what has been reported on what I regard as reliable sources, i.e. things that I don't regard as fake news. Some of the source materials for this are two documents, both available on the web, although one is classified. The DoD Cyber Strategy, which is not classified, was released in April 2015. That is the most authoritative statement to date about the US Department of Defense's cyber strategy. The other document is a memo titled the "US Cyber Operation Policy", which was leaked by Edward Snowden. It is still classified as top secret in the US, so unless you have the right security clearances, you are not allowed to read it. If you know where to look, though, you can find it.

Under Obama, the US view of on offensive cyber operations said that offensive cyber operations has a definite role as instruments of US military power. It was the first time that the United States had explicitly and overtly made such a claim. It said that such operations will be conducted in accordance with the laws of war. It specified the kinds of targets that were regarded as legitimate, meaning military targets. It specified when these might be used: "offensive operations may be conducted during periods of heightened tension," which basically means before anyone actually starts shooting. That is an interesting point, because it says that if you are going to use cyber operations before shooting, that is an indication of intent; that is the thing you are going to do first. In a tense situation, it is the closest thing to using force that you can do, without actually using force. Moreover, the strategy clearly warns against the fact that any operation against military-related infrastructures raises the possibility of dual-use, meaning it might also hurt civilian interests. Additionally, the

DoD report states that the DoD Cyber Command is not precluded from providing cyber capability for the intelligence community. In other words, the report says that the DoD will work with and assist any appropriate government agency in the US regarding all cyber related things. That is a nod to the fact that the DoD often provides support, personnel and equipment to help conduct operations that are undertaken by the intelligence community.

The PPD-20, the classified document mentioned above, also raises some interesting points. It directs relevant agencies to start assembling a list of "potential targets of national importance." The only other situation, that I know of, where there is a national target list like that, is the target list for nuclear weapons. That said, the PPD-20 directive requires specific presidential approval; the President has to give the "okay" for operations with "significant consequences", including loss of lives, significant damage to property, and so on.

Another source we might want to consider is the *World Threat Assessment* report by the US Intelligence Community. This report made a point, which many of us have been trying to make for a long time, that we don't believe in a cyber Armageddon anymore. We mostly believe in something that imposes a "long going series of low to moderate level cyber attacks, from a variety of sources over time, which will impose cumulative costs." Some might call it a death by a thousand cuts.

During his campaign, and shortly thereafter, here is what Trump said about cyber: "I will make certain that our military is the best in the world on both cyber offense and defense […] the united states must possess the unquestioned capacity to launch crippling cyber counter-attacks […] cyber security is not just a question of developing defense methodologies, but offense as well. We have to turn cyber warfare into one of our greatest weapons against the terrorists." On the White House website, a statement uploaded on inauguration day said: "we will make it a priority to develop defensive and offensive capabilities at our US Cyber Command," and his "America First" foreign policy includes the following text: "The Trump administration will work with international partners to […] engage in cyber warfare, to disrupt and disable propaganda and recruiting." This indeed resulted

in a joint initiative between the US and Israel, which discussed denial of cyber space to adversaries and so on.

What could be the meaning of "more aggressive use of cyber operations"? It is clear from the statements above that there is an attempt to use cyber operations more aggressively. The President doesn't mean building thicker and thicker walls around you and around systems. He means taking the fight to the adversary. Here is what I speculate. First, the requirement for specific presidential approval for offensive cyber operations with significant consequences will relax over time. This means a greater delegation of cyber authorities to the theater commanders, heading the six or seven commands, or theaters, in which the United States organizes its military forces around the world. Each one of these commands has a commander-in-chief, a four-star admiral or general, and my prediction is that they will get more authority to conduct cyber operations to support their efforts. It also means elevation of the Cyber Command from a unified subcommand underneath the Strategic Command, to a full unified command on its own. This, in turn, means it won't have to ask Strategic Command for permission to do anything anymore. There will be less oversight. I think it is supposed to happen soon, and legislation about this has already passed.

Another thing we will see is that cyber will become the instrument of first military use, when diplomatic, economic, and law enforcement measures fail. I am not here to say that the new administration is going to resort to force, even cyber force, as a means of first resort. However, in the event that diplomatic, economic and law enforcement efforts don't succeed, cyber is going to be the first military force that is to be applied. This has several implications. One is that there will be continuing or expanded efforts to establish ubiquitous presence on all potential cyber targets. One of the realities of cyber attack is that if you are going to attack a target today, you need to have infiltrated it two years ago so that you have access to it. You have to scout out all the vulnerabilities, get all the access paths, maintain them, and so on. And since you didn't know, two years ago, what you want to attack today, you have to do it everywhere you can possibly do it. I also suspect that there is, in here, an implied greater willingness to use

active defense measures, which are disruptive or destructive. In other words – hack back. Next, there is something called the "vulnerabilities equities process". When the United States identifies a vulnerability in a system, there is a dilemma: do we stockpile it for our own use, or do we tell the vendors about it, so they can fix it? Right now, the bias in the process is, allegedly, towards disclosing it to the vendor. That is the claim, anyway. However, I think that this process is going to change, and be more biased towards keeping them around for future use.

Another thing we can look at is that Trump, during the campaign, came out very much against Apple, and its decision to not cooperate with the FBI in unlocking a terrorist's phone. That debate has gone on for a while, and Obama tried to find a middle ground. Eventually he publicly said that he was on the side of law enforcement on this issue, but he was reluctant. Trump will go down that path enthusiastically. In May 2017 Trump signed the Cybersecurity Executive Order, citing the need for "strengthening the cybersecurity of Federal networks and critical infrastructures." It only had a few things that were not about generation of reports: holding agencies accountable, managing cyber security risk as an executive branch enterprise, usage of the NIST framework, and making sure that everybody has a preference for Cloud. This is not bad, but a sort of "modest efforts" type of order.

Lastly, I want to touch some outstanding issues, that we don't know how to react to. How do we interpret presidential tweets, on cyber or anything else? What do we need to make of it? The cyber strategy, with respect to Russian operations in particular, is an interesting question. If we think about the election hack, it wasn't really a hack of an election. Yes, they got emails, but people have been getting emails all the time. What the Russians were able to do is to leak and distribute information selectively, and amplify information. Most of these things are perfectly legal. The only illegal thing they did was to attack the DNC's and John Podesta's email accounts. Is the US going to respond to Russian interference? The President says no, they didn't interfere. It is unclear what we are going to do in this matter, and it will be interesting to see.

The significance of resolving details of cyber operations at lower levels of authority is, at the very least, questionable. Are military

commanders really the best parties is to decide on diplomatic, law enforcement, and economic equities? I don't know. We also need to think about the possibility that somebody at the top might be an impulsive decision maker. Imagine that instead of a cruise missile strike on Syria, there was a large scale of cyber attack order, with the same kind of careful thought and deliberation that went into that.

All of these are interesting questions, ones that that we don't know what to make of. What I can do, what anybody can do at this point, is guess.

Where Does US CYBERCOM Fit?

Martin Libicki, Professor, US Naval Academy, USA

Since its creation in 2009 and up to the present, CYBERCOM is a sub-unified command that reports to the US Strategic Command. The latter, frankly, could not care less about cyber, and would not have done anything about it if it wasn't in the job description. What they really like is to play with nuclear and space stuff. It is also joint at the hip to the NSA, which makes some people at NSA happy and some people at NSA less happy. There are proposals out there to separate CYBERCOM from NSA, or in other words, have a separate director for CYBERCOM and a separate director for the NSA, and to take CYBERCOM out from the US Strategic Command. If you look at the reasons for this, they are mostly bureaucratic, and I won't go over them. I am, however, going to propose where CYBERCOM should be; how it aligns to other instruments of military power, and has everything to do with what you would expect for delivering its primary focus. CYBERCOM has offensive and defensive responsibilities, but I am going to concentrate purely on the offensive ones, and I am also going to assume that offensive cyber operations can actually do something efficacious.

There are four possible options for offensive cyber operations. In going through each of these various options, I am going to ask the question: who are the hackers playing with? If it turns out that the hackers are playing with, for instance, the special operators, then you have a prima facie case for thinking about an organization that puts the hackers and the special operators together. However, not all of these missions are going to involve the same colleagues.

The first option is Tactical, supporting warfare and carrying out what at the US is called kinetic operations, the use of force, violence, etc. Its focus is concentrated, to a large extent, on taking apart opposition commanding control. The first case I would like to make is that if you are serious about CYBERCOM in a tactical position, you have to think about aligning it with electronic warfare. A command and control system is a graph, which has nodes and links. Nodes can be

taken care of through cyber warfare, and links can be taken care of by electronic warfare. More controversially, cyber operations would be aligned with signals intelligence, which is to say, cyber espionage. If you take a look of the structure of US law, we have Title 10 about the military, and Title 50 for espionage. We tend to think of them as two separate operations. That said, I read a story by Fred Kaplan, who has also written a book about cyberspace, which may or may not be true. In this story, hackers got into the Islamic State's computers, and faked an order, telling the Islamic State warriors to meet at a particular position. When they got there, US operators were there ahead of time, and were able to conduct a very successful military operation against them. Now take a slightly different scenario: instead of faking commands, the US operators were listening to the commands, and when they found out there was going to be a meeting of fighters, they got there a few hours ahead of time, ending with the same results. There are some subtle differences between the two, but making things happen, and listening to things that are about to happen, are not so different.

Here is another example. There isn't a lot of information available regarding the Russian cyber operations in Ukraine. There was, however, one report by Crowdstrike, which is a US cyber security company, that made the argument that the Russians had developed a piece of malware for Android phones, that was triggered when these phones were used for artillery targeting. The Ukrainian army used their phones for targeting, because the targeting system could use the GPS of the phones, which led to greater accuracy. However, the Russian malware on the phone changed the behavior of the targeting system, such that it broadcasted the location of the Ukrainian forces, using the phones, to the Russian forces, which then went ahead and took military actions against them. Technically, espionage; practically, probably more effective than simply blowing up their phones, which would have been an act of cyber war. The distinctions between collecting information and destroying information or altering information, are distinctions that make more sense in law than they actually do in the field.

If you want to put tactical cyber espionage together with the cyber command, what do you do with strategic espionage that doesn't belong in cyber command? What do you do with those guys who listen to

foreign leaders and tell the US President what the President of Russia is doing, for instance? Where do you make that split? Furthermore, even though cyber operations tend to be logically targeted after military targets, it is a fact of life that military targets are generally harder than civilians ones, which means you may have to do more investment, and you may not get good results.

The second option for offensive cyber operations is the "gray zone" focus, otherwise known as "hybrid conflict". If you are debating between a cyber attack or a kinetic attack, there are many reasons to favor one over the other. However, there may be situations under which cyber attacks are possible and kinetic attacks are simply off the table. In other words, when you are in a hybrid warfare situations, the relative attractiveness of cyber attacks increases, and that may be where you want to concentrate your cyber attacks on, and therefore concentrate your CYBERCOM capabilities on. In this option, cyber attacks are covert and even clandestine, which might be a good basis on which to organize a cyber command, but the problem is that we already have a clandestine service, and we know that it uses cyber attack tools.

The third option, "informational operations", is a little more complicated. The idea is that many things that used to be lumped under informational operations are actually fairly similar to cyber. Cyber attacks are similar to electronic warfare, psychological operations, weaponized surveillance, device hijackings, and more. I would argue that they fit in the same niche because they have many similar characteristics: they are non-lethal, their effect is far too unpredictable for combat use, they are often ambiguous in origin, and they have a persistent generation capability; hackers can survive the demise of a country's government, for a certain amount of time. Because they share the similar niche, one can understand why you might want to use them under the same circumstances, which leads me to the next question: if you are going to use them under the same circumstances, why don't you do the planning of their use by the same organization? This is the argument that was written up in the *Strategic Studies Quarterly*, about six months ago, in a paper about the logic for thinking of cyber attacks and cyber espionage as part of an integrated operation. There

are issues with this. First, Propaganda and informational warfare come more naturally to some countries than it does to others. Second, if you actually have an overall information operations orientation of CYBERCOM, you probably want to put somebody who understands psychological operations as the head of it. This, however, will cause other problems as the rest of the organization is highly technical, and requires management that can handle technical topics, meaning there is a need for two different profiles here.

Finally, there is the strategic focus, which is where you use cyber attacks as part of your overall deterrence and compulsion policy. In other words, if we want to influence a nation's behavior, we can do it by presenting them with an overall integrated threat. This can go all the way from nuclear warfare on one side, to informational warfare on the other side, depending on the circumstances under which you want to use the threat. For instance, if you look at the Russian use of tactical nuclear weapons in threat postures, you really see this logic taking hold. In the US, however, we tend to be much quieter about our use of nuclear weapons. We generally don't have as overt a compulsion policy, using nuclear tools as we could have, and as other countries do. When we do want to compel other countries to do what we want them to do, we generally use civilian agencies; we give them the threat of sanctions, we give them the threat of judicial action, and at least in the Obama administration, that has been the primary instrument. In this case, you are left with the dilemma of whether you put CYBERCOM under Treasury or under the Justice Department, and that is an absolute non-starter, neither which is a viable option.

To conclude, there is no obvious basis for the organization of CYBERCOM, because a lot depends on the relationship between how you want to use cyber attacks in the pursuit of national policy, and the particular vectors through which national policy operates. If I had to make a guess, I would say that the tactical missions of CYBERCOM should have the highest importance, followed by information operations, Gray Zone operations, and strategic operations. These subjective circumstances are also subjected to the character of the commander-in-chief, because different commanders-in-chief will have different instruments for expressing national power. That said, the first three

options tend to put electronic warfare, tactical cyber espionage and cyber operations under the same roof. If you take a look at other countries, you can see them evolve to that standard. These countries, as best as I understand, include China, which has formed the Strategic Support Force. From what I understand, I get the impression that Germany is going down that route, and I also get the impression that there is a strong similar element in how Israel organizes its army forces.

I advocate that before we rush in and say that we, meaning the US, made up our mind about the place of cyber attacks, perhaps some fundamental thinking about what we want to use them for, will not be missed.

Global Reactions to Endogenous Growth

Prof. Steven Weber, School of Information and Department of Political Science, University of California, Berkeley, USA

I want to ask you, the reader, to do something that might be a little bit unusual, and put yourself in the shoes of Mohammad bin Salman, the Crown Prince and the main driver of economic development and economic growth in Saudi Arabia. He has been given a task of diversifying the economy from oil, and I want you to imagine that you're in the position of trying to figure out what his vision for the 2030 economic growth plan will look like. The reason I'm asking people to do this is to get people thinking a little bit about what is called High Development Theory, for this data enabled economy. I want them to do that not only in the context of the awareness of the extraordinary economic growth potential that is embedded in all that, but also the awareness of all the risks and vulnerabilities of which we have been hearing lately. How would you think about that from the perspective of a wealthy, but non-diversified economy? Another question, if you prefer a different country, is how would you think about that from the perspective of a middle-income country, which is behind in the data and cyber aspects of growth.

There are big ideas for development and economic growth that used to be quite popular. In the 1960s people talked about import substitution industrialization, and the idea was that in order to grow quickly, you needed to have an entire supply chain inside your own national borders. That has largely been discredited today; it doesn't make a great deal of sense when you think about the global supply chain for digital goods or for data. In the 1980s and 1990s there was something called "The Washington Consensus", which was a big development idea which sort of turned the previous idea upside down, saying that the way you get growth is by joining global supply chains, and putting yourself and your economy in a position where global suppliers want to be in your location. The whole macro economic policy dimension around that was about making yourself an attractive

place for an investment. That has been a little bit discredited as well, largely as a result of the global financial crisis.

These days, or in the near future, with the extraordinary capabilities that are inherent in data economy, and particularly in the context of machine learning and artificial intelligence systems, we are asking ourselves, what do we do next? The core argument that I wish to make here is that the answer to this question is complicated, and also enriched by the following extraordinarily simple insight: the data economy creates very strong positive feedback loops, that link large scale data sets and raw data with rapidly improving machine algorithms. Those things together create a kind of semi-natural monopoly dynamic, where, in very simple terms, the best get better faster than anybody else, and the further you are ahead in the race the better you are, and the losers fall behind at an increasing rate.

I think that intuition is starting to take hold in competition policy concerns inside the United States, when looking at the very large data enabled machine learning companies, sometimes known as FANG (Facebook, Apple, Netflix, Google). Others have used other analogies, but the story is very simple: if you sit at Google headquarters, you get more data, and so you create better data products. The better data products you have, the greater market penetration you get, then more people use them and that gives you more data. This is, in a sense, the epitome of a learn-by-doing system, and there are many complimentary growth effects that start to kick in. Naturally, the best data scientists want to work with the best data sets, and they are going to be attracted to places that have access to the largest and most interesting data sets, which leads to further positive feedbacks. These are the kinds of dynamics that lead to winner-takes-all economies, or winner-takes-all companies, in case of a national economy. The fancy term is "endogenous growth theory", but winner-takes-all is simpler.

That is why I think, in the context of that question of how do we enter this next growth stage, it actually really matters that the vast majority of the leading data Internet-enabled platform companies are currently located in the United States. They are American firms, even though they don't like to talk about themselves in that way, and there is something to worry about – they have an accelerating

advantage over the rest of the world in the highest-end data products and machine learning algorithms that go with them. In some kind of an imaginary world, we can take that to an extreme; we can see a point of view in which, even if not only in the United States, the majority of data-intensive business takes place in one or few countries, and is not spread even remotely close to equally around the world. Those countries, then, own the very nice upside of data enabled endogenous growth. They are able to make very significant investments in human capital; they have higher rates of growth; that may spill over into other sectors, including the advanced military technologies that are enabled by similar kinds of growth; it becomes much harder, over time, for other companies or other countries to catch up; and you get a feedback loop.

It is almost obvious that no feedback loop like that goes on forever, but right now it is quite hard for me to figure out where that feedback loop would break down. I think it is also empirically hard to see, based on the evidence of global economic growth. It will break down at some point, but we don't know when it will hit its limit. Therefore, one can think of that as a data economy version of new growth theory. Over the years, Paul Romer has made an argument about ideas as a kind of core production factor, driving innovation and economies. The particularly interesting thing about the data machine learning loop is that, in many respects, it is the core engine of economic growth. This is because it distinguishes good ideas from bad ideas more quickly than human beings can do, and that is, in endogenous growth theory, what actually makes economies grow.

There is another element to that which is worth mentioning, and that is how economies like that are best governed, and what are the right legal restrictions on various kinds of activities. In endogenous growth theory, that is called "meta institutions" or "meta ideas", which are ideas about how to govern other ideas. But theory aside, you would naturally conclude that countries or companies and companies within countries, which are ahead in the game of producing those products themselves, will probably be ahead in figuring out what those meta ideas and governing ideas would have to be. These are the things that are needed in order to try and keep those positive feedback loops

going, and they are more likely to develop in places that are already ahead and advanced in the data economy; it reinforces this winner-takes-all problem.

I should add, though, that this is not a conventional view on data-enabled growth economics. If you try to search the academic and semi-academic literature for views on how to think about that, for the most part you will find what I think is an almost naive "a rising tide lift all boats" kinds of view, which are similar to the old arguments about the Washington Consensus. The McKinsey Global Institute has made a very strong set of arguments, saying that if you want to grow in today's economy, what you need to do is to get involved in data flows, in the way it used to be about global supply chains. It doesn't really matter what your position in those data flows is, as long as you are in the middle of it. That point of view has been largely discredited by what we have learned about global value chains in a previous era. There are some views that say that data is completely different in all respects, but I think it is not as different as people think. We have had those arguments about Intellectual Property in the past. Finally, in the United States these days, there is a particularly shocking point of view about how people ought to respond to this kind of challenge, and the best way to describe it is as the "get over it" argument. This point of view is basically telling people to stop complaining about the fact that American firms are winning in that race and start competing more aggressively. I think that is willful blindness of hegemonic powers, assuming that what is good for them is good for the rest of the world. There is a definite lack of a compelling response for countries who are behind in that race, and wanting to really play that game.

I will conclude with a few of thoughts about what that means. The first is that we absolutely need better metrics for how data flows actually move around the world, and how they affect economic growth patterns. All of these metrics are imperfect, but the goal is to do better than we currently do. I think that what we will eventually need, and what governments will be looking for, are data current account imbalances, the same way we talk about current account imbalances in regular trade flows. I think it would matter whether a country is importing or exporting raw data products, as compared to very sophisticated

data products with high value-add, in the same way it does matter if you're exporting potato chips or computer chips.

Now, think about the options that present themselves to Mohammad bin Salman, or someone like him, if he were to accept this argument, even provisionally. If you put yourself in the position of trying to understand what is the scenario that he worries about, the answer is domination by value and global data firms, which have absolutely no presence and create no value in Saudi Arabia. His first option is to try to join that American dominated data value chain, and bargain for leverage, even though he doesn't really have much leverage, because Saudi data is not that important to Google. Second, he could consider placing his bet on the only other major alternative value chain that is emerging in the world today, which are data value chains emerging around Chinese platform companies. A third option would be to try to do a little bit of both, and play them off against each other. Fourth, for which I think we're starting to see some elements of behavior that correlate to it, is that he could try to disconnect from the global value chains, and create his own, whether nationally or regionally. In earlier days we would have called people like that national industrial champions, today we might call them national data champions. For better or worse, one can justify that kind of economic protectionist behavior by pointing to all sorts of other related concerns like privacy, encryption, espionage and so on and so forth. Edward Snowden, in other words, has provided a great rationale for doing that, and I think that this is the most likely choice for Saudi Arabia. The way to do it is complex, but in essence it would be a return to a nationally-based import substitution model, where the existing currents that we sometimes refer to as the "Balkanization of the Internet", would be reinforced by economic development trajectories and objectives.

When you scale that up, that is a very different landscape of global political economy than we are used to or expect to see. It is different, obviously, for some direct security consequences, if you are trying to build, buy or make, for example, lethal autonomous weapon systems that depend on machine learning. Perhaps even more interesting are the indirect sort of consequences; it is actually really hard, for me at least, to see how that kind of a global economy sits in a compatible

relationship with traditional alliance structures, like NATO, or some of the nascent arrangements that are developing in Southeast Asia.

We don't know yet how regions such as Saudi Arabia and the Balkans, which are currently aren't able to compete in the global data flows market, will react. We do, however, know that staying put is not an option for them, so whatever choice they make will have a global impact. What that would be remains to be seen.

Data, Power and Sovereignty

Prof. Frédérick Douzet, Professor of Geopolitics,
Paris 8 University, France

The research I would like to present in this paper began following a conversation me and my colleague, Stéphane Grumbach, had with Steven Weber, where we realized that the issue of data and sovereignty was really important in Europe, but wasn't really an issue in the US. The belief is that the data revolution is a game-changer for economics, but also for geopolitics and in particular for states sovereignty. The problem is that we know very little about this, and we need new metrics. In truth, I have more questions than answers.

My research tries to offer a glimpse into the geography of data flows that are linked to social platforms, and what options we have in that area. These are data flows that are connected to the most powerful actors that have emerged from the data revolution, platforms like Facebook, Google or Baidu, that operate across borders. We call them "intermediation platforms", because they facilitate communications and transactions between Internet users, and therefore they intermediate. They offer services that disrupt traditional economies, they have economic models that challenge conventional powers, and that tend to raise very complex social, political, legal and even security issues. As we all know, these platforms collect and control enormous amounts of data, but their geographic distribution is not uniform. This generates major imbalances in data flows that are strategically significant. The problem is that we don't have a well-measured understanding of their geography, and we don't have a very good conceptual framework to discuss their broader geopolitical impact.

We can give many examples of how this affects the strategic field, and I will give here three very well-known such examples. In June 2013, the Snowden revelation came as a wakeup call for European nations, because that is really when they realized their extreme dependency on foreign platforms, and the fact that there might be a strategic advantage for the US hosting them, at least at the time. In February 2015, just after the terrorist attacks in Paris,

the French interior minister toured the Silicon Valley to ask the very same companies for their help in fighting terrorism. The terrorist attacks brought to light the helplessness of the French government, which had no authority over these companies. It completely changed the dynamics of how you can fight terrorism, because you have to cooperate with these platforms. Then, in February 2016, Apple refused to help the FBI unlock an iPhone in the name of consumer's data protection. This shows that today, even the US governments has to struggle to obtain critical data from the big private companies. I think these three stories somewhat illustrate some of the most fundamental challenges to national sovereignty in the digital age. We see that data transforms power relationships, both between nations with this issue of data imbalances that can matter, but also between nation states and the private sector. These are changes we can obviously observe, but we don't know yet how to measure and analyze them.

The goal of my research is to try to understand how these platforms, and how these imbalances in data and data flows, can transform the geopolitical landscape, and how they transform the exercise of sovereignty by nation states. What we are trying to do is to develop measures, and to map the data flows link to platform activities. The difficulty is to select the data flows that are strategically important for national power, in order to later assess their impact on the balance of power between nations. Right now we don't have any satisfying way to measure platform activity, because most companies like Alexa, which we used for our first map, or Traffic Estimate, or the Israeli SimilarWeb, capture only a small sample of the traffic. We could aim for companies that capture a big sample of traffic, like TeleGeography, which has about 50% of the traffic, but the problem there is that the data is not directional, and you don't have a breakdown by platform. The result is that no single source will give us a full, accurate picture. We probably need to combine several sources to get this process started, and then eventually encourage companies and maybe states to consider the strategic importance of these data flows. Ideally, we would like for them to start collecting the data, or organize the data they collect in a way that we can use.

We started doing that, based on previous work done by Stéphane Grumbach and a few other authors, with a selection of 30 countries. This shows that there is a huge concentration of data traffic in a few platforms, and if you just take the top 25 platforms that attract the most visits, you have half of the traffic for the top five hundred platforms. Of course, Google is the first, by far, followed by Facebook. In all countries, except China, Google directs more than 25% of the traffic to the top 25 websites, meaning that Google does have a global influence. If you look at the top platforms, the Chinese platform Baidu is ranked fifth. However, it ranks in the top 25 only in four countries; it is huge because of the number of visits, but its influence is not global, it is very regional. Furthermore, not only did we find a huge concentration in a few platforms, we also found and huge concentration of platforms in a few countries. Out of the 30 countries, most of the traffic goes to the US, about a third of the traffic goes to national platforms, and then the rest goes to the other countries. What we decided, therefore, to do for this project, is to map the data for Europe.

We started by using Alexa and Traffic Estimate, and for our first map we counted visits to the top 25 websites from each of the 28 countries in Europe. The data clearly showed that most of the data from those countries goes to US platforms. It actually shows number of visits, but this is a proxy for the data that is harvested at the same time. In France alone, foreign platforms captured 78% of all the visits. This is interesting because if you take other indicators, such as economic powers, the European Union and the US are approximately the same, but when you look at data exchange, the data going from the US to Europe is next to nothing. One can say that this is a business model, that the world is full of imbalances and everybody is free to compete; Europe just did a poor job at developing platforms, maybe we should just stop complaining and start developing. However, that is not the way it is perceived in Europe. It is perceived as a critical sovereignty issue, and now there are calls for digital sovereignty, even though we are not fully sure what that means. The real questions are: how do we jump in, how do we bridge that gap, and what are the implications?

Another country we looked at was China, which developed its own platforms, and is actually a leader in the digital platform industry.

We used the same type of Alexa data, and created a similar map with China. What we found out was an entirely different picture. Most of the traffic coming from China is internal, going to domestic platforms. Because the data is not reliable, and because we only have a small sample with Alexa, we tried another measure. We managed to get fifteen minutes of extraction of all the DNS requests for China, about 150 billion records. This is a very large set of data, so instead of counting visits to websites, this time we counted the number of DNS requests to specific IP addresses with a location. We got a slightly different picture, but not significantly different. We still had the same trend, since about 73% of all the traffic goes to Chinese websites, while 24% goes to the US and the remaining 3% goes to all other locations. We did find one very interesting thing. Like most websites, the Chinese websites we were looking at have web trackers, which allow for the analysis of who is doing what in the website; this is mostly used for doing web analytics for commercial purposes, when you want to target consumers with publicity. If we look at that traffic, then we have a completely different picture. Most of these trackers are operated by US companies, meaning that a major part of that data is actually going to the US. This, in turn, shows the limits of the Chinese strategy in getting control over its data.

As I said earlier, there are many more questions than answers at this point. We need to continue this research because we don't really know what type of interconnections there are between these platforms, which might affect this geography; we don't know which capital shares there are between the different platforms; we don't know who exchanges data with whom, who is going be able to cross what large sets of data with other large sets of data; and we don't even know what data is of potentially strategic value. For example, do I care about what goes through Netflix? Not really. Do I care about the data collected by Pokémon GO? Probably yes, because it can tell me what everybody is doing on a daily basis, where are they going and what pictures they are taking, and in general it can allow me to collect a lot of personal data that could have strategic importance.

All of this is raising very important questions. The first question, of course, is: how do these platforms change the balance of power

between states? We can always argue about the potential influence of our social practices, like over-usage and the promotion of certain values. There is also the whole issue of control over data, because you can argue that the data falls into the states' privacy and security legal framework of these platforms, although this tends to be contested now. This is especially true in the wake on the Snowden revelation, because platforms tend to develop data localization in order to escape from the authority of their country. They argue that they can't necessary know where specific data is located, or they can switch it very fast. This is not a question that is completely settled from a legal point of view, but it is an important question, as it can facilitate potential access by government, and it can empower economic growth and business development. The other set of questions relates to how this data geography changes the exercise of sovereign powers by states, even the states where these platforms are located. We see that the ability for these platforms to collect, process and cross immense amounts of data provides them with a form of power that can challenge state sovereign powers in many ways. Alternatively, if these platforms can be used or incentivized or coerced to cooperate with the government, they can contribute to the reinforcement of its powers.

We are in a situation where, from a national sovereignty perspective, these platforms are becoming both competitors and essential partners in the exercise of sovereign powers. We see that in the security area, and we also see that in the ability of these platforms to provide services that are adopted by citizens and become almost basic utilities for them. The more data that is out there, the more they are able to provide services that will potentially replace public services, and we have no idea what are the implications of this. We already see that in some areas that used to be very critical for states, such as cartography, now these platforms can do much better. We can even see that in some areas, especially counterterrorism, these platforms end up performing decisions that used to be the exclusive prerogative of these states.

One possible global response to these processes is the idea of protectionism, meaning that each country or set of countries will build their own isolated platforms and services. Russia has taken the opportunity of the Snowden revelation, and the fact it had already

developed their own platforms and an ecosystem, to develop a discourse about state sovereignty. Then, in turn, they used that discourse and turned it into a law that forces companies doing business and using data from Russian citizens, to localize data on its ground. Just because you don't have measures of what happens, it doesn't mean that you cannot have policies that, in turn, are going to have a strong impact on the digital sphere. We measured the proliferation of data centers in Siberia, and it is very interesting. Because of the low energy cost due to the very cold climate, and because it has all the infrastructure from the Gulag time, Russia will potentially be able to attract a lot of Chinese data on its territory, and reinforce its domination on central Asian republics. This is supposedly a simple issue of data, but it is strongly intertwined with the issue of sovereignty, and is a perfect example of why we need to do more research, and understand what the strategy implications are.

Bonnie Brinton Anderson, Professor of Information Systems, Brigham Young University, USA

You may have heard this quote before: "given a choice between dancing pigs and security, users will pick dancing pigs every time." This is mostly an unfavorable view of security users, or users in general. It assumes that they are indifferent or careless or lazy. But we don't think that is the case, and what we are trying to do in our lab is to look at the neuroscience tools and see if there are ways that we can make security more usable for users. Some things that happen in our brains cause us bias towards certain behaviors, and we try to work with those biases instead of against them.

I am going to discuss three specific areas that we have been working on in neuro-security. The first is dual-task interference, or as some people call it: multitasking. The second is about habituation, which you may know as "tuning out" or warning fatigue. The third is generalization, when images start to look similar, even though they are different.

The best way to demonstrate the dual-task interference problem is by watching a video. Studies have shown that while we are concentrated on the main plot of a video, we tend to miss insignificant parts happening in the background of the video, or outside the designated video box, if we watch it online. For example, if a pop-up message comes up on the side of the screen while we are busy watching the video, chances are that we will miss it completely. Most of us multitask all the time, and we generally feel like we are capable of doing everything at once, but the truth is that if we eat a hamburger and text while driving a car, there is a good chance that we will hit the kid on the bicycle. When we are doing stuff on our computers, and an important security message pops up, we almost never think it is a good time to stop what we are doing and address the security concern or notification or whatever it might be. Most of the time we actually think this a really bad time.

We wanted to see how bad this problem is, so we conducted a study. We scanned people with our fMRI machine at our lab at BYU and watched what was going on in their brains while we are

interrupting them with security messages at bad times. We started by establishing a baseline of their brain during rest, and then added a memory task, asking them to memorize a seven-digit number in five seconds. This task is utilizing the subject's working memory, which is supposed to simulate their brain while they are doing something on their computer. Then we gave them a security test. We showed them a set of permissions a certain weather app is asking for during installation and asked them to decide whether the permissions are reasonable or not. For example, a weather app asking to get access to all data on your device should be considered a bad thing. Finally, we did a high DTI test, which is a high multitasking or high dual-task interference test. We had them memorize a number, then asked them if the app is requesting reasonable permissions, and then asked them to recall the number they memorized. We found out that the security test interferes and inhibits the medial temporal lobe, which is where the Amygdala and the Hippocampus – both associated with working memory – are located.

When the subjects were asked to just look at an app that was asking for potentially scary permissions, they made mistakes, on average, 7% of the time, but when they were doing it in a bad time, this increased to almost 23%. We also wanted to see what will happen if we keep the tests but change the order. We asked them to memorize a number, recall it, and then deal with the security issue. The result of this low-DTI test was that mistakes were made 9% of the time, almost as low as if the only thing they were doing was dealing with the warning. And so, our recommendation is to do things at better times, if you can. We work with the Chrome security team, as they fund some of our research, and together we brainstorm to try and find what are good times, what are bad times, and if there is a security message that could be delayed. Some messages are critical, and so they have to be prominent and make the user stop what they are doing. Others, like a message telling you that some problem has occurred, and Chrome is suggesting to fix it for you, can be adjusted. We had about 1,000 subjects watching a video or doing some other task, and we would alternate when the message would come up. We saw that some of the better times were right after the video ended, or when

they were waiting for a page to load and had nothing else to do. We sometimes artificially added the rotating sign that signals the user to wait and saw that more of them noticed the message then. In terms of numbers, on average, during a good time or low interference, the subjects disregarded the message 36% of the time, but at bad times, they ignored it 80% of the time. The bottom line is that our brain is not good with handling interruptions, so if you can delay them to a more opportune time, you will get better results.

The second part is Habituation. We see many notifications of security updates, and eventually we start tuning those out. To test this, we exposed subjects to 40 security messages, while doing an MRI. We looked at the parts of our brain that are associated with virtual processing, like the Occipital lobe, and saw that the more times you see security messages, the less your brain is responding to them. Our brain is relying on memory instead, like a computer does with cache; it is faster to pull from cache than to reload data. Another thing we saw are brain activities associated with inner tension or boredom. The more we see something, the less we care. To mix it up a little, and to see if we could improve this, we changed how messages looked. The text remained the same, but we would put a border around it, or highlighted it, or we had some animated ones like a little fidget or zoom. This showed significant improvement in results. By constantly changing the way messages looked, we got the brain's attention. It is like we were able to tell the brain "this is important," and although we know it is saving energy by trying to not load this from scratch, this time it has to. We could see that the attention was higher, and boredom was lower, when we went through those variations.

This led to another question. We understand that if we tweak a message then our brains pay attention, but how often do we have to tweak it, and how much does it need to be varied? We did a longitudinal study, which is groundbreaking; no one else has done this. We also bought an eye tracker that goes at the end of the fMRI, so we were doing eye tracking, seeing exactly where the subject looked and what was going on in their brain, matched with blood flow to certain areas of the brain and the answers they gave. We did it for five days in a row per subject, and we found that the results were consistent; even

over multiple days, those receiving the varied messages held their attention better than the ones receiving the non-polymorphic ones. But in this study, we also looked at eye gaze fixations, where and how the subjects were looking at the messages, and we could see that the subjects' gaze was more focused on the alternating messages.

However, that wasn't enough for our reviewers at the top journals. They were not content with a lab setting and asked that we also do a real field study. And so, we did a longitude field study on mobile devices, using android, and this time we measured adherent behavior, actually capturing how the subjects behaved, not what was going on in their brains. What we did was to present our subjects each day with ten apps from a random different category every day, like entertainment, business, sport and so on, and they had to pick, download and install their three favorite apps. However, they were told to be cautious about what the download. When they wanted to download a risky application, wanting undue or risky permissions, they got a warning. One group received the same type of warning for all risky applications, while another group got variations of the message. It was not surprising that the people of the variations did better over time. The takeaways are that the human brain is wired to tune recurring things out, but by updating them, we can do better.

I also want to quickly mention generalization of habituations, as this research is just in its beginning. We found out that if different things have the same look and feel, we tune them out as well. We stop differentiating between security messages and general notifications, and ignore them all. If there is a beep, or a popup message, we just ignore them or swipe them away without reading them. We get too many notifications, about incoming emails and traffic and the best route to get home, and so we just don't pay attention any more. What we need to do is either limit the notifications or change the way they look. We are working with the Chrome security team on this as well, to change the way some things look or the timing of those messages, and we call out to any other software designers and announce that we can, and we do more research in this area as well.

To summarize, we are studying the three main biases of the human brain – dual task interference, habituation, and generalization – using

neuroscience. We mostly do fMRI and eye tracking studies, but we also have EEG, mouse cursor tracking, and cortisol tracking studies. We also found out that women generally have a higher surprise reaction to security notifications, so perhaps in the future we will examine how people react if we tone down the messages. In the future we would also like to see if there are differences in the reactions between regular people and security experts. We think neuroscience holds the future for making security more usable, so that we won't have to be the weakest link anymore.

The "Take No Action Paradox": The Media Dependency's Influence on Post-Attack Protective Behavior

Prof. Anat Zeelim-Hovav, Korea University, South Korea

The research I'm about to present here, at this point, is a bit theoretical, and not entirely based on studies. The baseline for this theory and my research are two areas I have been working on for quite some time. I have done a great deal of work in in the area of Behavioral Security. Most people know my work on internal misuse, which started a long time before Snowden came around. Others know my work on the economics of cyber security. This new research is an intersection between these two areas.

When you read industry reports, everybody runs around screaming that security is really problematic, and that it costs too much money. Quite often, especially if there is an issue with private information, one of the reasons that organizations lose money is because individuals tend to react. The basic assumption is that if there was a breach, and 300,000 credit cards were exposed, then the organization is going to lose money, partly because people will get up and leave. However, academic studies that look at the market reaction to cyber-attacks found that, actually, the market doesn't care. Most often, the market may react negatively for about a day, and then it will go back to normal, and nobody will care.

Interestingly enough, nobody went around to ask: what do people actually do? That is what my team and I decided to do. Industry reports claim that if the media talks negatively about an attack, there is a bigger chance that people would react. We didn't believe these reports and wanted to test them, using a theory called Media Dependency Theory. Without going into too much detail, this theory says that the more the media talks about something, the more people react, which is very similar to what we wanted to study. Within Media Dependency there are three measures, and we used two of them: media understanding, which addresses the question "does the media provide you some

understanding of what happened?" and media orientation, which is more reactive, and addresses the question "does the media tell you what you need to do in order to get out of the jam?"

Within Media Dependency research, there are studies saying that media messages could have several effects on people. One of them is what we call an affective effect, which means that people will change their attitude, and one is a behavioral effect, which means that people will change their behavior. Of course, from the cyber security perspective, we would like to see behavioral change, meaning that we would like to see people actually do something if they are hacked.

Within Media Dependency theory, there is a subset of studies that look specifically at Risk Communication, i.e. media messages that talk about risky things. The current body of work on this topic, and specifically articles published in 1979 by Kahneman and Tversky from the Hebrew University in Jerusalem, says that most people perceive risk a little bit different than they perceive positive messages. The idea is that if the consequence of an event is low, it doesn't matter to people what is the probability that it would happen. And that is what we are seeing with cyber security right now. People are assuming that until someone is killed by a cyber attack, the consequences are low, and therefore not truly interesting. That is really how risk communication works.

There is also an issue of centrality – if something is central to you, as opposed to it being in your periphery, you tend to analyze it differently. In short, the studies indicate that risk communication can increase anxiety, meaning it has an affective effect. Many studies have looked at anxiety, but no studies have actually looked at behavioral changes. There are some studies that look at an intention to change behavior, but not at actual behavioral changes.

The last item I wish to mention in the context of this model has to do with trust in the media. We all know about "fake news" and lack of trust in the media, so of course we had to measure trust in media messages.

We took all of these things mentioned above, and built a theoretical model, saying that if the media tells you that something really bad happened, it should increase your anxiety, and it should increase your

risk aversion, which means you should change your behavior. I needed to check my theory, and in a certain way I was really lucky. There was a major cyber attack in Korea, where hackers have attacked three major credit card companies, and over 60% of the adult population had their credit card information stolen. It was a wonderful opportunity to actually examine our theories. Immediately after the attack we developed our instruments, and conducted our research. We surveyed people that were attacked, as well as people that were not attacked but heard about it, which was almost everyone, since it was a very highly publicized case. We compared the two groups, and found out that for people who were not attacked, but only heard about the incident, the model worked very well. In Korea we use a central Public Key Infrastructure, or smart ID, for practically everything, so an attack on one major infrastructure may affect all other major infrastructures. In light of this, we asked them: "knowing what you know about the attack, if you were attacked now, would you change your behavior?" In this case, a change in behavior could be many things, such as changing their password or their PKI credentials, changing their credit card or changing their bank entirely.

We found that a significant part (43%) of the 40% who were not hacked intended to make a change in their behavior. Statistically speaking, these are very good results. However, when we asked the people who were actually attacked if they made any changes following the attack, the answer was an overwhelming no. We also collected qualitative data, indicating that there are many different reasons as for why people did not change their behavior. The most recurring reason was that it would be too costly, or too time consuming. Another one had to do with self-efficacy; we don't know how to do it, we don't have to ability, we don't have the knowledge, etc.

We also started looking at other things that we can research in the future. One of them is called the "numbing effect", which is what happens when people are given too much information, and just give up, saying: "I'm just going to sit here and do nothing." Some of future topics have to do with centrality, anchoring points, and prospect theory, where people feel that if the cost or damage is not that high, if the consequences are not that significant, then they can just forget about

it, as it isn't worth dealing with it. We would also like to conduct a longitudinal study, because we would like to know exactly at what point the anxiety was high. The problem is that we are not entirely sure how that could be measured. Perhaps we can do MRI scans, but we don't know yet how we can actually measure anxiety over time, and see if that makes any difference.

There are other theories, of course. A theory that one of my students proposed was Stockholm Syndrome, which in criminology is the idea that you feel close to your attacker. It isn't exactly what happened in Korea, but many people in the qualitative data that we collected said that the attack was not really the bank's fault. They felt that if they left the bank, then the employees would suffer because maybe they will get fired, even though it wasn't really their fault.

Our study has many potential implications for the industry and policy making, but I have to preface it with the fact that this study was conducted in Korea, so the results would have possibly been very different if it was done elsewhere. Additionally, this was a statistical study. We only surveyed several hundred people from each group, so if we had asked more people, maybe there would have been different results, although I doubt that. And so, these results might explain why the market doesn't particularly care when companies are attacked. If the costumer doesn't care, then why should the market care? We need to ask ourselves, how do we give some incentives to customers so that they would react, to go and change their password or behavior? Following the incident, the banks sent many emails to their clients, encouraging them to change their passwords, but we know most of them didn't. Currently, there are no real incentives for the clients to do anything. They don't receive penalty for any malicious usage of their credit cards following the attack, so they don't feel obligated to react.

Industry reports are propagating this notion of fear; that if there is a serious cyber attack on a company, then the result is a big wave of customers leaving it, but my study shows otherwise. If you are the CEO, and you know that your costumers are not going to do anything following an attack, what are the chances that you are going to invest billions of Dollars in putting extra security layers? Naturally, the companies are not going to tell their customers to throw away

their credit cards, because it costs them a hundred Dollars apiece to replace it. Therefore, what we need to figure out is a message and a messaging system that is effective, that is cost-effective for both sides, and let both sides win.

Who Should Pay for the Damages From Cyber Attacks?

Prof. Joachim Meyer, Tel Aviv University, Israel

I would like to describe the research work that I have been doing, alongside with Ronen Avraham from the Buchman Faculty of the Law in Tel Aviv University, and one of my students, Mr. Yehoshua (Shuki) Cohen. Our funding for the project comes mostly from the ICRC, as part of the interdisciplinary project. The question we ask ourselves is who should pay for the damages from cyber attacks, and I think that our research exemplifies the interdisciplinarity of the field of cyber security. In the academia we see research work on technology, on policy, and on legal aspects, and we do work on human behavior, and all of these play a part in the overall question of how we should cope with cyber threats.

The threats I am referring to are threats that involve users to some extent, and many of the recent attacks are somehow related to human behavior. In 2016, for instance, we saw a significant increase in spam messages that had malicious payloads, mainly related to ransomware. These messages require people to open them and take an action, and if they do this then their computers end up being encrypted. Companies, on their part, very often say that they are willing to offer some protection to people, meaning protect the users, even if this is not their core business. McAfee, obviously, sells security as their product, but even search engines such as Google, now give you some warning about possible phishing. Israeli readers probably knows Bezek Benleumi's "super parrot", who promises that the company, a prominent Israeli ISP, will protect you, even though this is not their core business. We get services from companies, and these companies tell us that they somehow protect us against bad things that could happen. This, of course, begs the question: what can these companies do in order to protect us, and how does it relate to what people do?

Ideally, of course they should eliminate vulnerabilities. If Microsoft publishes a patch for some vulnerability in the operating system then

applying it is the best course of action, as it fixes this specific problem entirely. Additionally, the patching process can be done with minimal user intervention, or in some cases, without any user intervention at all. However, in most cases, actual interference in the end-point is impossible. More often, what happens is that companies block actions; for instance, they filter incoming messages, and they block messages that are identified as spams or as malicious. As before, blocking can be done entirely without user intervention, or with some user intervention in the sense that if people want to override the blocking, then it is, to some extent, possible. And finally, we have alerts, which is the most common and the most frequent way companies can help us secure ourselves. We have this in many products and in many contexts, and this is also what I will be discussing here. It is a field in which I have been interested for a long time, not only in the context of cyber security, and so I've done quite a bit of work there.

If we have an attack, and we are subscribers of a company that promises us some sort of protection, who should be responsible for the damages caused by this attack? Who should pay for it? If my computer was encrypted by a ransomware of some sort, and I have this service from a company that told me it will protect me, who should pay for the damages I suffered? One possibility is to say that the user is responsible. Obviously, this has the advantage that users will be more cautious. They will take more care with what attachments they open, which is good. However, users may also be reluctant to use certain services or products. That is the reason why credit card companies do not, most of the time, charge users for fraudulent usage of their cards. That said, credit card companies are now trying to change this. If you are using a debit card, or if you subscribed to a two-step authentication process, you will be responsible to whatever charges you make and you are likely to be charged.

Another possibility is to have the company pay for it. This helps us in the sense that the companies will try to minimize the damages, but the problem is that we may be flooded by all kinds of alerts, where the company repeatedly warns us against certain dangerous things, to the point that it makes these alerts completely irrelevant. And so it seems that the best possible solution is to find some kind

of a mix, where we divide the responsibility for damages between the user and the company. Then the question becomes, what kind of a mix should we aim for?

What I would like to present in this article is a more economic, mathematical analysis of this question. In parallel, I should mention that we also ran an experiment in which we examined the behavior parts, and as I said before, it is a joint work with Ronen Avraham, who is looking at the legal aspects, as this is really an interdisciplinary problem. The basis of this analysis is something called Signal Detection Theory, which is the tool I am using. It is a very simple idea; essentially it means that we have two states, a message can be either malicious or not malicious, or in the professional lingo, can be either noise only (non-malicious) or signal + noise (malicious). We have some variables we observe, and if we look at all malicious messages then we have a normal distribution of values. This is true when looking at all messages, both malicious and non-malicious. The means of the two distributions are somewhat shifted, and the difference between them is the sensitivity of the system or the detector; the larger the difference, the better we are able to distinguish between signal and noise. How do we decide that something is malicious? We get some observations on its place on the x-axis, and compare it to a certain threshold value. If it is above the threshold, we say that it is malicious, and if it is below the threshold, we say that it is non-malicious. We have two parameters that we find interesting. One is the sensitivity, or prime, and the other is the response criterion, which is essentially the threshold.

At this point, we can to look at the question of dividing the responsibility, and I will describe a framework for analyzing this problem. We have two types of players in this "game". One is the system, meaning the company or service provider, and this player provides objects to the user, be it a service or a digital item. An example for this is a service that forwards emails to the customer. This service may issue alerts if it detects a possibility of malicious content in any item it forwards. The second type of players are users, and they can access objects. An object will have some positive value, some benefit for the user and for the system if it is non-malicious, and it can have

negative outcome if it is malicious. The goal is to maximize access to non-malicious objects and to prevent, to the extent possible, access to malicious objects. Also, if a user opens a malicious object then the damages are not only on the user, or does not entirely fall on the user, but rather the damages are divided between the user and the system, by some proportions we note as R. If R=0 then the damages will be paid entirely by the user, and if R=1 it is entirely on the system. If R=0.5, then the damages are divided equally. It is also possible that in cases where the system issued an alert, and warned the user, the proportion the system carries will be diminished (using a modifier with the notation r). That means it is possible to have an equation where R is close to 1, but still the final result will be 0. In other words, the system is responsible for damages if they didn't take any precaution, reaction, or warned about it.

In our study we developed an experimental program that helped us simulate this type of setting in our lab, or even on mobile phones, where people used their own devices to access the experiments. The program used a set of items, and these items have a certain probability (P_M) of being malicious. We have a system that has a set of threshold (β), and also a sensitivity factor (D'_s). Each item is compared to this β, and once it gets through, it gets into a list of incoming mails for the user. The user can now click on a mail, and if it is identified by the system as potentially malicious, the system will issue a warning. The user will see it, and will also be able to see information about the maliciousness level of this particular email message. The user has two sources of information: the information from the alarm system, and the independent evaluation of the message. Using these, the user will need to make a choice of whether or not to trust the email. For instance, if you get an email message saying that a Nigerian prince has $3.5B that he wants to share with you, you will probably think that this is unlikely to be a correct or a legitimate message. You will not transmit your money. Given the information we get from the warning and the private information, if we download it there are two possibilities. One is that it is non-malicious, and we have benefit for the user and benefit for the system. The other option is that it is malicious, and then we have the cost of it being malicious. The cost

for both the user and the system depends on the values or R and rR. For the user, the equation for damages (D) is $D = Cost * (1 - R)$ if there was no alarm, or $D = Cost * (1 - r * R)$ in case there was an alarm, and the outcomes will depend on the parameters we set.

Our main challenge was to find out what should be the optimum values of both R and r, and we acknowledge that in all probability, there is no single set of values that will always fit. The more accurate problem, then, would be to identify the criteria by which we should choose these values. The answer is that we should choose them so that the behavior we expect the system to choose should be as close to optimal as possible for the user. If we choose a value of R according to which the system will choose a threshold that will lead to very negative outcomes for the user, then that is not a good situation. We would like to have a situation where both parties get similar benefits from it. With that in mind, we computed the payoffs for both the user and the system, for different combinations of parameters. For instance, suppose we have a user with a $D'_{User}=1$, and a system with $D'_{system}=3$, meaning we have a system that is very able to distinguish between malicious and non-malicious items, and a user that is less able to do so. Also, we have r=0.5. What we did was to draw lines for different values of R, from 0.1 (the system is responsible for 10% of the damages) to 0.9 (system is responsible for 90%). Then we look for the point on both curves (user payoff and system payoff), where payoff is optimal for both. In this case, we found out that if R=0.3, and if the system chooses a threshold of 0.73, we get the best results. This means that there exists a single system payoff that would be best both for the user and for the system. We have shown that the possibility to choose something like this exists. From here, we can ask if this is something we would like to have enforced, a regulation of sorts, that defines the right way to divide responsibility between users and systems.

There are, of course, other cases. Suppose now we have a smart user with a $D'_{User}=3$, and a system with a $D'_{system}=1$. The curves will look completely different, and in fact, there is no optimal threshold. We can choose one of two different options here. One is that we go with the 0.1 division, where the payoff for the system and the payoff

for the user are more or less constant, no matter what threshold we choose. The other option is that we go with a high division, i.e. 0.9. What will happen here is that the system would have to choose a very low threshold, significantly increasing the number of alerts. For the user, this is not the optimal threshold, and it will become an annoyance, but it is still better than the payoff that they would get with the R=0.1.

To conclude, the system and user interests are more or less aligned, depending on the parameters, and we can specify a level of R that creates better alignment between the system and service provider user. If we implement such a mechanism, we can generate situations in which the user will optimally cope with threats, and that is essentially what we would like to do. However, implementing such a mechanism requires us to deal with major technical, legal and organizational challenges. Naturally, what I presented here was a simplified version, but in general we study any situation where we have some kind of information source that tells users something about the risks involved in taking a certain action. This can take the form of solutions like an IDS or similar solutions, a single service provider like the one who filters your spam, or something that is much more complex. What we say is that there should be rules that define how the information is provided to the user, and the relative responsibility for outcomes. Because if there is no responsibility, whoever issued this information will act differently.