



KHAN




SANDHU



HAGEROTT



CARLISLE



SHI

Roundtable on Security Issues in the Cloud-Assisted Internet of Things

Samee Khan, North Dakota State University

Ravi Sandhu, University of Texas, San Antonio

Mark R. Hagerott, North Dakota University System

Martin Carlisle, US Air Force Academy

Weisong Shi, Wayne State University

This roundtable focuses on the security issues related to the cloud-assisted Internet of Things. Participants represent a good mix of domain experts from across the United States who are actively engaged in research related to the cloud-assisted Internet of Things.

Samee Khan: First, let me thank everybody for being here. *IEEE Cloud Computing* was started with the intention to reach out to a very diverse set of individuals. They might not only work in academia, but also industry, or they might be up and coming students who want to work in this field. So with each issue, we try to bring in folks who have expertise in the field, have worked in the field in the past, or are currently in a position where they can comment on the policies and procedures. Let's start by having each of you make a general statement about the state of the art of the cloud-assisted Internet of Things and how important security is in that particular system.

Martin Carlisle: We're living in something like the Wild West. We don't have good regulation. We don't have good security. Security is essential in the Internet of Things. We now have all sorts of devices in our homes that are connected to the Internet and made by people who have no experience developing Internet-connected software. If your refrigerator is now on the Internet, you've introduced all sorts of new vulnerabilities to your home network. Perhaps criminals can now use your refrigerator to tell when you're not home and burgle your house. Or maybe they can use it to gain access to your home network and steal your bank account information from your computer.

We really have a long way to go in this area.

Mark Hagerott: That was a great lead-in, so I will build on what Dr. Carlisle was saying, why this is so significant, and how to convey to your readers, how it's so much more significant than they might otherwise observe. Before, when you had the security of the cloud, databases, banking information—you could have backups. You could have hard drives. You could have alternate sites. You could have insurance companies back up your data value if there was a space between what was being processed and what went down. You know, it's all recoverable in a way.

But now that you're moving into the Internet of Things that exist in what I call human space and they're becoming ambulatory—from the Roomba vacuum cleaner, to the thermostat, to the furnace that actually acts on the environment by bringing gas in, igniting it, burning it safely. In contrast, an actor

could hack the furnace, shut off the pilot light, send gas into the room, and then you turn it on. You've now acted on a physical environment that no insurance policy will bring back your house or your life when you're in it. Consider Google driverless cars, and so on. So what's different is the cloud-assisted Internet of Things is now the downside, in many ways, of the Internet, applied to things that move in human space.

This is new—fundamentally new for most Americans.

Ravi Sandhu: To follow up on the previous two speakers, I think what is fundamentally different from a security perspective is that there is physical risk coming into the picture. You know we've worried about security of Web-based systems for the entire life of the Internet. But ultimately, that's about money. Money is kind of easy to deal with. It's a common denominator. Its value is well understood. And if your loss can be compensated with money, it's just typically the mechanism today. We seem to have gotten somewhat of a handle on it. But when it comes to physical damage, which could range from being somewhat minor to life-threatening or certainly threatening parts of your body, I think we're going to see a qualitative change in what happens. I'm not sure how the security research community should address that change, but at least we have to recognize it's there, and think about what it means as to the nature of the research we do.

Weisong Shi: I agree with all that has been said. I think that the IoT combined with the cloud, no matter how we think about its security, it's going to be there. However, while we're enjoying the type of convenience the cloud-assisted Internet of Things brings to us, it also brings a huge amount of risk from both a security point of view—access control kinds of things—and privacy issues. For example, recently Shodan (www.shodan.io) was in the news. It's basically the search engine for the IoT equipment. If you join at Shodan, you probably can see a lot of things that are not intended to be seen.

I think this has created a lot of work for academia. It's an opportunity for researchers, but it also brings huge changes to us as human beings.

Samee: Thank you so much for starting off this conversation. Let me dive into the first question, which happens to be the most difficult one. How do you define cloud-assisted Internet of Things? And how is it different from a nonassisted Internet of Things? Or do you see any difference?

Mark: Well, for a general reader, that would be quite hard. My dad was a Navy technician in the 1950s. He operated mainframes called General Electric Computers and they linked the banks in their region together. That was kind of a cloud. So the question is kind of splitting hairs a little bit. I would say nonassisted Internet of Things is something that is running, it has the processing power to run and do its function, and it happens to be connected to the Internet to allow communications—such as with the

experience that as consumers we already have—things like wearable devices and so on. Unless the data is aggregated in the cloud and managed in the cloud, the value of the device isn't that compelling. So it's almost intrinsic, in my opinion, that data will have to be exchanged between the physical and virtual counterparts in real space and cyberspace. The current situation there is that these things are probably going to start being developed as proprietary islands.

For example, I have an ecosystem of a particular company, like a Fitbit, and an ecosystem of a competing company, and so on. But they're going to be fairly siloed. Eventually the vision would be that somehow they would need to collaborate or exchange data with each other. I think it's going to play out, and we end up with a totally siloed world, structured in sort of corporate interests, or we come up with something more open. I think it's a very different kind of world, depending upon how that technology works, if it works.

Weisong: I would say it is the cloud-assisted IoT that makes it really meaningful for the IoT. If an IoT device isn't connected to the Internet, then it's not really an IoT. I would say cloud-assisted

IoT is maybe the IoT becoming real IoT. But if you're looking for cloud-assisted, I think that could make a lot of things happen. Because before the IoT is not a command or without intelligence from the cloud. The core thing here is, with the cloud, you can combine multiple heterogeneous data sources together, and then it can help you to make more decisions. That's the power brought by the cloud-assisted IoT here.

As an example, in Detroit, there's a new app called ParkDetroit, which is really powerful. For example, my wife can drive to downtown Detroit, and she doesn't need to pay anything [for parking]. She just calls me and I can use my app, even, for example, when I am traveling in China. She's going to see a zone number there, and I can make a payment for the car and the particular location in Detroit. So now this is super convenient for drivers, because you don't need to find coins. I think this is one of the typical examples for how the cloud-assisted IoT works.

The app for the parking authority personnel is even more cool. What they do is they drive a car with a camera. As they pass by, they do a real-time analysis of each of these plates. If you already paid, they go. Otherwise, a ticket is generated immediately to you.

How do you define cloud-assisted Internet of Things? And how is it different from a nonassisted Internet of Things?

furnace example, where someone could tie into and change the thermostat.

Cloud-assisted—I would say it starts to now interact with the cloud in more of a network of things, learning behavior, other data. For example, using the furnace case, you're expecting an ice storm, so the furnace picks up the weather data and heats the house, because we know the power lines may come down, and you want a temperature boost before the power goes down.

So it's splitting hairs a little bit, but lends itself to more complexity and more potentially—and I guess the topic really here is security—more ways that trouble can start. Before it was just I'll connect to the Internet so I can watch the temperature. Now it's actually more of a control algorithm—artificial intelligence, low level—because the cloud is giving it advice and telling it what's best to do. That's my best swing at that, but I wouldn't spend too much time on that distinction.

Ravi: I think, first of all, this is relatively new, the cloud-assisted IoT, at least to me. I think of the IoT as existing at two layers—the physical layer and the virtual layer. The virtual layer is essentially in the cloud. And to get full benefit, even in the little

That's just one example to show that this cloud-assisted IoT can provide a huge benefit here. But later on when we're talking about security, I'm going to use this example again to mention the potential challenges here.

Martin: As was said, the IoT involves the Internet. If you don't have the Internet, you don't really have IoT. So when I was thinking about cloud assisted, if I could use the parking example—at midnight, we don't really need much computing capacity to support what's going on with the parking apps. But at noon we're going to need a lot more computing capacity. That's really what cloud computing is about. It's having a scalable computing infrastructure to support the devices. So, to me, the idea of the cloud-assisted Internet of Things means that we're using an on-demand scaled cloud computing infrastructure, rather than a fixed infrastructure, to support the devices.

Samee: Now that we have a good grasp on what is cloud assisted and what isn't cloud assisted, let me ask you this: what, in your opinion, is the primary security risk associated with the Internet of Things, cloud or non-cloud assisted?

Ravi: Risk involves security. I'll highlight two. From a researcher's perspective, what's interesting is really what is qualitatively new in this arena. And I think I already mentioned one aspect, which is the physical aspect of things actually causing you physical damage, which can happen in many different ways. It's fine to carry around a little wearable that gives you your heart rate, but if it reports a wrong value, an incorrect value, that can cause problems. And I have actually anecdotally known of people who have panicked because they had a bad reading on their wearable.

And so, the physical damage aspect, which can occur in more ways than first come to mind, is one thing. The second thing is the autonomy and the interaction between things, and the autonomy of the decision making. If you look at driverless cars, right now they kind of hand over to a human driver in a panic situation. There are some statistics about how often that happens with Google cars. But if you extrapolate that to a world where mostly we are in driverless cars, these cars will now have to make some decisions. Of course they will want to avoid accidents, but as a practical matter, there will be situations where accidents will occur, if only because one of these cars blows a tire, or the brakes fail, or something like that happens. They're not go-

ing to be perfect. And, there will be an issue about going through a calculation as to how to minimize the damage. Some serious ethical issues could arise here. For example, should they try to save the people in the expensive car or in the cheap car, to put it in terms that anyone can relate to.

Weisong: I think what most people worry about is the privacy issue. But I wouldn't say it's because you're collecting a lot of my information, and then this is potentially being seen by other people. For example, you're sending the data to the cloud, and you have no idea who is going to be accessing this type of data. As a data producer, you generate this data to the others and that is potentially a huge risk, and it's out of your control basically. That's one part. Since you rely on the cloud to make a decision for you, you have to really trust that all the security problems in cloud computing will be reinforced here.

Continuing with my parking example, if somebody is somehow controlling the City of Detroit's database, that person could easily give tickets to anybody if he or she wants. For this reason, for example, in Ann Arbor, Michigan, if you get a ticket, you can argue for yourself. If you can prove that, hey, I'm not in this location during this time, I can prove I'm not here at all, and send the ticket back to the parking authorities, they won't charge you.

With that said, the cloud-assisted IoT is actually open to more risks than the purely alone IoT equipment because you're disconnected. Another example you've probably heard about on the news, is people shutting down, stopping a car on the highway. It exists today. It's already doable, like GM's OnStar system allows you to start your car remotely. This is for a good reason. But if it was in the control of a malicious person, it could be a disaster. Somebody can open your car somewhere. This is already a huge risk for us. So I think that is a particular challenge for the so-called cloud-assisted IoT.

Martin: The big issue I see here is that we now have device manufacturers adding Internet connections to things, without concern for security or privacy, and we're repeating a lot of the mistakes that we made with personal computers, and then with mobile phones. A company might decide, wouldn't it be great if my refrigerator, my exercise bag, or my whatever talked to the Internet? And that company might be really good at making refrigerators or exercise bags, but they could have no idea what to do to protect their customers on the Internet. And they likely wouldn't have any plan for how to update the software on the devices if they discover a problem

later. I think that's the big issue. It's sort of a rookie hour on the Internet of Things.

Mark: These are all resonating. The framework that I used, that I briefed at several places (I was just at the National Convention of Chambers of Commerce) to try to put it in layman's terms is basically that what we're seeing today is two macro events. The massive big data we talk about where you can have so much predictive analytics and the whole cybersecurity implications of that, and then the movement of digital controllers into human space, which in the most extreme case, is autonomous machines, which we're testing in the military and up here in North Dakota, where you just completely turn it over. But the middle ground is human-machine integration, where a human operator has part of it, the algorithm has part of it, the cloud has part of it, and you get something that seems like it's running just fine. Like someone was saying—I know how to run a furnace. We know how to secure the cloud in itself. But then you put the cloud and computing into now an Internet-connected furnace, and you have something as simple as what happens if you have a power surge. Do you start the furnace automatically? Or does a human need to come in and verify that the pilot is there and running? And, quite frankly, depending on the age of the furnace, you have different procedures, but people back-fit things on here and don't think about the security.

So it's that intersection of human, machine algorithms, and the cloud that creates these unperceived combinations. And as someone said in the very beginning, people could really die this time. So it's not just like, oh, we had some data problems. So that's the biggest thing, is that interaction of different skillsets, of knowledge sets, in this gap that's now being filled in an imperfect way.

Samee: This is definitely a new era for us. In your opinion, are the conventional security techniques and standards sufficient and implementable to a cloud-assisted Internet of Things?

Weisong: When we're talking about conventional solutions, we're not just talking about cryptography techniques like public key and private key. I think the answer is no, because right now we don't have a good solution. Most of the security protocols today are used for specific applications. I expect that in the next five years or so, we will see multiple customized solutions for different applications. And then, at a certain point, maybe by 2021, people will sort of get a good understanding here, then they will

start to generalize, oh there's some things that can be applied here. Right now, society is at the stage where everybody is trying different kind of domains.

Martin: I'm not convinced that we have very rigorous conventional security standards. So it's not at all clear to me how we're going to transfer those to the Internet of Things. If I look at the computer on my desk, I regularly have to download patches because we are constantly finding all sorts of security vulnerabilities in the software. We really can't make that our model for the Internet of Things.

Mark: The two vulnerabilities I see that are instant problems for a doctoral student are, when you have big data, and big banking datacenters, and the Social Security Administration—they could have big IPS [intrusion prevention systems], IDS [intrusion detection systems], the latest patches with teams of people watching the patches come in, to build on what Dr. Carlisle said. They could have the strongest encryption. Now with the Internet of Things, you have some small mobile device that has this very limited processor, and has to run on limited memory capacity. So it's got weaker algorithms, yet through the cloud, you can attack it. You can go to a vector where you come in and attack that component, or you break it into a local area network and attack that component, which lets you get into the cloud in some way. So the first one is, you have these things, these tentacles going out from the cloud that connect back to the cloud. But how much can they protect themselves directly?

And then secondly, you have the problem of physical capture. If Morgan Stanley's service center was attacked physically, it would be in the news around the world. The local police would respond. You'd protect the facility physically. If you've got a bunch of Internet things, and someone goes up to a relay station in the power grid and literally, physically bypasses the firewalls, not only physically bypasses them, and goes in, and they're in now, physically. So now we have the Internet of Things ambulatory, out in space, human space, that can be captured and controlled, physically accessed and penetrated, and now a virus enters in and it goes back out.

So again, will that stop the world running? No. And as Martin can probably relate to, there are lots of theories of warfare that have become more and more relevant. We might be just seeing attrition warfare. Cyber warfare could be attrition warfare—good guys, bad guys, narco terrorists, rogue nation states—and it's just a war of attrition. But the Internet of Things is coming. It has just too many good

efficiency benefits, safety benefits to stop. But these are hard problems, and probably tamper-proof technology is going to be a great field. Build a tamper-proof technology.

Ravi: The conventional security standards have hardly proven to be sufficient in the current cyberspace we're already in. So clearly, the answer is that we are going to need something radically different. In particular, to a large degree, we've relied on patching. You let a few people get damaged, then you discover the malware, and you create signatures that help you protect other people from it and other computers from getting the same malware.

The update cycle isn't always feasible for the Internet of Things, in general. You're not going to be able to update these things like you update your PC or even your smartphone. That part of it has to be rethought.

Samee: What do you all think would be a method where one can monitor the end-to-end capabilities of the Internet of Things? With billions of "things" spread across the globe, how do you make sure that the system is still secure?

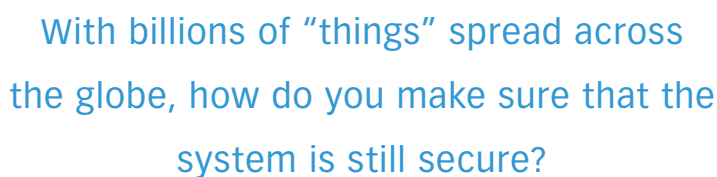
Martin: That's obviously an incredibly hard problem. I think one of the things that we're going to have to do to address the security issues is use strong encryption between the devices and the cloud. We'll need to have some sort of—probably artificial intelligence—techniques to say, "this device is now behaving differently from how I've seen it behave in the past, so maybe something is wrong here." But this is a million-dollar question, and I'm not sure I have the million-dollar answer.

Mark: I find myself agreeing with the Air Force. You keep letting him go ahead of me.

There could be advantages to reinserting uncertainty, heterogeneity into devices you think should be mass produced and all look the same. It's kind of funny using the biological analogy that when smallpox hit, some people made it and some people didn't because of genetic variation. In the military, you have gaps in systems, so you have the air gap type of stuff. Maybe you call it the human gap or the analog gap that certain systems just will not happen and actuate unless they're there, or maybe unless two separate cloud servers concur that's the right thing to do. Similarly, we can consider the Air Force tradition of two-person control of critical systems, particularly nuclear. In nuclear systems, two

people have to both agree and turn the key at the same time. Some analogies of biology and bureaucracy might be brought into automated machines, so one vulnerability doesn't take a whole bunch of them down. This might create some insufficiencies but save costs in long term when insurance companies will rate the cybersecurity of these systems, and you'll get a better rate if you have a lot of these built-in things that otherwise would seem inefficient. I'm sure that's already happening; cyber insurance is becoming quite the thing.

Ravi: I would just add one more consideration that is probably going to be qualitatively different with the Internet of Things, and that's the security of the supply chain or the integrity of the supply chain. You're going to have, just like you have malware, you're certainly going to have counterfeit devices.



With billions of "things" spread across the globe, how do you make sure that the system is still secure?

And you're going to have devices with malware or at least malicious intent embedded in them. We are already seeing some of these issues in our current environment because hardware is implemented all over the world and there are many opportunities along the way to play mischief with it. With the IoT, that's another major consideration that's going to come up.

Weisong: I agree that this is a very new problem here—how do we really monitor the end-to-end? I would like to point out that there is a movement here that we coined as edge computing, also called fog computing by the industry. Different from the traditional end-to-end argument that we put everything on the end, given that the IoT equipment does not have enough computing power and is resource constrained, I think that there might be a direction that leverages these edge devices, which, for example, could be your cellphone for body area networks. It also could be like at home, if you have a smart home, you have a gateway, but it is more powerful. You know the gateway has a home operating system running there. So this kind of end-to-end will be leveraged on the devices and the edge. I think

that edge computing and the edge of the network is probably one of the platforms that could potentially be used for people to design the new monitoring tools to solve these new issues.

Samee: It's certainly a great direction, so let me build on this because, if we are going to use edge devices, cellular phones, of course there have to be some regulations. Is there any need for the government to be involved in IoT security? If it happens to be that the government is trending toward using the IoT in the future, what type of measures can they take to secure their infrastructure?

Mark: I might just answer the first one so the Air Force has something they can say, because I'm sure they'll agree with me. But I will say the answer is yes. And, again, using the physical analogies, we had a time in human society, the feudal society, when basically individuals and businesses—if you want to

cloud services for its people. I just think it's inevitable that government regulation is going to come into play here, and as much as Amazon and Google and everyone else don't like it, that there's going to be more and more.

Eric Schmidt, in his book *The New Digital Age*, talks about this, that more and more restrictions—they call it the Great Firewall of China—but there's some wisdom to the idea that sovereignty extends into cyberspace. And, again, the legal people will tell you now, there's the whole issue of where's the data residing. The Germans are very anxious about their data residing in American companies. I'm not saying anything we haven't all heard about. But in answer to your question, yes. The government is going to have to be in there, not only thinking about where data is stored, but also standards, which is why we have NIST [National Institute of Standard and Technology] and other places that set standards. So, yes is the short answer.

Clearly there has to be a role for the government, but we don't want to slow down innovation.

Ravi: Clearly there has to be a role for the government, but we don't want to slow down innovation. So there's this tricky balance between having sufficient openness, lack of regulation, to allow innovation to take place, versus laying down regulations to control things. Historically, such regulations in cyberspace have been often completely out of date. We are currently seeing debate

call them that because they were agribusinesses, you know, the castle, the fiefdoms—they provided security. And it was called the Dark Ages, okay. It was a dangerous time. When you had obligations to protect the life and limb of its population, in return for taxes, and agreeing to regulations and laws, you had a safer environment to do things. So if all the science fiction writers are right that our identities, our very essence, our privacy, our deepest secrets are going to be going online—and I watch my kids putting more and more online all the time—then the elected leaders and the governments are going to have to have a role.

Another way to think about this: think about data residing in places you don't trust. Can you imagine, right now, the Iranians, who have a private company maintaining security for their people, that gets bought out by the Saudis, and the Iranians being okay with that. The Saudis bought out a company, and now they're going to run cloud services for the Iranians, for most of its people. You can see that case where the answer would be "no." The Iranians want an Iranian-regulated company providing

about encryption—privacy versus security issues. This is an important debate. It's about privacy, it's about innovation, and it's about national security, and our current government governance methods are simply not set to the same speed. If they don't work at Internet speed, they're not going to work with Internet of Things speed. This is a big dilemma. I don't have any good solutions. Somehow if our politicians could be made wiser than they currently are. Short of that, I'm not sure what we could do until we get some really thoughtful people who really care about the country and not just about their personal situations.

Weisong: When we're talking about the government's role in this kind of new security of the IoT—let's take a look at the government roles, for example, on the energy consumption. For example, when you walk into any appliance store, you see all of these home appliances have an Energy Star value, right? I think why the government can be easily involved is they have a well-defined metric. Similarly, we have a fuel efficiency metric for the car industry. However,

the security is very different here because security, right now, doesn't have a metric.

I think the government will need to go to a customized solution. Say if you're working on a smart grid, the government can put in a certain amount of regulations there in terms of how data is collected and used.

Martin: I strongly believe there is a role for the government in the security of the Internet of Things. If you look at bridges, once upon a time, anybody could build a bridge. After enough bridges fell down, we decided that we needed professional engineers. Only professional engineers could sign off on designs for bridges. Or if you look at toasters, we have the Underwriters Laboratories. You can't buy a toaster if it hasn't been certified by the Underwriters Laboratories. We don't really have anything like that for software.

We're reaching the point now where we've had enough failures, enough loss of security, enough loss of money, that we need to start doing something like that. The FAA [Federal Aviation Administration] has very strong requirements for software that's used in flight. We need to think about having requirements for software that's used in devices that consumers are going to put in their homes. Another big step would be providing legislation for liability for the manufacturers, so that if somebody puts a refrigerator in my home that has software, and that software enables people to spy on me through the refrigerator, I can bring a lawsuit against that manufacturer.

Samee: At this moment, I would like to thank all of you for participating in this roundtable, which was enlightening and beneficial on many levels. Thank you all.

The topics discussed in this roundtable on the security issues pertaining to the cloud-assisted Internet of Things are relevant to cutting-edge academic research as well as for the governmental entities to take notice on the policies and procedures related to the cloud-assisted Internet of Things. ●●

SAMEE KHAN is an associate professor at North Dakota State University. His research interests include optimization, robustness, and security of cloud, grid, cluster, and big data computing, social networks, wired and wireless networks, power systems, smart grids, and optical networks. Khan has a PhD in


computer science from University of Texas, Arlington. Contact him at samee.khan@ndsu.edu.

RAVI SANDHU is executive director of the Institute for Cyber Security and an endowed professor of cybersecurity at the University of Texas, San Antonio. His research interests include cybersecurity models and systems, particularly in authorization and access control. Sandhu has a PhD degree in computer science from Rutgers University. Contact him at ravi.sandhu@utsa.edu.

MARK R. HAGEROTT is chancellor of the North Dakota University System. A certified naval nuclear engineer in power generation and distribution, who also served as chief engineer for a major environmental project involving the defueling of two atomic reactors, Hagerott changed from engineering to a career in higher education in 2005. His research interests include the evolution of technology, education, and changes in career and the workplace. Hagerott has a PhD in history from the University of Maryland. Contact him at mark.hagerott@ndus.edu.

MARTIN CARLISLE is a professor of computer science at the US Air Force Academy and the director of the Academy Center for Cyberspace Research. His opinions expressed here are his own, and do not necessarily state or reflect those of the US Air Force or the United States government. His research interests include computer security, formal methods, and computer science education. Carlisle has a PhD in computer science from Princeton University. Contact him at carlisle@acm.org.

WEISONG SHI is a professor and Charles H. Gershenson Distinguished Faculty Fellow at Wayne State University. His research interests include big data systems, edge computing, mobile, and connected health. Shi has a PhD in computer engineering from Chinese Academy of Sciences. Contact him at weisong@wayne.edu.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.