

**THE ENERGY EFFICIENCY OF BLOCKCHAIN
AND SIMILAR TECHNOLOGIES AND
THE CYBERSECURITY POSSIBILITIES
OF SUCH TECHNOLOGIES FOR
ENERGY INDUSTRY APPLICATIONS**

HEARING
BEFORE THE
COMMITTEE ON
ENERGY AND NATURAL RESOURCES
UNITED STATES SENATE
ONE HUNDRED FIFTEENTH CONGRESS
SECOND SESSION

—————
AUGUST 21, 2018
—————



Printed for the use of the
Committee on Energy and Natural Resources

Available via the World Wide Web: <http://www.govinfo.gov>

—————
U.S. GOVERNMENT PUBLISHING OFFICE

WASHINGTON : 2019

COMMITTEE ON ENERGY AND NATURAL RESOURCES

LISA MURKOWSKI, Alaska, *Chairman*

| | |
|-------------------------------------|--------------------------------|
| JOHN BARRASSO, Wyoming | MARIA CANTWELL, Washington |
| JAMES E. RISCH, Idaho | RON WYDEN, Oregon |
| MIKE LEE, Utah | BERNARD SANDERS, Vermont |
| JEFF FLAKE, Arizona | DEBBIE STABENOW, Michigan |
| STEVE DAINES, Montana | JOE MANCHIN III, West Virginia |
| CORY GARDNER, Colorado | MARTIN HEINRICH, New Mexico |
| LAMAR ALEXANDER, Tennessee | MAZIE K. HIRONO, Hawaii |
| JOHN HOEVEN, North Dakota | ANGUS S. KING, JR., Maine |
| BILL CASSIDY, Louisiana | TAMMY DUCKWORTH, Illinois |
| ROB PORTMAN, Ohio | CATHERINE CORTEZ MASTO, Nevada |
| SHELLEY MOORE CAPITO, West Virginia | TINA SMITH, Minnesota |

BRIAN HUGHES, *Staff Director*

KELLIE DONNELLY, *Chief Counsel*

ISAAC EDWARDS, *Special Counsel*

ROBERT IVANAUSKAS, *FERC Detailee*

MARY LOUISE WAGNER, *Democratic Staff Director*

SAM E. FOWLER, *Democratic Chief Counsel*

JOHN RICHARDS, *Democratic General Counsel*

ELISABETH OLSON, *Democratic FERC Detailee*

CONTENTS

OPENING STATEMENTS

| | Page |
|---|------|
| Murkowski, Hon. Lisa, Chairman and a U.S. Senator from Alaska | 1 |
| Cantwell, Hon. Maria, Ranking Member and a U.S. Senator from Washington | 2 |

WITNESSES

| | |
|--|----|
| Skare, Paul, Chief Cyber Security Program Manager, Pacific Northwest National Laboratory | 18 |
| Golden, Thomas A., Program Manager, Technology Innovation, Electric Power Research Institute | 42 |
| Henly, Claire, Managing Director, Energy Web Foundation | 72 |
| Narayanan, Dr. Arvind, Associate Professor of Computer Science, Princeton University | 77 |
| Kahn, Dr. Robert E., President and CEO, Corporation for National Research Initiatives | 86 |

ALPHABETICAL LISTING AND APPENDIX MATERIAL SUBMITTED

| | |
|--|-----|
| Cantwell, Hon. Maria: | |
| Opening Statement | 2 |
| Golden, Thomas A.: | |
| Opening Statement | 42 |
| Written Testimony | 44 |
| Responses to Questions for the Record | 172 |
| Henly, Claire: | |
| Opening Statement | 72 |
| Written Testimony | 74 |
| Responses to Questions for the Record | 175 |
| Kahn, Dr. Robert E.: | |
| Opening Statement | 86 |
| Written Testimony | 89 |
| Responses to Questions for the Record | 181 |
| Murkowski, Hon. Lisa: | |
| Opening Statement | 1 |
| Narayanan, Dr. Arvind: | |
| Opening Statement | 77 |
| Written Testimony | 79 |
| Responses to Questions for the Record | 177 |
| Public Utility District No. 1 of Chelan County (Washington): | |
| Statement for the Record | 4 |
| Skare, Paul: | |
| Opening Statement | 18 |
| Written Testimony | 20 |
| Responses to Questions for the Record | 168 |

**THE ENERGY EFFICIENCY OF BLOCKCHAIN
AND SIMILAR TECHNOLOGIES AND THE
CYBERSECURITY POSSIBILITIES OF
SUCH TECHNOLOGIES FOR
ENERGY INDUSTRY APPLICATIONS**

TUESDAY, AUGUST 21, 2018

U.S. SENATE,
COMMITTEE ON ENERGY AND NATURAL RESOURCES,
Washington, DC.

The Committee met, pursuant to notice, at 10:08 a.m. in Room SD-366, Dirksen Senate Office Building, Hon. Lisa Murkowski, Chairman of the Committee, presiding.

**OPENING STATEMENT OF HON. LISA MURKOWSKI,
U.S. SENATOR FROM ALASKA**

The CHAIRMAN. Good morning. The Committee will come to order.

We welcome everyone. Back here in August, back for another week of work. We have a hearing today, a Subcommittee hearing tomorrow and hopefully a business meeting sometime this week. So we are working.

This morning, a topic that I think has generated a great deal of interest, not necessarily within this Committee, but certainly when you think about the implication to our energy grid overall and just energy more broadly, the topic this morning is one of considerable interest. We are going to delve into whether or not blockchain and related technologies will soon have a transformative impact on energy infrastructure.

While not everyone knows what 'blockchain' is, I think most people have heard of cryptocurrencies, like bitcoin. Blockchain is the way the bitcoin system stores data.

I feel like I am doing a little bit of Introductory 101, but having had this conversation with my family members at Christmas a couple years ago where it was confirmed that none of us knew what we were talking about—

[Laughter.]

—I think it is helpful to give a little bit of background.

Senator CANTWELL. Are you sure your sons did not know what they were talking about?

[Laughter.]

The CHAIRMAN. They professed to. They claim to be the experts. And in fairness, I listened to them more than any of the more mature adults in the conversation.

Electronic transactions are stored as blocks that are linked together to form a chain. The more transactions recorded, the longer the chain. The chain is stored in numerous locations simultaneously so the system is decentralized.

The verification needed for this data has created an entire new industry. So-called ‘miners’ are paid by some blockchain applications to verify data blocks as trustworthy. As a result, entire warehouses of computers have been set up to verify this kind of data.

Now obviously, this type of computer-driven industry needs electricity and a lot of it. Miners have flocked to places with the cheapest electric rates. I know, Senator Cantwell, you have certainly seen the impact in your state, but an overnight demand for more power can cause serious stress on a local utility and impact the grid. There is also the question of how long this new load will need to be served.

Some areas are starting to respond. The State of New York recently authorized its municipal utilities to charge cryptocurrency miners higher electric rates than other consumers. Hydro Quebec has proposed new rules that would require cryptocurrency miners to bid for electricity and quantify their community impact in terms of jobs and investments.

At the same time, utilities are looking at blockchain as a way to boost both consumer engagement and grid efficiency through secure energy transaction platforms. Puerto Rico is looking at this very concept, where the effort to rebuild in a more resilient way has focused on microgrids, and the use of blockchain technology to trade power among the companies that operate the microgrids.

Now finally, our hearing will examine any cybersecurity advantages that blockchain and similar technologies might offer over other ways of securing our energy infrastructure. That is something that is always at the forefront of the minds of many of us on this Committee.

We are fortunate this morning to have a very impressive panel of experts who are here today to help us understand these issues.

Including Dr. Arvind Narayanan, am I pronouncing that right? Narayanan? He is an Associate Professor at Princeton who literally wrote the book on bitcoin.

As well as Dr. Robert Kahn, who invented the fundamental communications protocols which are at the heart of the internet. It is truly a pleasure to have you here. I think it is recognized that Dr. Kahn is called one of the true “fathers of the internet.” We are very fortunate that he is here to discuss this technology and the issues surrounding its deployment, along with the other esteemed members of our panel this morning.

I am looking forward to today’s testimony and the opportunity to have an exchange with you on this important issue.

Senator Cantwell, I welcome your remarks this morning.

**STATEMENT OF HON. MARIA CANTWELL,
U.S. SENATOR FROM WASHINGTON**

Senator CANTWELL. Thank you, Madam Chair, and thanks for scheduling this hearing on the emerging technology in the energy sector of blockchain.

When many people hear blockchain, as you just mentioned, they think of bitcoin, but it is important to note at the outset that these two terms are not synonymous. The cryptocurrency bitcoin is one application of blockchain technology, and bitcoin mining is an issue of significant importance to the State of Washington and one which I will address shortly.

Nevertheless, I see great potential in blockchain technology to have a dramatic impact on the development of a more clean energy economy. At its most basic level, blockchain refers to the ability of individual actors to use independent computers to record and verify digital transactions without the involvement of centralized authority and with very low risk of alteration of that data.

In the energy sector, these attributes of blockchain enable peer-to-peer energy transactions using data-brokered calls that resist manipulation by bad actors which allow electricity consumers to purchase power from specific preferred sources. For instance, neighbor A could buy excess electricity generated by neighbor B's solar PV cells at a preset price. Obviously all the implications for distributed energy and driving down costs are great.

Blockchain technology will handle the transaction, verify the validity of the terms, accurately report to both parties and regulators without the need for a third party. A private investor interested in expanding electric vehicle deployment can install charging infrastructure using blockchain technology to enter into contracts with EV owners for payments for electrons used without having to negotiate into a business relationship. So these are very interesting applications.

Blockchain technology does present other challenges though. For instance, in the State of Washington we are experiencing a tremendous increase in electricity demand attributed to mining of bitcoin. These activities using blockchain processes to earn increments of cryptocurrency is, let's just say, very popular right now. It means that computers and servers churn around the clock and these server farms need a constantly increasing amount of electricity to run and cool the processors.

Because of inexpensive hydropower in Washington, we find ourselves at the forefront of dealing with this issue as our utilities deal with it. To protect against miners driving up the cost and negatively impacting reliability, the central Washington utilities are taking matters into their own hand. I would like to enter into the record a statement from Chelan Public Utility District, so we can have that as part of today's hearing.

The CHAIRMAN. It will be included.

[The information referred to follows:]

Public Utility District No. 1 of Chelan County
PO Box 1231
Wenatchee WA 98807

STATEMENT FOR THE RECORD

Submitted by Steve Wright, General Manager



ENERGY EFFICIENCY OF
BLOCKCHAIN AND SIMILAR
TECHNOLOGIES HEARING

Tuesday, August 21, 2018

ENERGY AND NATURAL RESOURCES COMMITTEE

Introduction

Madame Chairwoman, Ranking Member Cantwell, and members of the Committee, thank you for the opportunity to provide testimony today on experience in Chelan County regarding cryptocurrency mining and blockchain operation. I applaud the Committee for recognizing that the increasing adoption of cyber transactional tools have significant implications for our nation's electricity systems.

My name is Steve Wright. I am General Manager of the Chelan County Public Utility District in north central Washington. Chelan owns and operates roughly 2000 MW of hydropower, serving approximately 50,000 customers. I am also a Board member of the American Public Power Association and the Alliance to Save Energy. APPA is the national service organization representing the interests of the Nation's 2,000 not-for-profit, community-owned electric utilities. They collectively serve over 49 million people and account for 15% of all sales of electric energy (kilowatt-hours) to ultimate customers. Public power utilities are load-serving entities, with the primary goal of providing the communities they serve with safe, reliable electric service at the lowest reasonable cost. This orientation aligns the interests of the utilities with the long-term interests of the residents and businesses in their communities. The Alliance is a non-profit, bipartisan coalition of business, government, environmental, and consumer-interest leaders that advocates for enhanced U.S. energy productivity to achieve economic growth; a cleaner environment; and greater energy security, affordability, and reliability. The Alliance thanks the four members of this Committee serving on our Honorary Board of Advisors, including Chairwoman Sen. Lisa Murkowski (R-Alaska), Sen. Rob Portman (R-Ohio), Sen. Lamar Alexander (R-Tenn.), and Sen. Ron Wyden (D-Ore.) for all they do to advance federal energy efficiency policy.

ENERGY AND NATURAL RESOURCES COMMITTEE**Background**

Back in 2014, we began to see large shipping containers showing up in Chelan County that were filled with racks of computers asking to hook up to our electric system. We soon came to learn this was the beginning of a wave of entrepreneurs seeking to mine cryptocurrency mostly in the form of bitcoin. This placed a tremendous strain on particular locations in our distribution system. It also raised questions about how much energy this new business might consume and how soon. This led us to put in place a moratorium and spend roughly 2 years defining policies and rates that would help to assure a neutral-to-positive impact to our existing customers.

In early 2017, we completed policies and rates to apply to these new businesses and lifted the moratorium. Later that year, there was a tremendous run up in the value of cryptocurrencies. That led to a new influx of cryptocurrency miners arriving in Chelan County. Some of these were the same size as what we had seen previously. But we also witnessed a flood of miners operating out of residential properties - and there were a few operations that were seeking very large amounts of power. That led us earlier this year to put in place another moratorium to address the changing portfolio of mining operations we were witnessing.

ENERGY AND NATURAL RESOURCES COMMITTEE

Challenges

You might ask: Why it is perplexing for an electric utility to serve growing loads as we do this all the time? There are five reasons why cryptocurrency mining creates a unique set of challenges for an electric utility.

FIRST

The energy intensity or usage per square-foot is staggering relative to nearly all traditional loads on an electric power system running roughly 10 times that of the highest use local loads on our system. This is exacerbated by the staggering number of requests for service we were receiving such that our loads that had developed over 75 years could have doubled in a short time frame. This alone would not make cryptocurrency unique, as other large industrial loads are also electricity intensive. For example, we have had an aluminum plant in Chelan County for more than 50 years.

SECOND

The portability of the loads is highly unusual. We are accustomed to buildings being built and then staying in the same place. Electric power infrastructure is generally placed where the loads are. With cryptocurrency mining the loads can and do move on a moment's notice.

THIRD

Cryptocurrency miners can morph in size at a moment's notice. They can decide they want 100 MW in one location, and then change that same request to 10 MW in 10 locations, and then change again to 1 MW in each of 100 locations. Again, this increases the level of complexity and risk for

ENERGY AND NATURAL RESOURCES COMMITTEE

planning and infrastructure investment well beyond anything that we are accustomed to.

FOURTH

The value of cryptocurrency is subject to price volatility that is unprecedented for the businesses with which we are accustomed to dealing. Our aluminum producer operates in markets that see regular price volatility. But the price volatility in bitcoin mining in the last year has been roughly 10 times that experienced in aluminum markets. As we have learned with aluminum, the power supplier is in effect a partner with the business and must take into account the potential for business disruption when commodity prices decrease.

FINALLY

We have also noticed that cryptocurrency miners have an underwhelming understanding of how mining impacts the electric power system. This is an industry still in its infancy, and it shows. We have repeatedly found cryptocurrency miners, using various levels of electricity, who assume that they only need to plug-in to the wall or the distribution system. They expect that the electrons will flow without an understanding that they could be overloading circuits and creating a safety risk. The most troubling example was a miner that placed computer racks in a third floor apartment, increasing the usage from 500 KWh to 22,000 KWh and creating a fire risk for all inhabitants in the building.

The impact on our community is not entirely unique. APPA notes that there are public power systems in New York, Oregon, Idaho, Missouri, Colorado as

ENERGY AND NATURAL RESOURCES COMMITTEE

well as other public power utilities in Washington state that have experienced cryptocurrency mining requests.

ENERGY AND NATURAL RESOURCES COMMITTEE

Community Reaction

As a public power utility, our focus is on what enhances the quality of life in our community. We have created many opportunities over the last few years to hear from the public regarding their views about serving cryptocurrency mining loads. We have heard the following:

- There has been substantial concern about the potential impact of these loads on the need for infrastructure investment and potential for rate increases.
- Associated with the concerns about cost/rate implications, there is doubt about the value created. There appears to be very few jobs and not much so far in the way of investment that translates into an increased tax base.
- There has been substantial anxiety about the need to address health and safety risks.
- There is concern about the morality of cryptocurrency mining. The value creation driving cryptocurrency mining has been described in many ways, much of which appears opaque to the public. Assertions have been made about cryptocurrency being used to hide illicit activities.

On the other hand, there are those in our community who view cryptocurrency as having the potential for being on the cutting edge of a technological revolution, creating a long term economic growth opportunity. This perspective is held even more strongly and deeply when the conversation is expanded to blockchain. The potential for blockchain technology to substantially improve efficiency in industries that have a strong local presence (i.e. the potential for tracking personal medical

ENERGY AND NATURAL RESOURCES COMMITTEE

records in the health care industry or food safety protection in the form of stem-to-shelf tracking for the tree fruit industry) can be an attractive vision for many. To this point though, the translation of how providing electric power to blockchain entrepreneurs will lead to local job creation has remained opaque.

Chelan PUD has translated these community views into an evolving perspective that seeks to create opportunity for cryptocurrency/blockchain entrepreneurs while assuring existing PUD customers are left financially neutral (or even positive) as a result of their decision to locate in Chelan County.

ENERGY AND NATURAL RESOURCES COMMITTEE**Risk Mitigation Steps**

To ensure public health and safety, we have established strong standards with associated charges to assure cryptocurrency mining is self-reported to the PUD and that the infrastructure is adequate to avoid fire risks. The amount of the charges is tied directly to the costs that the District incurs locating and monitoring unauthorized or unsafe loads.

To protect our existing customer's interest, we implemented - and now have proposed updating - upfront charges to assure infrastructure investments we make will be paid in advance, reducing the financial risk of loads that are here today but gone tomorrow.

We have also proposed energy charges that cover all costs, including market purchases, that can more easily be tailored to address the commitment timeframes cryptocurrency miners are willing to make.

ENERGY AND NATURAL RESOURCES COMMITTEE**National Policy Considerations**

Cryptocurrency mining and likely blockchain technology are huge devourers of electricity. Electricity policy for both economic and environmental reasons has been an important focus of the national policy debate for decades. So it is sensible to ask about the energy national policy implications raised by cryptocurrency mining and blockchain. I would offer the following thoughts for your consideration.

1. What is the value of cryptocurrency? Over the last year, bitcoin prices have varied from hundreds of dollars to as high as \$20,000. At the low values, cryptocurrency mining is likely only attractive where there is very low-cost electricity. At the higher values, mining could be attractive anywhere in the United States. How to value cryptocurrency remains a subject of great debate. How the market defines this value is going to determine whether cryptocurrency mining is a niche regional issue or a significant national issue.

If cryptocurrency or blockchain value is defined as high, the load growth potential is enormous, creating needs for transmission, distribution and generation that far exceeds current plans. This portends substantial discussion as to whether the social value of more efficient transactions is worth the environmental cost of an expanded electric power system.

2. Bitcoin mining is based on a platform described as proof of work. It is extraordinarily electricity-intensive. There are other platforms for mining cryptocurrency. For example, proof of stake is used for Ethereum and uses roughly 15% of the electricity to produce a coin as

ENERGY AND NATURAL RESOURCES COMMITTEE

a proof of work platform. The choice about how cryptocurrency will be mined in the future will have a significant impact on the amount of electricity used.

3. The dramatic increase in cryptocurrency prices has led to a gold rush mentality. The financial incentives and sense that acting now is necessary to capture benefits has led to unfortunate actions on the part of some miners who are focused only on a short-term perspective. In some cases, unrealistic and impassioned requests for rapid responses for service and infrastructure investment without concern for impacts on communities have tarnished the cryptocurrency industry's reputation. Even worse, plugging in without understanding power system impacts has created unnecessary public health and safety risks in the desire for financial gain. A measured approach to cryptocurrency mining development is necessary to protect the public interest.

4. Legitimate questions remain about whether cryptocurrency mining will help stabilize or destabilize grid operations. Thoughtfully planned development can be added to the grid without risk to reliability. Proposals have been made by some cryptocurrency miners that the loads could be interruptible, filling in the famous "belly of the duck" while reducing loads during the ramp that stresses the grid in the afternoon as the sun sets and solar output declines. To this point though, we are not witnessing operations that have the sophistication to crossover between cryptocurrency mining and grid operations that would allow such operations to proceed. Moreover, as discussed earlier, the early stages of cryptocurrency mining development we have experienced has been predisposed toward rapid and unplanned

ENERGY AND NATURAL RESOURCES COMMITTEE

development that creates risk of grid destabilization and uncompensated cost risk.

5. An interesting phenomena regarding cryptocurrency mining is a fundamental question as to whether it should occur in a centralized versus decentralized model. Many miners have been attracted to the potential for operating a small number of machines connected to the larger blockchain ledger, but while being able to operate as an independent business. From a utility perspective, the decentralized model is likely to be the more costly way to proceed. Pulling miners into more centralized cryptocurrency enterprise zones allows the most efficient planning and use of infrastructure. In a time when there is substantial opposition to new distribution and transmission, public policy is likely to play a significant role in how this industry evolves.

6. Blockchain seems likely to be an increasingly important tool used to facilitate transactions of all kinds in U.S. energy markets including energy efficiency. Wholesale and retail power transactors will likely be the first to harness blockchain because of the potential monetary benefits of incremental transactional “efficiency” gains. Energy systems are likely to become more transactive as “smart” technologies proliferate and homes and commercial buildings become increasingly “grid-enabled”. The optimal role for the federal government is uncertain. But Congress should conduct careful oversight to ensure that opportunities for energy efficiency are realized while security and privacy concerns are properly addressed.

ENERGY AND NATURAL RESOURCES COMMITTEE**Conclusion**

I can't tell you with certainty whether cryptocurrency and blockchain are the next big thing, although the efficiency gains from blockchain in particular seem likely to attract new uses across our economy. What I believe is that the electricity impacts on local communities can be managed through pricing and policies if the issues are recognized early enough and proactive measures are taken. The potential for the total amount of national load growth needed for generation supply and impacts on the grid associated with crypto transactions are worth monitoring. Just like at the local level, with the right kind of proactive policies there is the potential for consumer benefits from the adoption of these technologies. Managing the pace of development will be important to insure we can learn and apply our learnings for the benefit of all consumers. I believe the Committee has been wise to hold this hearing, collect expert testimony, and consider how to respond.

I want to say how much we appreciate Chairwoman Murkowski's and Senator Cantwell's commitment to enacting the next generation of comprehensive energy legislation on a strong bipartisan basis. When Congress next addresses U.S. energy policy, I encourage this committee to ensure that energy efficiency remains a top priority. We also appreciate your work on hydropower licensing reform, and look forward to the Senate and House agreeing on a legislative path forward.

Senator CANTWELL. To put this into context, a recent estimate found that a single bitcoin transaction uses as much electricity as an average household in the Netherlands uses in a month. Needless to say, there are some issues here that I think our state is sorting through.

But we also know that blockchain has other great applications. We know that the grid is under near constant cyberattack, and blockchain technology which is relatively resistant to hacking could provide higher levels of cybersecurity than other means in our current electricity system. Blockchain applications may help accelerate clean energy and utility investment, and a recent report by the Energies Future Initiative estimates that the global investment in digital power sector infrastructure has increased 20 percent since 2014 and reached \$47 billion.

We know that clean energy innovators are expanding the use of blockchain applications across multiple sectors. I mentioned electric vehicles, where blockchain providers are developing incentive to bring more charging stations online, microgrid applications, enabling homeowners to use excessive power from other sources and grid edge technologies for blockchain transactions, optimizing smart technologies like meters, thermostats, and appliances that will allow most of these technologies to help develop with third parties.

So I find this hearing of great contrast, Madam Chair, to the President's continued insistence on trying to make coal the only reliable source of electricity. I guarantee you that what we need to be doing is upgrading our cybersecurity and making sure that we are not going to charge consumers more. This is the kind of technology that could help drive down costs for the future.

I look forward to hearing what the witnesses have to say in today's discussion. Thanks for scheduling this hearing.

The CHAIRMAN. Thank you, Senator Cantwell.

We will now turn to our panel.

We will ask that you try to limit your comments to about five minutes. Your full statements will be incorporated into the record.

We will start with you, Mr. Skare. Scar?

Mr. SKARE. Scaree, thank you.

The CHAIRMAN. Mr. Skare, I'm sorry.

Mr. Skare is the Chief Cyber Security and Technical Group Manager at the Pacific Northwest National Laboratory (PNNL). We welcome you to the Committee.

He will be followed by Mr. Thomas Golden who is the Program Manager for the Electric Power Research Institute, EPRI. Welcome.

Ms. Claire Henly is before the Committee this morning. She is the Managing Director for the Energy Web Foundation. We thank you for being here.

I mentioned Dr. Arvind Narayanan earlier. He is Associate Professor, Department of Computer Science at Princeton University. We welcome you.

And of course, Dr. Robert Kahn, who is the President and CEO at the Corporation for National Research Initiatives.

We welcome you all.

Mr. Skare, if you would like to lead off.

STATEMENT OF PAUL SKARE, CHIEF CYBER SECURITY PROGRAM MANAGER, PACIFIC NORTHWEST NATIONAL LABORATORY

Mr. SKARE. Good morning.

Thank you, Chairman Murkowski, Ranking Member Cantwell and members of the Committee for this opportunity to appear before you today to discuss blockchain as it relates to U.S. electric infrastructure.

My name is Paul Skare, and I lead the grid cybersecurity research at DOE's Pacific Northwest National Laboratory, located in Richland, Washington. I worked in grid cybersecurity for over 20 years both in private industry and at PNNL.

In my written testimony I've included more complete descriptions of these issues that we're discussing today. But for now, I'd like to cover the following points.

First of all, cryptocurrency mining. One particular application that includes blockchain technology is the general ledger. This is having localized impacts on the U.S. power grid, especially where energy costs are low. But most of our understanding of mining's impact remains anecdotal. It's unclear how long-term and widespread this issue will be for U.S. electric infrastructure, but blockchain is just one tool that PNNL and others are exploring to help secure the grid.

First, I'd like to get into the difference between blockchain and cryptocurrency and the associated energy intensive computing. Blockchain technology is essentially a business ledger, electronically distributed that securely captures transactions of value without the need for a centralized authority or intermediary. Computers in a blockchain's network all evaluate the transactions in parallel and entries in the ledger cannot be altered without getting consensus of the computers in the network. Cryptocurrencies are an example of an application that uses public blockchains which are open to anyone but require volunteers to serve complex digital puzzles to support new blocks being added to the chain. Volunteers are rewarded for their contribution of computational work with small amounts of cryptocurrency, a process known as mining.

The energy used in cryptocurrency mining has been compared to the total energy usage of states and even countries. Miners require increasing amounts of computational power and therefore, energy, to capture their cryptocurrency rewards. Thus, the practice is most profitable wherever electricity prices are low such as central and eastern Washington.

While there have been media coverage of the impact that large cryptocurrency mining loads can have on local utilities, including some utilities declaring moratoriums on this activity, I'm not aware of any quantitative studies specifically on cryptocurrency mining impacts on the grid.

Furthermore, it's unclear how the demand for cryptocurrency in this energy use for mining will respond to the fluctuating value of cryptocurrencies themselves. Bitcoin alone has dropped more than 50 percent in value this year.

While cryptocurrency use is a public blockchain and mining to control access and verify blocks, one can also use a private

blockchain which is not open and does not use mining, and thus does not require energy intensive computation.

At PNNL we're exploring the application of blockchain to grid cybersecurity with support from the Cybersecurity for Energy Delivery Systems program within the DOE Office of Cybersecurity, Energy Security, and Emergency Response.

At PNNL we take a holistic approach to securing the power grid, from stewarding operational capabilities, like the cyber threat monitoring program called CRISP, the Cybersecurity Risk Information Sharing Program, to developing entirely new technologies that keep our defenses at the forefront.

At PNNL's blockchain project, we're applying private blockchain solutions to a variety of use cases, including maintaining supply chain, chain of custody, ensuring integrity of control signals and managing distribution of software patches, among others.

Using a private blockchain has the potential for power system applications to add items to the blockchain every second and verify data upon the blockchain within the next second to alt scale. This quick update in capability is essential to handling increasing data requirements of a modern power grid and much more difficult to achieve with public blockchain approaches.

Blockchain and other distributed ledger of technologies, in fact, have many properties that make them well suited to facilitate more efficient and decentralized energy transactions but these properties also come with some potential challenges.

My written testimony discusses many of these challenges, but one I'd like to highlight here is endpoint security. No matter how secure the blockchain aspects of the solution are, the endpoints, those parts of the solution on either end of the blockchain, remains open to vulnerabilities as any other software.

Realizing the potential of blockchain for the grid, we're requiring studying in addressing these challenges in applying blockchain to the grid, alongside other technologies within a broader cybersecurity framework.

With all the potential for security and control systems that industry and DOE are working toward, it is important to keep in mind that blockchain is just one of a broad set of tools we must develop as we work to secure our power grid.

I appreciate the opportunity to discuss this important issue with you today and I'm happy to answer your questions.

Thank you.

[The prepared statement of Mr. Skare follows:]

Statement of Paul Skare
Chief Cyber Security Program Manager
Pacific Northwest National Laboratory

Before the
United States Senate
Committee on Energy and Natural Resources

August 21, 2018

Good morning. Thank you, Chairman Murkowski, Ranking Member Cantwell, and Members of the committee. I appreciate the opportunity to appear before you today to discuss blockchain as it relates to U.S. electric infrastructure issues and opportunities.

My name is Paul Skare, and I lead the Grid Cybersecurity Research Program at the Pacific Northwest National Laboratory (PNNL), a Department of Energy (DOE) National Laboratory located in Richland, Washington. I also support the Security and Resilience team in DOE's Grid Modernization Laboratory Consortium, a team of 14 National Labs that, along with industry and university partners, supports the Department's Grid Modernization Initiative. The consortium members include PNNL, the National Renewable Energy Laboratory, Argonne National Laboratory, Brookhaven National Laboratory, Idaho National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, the National Accelerator Laboratory at Stanford, National Energy Technology Laboratory, Oak Ridge National Laboratory, Sandia National Laboratories, and Savannah River National Laboratory. I have worked in the power industry for 38 years, starting at Northern States Power in Minneapolis, MN, Siemens Energy in Minnetonka, MN, and. I started working on cybersecurity for the grid 20 years ago. Today I will address these points:

1. Blockchain technology can be thought of as an electronic general ledger, securely capturing transactions without the need for a centralized authority, and cryptocurrency is *just one application* that uses this technology.
2. Cryptocurrency mining is having localized impacts on the U.S. power grid. However, most of our understanding remains anecdotal, and it is unclear what the long-term impacts will be as cryptocurrency prices fluctuate.
3. Grid cybersecurity is a multi-faceted issue with threats coming from many different directions, and blockchain is just one tool that PNNL and others are exploring that can help secure the grid. *But there is no silver bullet to securing our power grid and other critical infrastructure.*

Background

For more than two decades, PNNL has supported power system reliability, resilience and innovation for the State of Washington, the Pacific Northwest, and the nation. Over this period, the laboratory has:

1. Helped the North American Electric Reliability Corporation (NERC) and DOE design and implement the series of national grid cyber exercises known as *GridEx* which linked industry with government and law enforcement agencies and allows participants to practice their incident response plans. *GridEx III* engaged over 400 organizations and 4000+ participants in scenarios designed and operated with support from PNNL.
2. PNNL helped DOE develop the Electricity Subsector Cybersecurity Capability Maturity Model (ES-C2M2). This allows for utilities to assess their business practices supporting cybersecurity and learn where their business could invest more to meet their business's goals for cyber. Cybersecurity insurance companies have used this to influence the rates for insurance. PNNL has expanded on ES-C2M2 to create models for Buildings, the National Institute of Standards and Technology (NIST) Cybersecurity Framework, and Secure Design and Development Principles.
3. Led DOE-industry collaborations in developing and deploying synchrophasor technology to help avoid blackouts. Phasor measurement unit networks are designed to enhance situational awareness of wide area systems. This new grid tool has demonstrated value by detecting impending system control and equipment faults for system operators, thus avoiding major outages.
4. Led a public-private collaboration with utilities and vendors to develop and demonstrate transactive control concepts on the Olympic Peninsula in Washington and for the Pacific Northwest Smart Grid Demonstration project—the largest of its kind—to validate smart grid benefits and new control approaches that engage demand and distributed resources at scale.
5. Delivered the first applications of high performance computing to grid tools such as interconnection-scale contingency analysis, reducing run times from days to under two minutes. PNNL also applied high performance computing and phasor measurement unit data to deliver the first real-time dynamic state estimation to open the door to the future world of predictive grid tools. This parallelized state estimator tool enabled PNNL to deliver assessments of system risk at the interconnection scale (the Western Interconnection) in less than 2 minutes versus the traditional 24 hours.

These examples illustrate the high return on investment possible by utilities and National Labs across the country when combining advanced electric infrastructure technology innovation with public-private validation and deployment.

Blockchain: a new technology with many potential applications

Blockchain is most widely known as a technology that cryptocurrencies – Bitcoin being the most famous – use to secure digital transactions. The underlying blockchain technology, invented in 2008 as part of the Bitcoin cryptocurrency, consists of a database distributed across all computers in the blockchain network that maintains a continuously growing list of records – called blocks – that are protected from tampering and revision. Each block contains a timestamp

and a link to a previous block. The use of blockchain allows users to exchange value without intermediaries acting as arbiters of money and information, as the blockchain serves as a general ledger for transactions whose authenticity is established by the network itself.

Cryptocurrencies are an example of the ‘public’ use of blockchain (sometimes referred to as ‘permissionless’ or ‘open’), meaning anyone connected to the network can access the blockchain without restriction. This makes the transactions transparent and permanent. In these public blockchains, volunteers must solve complex digital ‘proof of work’ puzzles to make new entries in the blockchain without any other access control or identity verification. Solving these complex puzzles to verify the new block entries produces cryptocurrency as the reward for supplying computer calculations. This is known as ‘mining’ and is the source of unexpected energy usage electric utilities (especially those with lower energy prices) are facing today.

For other use cases beyond cryptocurrencies, ‘private’ blockchain is often used –also sometimes called ‘permissioned’, ‘consortium’, or ‘hybrid’ blockchain. This form does not allow public access to the blockchain, so mining is not necessary for access control and verifying blocks.

The key to the security of the blockchain transaction is the distributed, so called ‘Byzantine’ fault tolerant architecture used to validate each transaction. In this scheme, numerous computers (nodes in the blockchain network) are evaluating each transaction in parallel, independently looking at each transaction, and then comparing results with the other nodes. These blocks cannot be altered without altering all the blocks AND having a consensus of the nodes. In addition, access to the blockchain uses public-key cryptography to identify an address on the blockchain. This can be thought of as ‘secure by design’ in our current technology. It is fair to note here that there are no 100% secure systems or solutions. Please refer to a paper I cowrote with others on the application of Byzantine Architectures to secure control systems, ‘Survivable SCADA Via Intrusion-Tolerant Replication’ (IEEE TRANSACTIONS ON SMART GRID, VOL. 5, NO. 1, JANUARY 2014).

Cryptocurrency mining is straining localized parts of the US power grid

The energy used in cryptocurrency mining has been compared to the total energy usage of some states and some countries. From an electric utility’s perspective, this can appear as an unexpected energy load on their distribution grid. A qualitative analysis shows that this strain is typically appearing in areas with low energy costs and low fixed costs (such as rent).

Cryptocurrency miners must complete increasingly complex calculations, requiring increasing amounts of computing power and therefore energy, to capture returns in the form of cryptocurrency. Thus, the practice is most profitable wherever electricity prices are low, such as the Columbia River Basin in Eastern Washington. When large mining operations move into a particular location, the increase in load can cause the utility to increase its generation, buy power from others, or experience overloaded distribution circuits which could lead to localized power outages. A number of cities have placed moratoriums on new high-density load hookups to give staff time to develop a plan for dealing with the demand for electricity from digital currency miners (<https://news.bitcoin.com/washington-utility-increases-security-amid-crypto-mining-moratorium/>). While there has been no shortage of popular press coverage of the scale and

impact of the growing cryptocurrency mining community on the US power grid, I am not aware of any quantitative studies of cryptocurrency mining impacts.

Today, what we know comes from press stories and anecdotal evidence is that some utilities have knocked on doors to investigate unexpected loads and found rows of computer racks doing cryptocurrency mining in both residential and commercial areas. This year, Bitcoin alone has dropped more than 50% in value, and it is unclear how mining operations will respond to the fluctuating prices. Understanding the elasticity of demand for cryptocurrency is necessary to understand any potential long-term impacts on the US power grid.

Grid cybersecurity and how blockchain fits into the landscape

The U.S. power grid is rapidly changing from an earlier, simpler era with large generation stations and passive energy loads to a much more dynamic grid with growing distributed energy generation resources and much more active, connected, “smart” loads. With the increase in renewables on the grid, and the retirement of some older larger generators, generator inertia - which is crucial to the reliable operation of an AC power system - is strained. In addition, availability of black start generation - needed to restart generators after an outage - is also strained at some locations.

Blockchain technology shows potential in securing energy delivery systems (EDS) that have transactional attributes. This is important as these control systems require unprecedented levels of security and trustworthiness to verify integrity of data and manage complex demand response and market system exchanges. Improving the ability to identify, control, and secure grid devices with blockchain technology may increase the security and trustworthiness of real-time energy transactions without adding prohibitive costs, latency, interoperability or scale issues. The wide range of potential applications of blockchain to EDS has made the technology a priority for DOE.

In fact, blockchain and other distributed ledger technologies have many properties that make them well suited to facilitate more efficient and decentralized energy transactions, but these properties also come with some potential challenges:

- **Distributed consensus mechanism:** This supports decentralization of authority from single points of failure or compromise. This is a great advantage, but the challenge associated with purely distributed proof of work type consensus model is its vulnerability towards a 51% attack of the nodes. In such an attack, enough nodes can be compromised to approve a transaction. Such situations could be avoided by using permissioned type consensus model. In permissioned consensus models, privacy controls can also be implemented and customized as per the need of the application.
- **100% up time:** Blockchain provides a reliable, fail-safe logically centralized, physically distributed persistence mechanism. Bitcoin failures were focused on the application layer where there has been theft and loss of Bitcoins when users lose their private key required for signing a transaction or data content.
- **Strong immutability:** Even blockchain technology has proven nothing is immutable, with examples of mutations such as forks and or blockchain hacks that required rolling back

the blockchain. Blockchain technology does provide an atomically variable time stamped cryptographic signed electronic transaction that has proven to be very difficult to change.

- Immutability challenges: Immutability can lead to a number of challenges. Recently it was found that illegal images were saved in the Bitcoin blockchain. When undesired data are saved in the blockchain, it can prove very difficult to change. As discussed above, another way to change the blockchain is to control or compromise 51% of the nodes needed to reach a consensus.
- Big data management blockchain: Blockchain facilitates the distribution of prodigious data sets between organizations. Data can be synchronized and archived between multiple parties. The challenge here is at some point the blockchain might be overwhelmed by the amount of data being stored. It is important to address this by ensuring data is stored in efficient formats that minimize overall data volumes stored in the blockchain.
- Endpoint Security: No matter how secure the blockchain aspects of a solution are, the endpoints – parts of the solution on either end of the blockchain technology – remain as open to vulnerabilities as any other software.

PNNL is leading the way in Advanced Grid Cyber Security Approaches and new Technologies

PNNL is a leader in developing the foundational understanding and technologies for security of our power grid. We take a broad approach to this critical national need – from stewarding operational capabilities like the cyber threat monitoring program called CRISP to developing entirely new technologies that keep our defenses at the forefront.

CRISP – the Cybersecurity Risk Information Sharing Program – uses data shared by utilities to perform an intelligence-informed analysis that identifies threats that neither utilities alone, nor private cybersecurity firms, can identify. CRISP provides a strong complement to what utilities and private cybersecurity firms provide. Utility participation in CRISP will soon provide complete coverage of the continental United States. This program is able to identify traffic from malware that uses blockchain technology.

PNNL was recently awarded a project by the DOE Office of Cybersecurity, Energy Security, and Emergency Response's (CESER) Electricity's Cybersecurity for Energy Delivery Systems (CEDDS) program (formerly part of the Office of Electricity Delivery and Energy Reliability) to evaluate uses of blockchain in the electric grid. In this program, we are applying private blockchain solutions – so one that does not require the energy intensive mining process – to a variety of use cases in the power grid.

By using a private blockchain, this approach has the potential for power system applications to add items at scale to the blockchain every second, and to verify data from the blockchain within the next second. This quick updating ability is *essential* to handle the increasing data requirements of the modern power grid.

Use cases for this project:

- Supply Chain ‘chain of custody’ – the ability to trace products and components from origin to destination – this will allow utilities to see and track the source of all the components in a system, allowing for better understanding of risks due to potential vulnerabilities and patches.
- Device Integrity – verification of control signals
- Enhanced cybersecurity controls – trusted zone of nodes using verifiable digital signatures and signed messaging
- Patch Management – allows an asset owner to track and trace a patch from the vendor to their system.
- Supply and Demand transactions between microgrids

PNNL’s current program only scratches the surface of where we see potential for blockchain applications on the grid. We see many other potential use cases, including:

- EV Charging: Use blockchain to enable EV charging and billing interactions with the EV owner. This amount in kWh is subtracted from the smart meter read or billing to determine data for the prosumer (Prosumers are amateur advocates for products – in this case these are people who both produce and consume energy).
- Meter Data Access Management: Use blockchain technology to work with a central meter management system to allow consumers to manage who is allowed access to their meter data.
- Asset Lifecycle Management: Use blockchain to manage beginning-to-end lifecycle of assets’ parts and/or components (construction, operations, maintenance, disposal). This can also include the chain of custody for the supply chain.
- Distributed Energy Resources (DER) Transaction Processing: Use blockchain to process any transaction involving a DER asset, e.g., storage, solar photovoltaic (PV), EV, micro-combined heat and power (micro-CHP).
- Peer to Peer Trading of Distributed Energy: Use blockchain for bilateral trading of distributed energy generation (similar to the Brooklyn Microgrid project).
- Markets, Energy Trade Settlement: Use blockchain to support electronic trades at energy exchanges or for direct agreements/trades between market participants.
- Supplier Switching: Use blockchain technology to track supplier switching.
- Emission Certificates: Use blockchain to generate, own, and trade emission certificates related to energy generation.
- Energy Supply Chain: Supply chain reconciliation (energy delivered, technical/non-technical losses, consumption, etc.) spanning all measurement points all the way through from generation to consumption for commercial settlement.
- Blockchain based Metering: Use blockchain to augment smart meter for recoding energy use of appliances (EV, heating).

In addition to our program evaluating blockchain for grid cybersecurity, we are also completing many other projects for the CEDS program, including:

- MEEDS – the Mitigation of External-exposure of Energy Delivery Systems –This project works with an existing cybersecurity search engine tool called Shodan to support a private review of control systems connected to the internet for a particular utility without publishing results publicly.
- SSASS-E – Safe, Secure Autonomous Scanning Solution for Energy Delivery Systems - This project supports continuous scanning of control systems without negatively impacting the performance of the control systems and devices.

On the technology front, PNNL has begun a new program with Lab Directed Research and Development funding, called Proactive Adaptive Cybersecurity For Control, or PACiFiC, which includes two focus areas:

1. Deception:
Adaptive Cybersecurity Controls are being explored from many perspectives. Traditional approaches have put us in an asymmetric disadvantage against our adversaries in defending our systems. Adaptive Cybersecurity Controls can be a way to provide a more level playing field by adjusting control system environments on the fly to confuse, obfuscate, and mislead adversaries as they work their way through a system, increasing the effort and knowledge needed to get through the defenses, while also giving a better chance for detection solutions to be effective.
2. Secure Design and Development Principles:
While there are many documented methods to secure operation systems, there are no documented ways for a vendor to create a secure product that they can control and implement. This body of work provides a set of over 600 best practices that encompasses the entire product lifecycle.

Conclusion

Industry and DOE have partnered over the past several years to significantly advance our grid cyber technologies, and new projects are breaking new ground in leading edge technology such as blockchain. To see some potential of blockchain, look at Estonia – the first country to face a nationwide cyber-attack. As a result, ongoing investments have led to public services being digitized and accessed via secure digital identities provided to every citizen and resident. Integrated into the digital services is blockchain technology.

In parallel to “better securing” the grid, we need to leverage these same foundational science and technology tools of high-performance computation, analytics, deep learning and control theory to develop more resilient system designs for networks, data and grid control systems. These will enable the system to better resist inevitable attacks, better defend and ultimately recover quickly.

The DOE investments in fundamental science, applied technology and public/private partnerships in grid cybersecurity are essential elements of an effective, integrated national cyber readiness strategy for the U.S. electric power system and its related infrastructures. Securing our electric grid is a long-term endeavor that will require a range of strategies and new technologies;

there is no one silver bullet. Blockchain is just one of a set of tools we must develop as we work to accomplish the goal of securing our energy systems.

I appreciate the opportunity to discuss this important issue with you today, and I am happy to answer your questions.

Thank you.

CEDS-Supported Projects in Washington State (Active)

| Prime Performer | Project Title |
|--|---|
| Pacific Northwest National Laboratory (PNNL) | Automated, Disruption-Tolerant Key Management System |
| Pacific Northwest National Laboratory (PNNL) | Enabling Situation Assessment/Awareness for Utility Operators and Cybersecurity Professionals |
| Pacific Northwest National Laboratory (PNNL) | Universal Utility Data Exchange (UUDEX) |
| Pacific Northwest National Laboratory (PNNL) | Keyless Infrastructure Security Solution (KISS) |
| Pacific Northwest National Laboratory (PNNL) | Software Defined Networking for Energy Delivery Systems (SDN4EDS) |
| Pacific Northwest National Laboratory (PNNL) | Research Exploring Malware in Energy Delivery Systems (REMEDYS) |
| Pacific Northwest National Laboratory (PNNL) | Safe, Secure Autonomous Scanning Solution for Energy Delivery Systems (SSASS-E) |
| Pacific Northwest National Laboratory (PNNL) | Mitigation of External-exposure of Energy Delivery System Equipment (MEEDS) |
| Schweitzer Engineering Laboratories, Inc. | Secure Software Defined Radio |
| Schweitzer Engineering Laboratories, Inc. | The Alliance Project |
| Schweitzer Engineering Laboratories, Inc. | Tempus Project Time Synchronization Platform for GPS Spoofing |
| Schweitzer Engineering Laboratories, Inc. | Chess Master Project |

1. Pacific Northwest National Laboratory (PNNL): Automated, Disruption-Tolerant Key Management System

Partners: Arizona Public Service (APS), Lawrence Berkeley National Laboratory (LBNL)

Project Description: The project is working to design a standards compliant and interoperable system, implement a prototype key management and field device services, and evaluate and compare the performance and effectiveness of the prototype against existing key management systems for the energy sector. This effort is improving security and the efficiency of operations by providing a new key management architecture suited to the unique requirements of EDS.

2. Pacific Northwest National Laboratory (PNNL): Enabling Situation Assessment/Awareness for Utility Operators and Cybersecurity Professionals

Partners: Idaho National Laboratory (INL), GE (Alstom Grid), Peak RC, Total Reliability Solutions, WAPA

Project Description: Utility operators are bombarded with data from differing sources and systems and struggle to derive meaning from the data. To enable operators to make informed decisions in the finite amount of time available, operators need a cognitive system that displays the associated data to enhance situational awareness. This project will develop visualizations that power system operators and/or cybersecurity professionals can use to make fast, accurate assessments of situation, enabling them to maintain situation awareness during unfolding events.

3. Pacific Northwest National Laboratory (PNNL): Universal Utility Data Exchange (UUDEX)

Partners: MITRE, OATI

Project Description: The project team will develop a secure and flexible data-exchange approach to replacing key communication between control centers, including Inter-Control Center Communications Protocol (ICCP) data exchanges, threat information, synchrophasor, Reliability Coordinator Communications Information System (RCIS), and incident information. ICCP will be replaced with a modern model-driven data-exchange architecture and protocols, taking advantage of current methods of data transport and configuration.

4. Pacific Northwest National Laboratory (PNNL): Keyless Infrastructure Security Solution (KISS)

Partners: Avista, Cisco, Guardtime, Rocky Mountain Institute, TVA (Utility Advisor), Washington State University, OATI

Project Description: The project team will develop a Keyless Infrastructure Security Solution (KISS) to increase the trustworthiness, speed, integrity, and resiliency of EDSs responsible for complex grid-edge energy exchanges and integration of distributed energy resources, by developing a prototype of a secure and trustworthy blockchain energy platform.

5. Pacific Northwest National Laboratory (PNNL): Software Defined Networking for Energy Delivery Systems (SDN4EDS)

Partners: AECOM, CAISO, Dispersive Technologies, Juniper Networks, National Renewable Energy Laboratory (NREL), Sandia National Laboratory (SNL), Schweitzer Engineering Laboratories, Inc. (SEL), Southern California Edison (SCE)

Project Description: The project team plans increase the adoption of SDN technologies and improve security for local area networks (LANs) and wide area networks (WANs) components in the energy sector.

6. Pacific Northwest National Laboratory (PNNL): Research Exploring Malware in Energy Delivery Systems (REMEDYS)

Partners: ORNL, MIT, ANG Consulting, James P. Fama, Nevermore Security

Project Description: The project team will develop, evaluate, and refine organizational structures that could be used to coordinate the nation's multiple energy sector stakeholders in the rapid research, development, and distribution of mitigations that reduce the risk of an imminent or emerging malware cyber-attack that might otherwise disrupt energy delivery.

7. Pacific Northwest National Laboratory (PNNL): Safe, Secure Autonomous Scanning Solution for Energy Delivery Systems (SSASS-E)

Partners: Chelan County PUD, National Rural Electric Cooperative (NRECA), Tenable Security, University of Illinois at Urbana-Champaign (UIUC)

Project Description: The project team will develop, validate, and verify innovative safe scanning methodology, models, and architectures, and produce a prototype to transform Tenable's IT/OT platform, the most widely deployed vulnerability scanner in the IT space, to secure operational technology (OT) installed in critical energy infrastructure.

8. Pacific Northwest National Laboratory (PNNL): Mitigation of External-exposure of Energy Delivery System Equipment (MEEDS)

Partners: Shodan, LLC, National Rural Electric Coop, Tenable, Chelan PUD

Project Description: MEEDS is a user-friendly web application for utilities built upon the existing Shodan technology for performing continuous monitoring of utilities' internal networks to detect and identify any EDS equipment that may be inadvertently exposed to the internet.

9. Schweitzer Engineering Laboratories (SEL): Secure Software Defined Radio

Partners: Pacific Northwest National Laboratory (PNNL), San Diego Gas & Electric (SDG&E)

Project Description: The Secure Software-Defined Radio Project (SEL-3070) is developing a flexible platform for secure wireless communications to utility distribution automation devices, providing capabilities not offered in cellular, narrow-band licensed, or other unlicensed-band radios.

10. Schweitzer Engineering Laboratories (SEL): The Alliance Project

Partners: Sandia National Laboratory (SNL), Tennessee Valley Authority (TVA)

Project Description: The Alliance project is developing a proximity card reader and controller that allows physical and cybersecurity access to be monitored, tracked, and controlled using a single system. The reader and controller consist of four easy-to-deploy components: an access terminal, an access control processor, enhanced firmware for the SEL-3620 and SEL-3622 security gateways, and a card enrollment solution.

11. Schweitzer Engineering Laboratories (SEL): Tempus Project Time Synchronization Platform for GPS Spoofing

Partners: San Diego Gas & Electric (SDG&E), Southwestern Research Institute (SwRI)

Project Description: SEL will research, develop and demonstrate the capabilities of a secure, modular, and customizable time synchronization platform that provides layers of protection from GPS spoofing attacks. The project will include the development of innovative algorithms and electronics that detect GPS signal manipulation for critical applications that use timing signals in the energy sector.

12. Schweitzer Engineering Laboratories (SEL): Chess Master Project

Partners: Ameren Energy Resources, Sempra, Veracity Security Intelligence

Project Description: This project will provide system operators with a global view of their operational network, enabling them to set and view field network security policy and validate operational adherence to those policies. The Chess Master team will build on the successful commercial release of utility rated software defined network (SDN) technology under the previous CEDS project, Watchdog.

Survivable SCADA Via Intrusion-Tolerant Replication

Jonathan Kirsch, Stuart Goose, Yair Amir, *Member, IEEE*, Dong Wei, *Member, IEEE*, and Paul Skare, *Member, IEEE*

Abstract—Providers of critical infrastructure services strive to maintain the high availability of their SCADA systems. This paper reports on our experience designing, architecting, and evaluating the first *survivable* SCADA system—one that is able to ensure correct behavior with minimal performance degradation even during cyber attacks that compromise part of the system. We describe the challenges we faced when integrating modern intrusion-tolerant protocols with a conventional SCADA architecture and present the techniques we developed to overcome these challenges. The results illustrate that our survivable SCADA system not only functions correctly in the face of a cyber attack, but that it also processes in excess of 20 000 messages per second with a latency of less than 30 ms, making it suitable for even large-scale deployments managing thousands of remote terminal units.

Index Terms—Cyber attack, fault tolerance, reliability, resilience, SCADA systems, survivability.

I. INTRODUCTION

SUPERVISORY Control and Data Acquisition (SCADA) systems form the backbone of many vital services, such as electricity transmission and distribution, water treatment, and traffic control. As key components of our critical infrastructure, SCADA systems must continue operating correctly and at their expected level of performance at all times. In practice, ensuring such continuous availability requires the capability to tolerate and overcome various types of faults that arise in large distributed systems, including “benign” faults (e.g., hardware crashes, power failures, and network partitions) and more severe faults, including potentially malicious cyber attacks.

Unfortunately, contemporary SCADA systems exhibit an *availability gap* that leaves them vulnerable to downtime. While today’s systems are able to withstand effectively many types of benign faults using hardware and software redundancy techniques (e.g., primary/hot standby [1]), their ability

to survive in the face of cyber attacks remains limited. Many SCADA systems were designed to operate on isolated, private networks, but this assumption of an “air gap” no longer holds in many modern deployments: interoperability goals and the need to provide access to more grid stakeholders mean that the SCADA system is often connected to enterprise IT infrastructure, inheriting the associated vulnerabilities. SCADA has also become an increasing target for cyber attacks [2], resulting in an arms race between attackers and SCADA vendors and operators. Furthermore, although today’s systems employ a defense-in-depth approach to security that focuses on *preventing* attacks, it is impossible to prevent all attacks; insider attacks, in particular, pose a growing threat to critical infrastructure [3].

This paper reports on our experiences to date designing and implementing the first *survivable* SCADA system.¹ By *survivable*, we mean that the SCADA system continues to *operate correctly* and with *minimal performance degradation* even if malicious attacks compromise part of the system. These twin properties are essential for maintaining high availability in the face of cyber attacks.

To achieve survivability, our system employs *intrusion-tolerant replication* [5], [6]. It runs, in parallel, several copies of the SCADA Master application: the copies collectively behave as a single *logical SCADA Master* that provides correct, timely service as long as less than a threshold fraction of the copies is compromised. Intuitively, intrusion tolerance allows an application to act as its own firewall, providing protection even if the system’s security perimeter is breached. A distinguishing feature of intrusion-tolerant systems is that they do not require prior knowledge of attack signatures and behaviors to provide their guarantees.

Intrusion-tolerant replication protocols have been well-studied in the distributed systems community over the last decade (e.g., [6]–[11]), and in this paper we build on this previous research. Specifically, we use the Prime replication protocol [5], [6] as a fundamental building block in our survivable SCADA system. However, we were confronted with two significant challenges when attempting to integrate Prime with a SCADA system. First, existing intrusion-tolerant replication systems, including Prime, implicitly assume that the application being replicated is client driven (i.e., the server application takes action only in response to unsolicited requests submitted by clients). By contrast, in a SCADA system, the SCADA Master application also processes solicited requests, which are pulled from field devices by a *server-driven* polling

Manuscript received December 05, 2012; revised March 18, 2013; May 10, 2013; accepted June 10, 2013. Date of publication August 07, 2013; date of current version December 24, 2013. The work of Y. Amir was supported in part by DARPA Grant N660001-1-2-4014. The content of this paper is solely the responsibility of the authors and does not represent the official view of DARPA or the Department of Defense. Paper no. TSG-00841-2012.

J. Kirsch and S. Goose are with Siemens Technology-To-Business Center, Berkeley, CA 94704 USA (e-mail: jonathan.kirsch@siemens.com; stuart.goose@siemens.com).

Y. Amir is with Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: yairamir@cs.jhu.edu).

D. Wei is with Siemens Corporation, Corporate Technology, Princeton, NJ 08540 USA (e-mail: dong.wi@siemens.com).

P. Skare is with Pacific Northwest National Laboratory, Richland, WA 99352 USA (e-mail: paul.skare@pnl.gov).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSG.2013.2269541

¹A preliminary version of this paper appeared in the Proceedings of the Annual Cyber Security and Information Intelligence Research Workshop, 2011 [4].

operation. This need to support polling creates an architectural mismatch between Prime and SCADA. The second challenge we confronted relates to performance: the replication engine must be able to provide low enough latency to preserve the real-time control and monitoring of the SCADA system, while being able to support a high enough throughput so that the system can scale to large deployments.

Our efforts to date have resulted in a prototype implementation of a survivable SCADA system for electricity transmission and distribution, where our Prime-based intrusion-tolerant replication engine is integrated with a real Siemens SCADA product. We developed the prototype system on this product because the results of a compromise of a critical infrastructure system such as the power grid would be particularly disruptive and would have significant adverse effects on today's society, and SCADA plays a vital role in the power grid. However, Prime and the protocols described in this paper are generic and could also be applied to other mission-critical systems (e.g., distributed control systems).

The novel contributions of this paper are as follows. i) We present the design of the first SCADA system in which the SCADA Master application is able to survive a partial compromise. ii) To address the architectural mismatch between SCADA systems and traditional intrusion-tolerant replication systems, we present the first scalable and intrusion-tolerant *logical timeout* protocol (to support the scheduling of polling events) and a *logical channel* protocol (to enable reliable and intrusion-tolerant SCADA communication in a replicated environment). iii) We present performance results demonstrating the suitability of our intrusion-tolerant replication engine for use even in large-scale SCADA deployments.

The remainder of this paper is organized as follows. Section II presents background on SCADA and intrusion-tolerant replication needed to understand the rest of the paper. Section III presents the design of our survivable SCADA architecture, as well as our attack model and assumptions. Section IV discusses the integration challenges we faced and presents the new protocols we invented to overcome them. Section V presents our performance results, and Section VI places our solution in the context of related work. Finally, Section VII concludes the paper.

II. BACKGROUND

A. Conventional SCADA Systems

SCADA systems are large and complex. In this section we describe the components of a SCADA system most relevant to the work in this paper. These components include:

- One or more **Remote Terminal Units (RTUs)**, which communicate with, and aggregate data from, local sensors in the field (e.g., within an electricity distribution substation). Some larger systems can have several thousand RTUs.
- A **SCADA Master**, which periodically polls the RTUs by sending messages over a wide-area network. The SCADA Master maintains a real-time database containing the current state of each RTU. It can also send supervisory control commands to the RTUs. For fault tolerance, many SCADA systems use a Primary/Hot Standby (HSB) configuration,

in which two similar but slightly different copies of the SCADA Master application run in parallel. Although both the Primary and the HSB receive incoming events, the Primary is responsible for controlling the system, and the output of the HSB is suppressed. The HSB monitors the Primary and performs a "take-over" operation to assume control if it believes the Primary has succumbed to a benign fault.

- One or more **Human Machine Interface (HMI)** workstations, which periodically query the SCADA Master so that the state of the system (e.g., the power grid) can be graphically displayed for a human operator.

B. Intrusion-Tolerant State Machine Replication

An intrusion-tolerant protocol [12] assumes that some of the participants may be *Byzantine* [13] and act in an arbitrary manner (e.g., because they are compromised and under the control of a malicious attacker). The protocol is designed to operate correctly regardless of how the Byzantine participants behave,² as long as no more than a threshold fraction of the participants (typically f out of $3f + 1$) is Byzantine [13]. Intrusion-tolerant protocols often use *proactive recovery* techniques [7], where participants are periodically rejuvenated to a clean state. This allows the system to survive more than f Byzantine failures over the life of the system, as long as no more than f are Byzantine at the same time.

In this paper we make use of an intrusion-tolerant replication system, where replication is achieved via the *state machine approach* [14], [15]. In this approach, the several application *replicas* begin in the same initial state, and they cooperate to agree on the order in which to execute any event (i.e., message or timeout) that might change the state of the application. The state transition caused by executing an event is assumed to be deterministic. Therefore, by executing events according to the agreed upon order, the replicas proceed through the same sequence of states.

It is important to note that although the replicas in a state machine replication system must be functionally equivalent, they are permitted to have different implementations, provided they all adhere to the same abstract protocol specification [16]. Indeed, the effectiveness of the state machine approach to replication depends upon using replicas that are unlikely to suffer correlated vulnerabilities, which can be achieved by using replicas with diverse implementations. Diversity can be introduced at various levels, including at the operating system (OS) level [17] and at the application level.

At the operating system level, one may introduce diversity by running each replica on a different OS (or on as many as are available). Garcia, *et al.* [17] studied over 15 years of known OS vulnerabilities and found that there exist sets of operating systems that exhibit sufficient diversity to avoid suffering common vulnerabilities. While this approach may be a viable option in some deployments, in others it may result in excessive management complexity or may simply not be possible (i.e., if an application is tied to a specific OS). Within a single OS deployment, address space layout randomization (ASLR) [18] can be used

²Usually subject to certain cryptographic assumptions.

to generate diversity. ASLR is used by many modern operating systems (e.g., OpenBSD, Linux, Solaris, Microsoft Windows, Mac OS X). It places the sections of a process' address space at random offsets. This mitigates certain types of attacks that rely on being able to predict addresses, because the addresses are likely to differ at each replica.

Traditionally, achieving diversity at the application level has required expensive techniques such as N-version programming [19]. However, newer approaches can automatically create software diversity during compilation [20] or (if source code is unavailable) through binary re-writing [21]. Such approaches require no additional development effort and have been demonstrated to have minimal performance impact. Encouragingly, the compiler-based approach of [20] has been used to diversify an entire Linux operating system. Our system also introduces diversity by deploying each replica with its own private key. An attacker that compromises a replica can cause it to send messages that appear legitimate but have invalid content only if the attacker can compromise the replica's private key. For this reason, in a real deployment it may also be prudent to protect the cryptographic keys using tamper-proof hardware.

As noted in Section I, we selected the Prime intrusion-tolerant replication protocol [5], [6] for our survivable SCADA system. Prime requires $3f + 1$ replicas to tolerate f Byzantine faults and was the first protocol to guarantee both correct operation and good performance in executions in which up to f of the replicas actually exhibit Byzantine behavior. Prime uses an elected leader to coordinate the ordering of events. We selected Prime because its service properties make it a particularly good fit for real-time applications, such as a SCADA Master. Specifically, Prime bounds the amount of delay that can be introduced by Byzantine replicas: assuming enough correct (i.e., non-Byzantine) replicas can communicate with one another in a timely manner, Prime ensures that any event submitted to the system will be ordered by the correct replicas within a bounded delay δ , where δ is a function of the network round-trip times (and their variance) among the correct replicas. To achieve this property, Prime runs a distributed monitoring protocol, whereby the replicas constantly monitor the performance of the current leader and quickly elect a new leader if they detect that the current one is performing too slowly.

III. SURVIVABLE SCADA: SYSTEM ARCHITECTURE

A. Motivation and High-Level Design

We believe the highest value asset in a SCADA system is the SCADA Master, because its compromise can have serious consequences for the control and monitoring of the entire system. Therefore, our focus in this paper is on improving the robustness of the SCADA Master by making it survivable.

In our survivable SCADA architecture, instead of running a Primary and a Hot Standby, we run $3f + 1$ peer replicas of the (primary) SCADA Master application, where f is the maximum number of replicas that may be Byzantine. Fig. 1 depicts a minimal configuration of the system, which runs four SCADA Master replicas (i.e., $f = 1$). Each replica links with a local

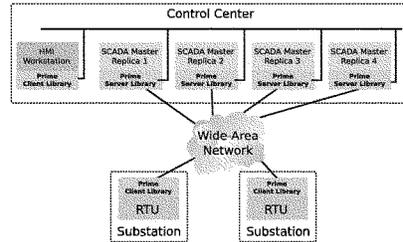


Fig. 1. A survivable SCADA system capable of tolerating the compromise of one SCADA Master replica.

copy of the *Prime Server Library*, the intrusion-tolerant replication engine that delivers to each replica the same events in the same order.

The collection of SCADA Master replicas forms a *logical SCADA Master* that behaves correctly even if f replicas are Byzantine; that is, upon processing an event, the logical SCADA Master makes the same state transition as an unreplicated, uncompromised SCADA Master application would make given the same event. Moreover, Prime's service properties bound the degree to which the Byzantine replicas can slow down the performance of the logical SCADA Master compared to a correct, unreplicated SCADA Master.

In order to interact with the replicated SCADA Master, the HMI and the RTUs each link with the *Prime Client Library*. This library provides functions for i) sending messages to multiple replicas and ii) voting on the messages that arrive from the replicas to determine the correct content.³ In some deployments one may not have access to the RTU source code in order to link it with the Prime Client Library. We addressed this by implementing an RTU proxy, which communicates with the RTU using its native protocol and uses the Prime Client Library to interact with the SCADA Master replicas. This "bump-in-the-wire" solution may also be suitable for legacy RTUs with limited computational power or memory.

Although our solution significantly improves the robustness of the SCADA Master and its communication with RTUs, we emphasize that achieving system-wide survivability requires taking a systematic approach to security at all levels and in all parts of the system. Intrusion-tolerant replication is complementary to more traditional host, network, and perimeter security technologies and should be deployed alongside them rather than being seen as a replacement for them.

In previous work we discussed how virtualization can be used to reduce the hardware cost of replication to that of a conventional SCADA system: four SCADA Master replicas can be run on only two physical machines (the number required for a Primary/Hot Standby deployment), while still maintaining the ability to survive the crash of either machine. We refer the interested reader to [4] for details.

³Since at most f replicas may be Byzantine, a client can act on a message when it receives $f + 1$ copies of the message from different replicas, indicating that at least one correct replica sent a message with the given content.

B. Attack Model and Assumptions

As described above, the logical SCADA Master is implemented by a set of $3f+1$ replicas, f of which may be Byzantine. Byzantine replicas may send invalid or conflicting messages to other replicas or clients, drop or delay messages, or otherwise attempt to disrupt the system. Digital signatures provide message integrity, authentication, and non-repudiation for all Prime messages.

We assume that Byzantine replicas cannot disrupt the communication between correct replicas; this can be achieved through network isolation techniques [11] or by using tamper-proof network cards capable of rate-limiting outgoing traffic (e.g., [22]). Such solutions insulate the correct replicas from denial-of-service attacks launched from within the control center. Mitigating external denial-of-service attacks is a difficult problem for which production systems typically rely on commercial solutions. Note that denial-of-service attacks may, for their duration, affect the performance and monitoring capability of the system but do not affect the consistency of the replicas and, therefore, the correctness of the supervisory control commands issued by the logical SCADA Master.

As in a conventional SCADA system, the ability of the logical SCADA Master to monitor effectively the physical assets of the system relies on the RTUs to report legitimate values. In the current version of our system, a compromised RTU may report incorrect values that will be replicated consistently by the SCADA Master replicas. This impacts the SCADA Master's ability to monitor the substation containing the compromised RTU but is unlikely to affect the monitoring of other substations. Note that although our focus is on making the SCADA Master survivable, intrusion-tolerant replication could also be used at the substation level to form survivable RTUs. In a real deployment, a utility would need to evaluate the costs and benefits of such an approach to determine the feasibility of using replication at this level.

IV. SURVIVABLE SCADA: CHALLENGES AND SOLUTIONS

Conventional SCADA systems are *server driven*: the SCADA Master (the server) periodically sends poll requests to the RTUs (the clients), and the RTUs respond with their current status.⁴ The periodic sending of poll requests is triggered by the expiration of *timeout events* at the SCADA Master. In contrast, existing intrusion-tolerant replication systems implicitly assume that the state machine of the server application being replicated is *client driven*: the application processes a client request as input, sends a reply to the requesting client as output, and then processes the next client request.

In this section we describe how this seemingly small difference has large implications on the functionality required from the replication engine in our survivable SCADA system. Indeed, addressing this architectural mismatch required the invention of several new protocols.

A. Scalable Intrusion-Tolerant Synchronization

1) *Motivation*: When polling an RTU, a SCADA Master takes action based on the passage of time: the expiration of

⁴Some SCADA communication follows the traditional client-server pattern, such as the request/reply protocol between an HMI and the SCADA Master.

a local timeout triggers the SCADA Master to send a poll request to the RTU. However, absent perfectly synchronized clocks, the passage of time will be observed in a non-deterministic way at the different SCADA Master replicas in our survivable SCADA system. As a result, if the replicas were to make state transitions based on their local clock values, they could become inconsistent with one another. Therefore, in our survivable SCADA system, the SCADA Master replicas use a *logical timeout protocol* to agree on the *logical time* at which a time-based action (such as generating a poll request) should be taken. This protocol must be intrusion-tolerant so that Byzantine replicas cannot disrupt the agreement or trigger the expiration of spurious timeouts. Moreover, since large SCADA systems may contain thousands of RTUs, each of which may be polled individually, the protocol must scale with the number of different logical timeouts being agreed upon.

Existing techniques for time-based synchronization in a Byzantine environment [5] are intrusion tolerant but not scalable, requiring a number of messages proportional to the number of logical timeouts set by the application. In [5], each replica sends a "vote" message each time its local clock indicates a logical timeout should expire (i.e., once per timeout), and the replicas act on a given logical timeout when they agree that $f+1$ replicas have sent corresponding votes.

To provide a solution that scales to large SCADA deployments, we developed a new protocol that only requires a constant number of messages to be exchanged, independent of the number of logical timeouts requested by the application. Our protocol makes no assumptions about the relative speeds of the replicas' local clocks and prevents Byzantine replicas from arbitrarily advancing or delaying the logical time at which a logical timeout expires. As explained below, our protocol achieves scalability at the cost of a (potentially) slightly lower timer resolution compared to [5].

2) *Protocol Details*: The Prime Server Library provides an API call enabling an application to *set* a logical timeout, (t, d) , where d is the number of seconds that should pass before t expires. For example, to poll an RTU a SCADA Master replica might set a logical timeout, t , with a duration, d , of 1 second. Approximately 1 second later, the replica will receive from the Prime Server Library a notification that t has expired. Upon receiving this notification (which is delivered to all replicas at the same logical time), each replica generates an identical poll request and sends it to the RTU.

Observe that since the replicas set each logical timeout at the same logical time, they can consistently map each logical timeout to a unique sequence number, where the i th timeout set is mapped to sequence number i . Although logical timeouts are set in sequence number order, they may *expire* out of order, because they can be set with arbitrary (non-negative) durations.

Each replica r periodically broadcasts to the other replicas a SYNC message of the form $(\text{SYNC}, \text{localClock}, \text{lastTimeoutSet}, r)$, where *localClock* is r 's current local clock value and *lastTimeoutSet* is the sequence number of the last logical timeout that r has set. The replicas use Prime to order the SYNC messages. Thus, each replica executes an (identical) ordered stream, S , of SYNC messages. Prime guarantees that if replica r sends SYNC message S_1 before sending SYNC message S_2 , then S_1

appears before S_2 in S ; SYNC messages from different replicas may be interleaved in S . Note that SYNC messages are sent periodically and are the only types of messages sent during the protocol. Thus, the protocol has a constant message complexity.

As explained below, the replicas agree upon the logical time at which each logical timeout should expire by using a deterministic, online algorithm that examines the ordered stream of SYNC messages, $S = S_1, S_2, \dots$, one message at a time, as each new message S_i is ordered by Prime. The examination of the SYNC message most recently ordered results in the expiration of a (potentially empty) set of logical timeouts. The algorithm proceeds in three steps.

Step 1: Identifying candidates for expiration. Let $M = \langle \text{SYNC}, c_r, i, r \rangle$ be the current SYNC message being examined by some replica s . From this message, s learns that M was originated by replica r at time c_r (i.e., r 's local clock value) and that the last timeout set by r had sequence number i . Since logical timeouts are set in sequence number order, M also implies that r has set all timeouts 1 through i .

At replica s , we say that a logical timeout with sequence number j becomes a *candidate for expiration with respect to* r the first time that s examines a SYNC message implying that r has set a timeout with sequence number j . In general, the examination of SYNC message M by s may cause several logical timeouts to become candidates for expiration with respect to r ; this stems from the fact that r generates SYNC messages periodically rather than each time it sets a logical timeout. As a concrete example, if s is currently examining a message M_i from r with $\text{lastTimeoutSet} = 5$, while the last message that s examined from r had $\text{lastTimeoutSet} = 3$, then the logical timeouts with sequence numbers 4 and 5 would become candidates for expiration when s examines M_i . Each replica maintains N *candidate lists* (N being the number of replicas), numbered 1 through N , where list r contains an entry for each logical timeout that has become a candidate for expiration with respect to r .

Step 2: Computing when each candidate should expire. Each candidate timeout t is stored along with its *local expiration time with respect to* r . This value is computed as $c_r + d$, where c_r is the local clock value contained in the SYNC message that caused t to become a candidate and d is the duration of t as requested by the application. Intuitively, c_r represents the “best guess” of replica s for when r set timeout t , so s believes t should expire when r 's local clock reads $c_r + d$ (i.e., d seconds later). Note that since r only sends SYNC messages periodically, c_r may be up to one period later than the actual time at which r set t . This error is the price paid by our protocol to achieve constant message complexity.

Step 3: Triggering the expiration of candidate timeouts. At replica s , we say that a candidate timeout t is *triggered with respect to replica* r when s examines a SYNC message from r indicating that r 's local clock has reached the local expiration time associated with t in candidate list r . Observe that a logical timeout that became a candidate upon examination of a SYNC message from r will typically not be triggered until a later SYNC message from r is examined; this is due to the fact that examining new SYNC messages from r is the only way in which s updates its estimate of r 's current local clock value. When a logical timeout has been triggered with respect to $f + 1$ dif-

ferent replicas, the timeout is said to *expire* and is delivered to the application.

3) *Discussion:* We make several observations about our logical timeout protocol. First, because the examination of ordered SYNC messages is deterministic, each logical timeout (t, d) expires at the same logical time at all correct replicas. Second, Byzantine replicas cannot cause t to expire before d seconds have elapsed on the local clock of at least one correct replica; thus, Byzantine replicas cannot cause t to expire “too soon.” This property holds because t does not expire until it is triggered with respect to $f + 1$ replicas, at least one of which is correct. Third, Byzantine replicas cannot delay the expiration of t : t becomes a candidate for expiration and becomes triggered with respect to a correct replica r based solely on the SYNC messages ordered from r , which cannot be influenced or delayed by the Byzantine replicas.

The “clock resolution” of our logical timeout protocol is determined by two configurable parameters: i) the rate at which SYNC messages are generated (since this dictates how quickly a replica's local clock value can be observed to advance), and ii) the rate at which SYNC messages can be ordered by Prime. The latter is determined by the network delay between replicas (which, on a LAN, should be small) and the rates at which certain periodic messages are sent in Prime. The resolution of the timeout protocol is therefore limited by the slower of rates i) and ii). As explained in Section V, our current implementation achieves a resolution of approximately 17 ms, with the limiting factor being the rate at which Prime orders SYNC messages.

B. Intrusion-Tolerant Reliable Channels

Many SCADA systems make use of a reliable transport protocol, such as TCP, to pass messages between the SCADA Master and the RTUs. By contrast, existing intrusion-tolerant replication systems tend to use UDP and implement their own reliability.⁵ In such systems, the replication engine passes application messages between clients and the server replicas. Unfortunately, since existing intrusion-tolerant replication systems implicitly assume that the application is client driven, they provide only limited support for reliable communication between a client and the server replicas. Most use a transaction-based protocol, where the client retransmits its request if it does not receive a response within a timeout period. This approach makes additional implicit assumptions, namely that the server application will always generate a response that can be used by the client as an acknowledgement, and that this acknowledgement message will be sent to the requesting client.

In our survivable SCADA system, messages may be originated by both the SCADA Master replicas and by clients. Moreover, the execution of a client message (e.g., a poll response) by the replicas may cause them to send a reply to an entirely different entity (e.g., the HMI workstation) or not to send a reply at all. Therefore, the system needs a more flexible approach to achieving reliability than the simple client-driven scheme just described.

⁵As noted in [7], TCP is poorly suited to systems with potentially Byzantine receivers because, by failing to send acknowledgements, the Byzantine receivers can require correct replicas to buffer an unbounded number of messages.

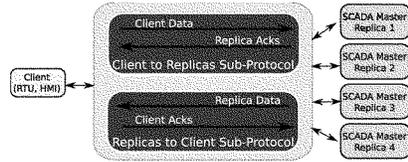


Fig. 2. Intrusion-tolerant reliable channel abstraction.

To address this issue, we developed two protocols that together implement the abstraction of an *intrusion-tolerant reliable channel* between clients and the SCADA Master replicas. Each protocol handles a different communication direction (see Fig. 2). The two endpoints of a channel are asymmetric: one is a client (RTU or HMI) and the other is the set of SCADA Master replicas. Using two unidirectional protocols allows us to take advantage of this asymmetry to optimize performance in each direction. The SCADA Master replicas may communicate with many clients, using a separate channel for each client. Each client is an endpoint of exactly one channel.

Applications interact with a channel using an API similar to that of a TCP socket. The API provides calls for establishing a connection, sending and receiving a message into and from the channel, and closing a connection. Despite the similarity of our API to the socket API, the fact that one end of our channels is replicated has several practical implications. First, when a client application sends a message m into a channel, the channel implementation actually sends m to multiple replicas ($f + 1$) to ensure its delivery (since some replicas may be Byzantine). Second, although it is a useful abstraction to imagine that a logical SCADA Master application sends a single message m to a client, the sending of m by the logical SCADA Master actually requires several physical messages to be sent, since multiple replicas ($2f + 1$) introduce m into the channel at the same logical time. Byzantine replicas may also introduce arbitrary messages into the channel at any time.

We refer to messages introduced into a channel by correct replicas or correct clients as *legitimate*; all other messages are *spurious*. The correctness requirements of our channel abstraction, which we now state, define the delivery properties of legitimate and spurious messages.

Let \mathcal{C} be a communication channel established between the SCADA Master replicas and a correct (non-Byzantine) client. We say that \mathcal{C} is an “intrusion-tolerant reliable channel” if, even when up to f of the SCADA Master replicas are Byzantine, it i) delivers legitimate messages without modification or unnecessary delay; ii) does not deliver spurious messages; and iii) prevents either endpoint from causing the other to consume excessive computational or networking resources. Achieving intrusion tolerance is challenging, because Byzantine replicas may attempt to insert spurious messages into the channel, delay or prevent the delivery of certain legitimate messages, attempt to deliver messages out of order, or otherwise attempt to disrupt the protocol.

Let \mathcal{C}' be a communication channel established between the SCADA Master replicas and a Byzantine client. \mathcal{C}' makes

no delivery or timing guarantees for messages sent from the replicas to the client. Messages sent from the client (which, by definition, are spurious) to the replicas may or may not be delivered, but if any correct replica delivers a message, then they all do. The channel also prevents the Byzantine client from consuming excessive resources at the replicas.

The following sub-sections provide a high-level overview of the operation of each reliable channel sub-protocol.

Client-to-Replicas Sub-Protocol: To introduce a data message into the channel, a client assigns it a sequence number and sends it to a set of $f + 1$ replicas. Since at most f replicas are Byzantine, this ensures that the message is received by at least one correct replica. Each client message carries a digital signature, which prevents Byzantine replicas from modifying its content. Upon receiving a data message, a replica places it into its *local window*. A replica introduces a message for ordering (via Prime) when all messages with lower sequence numbers have been placed in its window. Note that the same client message may be (legitimately) introduced for ordering multiple times (by different replicas), and that Byzantine replicas may introduce messages for ordering out of sequence number order. The replica’s channel implementation overcomes these issues and ensures that the client’s messages are delivered exactly once, in sequence number order.

Each replica sends cumulative acknowledgements containing the sequence number of the last client data message it has executed. Since the replicas execute data messages at the same logical time, all correct replicas construct identical acknowledgements. The client can slide its window forward (and thus send more messages) when at least $f + 1$ distinct replicas have acknowledged executing the data message at the front of the window. This prevents Byzantine replicas from causing the client to prematurely slide its window. Note that although the replicas send identical acknowledgements, they send negative acknowledgements individually (i.e., based on which data messages they have locally received). This separation enables faster packet recovery in the face of loss, because a replica need not wait until an out-of-order message is executed before requesting a retransmission.

Replicas-to-Client Sub-Protocol: Since the SCADA Master replicas construct data messages at the same logical time, each correct replica introduces identical data messages into the channel, in the same order. Outgoing messages are assigned a sequence number and sent to the client. A client’s channel implementation delivers a data message to the client application when the client receives $(f + 1)$ copies of the message (from distinct replicas) and when the channel has delivered all messages with lower sequence numbers.

Clients send cumulative acknowledgements containing the sequence number of the last data message they have delivered. A client also sends negative acknowledgements for those messages it knows to be missing. For efficiency, the client indicates, for each missing sequence number, from which replicas it has already received a copy of the message. Such replicas do not need to retransmit their copies.

The SCADA Master replicas explicitly avoid introducing for ordering (via Prime) client acknowledgements that they receive from the network. This significantly reduces the computational

overhead of the protocol, because it avoids the several rounds of message exchange and the corresponding cryptographic operations associated with ordering. However, an important implication of this design decision is that different replicas may process client acknowledgements at different logical times. As a result, replicas may disagree on which messages the client has received so far, and they may slide their windows asynchronously with respect to one another. To ensure that the replicas still proceed through the same sequence of application states despite the asynchrony that may occur within their channel implementations, the replicas *do* use Prime to order the limited number of key events that might cause the behavior of the application to change. These events include i) connection establishment messages, ii) connection termination messages, and iii) messages indicating that a given replica believes a connection should be closed due to a timeout or because the client is not reading fast enough. Therefore, the replicas always agree on whether the logical state of a channel is open or closed, so they still behave like deterministic state machines at the application level.

V. PERFORMANCE EVALUATION

We have integrated our Prime-based intrusion-tolerant replication engine with a real SCADA Master product for electricity distribution, and we have developed a proxy to integrate this survivable SCADA Master with an RTU. Before integrating our engine with this product, we benchmarked the engine in both fault-free and under-attack scenarios to verify that its throughput and latency could meet the performance requirements of a SCADA system. We also evaluated the performance of our logical timeout protocol to determine whether it could scale to large deployments. Benchmarking the engine in this way enabled us to assess its performance in isolation from the effects of any particular SCADA product or deployment. This section presents the results of our benchmarks.

A. Testbed and Network Setup

We used a cluster of four Dell Precision T3500 servers. The machines had 64-bit, 6-core, 3.47 GHz Intel Xeon processors, with 12 GB RAM and hyper-threading enabled. The machines were connected on a local-area network via a Netgear ProSafe 8-port Gigabit switch. All machines ran 64-bit Fedora 12 Linux. 1024-bit RSA signatures [23] provided authentication and non-repudiation. Each machine can compute an RSA signature in 0.389 ms (2570/sec) and can verify an RSA signature in 0.02 ms (49,683/sec).

For our benchmarks, we implemented a simplified SCADA Master and a simplified RTU. The SCADA Master links with the Prime Server Library and the RTU links with the Prime Client Library (see Fig. 1). The SCADA Master can be configured to be server driven (i.e., to poll one or more RTUs, driven by the expiration of logical timeouts) or client driven (i.e., to receive RTU state updates and send replies). In the experiments described below, we ran two SCADA Master replicas on each of two machines (for a total of four replicas), and the remaining machines ran RTU processes.

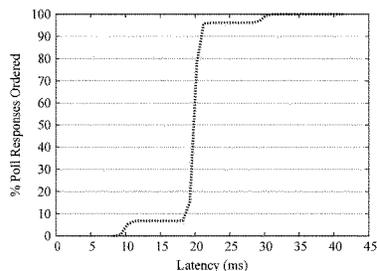


Fig. 3. Poll operation latency, cumulative distribution function.

B. Polling Scenario

Since the main operation in a SCADA system is the polling of RTUs by the SCADA Master, we first evaluated the performance of our engine in a **polling scenario**. We ran four replicas of the SCADA Master, and we ran 1000 RTU processes. The replicas polled each of the 1000 RTUs individually, at a rate of once per second. The time at which each RTU was initially polled was selected uniformly at random over a 1 second interval. Poll requests were 100 bytes long. Upon delivering a poll request, an RTU responded by sending a 100-byte poll reply. When the SCADA Master replicas delivered an ordered poll reply, they re-scheduled a logical timeout for 1 second in the future to poll the associated RTU again. At steady state, replicas set (approximately) 1000 logical timeouts per second, sent 1000 poll requests per second, and received 1000 poll replies per second.

Polling Latency: First, we measured the latency of each poll operation, as measured by SCADA Master replica 1 during an 8-minute run. The latency of a poll operation is computed as the time between the replica sending the poll request to a given RTU and executing the ordered poll reply from that RTU. Since the test was performed on a LAN with sub-millisecond link delay, the measured latency was dominated by the time required for the Prime Server Library to order (and subsequently deliver) the incoming poll reply. This enabled us to measure the amount of latency added by Prime to a typical polling roundtrip. In a real deployment, the SCADA Master replicas would be separated from the RTUs by a wide-area network, so the actual polling latency reported here would be scaled up by the network roundtrip time.

Fig. 3 shows a cumulative distribution function of the polling latencies measured during the run. The y-value of a point represents the percentage of poll operations whose latency was less than or equal to the x-value of the point. For example, the figure shows that about 96 percent of poll operations had a latency less than or equal to approximately 22 ms, and all poll operations had a latency of less than 43 ms. Our discussions with SCADA system architects suggest that, given the supervisory nature of SCADA, this latency is sufficiently low to be suitable for real deployments (in fact, even a latency added by Prime that was twice as high would likely be low enough).

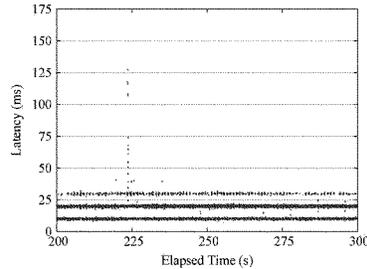


Fig. 4. Polling operation latency, under-attack scenario. At time 220 s, server replica 1 (the coordinator of the replication protocol) began delaying its outgoing messages by 100 ms. The other replicas quickly detected the attack so that subsequent operations were not affected.

Polling Under Attack: To demonstrate that Prime can mitigate performance attacks effectively, we re-ran the polling experiment, but this time we instrumented the coordinator of Prime so that it would begin delaying its outgoing packets by 100 ms after the system had been running for 220 seconds. Fig. 4 shows the latency of the poll operations initiated between time 200 and 300 seconds, as measured at SCADA Master replica 2. Each point represents the latency of a single poll operation. Before the attack, all poll operations had a latency of less than 30 ms, with most falling into two bands at 10 ms and 20 ms. The bands reflect the batching period of 10 ms used for certain messages within Prime. At time 220 s, there is a momentary spike in latency when the attack is triggered; poll operations initiated at this time were delayed by as much as 130 ms. The spike contains at most one operation for each RTU. That the spike is so “skinny” implies that the other replicas quickly detected the attack and reconfigured Prime to mitigate the attack (by electing a new coordinator) so that subsequent operations were not affected.

C. Scalability Scenario

To better understand the scalability of our replication engine, we ran a **scalability scenario** in which we measured Prime’s request ordering latency at different throughputs. To generate load, we configured the SCADA Master to be client driven: the RTU submitted to the SCADA Master replicas a 100-byte request to be ordered, and then the replicas responded with a 100-byte reply. Upon receiving a reply, an RTU submitted a new request. We ran between 1 and 30 RTUs, each of which had up to 40 outstanding requests at a time.

In our first, baseline scalability test, each RTU submitted each of its requests to $(f+1)$ SCADA Master replicas. Since at most f replicas may be Byzantine, this ensures that the RTU’s message is received by a correct replica the first time it is sent. The trade-off is that in fault-free executions (like the one tested), each message is actually ordered $f+1$ times (twice, in this case), reducing the maximum throughput that can be achieved. The throughput numbers that we report apply to the number of *unique* requests ordered per second.

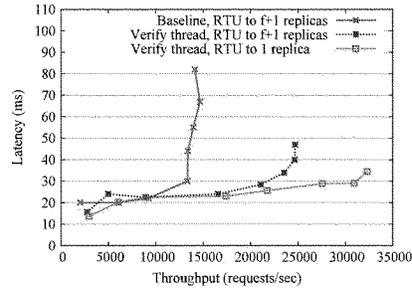


Fig. 5. Request latency vs. replication engine throughput.

Fig. 5 shows the average latency of a request vs. the throughput achieved by our replication engine. Latency was measured at the RTU and computed as the time between submitting a request for ordering and receiving the corresponding reply, averaged across all requests during the run. Throughput was measured at the SCADA Master replicas as the number of requests ordered per second. The performance of our baseline test is shown in the line labeled “Baseline, RTU to $f+1$ replicas” in Fig. 5. Requests experienced a latency of about 25 ms when 12 000 requests per second were ordered, and they experienced a latency of about 30 ms when 13 500 requests per second were ordered. For loads beyond this point, latency dramatically increased because the system was saturated and at its peak throughput.

In running our baseline scalability test, we observed that the throughput of the system was CPU bound, and that the performance bottleneck was the computation required to verify the RSA signatures contained in all Prime messages. Therefore, we re-engineered our engine to use a separate thread for the verification of digital signatures, allowing the implementation to better exploit multiple CPU cores. As seen in the middle plot in Fig. 5, using a verification thread significantly increased the peak throughput of our engine, resulting in requests experiencing a latency of about 27 ms when 20 000 requests per second were ordered and about 35 ms when 24 000 requests per second were ordered.

In the two tests just described, the RTU submitted its requests to $(f+1)$ replicas. A different strategy would be for the RTU to initially submit its request to only one replica, and if it does not receive a reply within a timeout period, it submits to a different replica. This strategy can result in higher peak throughput in fault-free executions, but in under-attack scenarios it can cause latency to be increased by the duration of the RTU’s timeout. To assess the fault-free performance impact of this approach, we configured each RTU to submit its requests to a randomly-selected replica. As seen in Fig. 5, this modification increased the scalability of our engine even further, enabling it to order 30 000 requests per second with a latency just under 30 ms. We comment that although this optimistic strategy may be suitable for some SCADA systems, others may be more sensitive to delay, and thus which strategy to prefer is deployment-specific.

TABLE I
LOGICAL TIMEOUT ACCURACY, SINGLE TIMEOUT

| Target Delay (ms) | Minimum Delay (ms) | Maximum Delay (ms) | Average Delay (ms) | Average Δ (ms) | Standard Dev. (ms) |
|-------------------|--------------------|--------------------|--------------------|-----------------------|--------------------|
| 0 | 8.8 | 26.5 | 16.8 | 16.8 | 4.7 |
| 10 | 19.3 | 33.1 | 30.0 | 20.0 | 1.4 |
| 100 | 109.7 | 122.4 | 119.5 | 19.5 | 3.3 |
| 500 | 512.1 | 524.5 | 514.7 | 14.7 | 3.1 |
| 980 | 993.6 | 1000.4 | 997.0 | 17.0 | 0.8 |
| 1000 | 1015.6 | 1020.7 | 1017.1 | 17.1 | 0.8 |

The above results suggest that the performance of our intrusion-tolerant replication engine will be sufficient for most SCADA systems, even large-scale systems with several thousand RTUs being polled approximately once per second.

D. Logical Timeout Performance

Accuracy: To measure the accuracy of our logical timeout protocol, we configured the SCADA Master application so that it set logical timeouts at various periods; upon delivering a logical timeout, each application replica re-scheduled a new one with the same duration. As shown in Table I, we ran the application in this scenario with several different target periods (Table I, column 1). In each run, the replicas set timeouts at the given period for a five minute duration. For each logical timeout, we measured the time between replica 1 setting the timeout and delivering the event indicating that the timeout expired. Columns 2, 3, and 4 of Table I report the minimum, maximum, and average delays measured by replica 1, respectively. Column 5 reports the average Δ , defined as the actual delay experienced minus the target delay, and Column 6 reports the standard deviation.

Table I shows that the average Δ for all measured target delays is between 15 and 20 ms, indicating that the application consistently experienced a delay 15–20 ms higher than it requested. As discussed in Section IV-A, this “error” reflects the clock resolution of the logical timeout protocol and is caused by the time required for Prime to reach agreement on SYNC messages. Since the Δ values are fairly consistent, the results suggest that an application should take the error into account when specifying its polling period. For example, Table I shows that an application wishing to poll at one second intervals should specify a polling period of 980 ms to compensate for the error. The 0 ms row in Table I shows the clock resolution directly when logical timeouts are set sequentially (i.e., the next one is started only after the current one expires). The data show that Prime can deliver roughly 60 sequential timeout events per second, reflecting an agreement time of roughly 17 ms.

Scalability: To evaluate the scalability of our logical timeout protocol, we measured the average Δ observed by the application when increasing numbers of periodic logical timeouts are set. We repeated this experiment for four different target periods: 10 ms, 100 ms, 500 ms, and 1 s.

Fig. 6 shows the intuitive result that the number of timeouts that can be set before the system reaches saturation (and the Δ values spike) increases with the target period. For example, the replicas can reach agreement on roughly 20 000 500 ms timeouts with a Δ of about 27 ms, and they can reach agreement on roughly 30 000 1 s timeouts with a Δ of 28 ms. Recall from Table I that the Δ values for 500 ms and 1 s timeouts when only

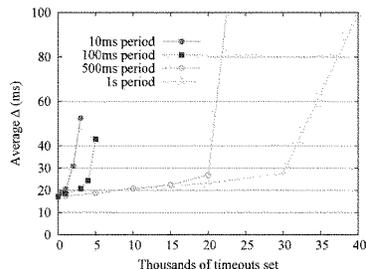


Fig. 6. Logical timeout scalability. Δ is the difference between the actual delay observed by the application and the target delay (timeout period).

one timeout was being set at a time were 14.7 and 17.1 ms, respectively. Thus, in the 1 s case, the number of timeouts being set increased by a factor of 30 000 while the average Δ increased by less than a factor of 2, reflecting the protocol’s scalability.

VI. RELATED WORK

The number of reported cyber attacks, especially insider attacks, continues to rise; McAfee reported more than 90 million unique pieces of malware in its database in its Q2 2012 threat report [24], up from 70 million just one year earlier. Therefore, SCADA systems have begun deploying standard IT security technologies to harden their defenses. Some representative technologies include firewalls [25] to police incoming traffic; intrusion-detection systems [26] to monitor network and system events to log and report suspicious behavior; and application whitelisting [27] to ensure that only known, trusted executables can run. Although these technologies significantly enhance the security of today’s SCADA systems, they focus on preventing attacks and do not protect the SCADA application if an attack compromises part of the system.

The field of *intrusion tolerance* [12] represents a different way of thinking about security and availability. An intrusion-tolerant protocol assumes that some of the protocol participants may be *Byzantine* [13] and act in an arbitrary manner. Over the last decade, using intrusion-tolerant protocols to achieve consistent global state (e.g., [6]–[11]) has been shown to be an effective technique for building highly available systems able to withstand partial compromises. Such protocols are known as *intrusion-tolerant replication protocols*. While earlier protocols guaranteed correctness (i.e., replica consistency) in the face of partial compromise, more recent protocols [6], [11], [28] also guarantee minimal performance degradation whilst under attack and hence meet our definition of survivability. Importantly, these recent protocols can also scale to support thousands of clients.

A different approach to applying intrusion tolerance techniques to critical infrastructure systems was presented by Bessani *et al.* [29]. Rather than integrating intrusion-tolerant replication within the SCADA system itself, one creates an intrusion-tolerant “firewall” (called a CRUTIAL Information Switch) that sits on the perimeter of the network and ensures

that only messages which adhere to policy are admitted into the system, even if some of the replicas implementing the firewall are compromised. Such an approach has the benefit that it does not require any changes to, or integration with, the SCADA Master. However, its effectiveness requires the policy to be specified and implemented correctly, and (unlike our approach) it does not protect the SCADA Master from attacks launched from within the system's security perimeter. Another important distinction is that CRUTIAL relies on trusted hardware components, which are assumed to be unable to be compromised. In contrast, our survivable SCADA system assumes that any machine may be compromised.

VII. CONCLUSION

In addition to the conventional challenges to availability, such as hardware crashes, power failures, and network partitions, SCADA providers must also anticipate the consequences of cyber attacks. Whereas conventional enterprise security technologies have sought to build increasingly sophisticated perimeter defenses, in this research we sought to answer whether it is possible to build a SCADA system that is able to operate correctly and with good performance even if a cyber attack was successful at evading these conventional defenses.

As the compromise of the highest value asset, the SCADA Master, can have potentially disastrous consequences, our work has focused on protecting this entity via intrusion-tolerant replication. In effect, intrusion tolerance allows the SCADA Master application to act as its own firewall, thus providing protection in the event of a security breach.

This paper reports on our experience designing and evaluating the first survivable SCADA system. We described the unique requirements imposed by the SCADA architecture and gave an overview of several new techniques facilitating the integration of intrusion-tolerant replication and SCADA. Our experimental results illustrate that our replication engine performs sufficiently well to meet the needs of even large-scale SCADA systems containing thousands of RTUs.

REFERENCES

- [1] N. Budhiraja, K. Marzullo, F. B. Schneider, and S. Toueg, "Primary-backup protocols: Lower bounds and optimal implementations," in *Proc. 3rd IFIP Conf. Dependable Comput. Critical Appl.*, 1992, pp. 187–198.
- [2] E. Byres, "Next generation cyber attacks target oil and gas SCADA," *Pipeline Gas J.*, vol. 239, no. 2, 2012.
- [3] D. S. Wall, "Organization security and the insider threat: Malicious, negligent, and well-meaning insiders," 2011, Symantec.
- [4] J. Kirsch, S. Goose, Y. Amir, and P. Skare, "Toward survivable SCADA," in *Proc. Annu. Cyber Security Inf. Intell. Res. Workshop (CSIRW'11)*, Oct. 2011.
- [5] J. Kirsch, "Intrusion-tolerant replication under attack," Ph.D. dissertation, Johns Hopkins University, Baltimore, MD, USA, 2010.
- [6] Y. Amir, B. Casan, J. Kirsch, and J. Lane, "Prime: Byzantine replication under attack," *IEEE Trans. Dependable Secure Comput.*, vol. 8, no. 4, pp. 564–577, 2011.
- [7] M. Castro and B. Liskov, "Practical Byzantine fault tolerance and proactive recovery," *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, 2002.
- [8] J. Yin, J.-P. Martin, A. Venkataramani, L. Alvisi, and M. Dahlin, "Separating agreement from execution for Byzantine fault-tolerant services," in *Proc. 19th ACM Symp. Oper. Syst. Principles*, 2003, pp. 253–267.
- [9] J.-P. Martin and L. Alvisi, "Fast Byzantine consensus," *IEEE Trans. Dependable Secure Comput.*, vol. 3, no. 3, pp. 202–215, 2006.
- [10] Y. Amir, C. Danilov, D. Dolev, J. Kirsch, J. Lane, C. Nita-Rotaru, J. Olsen, and D. Zage, "Steward: Scaling Byzantine fault-tolerant replication to wide area networks," *IEEE Trans. Dependable Secure Comput.*, vol. 7, no. 1, pp. 80–93, 2010.
- [11] A. Clement, E. Wong, L. Alvisi, M. Dahlin, and M. Marchetti, "Making Byzantine fault tolerant systems tolerate Byzantine faults," in *Proc. 6th USENIX Symp. Netw. Syst. Design Implementation*, 2009, pp. 153–168.
- [12] P. E. Verissimo, N. F. Neves, and M. P. Correia, R. Lemos, C. Gacek, and A. Romanovsky, Eds., "Intrusion-tolerant architectures: Concepts and design," in *Architecting Dependable Systems*, 2003, vol. 2677, Lecture Notes in Computer Science.
- [13] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [14] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. ACM*, vol. 21, no. 7, pp. 558–565, 1978.
- [15] F. B. Schneider, "Implementing fault-tolerant services using the state machine approach: A tutorial," *ACM Comput. Surv.*, vol. 22, no. 4, pp. 299–319, 1990.
- [16] R. Rodrigues, M. Castro, and B. Liskov, "BASE: Using abstraction to improve fault tolerance," in *Proc. 18th ACM Symp. Operating Systems Principles (OSPP'01)*, Banff, AB, Canada, 2001, pp. 15–28.
- [17] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Oliveira, "OS diversity for intrusion tolerance: Myth or reality," in *Proc. 41st IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN'11)*, 2011, pp. 383–394.
- [18] PaX Team, PaX Address Space Layout Randomization [Online]. Available: <http://pax.grsecurity.net/docs/aslr.txt>
- [19] A. Avizienis, "The N-Version approach to fault-tolerant software," *IEEE Trans. Softw. Eng.*, vol. SE-11, no. 12, pp. 1491–1501, Dec. 1985.
- [20] M. Franz, "E unibus pluram: Massive-scale software diversity as a defense mechanism," in *Proc. New Security Paradigms Workshop (NSPW 2010)*, Concord, MA, USA, Sep. 2010.
- [21] R. Wartell, V. Mohan, K. W. Hamlen, and Z. Lin, "Binary stirring: Self-randomizing instruction addresses of legacy x86 binary code," in *Proc. 2012 ACM Conf. Commun. Security (CCS'12)*, 2012, pp. 157–168.
- [22] *Secure Computing SnapGear User Manual, Revision 3.1.4* Aug. 2006, Secure Computing.
- [23] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Commun. ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [24] Z. Bai, T. Dirro, P. Greve, Y. Lin, D. Marcus, F. Paget, V. Pogulievsky, C. Schmugar, J. Shah, D. Sommer, P. Szor, and A. Wosotowsky, "McAfee threats report: Second quarter 2012," 2012.
- [25] R. Oppiger, "Internet security: Firewalls and beyond," *Commun. ACM*, vol. 40, no. 5, pp. 94–102, 1997.
- [26] D. E. Denning, "An intrusion detection model," in *Proc. 7th IEEE Symp. Security and Privacy*, May 1986, pp. 119–131.
- [27] J. V. Harrison, "Enhancing network security by preventing user-initiated malware execution," in *Proc. Int. Conf. Inf. Technol.: Coding and Computing (ITCC'05)—Volume II*, Washington, DC, USA, pp. 597–602, IEEE Computer Society.
- [28] G. S. Veronese, M. Correia, A. N. Bessani, and L. C. Lung, "Spin one's wheels? Byzantine fault tolerance with a spinning primary," in *Proc. 28th IEEE Int. Symp. Reliable Distrib. Syst.*, 2009, pp. 135–144.
- [29] A. Bessani, P. Sousa, M. Correia, N. F. Neves, and P. Verissimo, "The CRUTIAL way of critical infrastructure protection," *IEEE Security Privacy*, vol. 6, no. 6, pp. 44–51, Nov.–Dec. 2008.



Jonathan Kirsch received the B.Sc. degree from Yale University, New Haven, CT, USA, in 2004 and the M.S.E. degree from Johns Hopkins University, Baltimore, MD, USA, in 2007. He received a Ph.D. degree in computer science from Johns Hopkins University in 2010.

He is currently a Research Scientist at the Siemens Technology to Business Center, Berkeley, CA, USA. His research interests include fault-tolerant replication and survivable systems.

Dr. Kirsch has served on the program committee for the International Symposium on Stabilization, Safety, and Security of Distributed Systems, as well as the Cyber Security and Information Intelligence Research Workshop. He is an active reviewer for several journals, including IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING, IEEE TRANSACTIONS ON COMPUTERS, and ACM Transactions on Computer Systems.



Stuart Goose received the B.Sc. and Ph.D. degrees in computer science from the University of Southampton, U.K., in 1993 and 1997, respectively.

He held a Postdoctoral position at the University of Southampton. He then joined Siemens Corporate Research Inc., Princeton, NJ, USA, holding various positions in the Multimedia Technology Department where he led a research group exploring and applying various aspects of Internet, mobility, multimedia, speech, and audio technologies. His current position is Director of Venture Technology

at Siemens Technology-To-Business Center in Berkeley, CA, USA. He scouts for disruptive technologies from universities and startups, runs projects to validate the technical and business merit of technologies, and, if successful, the technologies are transferred to the relevant product lines within Siemens.

Dr. Goose serves as program committee member and reviewer for IEEE International Conference on Distributed Multimedia Systems and IEEE International Conference Multimedia Expo.



Yair Amir received B.S. and M.S. degrees from the Technion, Israel Institute of Technology, in 1985 and 1990, respectively, and a Ph.D. degree from the Hebrew University of Jerusalem, Israel, in 1995.

He has served as Professor of Computer Science at Johns Hopkins University, Baltimore, MD, USA, since 1995. Prior to his Ph.D., he gained extensive experience building CM systems. He is a creator of the Spread and Secure Spread group communication toolkits, the Backhand and Wackamole clustering projects, the Spines overlay network messaging

system, and the SMesh wireless mesh network.

Dr. Amir has been a member of various program committees including the IEEE International Conference on Distributed Computing Systems, the ACM Conference on Principles of Distributed Computing, and the IEEE/IFIP International Conference on Dependable Systems and Networks. He currently serves as an Associate Editor for the IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING. He co-founded Spread Concepts LLC (2000) and LTN Global Communications Inc (2008), and is a member of the ACM and the IEEE Computer Society.



Dong Wei (S'00-M'04) received his B.S. degree in electrical engineering from Tsinghua University, Beijing, China. He received his M.S. and Ph.D. degrees from New Jersey Institute of Technology, Newark, NJ, USA, both in electrical engineering, is a research scientist at Siemens Corporation, Corporate Technology.

He has worked at Siemens for more than 10 years. He has worked on factory automation systems, PLC, motion control, human-machine interface, drive system, and industrial communication networks for

more than 10 years. He has more than 20 publications, including book chapters and journal papers.

Dr. Wei works as Principal Investigator for several government-funded research projects. He is also an active reviewer of IEEE TRANSACTIONS ON SMART GRID, IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING, IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS, IEEE TRANSACTIONS ON SYSTEM, MAN, AND CYBERNETICS, *Computer in Industry*, *Journal of the Network and Systems Management*, etc.



Paul Skare is the Chief Cyber Security Program Manager in Electricity Infrastructure at Pacific Northwest National Laboratory, Richland, WA (PNNL). Programs that he manages include PNNL's work on the Department of Energy's Cybersecurity for Energy Delivery Systems (CEDSS) and the Cybersecurity Risk Information Sharing Program (CRISP). Previously he worked for Northern States Power for 4 years, and Siemens Energy for 26 years, including roles in power applications, R&D manager for SCADA, Product Lifecycle Manager for EMS

and substation automation products, and most recently he was the Director of Cyber Security. He has a patent published on cybersecurity for control systems. He is the Convener of Working Group 19 in IEC TC 57 and a member of WG 13 & 15 and former WG 7. He is a member of the IEEE PES and worked in IEEE P1689, P1711, and P2030. He has twice testified to the U.S. Congress on cyber security for control systems. Paul has been active in the NERC CNSWG, has been in numerous NERC working groups including Hydra and the GridEx cyber security exercises, and has been active in the DHS ICSJWG and CyberStorm III programs.

The CHAIRMAN. Thank you, Mr. Skare.
Mr. Golden, welcome.

**STATEMENT OF THOMAS A. GOLDEN, PROGRAM MANAGER,
TECHNOLOGY INNOVATION, ELECTRIC POWER RESEARCH
INSTITUTE**

Mr. GOLDEN. Thank you and good morning.

Chair Murkowski, Ranking Member Cantwell and members of the Committee, thank you for inviting me here today to discuss energy efficiency of blockchain and other similar technologies, as well as the cybersecurity possibilities of such technologies for energy industry applications.

My name is Thomas Golden, Program Manager for Technology Innovation, appearing before you on behalf of the Electric Power Research Institute, also known as EPRI.

The goal of both my written testimony as well as my testimony today is to provide this Committee with objective research findings to help inform your discussions regarding this potentially important technology.

As you know, many times there is a desire to think of blockchain and bitcoin as one and the same; however, bitcoin is not blockchain. Rather, it uses blockchain as the underlying technology platform. That being said, it is also important to note that much of the relevant early research to date has been conducted on bitcoin because of the relatively wide adoption of volatile value and popularity in the press.

There are several different types of blockchain architecture currently in use, including the proof of work, proof of stake, proof of authority, and tangle. Each of these architectures require varying levels of energy.

Bitcoin uses the proof of work architecture which is often most energy intensive. For that reason, combined with the Committee's hearing topic, I would like to share a few thoughts regarding the mining process for bitcoin.

Mining is a process of using computing power to solve cryptographic puzzles to validate new transactions in the blockchain. When bitcoin was first established mining was possible using a standard desktop PC; however, the cryptographic puzzle is made more difficult every 10 days to maintain an average of 10 minutes of solving new transactions. As these cryptographic puzzles become more difficult, the amount of computing power required to solve these puzzles increases, resulting in an increase in energy usage based on the computing power and cooling requirements.

Much like gold miners of the past traded their pans for pickaxes and their pickaxes for front end loaders, blockchain miners are constantly looking to gain efficiencies in both processes and energy requirements. There has been a transition from the standard desktop PC to something called the graphic processing unit and finally to the application-specific integrated circuit. Each transition has resulted in increasing efficiencies by either one or two orders of magnitude.

Similar to the mining route times of the past where miners congregated geographically, bitcoin miners have sought to locate themselves with unused high capacity electric grid connections, rel-

atively inexpensive electricity and a cold, dry climate. This is all intended to reduce their energy costs which have been estimated to be as high as 32 percent of overall operating costs.

When bitcoin prices hit record highs, many began to enter this new market and establish mining operations. Although the mining operations have become more sophisticated, this surge of new participants helped to drive up overall energy consumption.

Today, worldwide bitcoin energy usage is estimated between two to three gigawatts of power. To put that in perspective, this is approximately 0.1 percent of the total worldwide generating capacity, or more simply put, equivalent to the power required for nearly two million residential homes. This power consumption can be thought of as somewhat small in a global context but can be seen as very large in concentrated areas that are experiencing bitcoin boom towns.

EPRI is working with its members to understand the potential challenges associated with blockchain mining operations, including potential cost to customers. Our research will continue to examine a wide array of potential impacts this technology may impress upon the electric grid.

Additionally, EPRI has recently convened a member group called the Utility Blockchain Interest Group. This group of nearly 40 energy companies has been chartered to discuss research findings, level set technology intelligence and share results of early pilots.

Finally, it is important to state that bitcoin is not the only use for blockchain technology. Any transaction that requires trust and currently uses a third party to deliver that trust, will most likely be looked upon as a place where bitcoin can add value.

Many changes are underway in the electric grid. Gone are the days where consumers simply buy their electricity from their local trusted utility. Today we continue to see the installation of distributed energy resources such as solar panels on commercial and residential roofs. This presents an opportunity for what many are calling transactive energy. Rather than simply buying electricity from a utility, there exists a possibility where in the future you could buy and sell electricity in an open market with your neighbors and your utility. Many have theorized that blockchain technology may solve many of the challenges associated with setting up such a market. Additional research and testing is required before this theory can be truly vetted. EPRI is committed to this research and has created an initial version of the blockchain energy market simulator to test this theory.

In closing, I thank you again for the opportunity to testify before the Committee today. I look forward to discussing many of the nuances and potential use cases for this technology in the energy industry and the potential applications regarding cybersecurity.

Thank you.

[The prepared statement of Mr. Golden follows:]

Written Testimony

Hearing of the U.S. Senate Energy and Natural Resources Committee

**Thomas A Golden
Program Manager, Technology Innovation
Electric Power Research Institute**

“The purpose of the hearing is to consider the energy efficiency of blockchain and similar technologies and the cybersecurity possibilities of such technologies for energy industry applications.”

August 21, 2018

Chair Murkowski, Ranking Member Cantwell, Members of the Committee – thank you for inviting me to discuss the energy efficiency of blockchain and similar technologies as well as the cybersecurity possibilities of such technologies for energy industry applications.

The Electric Power Research Institute (EPRI) conducts research and development relating to the generation, delivery, and use of electricity for the benefit of the public. An independent, non-profit organization, EPRI brings together its scientists and engineers, as well as experts from academia, government, and industry, to help address challenges in electricity, including **reliability, efficiency, affordability, health, safety, and the environment**. EPRI’s members represent approximately 90 percent of the electricity generated and delivered in the United States, and international participation extends to more than 30 countries.

EPRI’s research into **blockchain** and its capabilities began recently (2016) as early interest in bitcoin led to questions on bitcoin mining energy usage and how other blockchain enabled technologies could impact energy industry processes and operations. EPRI’s early research efforts related to blockchain technology in the energy sector have revealed several pilots that have shown potential promise in the use of blockchain to enable transactive energy. The **GridWise Architecture Council’s (GWAC) Transactive Energy Framework** defines transactive energy as techniques for managing the generation, consumption or flow of electric power within an electric power system through the use of economic or market-based constructs while considering grid reliability constraints. The term “**transactive**” comes from considering that decisions are made based on a value. These decisions may be analogous to or literally economic transactions.

While innovation in the blockchain space is rapidly expanding blockchain capabilities, questions remain as to the standards, scalability, energy usage, and potential return on investment related to deploying blockchain-enabled technology into distribution and transmission networks. EPRI has

helped to raise awareness and provide information to the energy industry via our Technology Innovation research program and EPRI's Utility Blockchain Interest Group (UBIG). This UBIG is comprised of nearly forty energy companies and growing as the technology continues to generate great interest within the industry. EPRI has also begun developing a blockchain-based energy market simulator to explore how loads and renewable resources could work together using more granular market information.

Information & Insights

EPRI has published a whitepaper (attached) and what we term "Quick Insights" (*Blockchain: Early Activity for Utilities; Bitcoin Mining, Blockchain, and Energy Consumption* attached) to provide the public and energy industry with a high-level view of blockchain basics and potential impacts on industry capabilities if blockchain technology were to be adopted. These two documents provided much needed education to counter some hype that often surrounds emerging technologies. In addition to these early education efforts, EPRI's UBIG holds regular webcasts to share experiences and applications of blockchain among supporting members.

Blockchain and Energy Use

In response to questions around blockchain and energy use, EPRI published *Quick Insights: Bitcoin Mining, Blockchain, and Electricity Consumption*. Questions have been raised about data mining operations. These operations often seek out locations with a cool, dry climate, which reduce HVAC expenses and lower energy costs. The concern is that these operations may shutter if cryptocurrency mining becomes less profitable. The price of Bitcoin, one 'cryptocurrency', has seen a decline in recent months from a high of ~\$19,000 in December of 2017 to roughly ~\$6,000 today as shown in the figure below.



Figure 1 Bitcoin price from Nov 1, 2017 to August 16, 2018 Source: Coindesk

Blockchain Architecture

Blockchain is as an append-only file; data can only be added and verified. Once added, data cannot be changed or deleted. The various blockchain architectures of Proof-of-Work, Proof-of-Stake, Proof-of-Authority, and tangle all make different design trade-offs and hence, use different amounts of energy, have different security requirements, and differing performance characteristics. A blockchain company called Ethereum is already experimenting with using Proof-of-Stake¹ for some blocks in its chain to counter scalability and energy use issues.

The three primary characteristics that collectively make blockchain interesting -- security, transparency, and immutability -- don't fit all types of transactions. Security and control requirements vary due to design trade-offs inherent in the technology. It is important to find processes where these three characteristics add value. Some of the early uses EPRI is exploring in addition to transactive energy involve applications that may be subject to audits and safety checklists. For customer-facing applications the main value may be in increasing trust for those customers who would prefer an independent system capturing energy usage.

Transactive Energy

Blockchain is seen as an enabler for Transactive Energy. The challenge facing any transactive energy system is that it must run on disparate devices, through many levels of the grid; consumer, microgrid, feeder, distribution system operator, and transmission system operator to enable transactions. These transactions will be between the customer and the utility, as well as any willing buyer and seller (prosumer). Blockchain could potentially solve this challenge and provide a platform that handles exactly what is described above. Regardless of the type of device, if the market constructs are standardized, then the device would only need to be able to exchange price/energy data with the blockchain being used to enable the market. In the energy sector, nearly all the attention has been on blockchain's enabling capability for transactive energy or eMobility (e.g., payment platform for electric vehicle payments). However, there are regulatory barriers that currently restrict transactions to being between a customer and their local utility and questions about the cost, return on investment, and capability of the devices required to enable this infrastructure. Also, as in traditional smart metering, there are differences in geography and topology that impact the design of the required communication networks. What may be feasible in downtown New York with ubiquitous broadband connectivity, may not work in rural areas that are usually more limited to Power Line Carrier (PLC) or intermittent communication.

To better understand this environment EPRI has created an initial version of a blockchain energy market simulator. This platform was developed as part of the Information Communication Technology Security Architecture for DER research program. This platform will be expanded to simulate many more nodes, loads types, and combined with the EPRI smart inverter simulator,

¹ <https://www.coindesk.com/bitcoins-ta-proot-privacy-tech-is-ready-but-one-things-standing-in-the-way/>

and should provide robust simulation capabilities for loads, generation, and energy from solar panels. This simulation capability, built on an open platform, coupled with the projected deployment cost, may finally give some insight into the total cost of ownership required to enable transactive energy.

Concluding Remarks

Working collaboratively with other stakeholders, EPRI will continue to explore energy efficiency of blockchain and similar technologies as well as their cybersecurity implications.

EPRI is committed to developing science-based solutions to these difficult problems, and offers technical leadership and support to the electricity sector, public policymakers, and other stakeholders to enable safe, reliable, affordable, and environmentally responsible electricity.



RESEARCH QUESTION

What is blockchain and its associated capabilities and applications in the utility industry?

KEY POINTS

- Blockchain is an emerging digital technology acting as a distributed ledger to record transactions.
- The technology removes the need for centralized third-party intermediaries and supports cryptocurrencies that function similar to cash, which are exchanged immediately with no provision for money being returned.
- As the energy internet of things (EIIoT) evolves and connected devices proliferate, blockchain may facilitate payments and other information exchanges among an exponentially increasing volume of customers and service providers.
- The technology is in its early stages of development, with only a couple of utility-related proof-of-concept implementations, though engagement is starting to increase with companies developing the technology for various use-cases.

INTRODUCTION TO BLOCKCHAIN

Blockchain is a "distributed ledger" technology. Like a traditional ledger, it keeps a record of every transaction in a system. Unlike centralized ledgers, it is considered transparent because every participant in the peer-to-peer network has a copy of the ledger and can see the contents of every transaction. Blockchain is currently most closely associated with enabling cryptocurrencies such as Bitcoin, Ethereum, or zCash, but in addition to its uses as a currency, hundreds of use cases are being explored; everything from games to contracts.

Blockchain gets its name because it is a chain of data blocks, each containing a given set of transactions. Additionally, each block contains a mathematical algorithm called a *cryptographic hash* which is based on all the content of all the blocks in the chain to that point including a timestamp based upon the time of creation.



Each block has a "hash" that is based on the contents of all prior blocks, creating a chain.

While there are a number of kinds of cryptographic hashes, they all share the property that it is relatively easy to verify that a particular block of data matches a given cryp-

ographic hash, but that the reverse operation of creating a block of data that matches a particular hash is very difficult. This computationally difficult verification is called a *proof of work* in blockchain parlance. In Bitcoin, one of the more well-known cryptocurrencies based on blockchain technology, the “miners” (the entities that create a new block) are rewarded with bitcoins for performing this task. Other entities also exist in the cryptocurrency ecosystem, such as exchanges, which will exchange the cryptocurrency for something else (dollars, euros, etc.), and wallets, a mechanism that allows one to buy or sell using the cryptocurrency, without the energy or computational overhead of the entities that maintain the blockchain.

There are tradeoffs made with the distributed ledger design choice. For example, with Bitcoin, the blockchain is computationally expensive, and for Bitcoin miners (the entities that create the proof-of-work), only the largest organizations can afford to pay for the energy required to run the computers, whereas smaller entities may pool their resources and share the rewards. The size of the blockchain is also a challenge. It is estimated that for a Visa-scale transactional system (~3000 per second) the blockchain would grow at the rate of approximately 25 terabytes per month.

While a blockchain is inherently secure due to how blocks are created and the use of cryptographic keys, it is not without its challenges. While the chain is secure, the computers and devices that would participate are still vulnerable to hacking. However, to compromise the blockchain, a hostile entity would need to control more than half of the participating devices due to the nature of how participants need to “agree” on each block that is added to the chain; the majority “wins”.

PERMISSIONED VS PERMISSIONLESS BLOCKCHAINS: A MATTER OF TRUST

The Bitcoin implementation of blockchain is permissionless; that is, anyone can choose to participate and decide how much information they wish to reveal about their identity. In this way Bitcoin is “pseudo anonymous”—while identities may be hidden, some companies now offer services that perform analysis on the blockchain in an attempt to reveal participant identities. However, in a permissioned blockchain, entities are only allowed to participate if their identity has been verified.

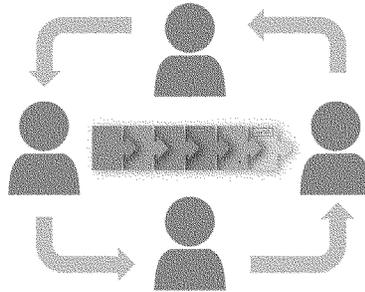
There are cryptocurrencies such as zCash that allow complete anonymity using what is called a “zero knowledge proof”, where a verifier can validate that the user knows something to a third-party without revealing the actual item in question. Zero knowledge proofs are probabilistic proofs rather than deterministic proofs. The cheater has some chance, albeit small, to make the verifier believe they satisfy the proof.

Due to current and potential emerging regulatory requirements, blockchains applied in the utility industry will likely be the permissioned variety.

APPLICATIONS OF BLOCKCHAIN

There are hundreds of use cases for blockchains in various stages of development (a list can be found at <http://slaps.altracasts.com>). These applications run the gamut from smart contracts (insurance, ticket purchasing) and keeping personal data or identity records, to unalterable constitutions or governance—now enforced by algorithm rather than humans.

Nominally, any transaction that the utility participates in that requires an exchange of currency or of paper could use a blockchain. When all parties can verify a transaction without requiring a third party for validation or confirmation to process or hold information, the cost and speed of transactions is improved, potentially reducing costs and benefitting society. EPRI’s investigation of blockchain as well, aligns with its public benefit mission. This may also be the answer to facilitating the



The blockchain is distributed to all participants on the peer-to-peer network

transactions required to enable transactive energy. This concept changes the relationship of customer and utility to one of prosumers. A prosumer can buy or sell, not just to the utility, but to any willing participant.

INDUSTRY ACTIVITY

The Cleantech Forum Conference (January 23-25, 2017) offered a panel discussion on blockchain technology which brought together international representatives from the utility, vendor, and startup communities. The startup companies see, and are pursuing, clear opportunities in the financial industry. They were looking to the utilities to outline potential use cases for blockchain in the energy sector, while the utility attendees were interested in potential applications and impact on their businesses.

There is a clear knowledge gap between technology developers and existing market participants, as well as legal, regulatory, and technical issues which will need to be addressed. However, three use cases were mentioned that may demand closer inspection:

- ◆ Transactive energy to support DER and their interaction with DER management systems (DERMS)
- ◆ eMobility – the ability to transact energy charging at stations in multiple service territories
- ◆ Customer contracts – removing the middleman from the retail energy market

NEXT STEPS/ONGOING EPRI RESEARCH

EPRI will continue to survey new technologies and the marketplace for blockchain-related capabilities and use cases that would present opportunities in the utility industry. EPRI will also engage utility leadership and thought leaders in this emerging industry to provide information and assess potential impact to the energy industry, and to inform related research activities.

CONTACT INFORMATION

For inquiries regarding the technical content of this brief or for general inquiries about EPRI's Quick Insight Briefs, please send an email to QuickInsights@epri.com.

Quick Insights are developed by EPRI to provide insights into strategic energy sector questions. While based on sound expert knowledge, they should be used for general information purposes only. Quick Insights do not represent a position from EPRI.

3002009889

February 2017

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA • 800.313.3774
 • 650.855.2121 • info@epri.com • www.epri.com

© 2017 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER ... SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.



RESEARCH QUESTION

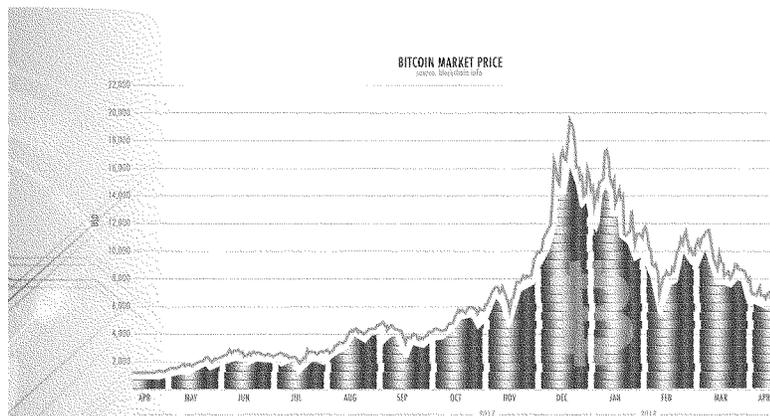
What is the energy consumption of mining cryptocurrencies such as Bitcoin, and how can utilities best interact with these customers?

INTRODUCTION

In 2009, Bitcoin became the first digital currency based on cryptography—creating what has become broadly known as *cryptocurrencies*—to provide a medium of currency exchange without a central authority and without backing by a physical commodity or nation-state. There are currently more than 1,300 similar cryptocurrencies using cryptography to secure transactions, control the creation of new currency, and validate the transfer of value. Cryptocurrencies are backed by blockchain technology, which employs cryptography to validate each transaction and create a permanent public record. Bitcoin mining requires large amounts of electricity, but its inherent volatility, decentralized operations, and uncertain future create challenges for electric utilities engaged in long-term resource planning.

KEY POINTS

- Bitcoin mining requires an estimated 1 to 3 GW of continuous electricity demand—representing less than 0.1% of global electricity generation capacity.
- It is difficult to determine the actual electricity use for mining Bitcoin at any given time because there is no central registry of miners. Similarly, it is virtually impossible to accurately predict future growth because the efficiency of mining equipment is changing rapidly, the difficulty of mining varies, and the revenue paid to miners is highly volatile.
- Given that the values of many cryptocurrencies have recently skyrocketed, any reporting that extrapolates current growth rates to project future electricity demand will likely inflate future predictions of consumption.
- The potential boom-bust nature of cryptocurrency mining and the risk of failure for this emerging industry may present a risk to electric utility cost-recovery or lead to stranded assets.
- Amid rapid Bitcoin mining growth in U.S. regions where electricity is inexpensive, local utilities have grappled with accommodating or banning this type of electricity load.
- The blockchain technology that underpins cryptocurrencies could eventually streamline the management of other transactive processes, but it is too soon to determine its ultimate impact.



WHAT IS BLOCKCHAIN?

Fundamentally, a blockchain is a series of digital blocks, each of which contains a set of transactions. A unique identifier represents the contents of each block and the combined value of all prior blocks in the chain. This linkage of unique identifiers, called a "cryptographic hash," ties the blocks together in the chain. Rather than having a centrally stored and controlled ledger like a traditional accounting system, the blockchain's "ledger" is distributed, with each participant in the peer-to-peer network holding a copy of the "distributed ledger."

Each block of transactions recorded in a blockchain requires a proof of work (PoW) to validate the block and securely append (and timestamp) it to the ledger. This creates a chain of blocks, hence the name blockchain.

A PoW is a cryptographic hash discovered by performing a computationally intense algorithm called *mining*. A hash function is simple to compute given an input value, but the inverse function—i.e. solving for an input given the output—can only be determined through brute-force trial and error. Because a PoW is required to validate each block of transactions that is added to the ledger, mining is necessary to support the use of the currency. In exchange for computing the hash, a miner earns a reward (typically a small amount of the cryptocurrency).

For more information on potential applications for blockchain technology, see EPRI Quick Insight 3002009889 [1] and EPRI white paper 3002010242 [2], which explain how blockchain technology could be applied to other utility transactional business operations.

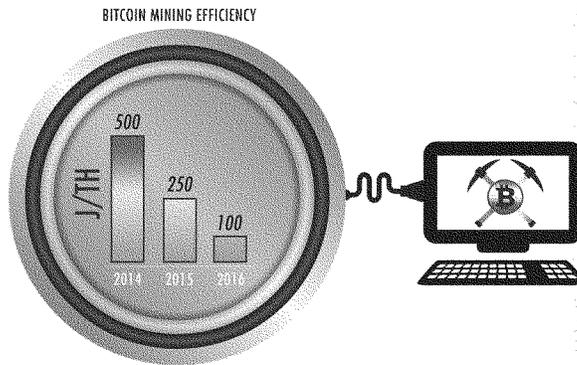
MINING FOR CRYPTOCURRENCY

When Bitcoin was first established, mining was possible using the CPU of a standard desktop PC. As more miners have joined the Bitcoin network, the global *hash rate* (overall number of hash functions that are solved by all miners on the network, now represented by exa hashes per second [EH/s, 10^{18} H/s]) has risen exponentially, to as high as 26 EH/s in March 2018.

The difficulty of the mining algorithm is adjusted roughly every 10 days to maintain an average creation of one block every 10 minutes. With the mining reward set to be halved about every four years, mining will become less profitable over time. As a result of continuously growing resource requirements—particularly the amount of electricity needed for processing and cooling—CPUs are no longer cost-competitive for Bitcoin mining. To improve efficiency, miners initially shifted from CPUs to graphics processing units (GPUs), which offer about an order of magnitude superior mining performance over standard CPUs.

Today's best-in-class mining hardware—which employs an application-specific integrated circuit (ASIC) specially designed for mining Bitcoin—performs two orders of magnitude better than GPUs. Due to advances in chip technology, reported mining efficiencies have roughly doubled every 12 months, from about 500 J/TH [joules per terahash, equivalent to watts per trillion hashes per second] in late 2014, to 250 J/TH in late 2015, to 100 J/TH in mid-2016. Recognizing that one of the largest manufacturers of mining hardware also operates one of the world's largest mining facilities, there may be preproduction mining machines in operation that surpass the commercially available 100 J/TH efficiency level.

Serious Bitcoin mining operations are not expected to reside in conventional data centers because the core business of data centers that house mining operations is to maximize the number of hashes computed for the lowest operating cost. With little concern for equipment availability, mining facilities do not employ the redundancy, fault-tolerance, or power conditioning equipment used in conventional data centers. In addition, many miners use "free cooling," relying on evaporative "swamp" cooling rather than mechanical (vapor compression) cooling. Some mining facilities have no mechanical cooling aside from fans that bring in outdoor air.



INDUSTRY ENERGY CONSUMPTION

Since December 2017, when the value of Bitcoin reached an all-time high of \$20,000 per bitcoin, numerous media outlets have reported on the growing energy consumption of the Bitcoin network. These reports cite the Bitcoin Energy Consumption Index (BECI) published on Digiconomist.net [3] which uses an economic approach to estimate the annual energy consumption of Bitcoin (more than 50 TWh as of March 2018). Note that any estimate of overall energy consumption must make numerous assumptions because there is no central registry of all active Bitcoin mining machines. Moreover, there is neither published data on the efficiency of mining machines in real-world applications, nor data on the number of machines in operation.

However, this widely cited estimate is fundamentally flawed: it assumes that 60% of mining revenue is spent on electricity, without providing a citation. One critic of this approach [5] suggests that the actual percentage of mining revenue spent on electricity may range from 6% to 32%, when accounting for capital recovery. In addition, the author of the BECI estimate presents a case study from an operating mining facility that found that the real-world efficiency of mining machines was less than rated efficiency—a finding attributed to the elevated operating temperature and failure rates seen in the real-world application.

Marc Bevand, a cryptocurrency researcher and entrepreneur, makes a more detailed evaluation of hardware efficiency on his website [4]. This approach took an in-depth look at the evolution of mining hardware efficiency over time, estimating the number of machines added in each hardware generation as a function of the increasing global hash rate. It makes the conservative estimate that only the least-efficient hardware available in each generation was added, so long as it was profitable to operate at \$0.05/kWh. On January 11, 2018, Bevand updated his estimate of the global Bitcoin mining network to be 2.1 GW of demand (upper and lower bounds of 1.6 and 3.1 GW) and 18 TWh of annual consumption (bounded by 14 and 27 TWh).

The results of this estimate and others suggest that Bitcoin mining worldwide is on the order of 2 to 3 GW. With global installed generating capacity totaling more than 6,200 GW as of 2015 [6], Bitcoin mining represents less than 0.1% of world generating capacity. In 2014, the annual energy consumption of data centers worldwide was estimated to be 194 TWh [7], roughly 10 times the annual consumption of Bitcoin mining estimated by Bevand [4].

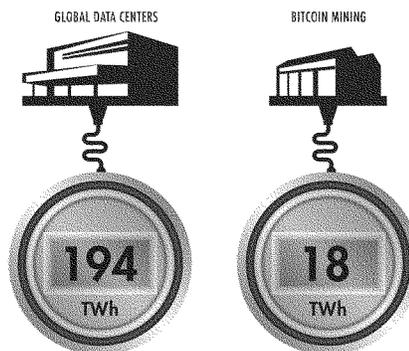
There may be valid concerns as to how the energy intensity of cryptocurrencies would scale if they were to handle the number of transactions supported by credit cards each year (roughly two orders of magnitude more). Any reporting that extrapolates current growth rates—when the value of the currency has recently skyrocketed—will likely inflate future predictions of energy consumption. It is virtually impossible to make an accurate prediction of growth in mining demand due to the number of unknown variables, including:

- ◆ Mining difficulty depends on the number of miners connected;
- ◆ The real-world efficiency of mining hardware is unknown, and reported hardware efficiency has improved at an unprecedented rate;
- ◆ The revenue paid to miners is highly unpredictable because it is determined by the value of the cryptocurrency, which has been highly volatile.

EVOLVING INDUSTRY

While it is impossible to predict the rate of future mining growth, there is evidence that this industry may continue to expand in the near future. First, demand for mining equipment has created a scarcity of best-in-class hardware, and prices have risen significantly. In addition, anecdotal reports from mining operations and utilities indicate that it is growing rapidly in certain parts of the United States.

Previously, the majority of commercial Bitcoin mining operations have been located in China—specifically, Inner Mongolia—due to its cool, dry climate and reportedly low cost of electricity. In January 2018, the Chinese government began to curb mining, in part due to its impact on demand for electricity. Since then, there have been several reports on mining operations seeking to make significant expansions in areas of North America where electricity and land prices are modest and the climate favors free cooling. For example, mining operations in Wenatchee, WA have been able to take advantage of low-cost electricity and unused distribution capacity left by shuttered industries (namely aluminum and logging). On the other hand, Plattsburgh, NY has banned any additional mining operations from locating there for 18 months due to the impact that these facilities may have on local electricity prices.



Even with low-cost electricity, the volatility of mining revenues may drive some mining companies to cease operations before electric utilities fully recover the costs of delivering service. With so much uncertainty in the longevity of this market, utilities may best consider these customers cautiously. Yet given the high load factor of these facilities, some utilities might consider them very attractive customers, if only for a limited time.

One aspect largely missed in the discussion of blockchain efficiency is the potential for the technology to make other transaction processes more efficient. Because blockchain eliminates the need for a central database manager or independent transaction validator, it could streamline transaction management in areas other than cryptocurrency. If used for transactional databases in other industries (e.g. the insurance, medical, real estate, and banking sectors), Blockchain technology has the potential to offer global (societal) efficiency gains (i.e. digitizing and streamlining the management of verified transactions) while increasing electricity use to validate the transactions. However, it is too early in the maturity of this technology and its deployment to predict the potential efficiency gains. This is an area of continuing research that EPRI is conducting under its Information and Communications Technology for Integration of Distributed Energy Resources program [8].

RESEARCH GAPS

- ◆ Can the actual energy consumption of cryptocurrency mining be more accurately estimated? What is the real-world efficiency of deployed mining machines?
- ◆ What is the risk to an electric utility of stranded assets or failure to recover costs?
- ◆ Can a mining operation follow time-of-use rates and only be active during periods of low electricity prices?
- ◆ Are there methods for making cryptocurrencies more efficient while maintaining security and validity?
- ◆ Can blockchain technology offer global (societal) efficiency gains (i.e. digitizing and streamlining the management of verified transactions)?
- ◆ What electric utility or other industry processes are best suited to blockchain technology?

REFERENCES

1. *Quick Insights – Blockchain: Early Activity for Utilities*, Electric Power Research Institute, Palo Alto, CA. Product ID: 3002009889, February 2017. <https://www.epri.com/#/pages/product/3002009889/>
2. *Blockchain: Technology Risk and Rewards for Utilities*, Electric Power Research Institute, Palo Alto, CA. Product ID: 3002010242, October 2017. <https://www.epri.com/#/pages/product/3002010242/>
3. A. de Vries, "Bitcoin Energy Consumption Index." [Online]. (Accessed: 15 March 2018). <https://digiconomist.net/bitcoin-energy-consumption>
4. M. Bevand, (2017, 10 March). "Electricity consumption of Bitcoin: a market-based and technical analysis," [Online] <http://blog.zarinaq.com/bitcoin-electricity-consumption/>
5. M. Bevand, (2017, 1 February). "Serious faults in Digiconomist's Bitcoin Energy Consumption Index," [Online] <http://blog.zarinaq.com/serious-faults-in-beci/>
6. U.S. Energy Information Agency, <http://www.eia.gov>
7. *Digitalization & Energy*, International Energy Agency, Paris, France, 2017.
8. *Information and Communications Technology and Security Architecture for Distributed Energy Resources Integration*, Electric Power Research Institute, Palo Alto, CA. Product ID: 3002009694, January 2017. <https://www.epri.com/#/pages/product/00000003002009694/>

CONTACT INFORMATION

For inquiries regarding the technical content of this brief or for general inquiries about EPRI's Quick Insight Briefs, please send an email to QuickInsights@epri.com.

Quick Insights are developed by EPRI to provide insights into strategic energy sector questions. While based on sound expert knowledge, they should be used for general information purposes only. Quick Insights do not represent a position from EPRI.

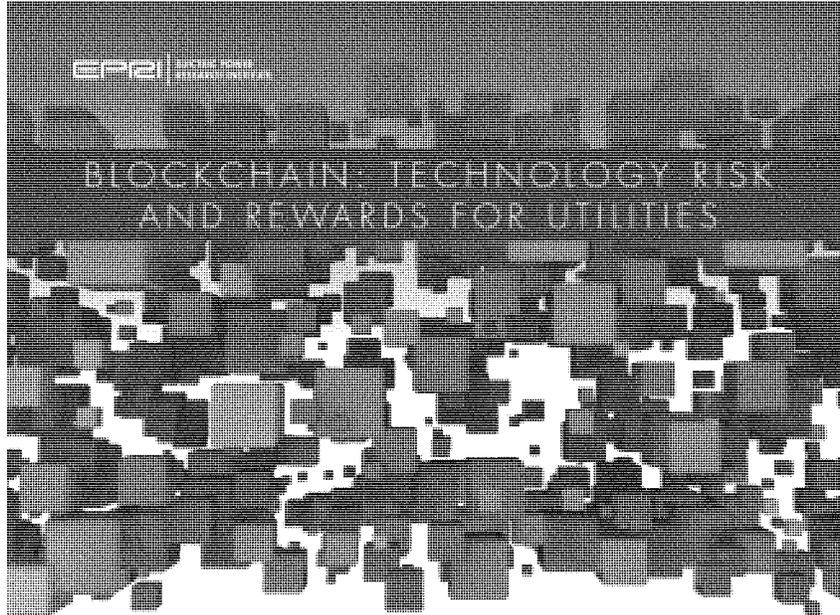
3002013910

April 2018

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA • 800.313.3774
• 650.855.2121 • tsa@epri.com • www.epri.com

© 2018 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER... SHARING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

**Abstract**

Blockchain is a potentially disruptive technology that will impact the way in which many business transactions are conducted in the future, including those used by the utility industry and its trading partners. While it is most commonly known as the technology behind cryptocurrencies such as Bitcoin, the greater impact will likely be with the implementation and automation of “smart” contracts that reduce costs by eliminating intermediaries. In this white paper, the characteristics of blockchain will be explored, providing insight into why this is a disruptive technology, the places blockchain is being used today, some of the potential applications of this technology in the utility industry, and the current challenges and limitations of the technology of which utilities need to be aware.

WHAT IS BLOCKCHAIN?

Fundamentally, a *blockchain* is simply a chain of blocks (hence the name), with each block containing a set of transactions. Within each block, a unique identifier is generated that represents the contents of that block, and additionally, the value of all the prior blocks in the chain. This linkage of unique identifiers, called a “cryptographic hash”, is what ties the blocks together in the chain. Additionally, rather than being centrally located like a traditional accounting system, with a ledger stored and controlled in a central database, the blockchain’s “ledger” is distributed, with each participant in the peer-to-peer network holding a copy, hence the term “distributed ledger”.

the community, such as different miners running different versions of the code. This protects the community by helping to ensure that each node in the blockchain P2P network runs the same version of the software. This prevents challenges such as differences in determining how consensus is reached and, hence, which are the valid blocks in the chain.

Large miners / Pool operators – In the Bitcoin ecosystem, miners are rewarded for generating a key and establishing a proof-of-work with bitcoins. In the early days of Bitcoin, a person with a single desktop computer had a reasonable chance to get such a reward. However, as the number of participants has increased, the reward for generating the proof-of-work has been outstripped by the energy costs to run the computer system. Only the largest miners, with the latest technology, optimized for mining and energy costs, can now reasonably expect to be rewarded with bitcoins. An individual operator has an increasingly smaller chance of getting such a reward. Hence, the evolution of pool operators. A smaller miner can join this pool, increasing the chance of a reward being earned by this group, and the rewards are spread across all the participants in the pool.

Users / Wallet providers – Users or customers, create new transaction requests, for example, maybe they would like to buy a hat. Software, called a wallet, passes requests to the P2P network, and these transactions are then packaged into blocks. A wallet allows a person or entity to generate a transaction without having to do the mining.

Payment processors – This software allows organizations the ability to offer payment services in Bitcoin, but then pay their clients in non-digital currencies.

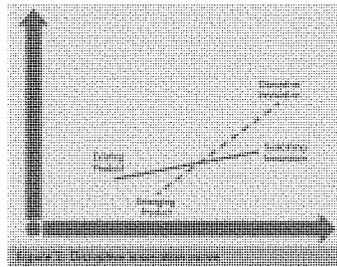
Exchanges – Just like traditional currency exchanges, this allows bitcoin to be exchanged with other currencies.

BLOCKCHAIN AND THE DISRUPTION CURVE

The disruption curve [Figure 3 below] illustrates how new technology can destroy incumbent players in established markets. Based on the work of Clayton Christensen, who discovered the phenomenon while investigating architectural changes in the hard drive industry and found that the curve repeated itself across many industries and technologies. Christensen makes the distinction between “sustaining innovation” versus “disruptive innovation”. This finding dispelled the notion that perhaps companies were not listening to their customers or investing in their product lines. Quite the contrary. They often did, but the investment

in existing products and listening to customers led to only changes in existing technology.

Disruptive technology is often not recognized when it emerges. It often underperforms existing technology. However, disruptive technology presents certain attributes that new customers prefer, and in fact, disruptive technology often must identify new customer to be successful. When the new technology matures, and begins to outperform the incumbent technology, then it begins stealing customers from the incumbent (Figure 3 below).



It remains to be seen if blockchain’s distributed ledger technology appears to be just such a disruptive technology.

Blockchain emerged in response to the financial crisis in the hopes of increasing transparency and to create a currency with specific attributes such as immutability, and limited supply (much like gold). At the time, these attributes were attractive to a small audience. As Bitcoin demonstrated the capability of the technology and became more accepted, innovators began exploring how distributed ledger technology could be employed in other types of transactions. While there are challenges with Bitcoin itself when it comes to issues such as scalability, the technology is being applied to a host of “smart contracts” and distributed applications. While any new technology must answer the question of, “How does this technology do something better than the existing technology?” to justify investment, it appears that it may be applicable anywhere contracts or exchanges are used today, with the added benefit of eliminating third parties to the transaction, which lowers costs [4].

The attractive advantages for new customers are reduced transaction costs, anonymity (to a greater or lesser degree), and immutability of the transaction. In this regard blockchain has been creating a new set of customers; however,

the destruction of incumbents, the key hallmark of disruptive technology, has not been observed yet.

Blockchain and Distributed Applications

While blockchain is more well known as the technology behind cryptocurrencies, its greater potential is in the use of distributed applications, or in blockchain vernacular, “dapps”. There are hundreds of dapps that cover everything from games, to insurance, to unalterable constitutions. Anything that can be contracted (essentially, any exchange), can be codified in a distributed ledger. The value-add of dapps is that they leverage the transparent quality of distributed ledger and payment is typically immediate. This immediacy of payment, which functions essentially like cash, is one of the other benefits of blockchain.

BLOCKCHAIN CHALLENGES

Like any technology, there are tradeoffs intentionally made in the design to favor certain characteristics over others. Looking at the Bitcoin implementation, there have been challenges with scalability, anonymity, and hacking [4][5] [6], albeit with some interesting twists.

Scalability – The Bitcoin blockchain is now 149 GB³. This is easily handled by computers of the scale that are typical in utility data centers. However, it is estimated that a “Visa scale” global network (~2000 transaction per second) would grow at a rate of 2.5 terabytes (TB) per month [7], which requires planning to accommodate such a file, and would, of course, outpace the ability of utility grid edge devices to be a peer on the Bitcoin blockchain network (although edge devices could run a wallet and transact with a peer). As it related to the Internet of Things (IoT) there are questions about the ability of cryptocurrencies to support the micro transactions that these devices might employ as they exchange services.

Anonymity – As discussed in the “Getting Technical” section, blockchains can be permissionless (anyone can join) or permissioned (only authorized entities may join). It is often assumed that Bitcoin is anonymous, when it is technically *pseudo-anonymous*. Parties to transactions do not have to give identifying information beyond their public key, but there are methods available to determine a given party’s identity (for example, some companies offer services wherein they examine the Bitcoin transaction to deduce a party’s identity). Alternatives to Bitcoin, such as

zCash [8] and Monero⁴, promise complete privacy for those engaging in transactions. There is an interesting dichotomy at play when Bitcoin’s acceptance has been reflected in increasing regulation. [9][10][11] Blockchains that are completely anonymous not only disrupt markets, but disrupt regulation as well (if you’re using it, no one can tell). Also, when the particulars of a transaction are stored or controlled by a “smart contract” [12], this requires less need for oversight because the “oversight” is encoded in the contract to which the parties agree.

Security and Control – Those that control the nodes, control the contents of the blockchain. This goes back to how consensus is reached on the blockchain. The nodes “vote” on the longest chain. If a single entity controls more than 50% of the nodes, then that entity can determine which block “wins”. This has implications for both permissionless and permissioned blockchains, where a consortium controls the nodes; participants need to be mindful of the 50% rule. This also has implications for claims that blockchain will solve IoT security challenges. This is because the devices themselves are still vulnerable. If more than 50% of the IoT devices are compromised, then the blockchain they transact on can also be compromised.

An IoT consortium, led by CISCO, Bosch, and others, is working to leverage blockchain to “secure and improve” IoT applications [13]. However, while the blockchain itself is secure due to the nature of the security employed to encrypt the transactions, this does not inherently secure the IoT devices themselves. For example, traditional device issues, such as not changing the manufacturers default password or inappropriately applying vulnerability patches, could allow bad actors to launch a distributed denial of service attack (DDoS) from these vulnerable devices [14]. If the devices are compromised, then these devices could flood an IoT-based blockchain with bogus transactions.

While blockchain technology has some technology challenges that need to be addressed as it matures, it still bears the hallmarks of a disruptive technology.

The Ethereum Hard Fork – One demonstration that blockchain technology is still maturing, was the “hard fork” (a change in the software that runs Ethereum) that was used to restore stolen funds [15]. This change to the software was required to return roughly \$40 million worth of ether that had been stolen from an account owned by an unknown

3. <https://kalinicharts.com>

4. <https://getmonero.org/learn>

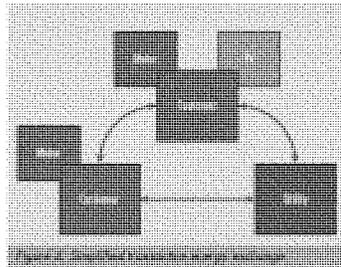
hacker. While the hack was remediated, the act of remediation raises questions about the supposedly immutable blockchain and the finality of any contracts based on said blockchain.

APPLICATIONS IN THE UTILITY INDUSTRY

Outside the exchange of Bitcoins used with financial applications, there have been few implementations of blockchain in the utility industry. In terms of simply enabling financial transactions via electronic data interchange (EDI), automated clearing house (ACH), or other means, Bitcoin is simply another currency that could be enabled to be used for payment just as dollars or Euros are used. However, there are several other applications in the utility industry that could also implement blockchain.

Transactive Energy

Transactive energy is a concept that refers to the "economic and control techniques used to manage the flow or exchange of energy within a power system"⁵. The example of transactive energy exchange is a bit more complicated because transactions are not limited to being between a utility and their direct customers. Customers can also sell to their neighbor – the "prosumer" concept, which is the notion that a customer might be both a buyer and seller of energy.



For example, in Figure 4, suppose the customer with the PV generates more energy than they need and offers to sell it to the other utility customer. In this future looking scenario, a meter still provides the measurement and verification (one wants confirmation that what one is paying for has been produced). The utility, in this scenario, is the distribu-

tion system operator and provides the wires for the energy exchange to take place, and one or both parties would have a transaction fee to the utility that pays maintenance overhead on this operation, in addition to the transaction fee that would be associated with the blockchain in use. Again, blockchain computational requirements exceed the capability of residential meters, so another platform would need to be provided to manage the peer-to-peer network, perhaps an additional meter behind the utility meter as we see in the LO3 configuration of the Brooklyn microgrid, or a device that can run a blockchain wallet, or there may be a distributed peer on the circuit run by the utility, e.g. a distributed DMS or DERMS that in addition to providing control function could also provide the mechanism for the blockchain transaction. Contrast this with how transactive energy is provided in the United Kingdom.

P2P Energy in the UK

The United Kingdom already supports a P2P energy trading program, albeit a traditional one. Good Energy⁶ allows customers to sign up to purchase power from renewable energy suppliers that come in two categories: a 10 kW – 100 kW provider with a pay in tariff (PIT) tariff, or for >100 kW suppliers, power purchase agreements (fixed, variable, and "cost for difference"). Buyers can choose from whom they wish to purchase and sellers get a predictable income.

Customers who sign up to the service are given access to an online portal where they can set preferences and priorities for their energy supply at certain points throughout the day. If a generator is available, the two parties are matched and the business will effectively pay that generator for the electricity it consumes. [1]

The difference is that Good Energy appears to be the normal third party handling the arrangements. In a blockchain-based P2P, once contract terms are met, buyers and sellers trade; there is no third party setting the market.

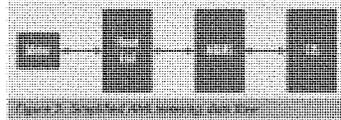
Metering

While blockchain could be used to secure transactions, there are some facets of smart metering today that do not lend itself to this application. Metering actors include the meter, the AMI Head-End, optionally a Meter Data Management System (MDMS) that may serve as a data ware-

5. http://www.epri.com/epri/products/whitepapers/transaction_energy.aspx

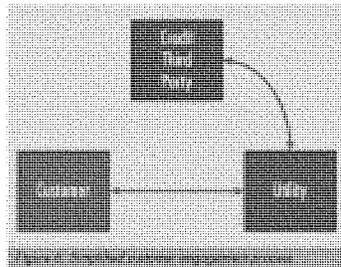
6. <https://www.goodenergy.co.uk/business/our-operations/energy-purchase-agreements.aspx/>

house for metering data, and the Customer Information System (CIS) which ties meters to customers and generates the bill. In traditional metering, there does not appear to be much place for blockchain to provide a "better" story.



The data flow for AMI metering is already secured, so there is not a value proposition for blockchain relative to that. In this case, the data exchange is only between the utility customer and the utility. There is no third party involved. Also, blockchain as currently available outstrips the ability of a standard residential meter's computational capability.

Move-in / Move-out / New Service / Prepaid metering



Where there might be a place for blockchain in this scenario is for the move-in/move-out process or to setup pre-paid metering. This is not the metering information path, but rather the "out of band" communications that goes on between a potential customer, the utility, and third parties that may confirm the creditworthiness of a customer. In this scenario, the customer indicates a desire to become a utility customer. The utility uses a third party to confirm the creditworthiness of a customer and may require the customer to provide a deposit before engaging in a service contract.

Now consider using blockchain in this scenario. The tokens used by blockchain such as Bitcoin function like cash. If the customer wants to engage in business, blockchain assures the merchant, in this case, the utility, that the cash (tokens)

are on hand before entering into a contract. The credit third party can be eliminated (saving the utility money) and speeding the execution of the transaction. The customer and utility agree to pay via a blockchain currency, with the costs limited to the transaction fee of the blockchain in use. If the utility and customer use a permissioned blockchain, the transaction has the added benefit that the customer identity is verified before the parties enter into the transaction.

Mobile Payments

Even as early as some of the first requirements gathering efforts were occurring for Home Area Networking (HAN), vehicle charging and the payment for those services had been considered as part of that development. If a customer has an electric or plug-in hybrid electric vehicle, this process is straightforward. One adds electric vehicle supply equipment (EVSE) to their premise, contacts their utility if there is an applicable tariff or program, and they can charge their vehicle.

The problem gets a bit more complex if the vehicle owner drive the vehicle to a different location and charges their vehicle there. If the location is in the same utility service territory, the driver simply needs to pay for the service locally, and if they wish to get credit within the utility program, then they need to identify themselves at that location.

It gets more complex when the electric vehicle user crosses territory. Potentially this customer would need to create an identity with every utility with which it desired to charge from. In the early days of HAN requirements, it was supposed that a national clearinghouse would emerge much like a VISA or American Express that would handle these transactions. But these clearinghouses operate on a per-

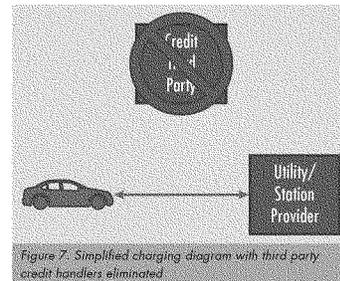
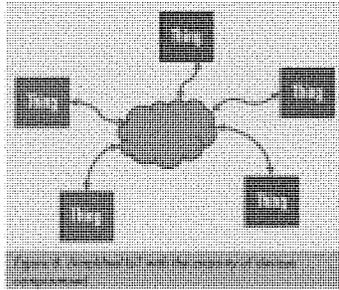


Figure 7. Simplified charging diagram with third party credit handlers eliminated.

centage of the transaction, 3% and 5% respectively. Given the low transaction amount for vehicle charging this does not provide much of an incentive to enter this market.

Now imagine that one is not only crossing utility service territory, but national boundaries as well. This is the situation with electric vehicles in Europe and the Ethon BloT charging stations. The Ethon BloT stations accept cryptocurrencies that can be paid via a blockchain wallet so the fees associated with the transaction are very low and no registration with a local charging provider is required because the identities and transactions are managed by the blockchain, and if the underlying infrastructure is used, it does not matter what local entity (utility, government, private) operates the charging station. For a similar case to work without blockchain, the existing charging station providers would need to agree to a common messaging infrastructure between all of the potential transaction participants, secure it, and then provide a mechanism to ensure transparency (or be audited), all of which would increase the operating costs. While the market is small today, the value proposition will increase as the numbers of electric vehicles and charging stations increase.



Internet of Things

One of the promised benefits of blockchain is that it will make IoT more secure. But again, there are two problems: 1) the computational capability of the device and 2) blockchain might secure the message, but it does not inherently secure the device. The promise of IoT as it relates to utilities (sometimes referred to as energy IoT or "eIoT") is that as more devices are deployed into the grid it gives utilities greater visibility and control.

When the home security cameras were compromised, the issue was not with the messages per se, it is that the devices were compromised in the first place, allowing the devices to participate in the distributed denial of service attack. With IoT, the messages might be secured and properly formatted, but if the devices themselves are compromised, if an entity gets control of the majority of devices, then that entity will determine the content of the blockchain. When a device uses a low-power protocol such as CoAP that is designed to run in as little as 10 KB RAM, that suggests it may not have the processing power to support blockchain, even if it is only the wallet (allowing the device to transact) and not a peer. Again, there are various strategies in place today to secure IoT messages (TLS/SLS, DTLS, depending on the messaging protocol), but the challenge is securing the devices. It remains to be seen if the design compromises required to get an IoT device to support blockchain would be more or less secure than the security protocols currently supported for the messages; device security remains a separate issue.

Asset Management

The story of the "lost" transformer is an issue as old as transformers themselves. (Transformer in this case is a surrogate for any asset deployed in the field). The utility orders a transformer; it arrives and is scanned into the system. When it is needed for a job, the transformer is scanned, removed from inventory, and loaded onto the truck. The transformer is deployed in the field, with its location theoretically noted by the crew. However, in the past days of paper-based paperwork, it was not uncommon for the paperwork to be lost or incomplete. As systems have been automated and barcodes or RFID employed to tag assets with tablets, laptops, and scanners employed to automate the transfer of custody, the barriers to adopting these leading practices have been reduced, though perhaps not eliminated, as the process still relies on a human to perform these functions.

The benefit that blockchain brings to this situation is that the record is immutable. It cannot be changed or forged (assuming some nefarious intent and not just "I forgot to scan the asset"). Although an asset still relies on a human actor for transfer, one thing will be known for certain, who the last person was to "touch" it.

“Smart” Contracts: Appliance Service Plan example

Another area where a smart contract might be employed is with Appliance Service Plans. Often utilities offer this service themselves or via a contractor. Customers typically sign up via a web site and the work is contracted to annually inspect and service equipment, and in the case of a failure, replace or repair the equipment that is included in the plan. One of the benefits of a smart contract is that it is executed once the terms of the contract are met. Again, assuming a known entity via a permissioned blockchain and tokens that are treated as cash (no need for a credit check or if funds are available), the contract could be executed immediately. Compare this experience to banking or real estate (or utilities) where signed paperwork still needs to be faxed before an agreement can be executed⁷.

European Utilities Get Involved

Other utilities in Europe investigating blockchain technology include: ENEL, ECN, EDF, Innogy, CEZ, Fortum, Vattenfall, Iberdrola, EDP, BSE, Eandis, ACEA and Allionder, as well as the British TSO National Grid and Eurelectric [21] that participated in a two-day workshop in Amsterdam. ENEL indicated that they see a need to review architecture impacts and the primary area of focus for them is low-medium voltage grid management, trading on the energy and commodities markets, and renewable energy, for facilitating payments within microgrids[25].

LOOKING FORWARD

Bitcoin has moved from nascent technology to being accepted, albeit with at times dramatic fluctuations, as a currency that is used on world markets. But the disruption associated with blockchain distributed ledger technology is not with the emergence of this cryptocurrency, but with dapps, the distributed applications that will replace how any exchange based on a contract (buyer, seller, consideration, and terms) are executed. Realizing this, several consortiums have formed to address the financial sector [22], IoT, Logistics, and a coalition in the energy sector called the “Energy Web Foundation”⁸. EPRI will also be exploring how blockchain may impact different facets of

the utility business, with emphasis on cyber security in the supply chain within the Information Communication and Cyber Security Program 183, and the *Information and Communications Technology and Security Architecture for Distributed Energy Resources Integration*.⁹

We have only touched upon a small sample of use cases here. While these show some promise, the “killer app” for distributed ledgers has yet to be identified. For those looking for potential applications of distributed ledger technology in the utility industry, there are two sources one might start with, either the American Productivity and Quality Center (APQC)¹⁰ Process Classification Framework (PCF) utility specific model, or the EPRI Business Architecture Service Repository¹¹. Both list hundreds of business processes that exist within utilities. These lists of processes and services could be reviewed for how distributed ledger technology either makes the process better (using the unique distributed ledger characteristics of immutability and transparency), or faster, by eliminating intermediaries.

In addition to utilities, it seems all the well-known technology companies, IBM, Google, Microsoft, in addition to a plethora of startups are emerging, all to create new blockchain applications. EnergyBiz cited a Navigant estimate that the spending to support these related technologies “will be about \$182.6 million in 2016, growing to around \$2.1 billion by 2025” [23].

Keep an eye on incumbent players. For example, while big banks are investing in blockchain applications because of the reduction in transaction costs, there will be emerging companies trying to displace banks for the very same reason.

Similar issues revolve around regulation. There are emerging attempts to regulate cryptocurrencies, but a permissionless blockchain where all the participants are completely anonymous has the ability to sidestep regulatory frameworks. Utilities will not want to be caught up in such schemes, but the potential disruption to regulatory agencies and their attempts to deal with this will bear watching.

As with any emerging technology, utilities should continue to monitor the capabilities as they evolve, but be wary of

7. Stokor, I. (2016). Good Energy formally launches peer-to-peer renewables trading service Solecity. Available [Online]: <http://www.cleanenergypartners.co.uk/news/stokor-joined-energy-formally-launches-peer-to-peer-renewables-trading-service-5133>

8. www.energyweb.org

9. <http://www.epri.com/#/pages/product/0000000030002070624/>

10. <http://www.apqc.org/pcf>

11. <http://www.epri.com/#/pages/product/0000000030002011034/>

start-ups that are not familiar with the utility business. Expect to see smaller companies go out of business, be acquired, or exit the space where they cannot compete. As the many pilots begin to emerge, it will also be important to be wary of products that might still be in beta, or otherwise not ready for full production mode, or companies that may outgrow their ability to support the customers they do gain. Gartner has blockchain at or near the peak of the hype cycle [24]. Be prepared for a shaking out period as blockchain enters what Gartner refers to as the "trough of disillusionment". Cryptocurrencies such as Bitcoin and Ethereum seem to have established themselves, but other applications of the technology have yet to do so and no "killer app" outside of these cryptocurrencies has really emerged.

Utilities should also be wary of pilots or applications that do not have a clear story of how blockchain technology improves or eliminates redundant processes, improves speed of delivery, or provides some other operational benefit. History has shown that in the IoT space, there have been overstated claims of how simply using blockchain solves the security issues related to the devices themselves, so utilities will need to dig into benefit claims and ask for demonstrated capability of blockchain related products and not chase blockchain because it is the latest silver bullet to hit the industry. There will be cases where blockchain provides a clear benefit, but these "can't miss" use cases are still in the process of being determined.

For more information

Please contact Dr. Gerald R. Gray, ggray@epri.com,
+1.865.218.8113

Appendix

GETTING TECHNICAL: THE DISTRIBUTED LEDGER

Blockchain is more than the technology behind cryptocurrencies such as Bitcoin. Due to the nature of distributed ledger technology that blockchain supports, the capabilities of smart contracts is likely to be the most disruptive feature of the technology. But first, readers will need to get technical for a moment while some of the pieces that make this technology work are explored to have a better understanding of the potential impacts. Also, when vendors come calling, it is important to be able to understand the basics of the technology so as to be able to validate any vendor claims about their products capabilities.

Public Key Cryptography

An important mechanism employed in blockchain is *public key cryptography* which is also known as *asymmetric key cryptography* to contrast it with *symmetric key cryptography*. With any symmetric encryption mechanism, the two fundamental operations are *encryption* and *decryption*. With encryption, an input message is rendered, often called the *plaintext*, into an equivalent message called the *ciphertext* which, ideally, is completely unintelligible. In symmetric key cryptography, which includes the well-known and widely used Advanced Encryption Standard (AES), the same key is used for both encryption and decryption. It is important to note that the security of such encryption is solely dependent on keeping the key secret. The fact that the algorithm is well known does not adversely affect the security of messages encrypted with such algorithms.

By contrast, public key cryptography employs key pairs. The pair of keys are related mathematically by one of a few so-called *trapdoor* or *one-way functions* in which is relatively computationally easy to compute a function, but for which the inverse function has no such algorithm. Of these key pairs, one must be kept secret and is therefore called the *private key*, but the other related key may be freely shared without compromising security and is therefore called the *public key*. The only difference between the keys is that one is kept secret. Mathematically there is no distinguishing characteristic to determine which is which, so for a given key pair, which one is made public and which one is kept private is an arbitrary choice. This feature allows some interesting uses that are much easier to accomplish with public key cryptography than with symmetric key cryptography.

Digital Signature

One such use is the *digital signature*. A digital signature is a way of associating a message with a private/public key pair without revealing the private key. The usual way that a digital signature is created for some message is to follow these steps:

1. Calculate the cryptographic hash of the message, called a message digest
2. Use the private key to generate a digital signature over the message
3. Transmit both the encrypted hash and the message

Using this, a receiver who knows the associated public key can verify the message by these steps:

1. Calculate the message digest of the message
2. Perform signature verification on the digital signature, resulting in a message digest
3. If the two message digests match, accept the signature.

There are two reasons that a cryptographic hash is used here. First, asymmetric cryptography tends to be much more computationally intensive (and therefore slower) than the correspondingly secure symmetric key encryption, so it is faster to encrypt a small digest than a large message. Second, the use of a hash enhances the security of the scheme for reasons not explained here.

Using these mechanisms provides the features of *authentication*, *integrity* and *non-repudiation*, but not *confidentiality*. Simply explained:

- Confidentiality means that only authorized parties may read the message
- Integrity means that one can verify that the message has not been altered either accidentally or maliciously
- Authentication, in this context, means that if the digital signature of a message is verified, the message was created by the entity possessing the corresponding private key
- Non-repudiation means that an entity cannot plausibly deny creating a verified digitally signed message

Note again, that these features require that the private key remains secret.

The Cryptographic Hash

A cryptographic hash is a mathematical algorithm that reduces an arbitrarily-sized block of data called the message into a shorter, fixed-size sequence of bits that is often called

a digest. There are several kinds of cryptographic hashes, but they all share the property that it is relatively easy to verify that a block of data matches a given digest, but that the reverse operation of creating a block of data that matches a digest is very difficult. This property allows for a simple method of verifying that the contents of a message match a given digest. One commonly used cryptographic hash function is called SHA-256. SHA256 is endorsed and used by the US Government and is standardized: FIPS180-3 Secure Hash Standard. It should be noted that the NSA plans to retire current cryptography standards and already recommends using at least SHA-384¹². Ensuring the cryptography standards stay ahead of hacker capabilities is an important facet of ensuring financial transaction security. The cryptographic hash is what ties the blocks in a blockchain together.

Wallets

A wallet is software that allows a user to transact with the blockchain P2P network, without the requirement of being a peer or doing a mining activity. When it is created, there is a seed function that will create an account and manage the secure keys that a user needs to transact with a blockchain. Additionally, the wallet is encrypted upon whatever device runs the software and the user's password for the wallet is used to both encrypt and decrypt the contents¹³. As noted in the "Transaction Basics" section, when a customer wants to buy a hat, it is the wallet software that connects with merchant system and a peer on the P2P network, authorizing the transaction (via the user) and sends the appropriate number of tokens for the transaction.

Smart Contracts

A smart contract is not the same as a legal contract, although the same parameters one would use in a legal contract can also be used in a smart contract. A smart contract is simply a digital conditional trigger that is configured and embedded in the blockchain coding, that then executes when the conditions are met. In fact, a smart contract does not require the use of a blockchain, it is that simply a smart contract, in conjunction with the blockchain, takes on the attributes that make the blockchain attractive in the first place: security and transparency.

DESIGNING A BLOCKCHAIN

If an organization did not want to use an existing blockchain but rather design a blockchain for a specific purpose or industry, there are several choices that need to be made. Each choice is briefly described below.

Who can access the network?

In cryptocurrency systems, such as Bitcoin, literally anyone can participate on the network (*permissionless*) [3]. All nodes can see the entire contents of the blockchain (which functions as a distributed ledger) and can participate in creating and verifying new transactions. However, there are other possible models. For instance, it is possible to create a private blockchain system in which only certain qualified nodes may participate (*permissioned*). It is also possible for different nodes to have different roles as further described below.

How are tokens created?

In some systems, including Bitcoin, tokens are created via a special "genesis block" (the first block of the blockchain) and then further tokens are generated during the process of mining. With Bitcoin, miners are incentivized to mine by being rewarded with Bitcoin. Eventually all Bitcoins will be mined and no further ones can be created (the design of Bitcoin is such that the limit of generated tokens will be ~21 million) although fractions of Bitcoins will continue to be used in transactions.

An alternative scheme is to give some nodes the ability to create digital tokens that are then used by the system. A genesis block is still required, but need not actually contain any tokens. Similarly, there may be a mechanism for destroying or "retiring" tokens.

How are transactions validated?

In the case of Bitcoin, a message is at least potentially valid if the digital signature of a transaction is verified. To prevent "double spending" of Bitcoin, the network must arrive at a new consensus view before the transaction is verified.

Alternatives could include matching the public key to an authorized list of nodes who can initiate that kind of transaction, though in this case, a mechanism would need to provide for a means to verify the identity of such a node.

12. <https://blog.cloudflare.com/ps-1512/>

13. <https://blockchain.info/wallet/how-it-works>

How does the network arrive at consensus?

Within blockchain, the mechanism for the nodes to agree on a valid block is referred to as consensus. There are several ways that the network might arrive at consensus. Bitcoin uses a "proof-of-work" (PoW) scheme that involves solving a mathematical puzzle involving a cryptographic hash of the block content. Because it is computationally expensive to solve this puzzle, the first mining entity to solve the puzzle presents it to the network and other nodes may easily verify it. In the case that two nodes solve the puzzle nearly simultaneously, there are temporarily two versions of the blockchain on the network, but this is generally resolved when the first solution to the next block is propagated; the network only accepts the longest chain as the correct one.

Another scheme is called Proof-of-Stake (PoS) which is less computationally expensive and uses significantly less energy. In this mechanism, nodes which have a larger balance have a higher probability of creating the next block. The cost to mine is significantly lower, and transaction throughput is increased and energy use is reduced versus PoW schemes. Lower mining costs, however, benefit both legitimate and malicious entities so the rate of introduction of fake blocks could be higher. There is also the potential that the entity with the largest stake might also be malicious. Various alternatives have been introduced to attempt to mitigate this risk including randomization and delegated PoS (DPoS) systems.

A class of algorithms generally called Byzantine Fault Tolerant (BFT) algorithms is also an alternative approach. It was originally described in terms of what is called the Byzantine Generals Problem¹⁴. Several Byzantine generals have, with their armies, surrounded a city. Some of them, but a minority, may be traitors. The problem is to find a way in which the generals may arrive at a plan of attack using only messengers between them such that all loyal generals agree on the same plan. There are a few variants, but essentially all of them have a maximum threshold of disloyal generals. The algorithm works if the actual number of disloyal generals does not exceed this threshold.

The Byzantine Generals problem applies to blockchain because one cannot assume that there are no malicious nodes in the network. In fact, if a malicious entity gains control of more than half of the nodes on a given network, that entity determines what block, and its contents, achieve consensus.

14. Lamport, L.; Shostak, R.; Pease, M. (1982). "The Byzantine Generals Problem" [PDF]. *ACM Transactions on Programming Languages and Systems*. 4 (3): 382-401. doi:10.1145/357172.357176.

REFERENCES

1. Peck, M. E. [2015, November]. Bitcoin needs to get its act together. IEEE Spectrum. Available [Online]: <http://ieeexplore.ieee.org/document/7335883/>
2. Christensen, Clayton M. (1997), The innovator's dilemma: when new technologies cause great firms to fail, Boston, Massachusetts, USA: Harvard Business School Press, ISBN 978-0-87584-585-2.
3. Allaby, D. [2016, October, 27]. The Trust Trade-off: Permissioned vs Permissionless Blockchains. www.fjordnet.com. Available [Online]: <https://www.fjordnet.com/conversations/the-trust-trade-off-permissioned-vs-permissionless-blockchains/>
4. IEEE [2016, June]. The blockchain has a dark side. IEEE Spectrum. Volume 53, Issue 6, p. 12 – 13. Available [Online]: <http://ieeexplore.ieee.org/document/7473136/>
5. Underwood, S. [2016, November]. Blockchain beyond Bitcoin. Communications of the ACM. Vol. 59, No. 11, P 15 – 17. Available [Online]: <http://cacm.acm.org/magazines/2016/11/209132-blockchain-beyond-bitcoin/fulltext>
6. Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2016, April). A fistful of Bitcoins: Characterizing payments among men with no names. Communications of the ACM. Vol. 59, No. 04, P. 127 – 140. Available [Online]: <http://dl.acm.org/citation.cfm?id=2504747>
7. Zohar, A. [2015, September]. Bitcoin: Under the hood. Communications of the ACM. Vol. 58, No. 9, P. 104-113. Available [Online]: <http://cacm.acm.org/magazines/2015/9/191170-bitcoin/abstract>
8. Peck, M. [2016, November 18]. A blockchain currency that beats Bitcoin on privacy. IEEE Spectrum. Available [Online]: <http://spectrum.ieee.org/computing/networks/a-blockchain-currency-that-beats-bitcoin-on-privacy>
9. Ito, J., Narula, N., and Ali, R. [2017, March]. The blockchain will do to the financial system what the internet did to media. Harvard Business Review. Available [Online]: <https://hbr.org/2017/03/the-blockchain-will-do-to-banks-and-law-firms-what-the-internet-did-to-media>
10. New York State Department of Financial Services. Part 200. Virtual Currencies.
11. Regulation (EU). No. 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the intermarket and repealing Directive 1999/93/EC. Available [Online]: <https://www.anacom.pt/render.jsp?contentId=985913#VWt0Rn85UvL>
12. Kouzmanoff, J. [2017, January, 19]. Blockchain technology: a hidden gem of innovation for the solar industry. Solarplaza.com. Available [Online]: <http://www.solarplaza.com/channels/finance/11648/blockchain-technology-hidden-gem-innovation-solar-industry/>
13. Ileria, I. [2017, March 29]. Bosch, CISCO, BNY Mellon, others launch new blockchain consortium. Reuters. Available [Online]: <http://www.reuters.com/article/us-blockchain-iot-idUSKBN15B2DZ>
14. Smith [2016, June 8]. IoT botnet: 25,513 CCTV cameras used in crushing DDoS attacks. NetworkWorld. Available [Online]: <http://www.networkworld.com/article/3089298/security/iot-botnet-25-513-cctv-cameras-used-in-crushing-ddos-attacks.html>
15. Castillo, M. [2016, July 20]. Ethereum Executes Blockchain Hard Fork to Return DAO Funds. Available [Online]: <http://www.coindesk.com/ethereum-executes-blockchain-hard-fork-return-dao-investor-funds/>
16. Besnainou, J. [2017, February]. Blockchain meets Energy: State of the Market. Cleantech Group. Available [Online]: <http://www.cleantech.com/blockchain-meets-energy-state-of-the-market/>
17. Higgins, S. [2016, November, 1]. Ethereum Energy startup awarded blockchain patent. Coindesk. Available [Online]: <http://www.coindesk.com/blockchain-energy-startup-patent/>
18. Blockchainfirst. [2017, January 18]. The first Multipurpose Blockchain enabled EV Charging Station. Available [Online]: <https://medium.com/@blockchainfirst/the-first-multipurpose-blockchain-enabled-ev-charging-station-d8265c1bcb38>
19. Engerati [2016, April, 11]. Blockchain transactive grid set to disrupt energy trading market. Available [Online]: <https://www.engerati.com/article/block>

- [chain-transactive-grid-set-disrupt-energy-trading-market](#)
20. Engerati (2017, January, 10). Australia – land of blockchain opportunity. Available [Online]: <https://www.engerati.com/article/blockchain-transactive-grid-set-disrupt-energy-trading-market>
 21. Burger, A. (2017, February 27). More than Half German Utilities Carrying Out or Planning Blockchain Pilots. Energy Central. Available [Online]: <http://www.energycentral.com/c/pip/more-half-german-utilities-carrying-out-or-planning-blockchain-pilots>
 22. Clancy, H. (2016, October 3). How the blockchain will disrupt energy markets. EnergyBiz. Available [Online]: <https://www.greenbiz.com/article/how-blockchain-will-disrupt-energy-markets>
 23. Khyati, K. (2015, September 17). World's 9 Biggest Banks to adopt Bitcoin's Blockchain, The Hacker News. Available [Online]: <http://thehackernews.com/2015/09/bitcoin-blockchain.html>
 24. Hype Cycle for Blockchain Technologies and the Programmable Economy, 2016, (2016, July 27). Gartner Research. ID: G00308190. Available [Online]: <https://www.gartner.com/doc/3392717/hype-cycle-blockchain-technologies-programmable>
 25. ENEL (2017, February 10). ENEL to the discovery of the blockchain. ENEL. Available [Online]: <https://www.enel.com/en/media/news/4201702-enel-to-the-discovery-of-the-blockchain.html>

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI members represent 90% of the electric utility revenue in the United States with international participation in 35 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; and Lenox, Mass.

Together ... Shaping the Future of Electricity

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94309-0813 USA • 800.313.3774
• 650.855.2121 • epri@epri.com • www.epri.com

© 2017 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER ... SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

The CHAIRMAN. Thank you, Mr. Golden.
Ms. Henly, welcome.

**STATEMENT OF CLAIRE HENLY, MANAGING DIRECTOR,
ENERGY WEB FOUNDATION**

Ms. HENLY. Good morning, Madam Chair Murkowski, Ranking Member Cantwell, Senators. Thank you for inviting me to speak today on this exciting and important topic.

My name is Claire Henly. I'm a Managing Director at the Energy Web Foundation which is the brainchild of a German technology start-up, Grid Singularity and a U.S. energy non-profit, Rocky Mountain Institute. Many of you will be familiar with Rocky Mountain Institute's founder, Amory Lovins, who sat in this chair often. The Energy Web Foundation builds open-sourced blockchain tools and technology for the energy sector in collaboration with our more than 80 affiliate companies, the likes of Duke, PG&E, Exelon, Sempra and many more globally.

Today I would like to leave you with three messages. First, leading blockchains are replacing energy intensive bitcoin mining practices with efficient alternatives. Second, blockchain presents a valuable opportunity for the U.S. electricity grid to improve security, increase efficiency and lower costs. And third, the U.S. is behind both Europe and Asia in advancing the frontier of blockchain research and development.

First, bitcoin uses significant electricity, as we've heard, roughly 10 to 20 terawatt hours per year in our median estimates which is enough to power one to two million U.S. homes. Bitcoin's electricity use is required by its security mechanism, proof of work, in which block validators, known as miners, work expending computing power and electricity to add blocks to the bitcoin chain. Bitcoin's energy use is a substantial concern, not only for you, Senators, but also for the industry. We know energy consumption on this scale will limit the potential of the technology to expand into and create value in energy and other sectors.

As a result, there are several alternatives that are under development. Two are leading the pack. One called proof of stake requires validators to deposit value on chain that is seized in the case of bad behavior. The other, called proof of authority, requires known, trusted and regulated entities to validate transactions. These alternatives have drawbacks, respectively, cost of capital and increased centralization, but they both also have the important benefit of using many orders of magnitude, less electricity which would lead to bitcoin consumption on the scale of a small office building, not a small country.

I don't mean to suggest that blockchain's electricity use will disappear overnight. There's currently little momentum in the bitcoin community to move away from proof of work. However, other leading blockchains are adopting the alternatives. Actually, the most widely-used blockchain, Ethereum, is in the process of switching to proof of stake and the Energy Web Foundation's blockchain is launching next year using proof of authority.

Second, beyond bitcoin blockchain presents a valuable opportunity for the U.S. energy sector. Bitcoin created a secure, distrib-

uted currency ledger, but blockchain has become, as we've heard, more than just bitcoin.

Subsequent innovations in blockchain have added the ability to execute code, turning the distributive ledger into a distributed computer. Features of this computer include automatic contract execution, no single point of failure, full data traceability and selective data permissioning and perhaps, most importantly, a common record of the state of the network held by all users.

What could this mean for the U.S. energy sector? It means a grid that is no longer only centrally controlled and vulnerable to grid operator attack; it means an energy market where customers can choose where they buy their electricity without fear of providence uncertainty or high broker fees; it means a grid where households, like large generators, can be accurately compensated for self-generation and efficiency; and it means an electricity system that is no longer unidirectional, but instead it supports local energy exchange, making for an overall network that is more dynamic, resilient and efficient. These are just a few examples. In the growing energy blockchain ecosystem there are dozens of companies who are actively working to develop applications, specifically for the energy sector.

Last, the technology is at an early stage. Important for you to know is that the U.S. is behind both Europe and Asia on research and development and the global hub for blockchain is not in San Francisco, as you might expect, but is in Berlin. The DOE's funding to explore blockchain's cybersecurity benefits is one good example of how the U.S. Government can support the technology but more is needed.

The path to the genesis of the internet, the fusing of Arpanet and TCP/IP in 1983 was not straightforward or without problems. Similarly and perhaps unsurprisingly, the first blockchains have flaws. But as an industry we are actively working to implement solutions.

Also, as in the early days of the internet, the current benefits of blockchain are not simple to grasp. Before email many believed the internet would have purely military applications. No one was dreaming of Amazon. But while the internet has allowed unprecedented information sharing, blockchain can create secure information agreement leading to open markets, distributed ownership and transformed institutions.

On behalf of everyone at the Energy Web Foundation, thank you again for inviting me to speak to you today. I welcome your questions.

[The prepared statement of Ms. Henly follows:]



Written Testimony of Claire Henly
Managing Director, Energy Web Foundation
 Committee on Energy and Natural Resources, U.S. Senate
 August 21, 2018

Summary

- *The blockchain industry is moving away from energy intensive versions of the technology.* Early blockchains such as Bitcoin and Ethereum are energy intensive, but there are two much less energy intensive alternatives in late stages of development. These alternatives are not without tradeoffs, but it's become clear to the industry that the energy intensity of Bitcoin and similar networks is a critical problem to be addressed.
- *There are valuable potential applications of blockchain in the energy sector.* Blockchain's novel functionality can render energy markets more efficient and more open.
- *There are still barriers to be overcome before blockchain can contribute to the energy sector at scale.* More work is needed—in particular in the US—to address these challenges and establish the value of blockchain in the energy sector.

The Blockchain Industry is Moving Away from Energy Intensive Versions of the Technology

With Bitcoin's Growing Popularity Comes Growing Energy Use

With its increasing popularity¹, estimates of Bitcoin's global electricity consumption have grown to between roughly 1 TWh and 32 TWh² per year, depending on methodology. The former is enough to power roughly 90,000 U.S. homes for a year³, the latter is on par with the annual electricity use of Denmark. Today a single Bitcoin transaction can consume as much electricity as an average American home does each week. And that's just Bitcoin; there are hundreds of other blockchains that, though they are less used, can be similarly energy intensive per transaction.

Not All Blockchain Networks Are Created Equal

Bitcoin's consensus protocol—the mechanism by which the computers in the network validate and agree upon transactions—is called “proof-of-work” (PoW). It is so named because miners work hard (i.e., devote real-world resources like computers and energy) to solve an equation and prove it to the rest of the network. Though inherently energy-intensive, PoW has the benefit of setting an extremely high bar for validating blocks, making it exceedingly difficult to manipulate or corrupt the blockchain.

PoW is only one way to validate transactions, and at least two alternatives hold promise as blockchain approaches with lighter energy footprints:

- *Proof-of-Stake (PoS):* Under a PoS system, network participants own a share of the system's digital currency and are selected to validate blocks in proportion to their share. This proportional stake—and the risk of losing this “deposit”—disincentivizes malicious actors. Since there isn't computing competition among all participants to solve an encryption problem, as in a PoW network, PoS blockchains use a fraction of the energy.

¹ [Wall Street Journal](#)

² [Energy Web Foundation](#)

³ [EIA](#)



However, PoS has a cost of capital for deposited money, creating economic cost and skewing decision-making power to the wealthy.

- *Proof-of-Authority (PoA)*: PoA systems rely on a trusted set of authorities to create and validate blocks. To ensure good governance, there are rules that regulate how authorities join the network and how transactions are validated. Such PoA networks are well-suited to regulated industries where entities responsible for maintaining the network (authorities) need to be known, rather than remain anonymous, as in PoW or PoS chains. Since only approved authorities are validating the blockchain, there is no competition amongst authorities to race each other, which means drastically less power consumption than PoW blockchains. However, PoA requires a set of trusted intermediaries, reducing the distributed nature of the blockchain.

The Industry is Moving Away from Energy Intensive Networks

Blockchain's energy use will not change overnight, but the most widely used chain, Ethereum⁴, is actively moving away from Proof-of-Work.⁵ Bitcoin may take longer to change, but the industry's overall energy usage will decline as existing networks move away from Proof-of-Work and new networks go live with PoS and PoA. The Energy Web Foundation, for example, was founded in part to address these and other technical limitations of blockchain. We are currently running a Proof-of-Authority test network that we will launch live in 2019.

Blockchain Has Novel Functionality That Could Transform Energy Markets

Blockchain's Functionality is Relevant to the Energy Sector

Blockchain has the potential to play a valuable role in energy markets, providing several novel functionalities:

- *Increased Market Access*: with smart contracts automating many of the functions necessary to bid, settle and participate in markets, blockchain can open up high-barrier-to-entry energy markets to smaller participants.
- *Enhanced Traceability*: by creating unique and trusted digital identities and allowing all users to work off a common ledger, blockchain can seamlessly track ownership of assets (e.g., electric vehicles) and data (e.g., smart meter data), increasing certainty of the origin of assets and electricity.
- *Direct Ownership*: through automated smart contracts, blockchain makes it possible to raise financing for an asset that directly represents an ownership stake and right to partial profit, allowing for enhanced execution of locally owned energy projects.
- *Asset Agency*: through unique and trusted digital identities, blockchain can enable assets like batteries to participate directly in markets without the need for a human intermediary, increasing grid efficiency and decreasing overall electric costs.
- *Data Sovereignty*: by creating unique identifiers for asset owners, assets, and the data produced by those assets, blockchain can create direct data ownership and selective permissioning, allowing for better customer data management and privacy.
- *Distributed Cybersecurity*: through its distributed ledger and distributed consensus mechanism, blockchain ensures that there is no single point of failure for grid control systems, increasing the robustness of the grid to certain types of cyber attack vectors.

⁴ The Ethereum main-net processes up to 1,300,000 transactions per day ([Etherscan](#)), more than the number of transactions processed by all other public blockchain (excluding Ripple) combined.

⁵ [Coindesk](#)



Blockchain Can Make Energy Markets More Efficient and More Accessible

One concrete example of a blockchain energy use-case is the Energy Web Foundation's reference application, *Origin*, a blockchain based tracking system for renewable energy credits. Current renewable energy credit markets have high transaction fees, are opaque, and are closed to small participants. Blockchain technology, through its trusted common ledger and automated transactions, allows for drastically lower transaction costs, higher functionality, and greater market access. This has the potential to create new sources of revenue for households with self-generation and allow all customers to buy renewable electricity at their discretion.

While renewable certificate markets are relatively small worldwide, a similar use of blockchain technology could apply to wholesale electricity markets. *Origin* is intended as a reference for commercial actors to build applications in global energy markets of all types.

There Are Barriers to be Overcome Before Blockchain Can Contribute to the Energy Sector at Scale

Barriers, Beyond Energy Use, Are Limiting Blockchain's Usability

The technology is at an early stage of development. Remaining barriers to deployment include:

- Enabling secure and seamless connections with real world assets
- Increasing transaction throughput
- Implementing low-cost data storage
- Educating IT professionals in the energy sector how to use the technology
- Educating energy sector regulators about the potential benefits of the technology
- Deploying a governance mechanism for actors to agree upon needed software upgrades

Energy Web Foundation (EWF) is Working to Address These Issues

EWF has assembled a consortium of 80 Affiliates ranging from large energy companies (such as PG&E, Duke, and Exelon) to small blockchain and energy startups (such as Electron, LO3, and Share&Charge), all focused on creating industrial-grade applications of blockchain technology in the energy sector. EWF's primary project is to develop and deploy the EW Chain, a public, open-source blockchain, purpose built to support applications in the energy sector. The EW Chain currently uses a proof-of-authority based consensus mechanism, enabling it to run with the energy footprint of a medium-sized office building, not a medium-sized country.

This Technology is at a Promising Early Stage; Government R&D Support is Needed to Establish the US as a Leader in this Space

Blockchain technology is at an early stage and more work is needed to address the limitations of the technology and understand its nuanced benefits for electricity grids and the energy sector. The DOE-funded research on blockchain's cybersecurity benefits is a good example of how the federal government can support the technology.⁶ However, Europe is far ahead of the US when it comes to blockchain demonstrations and expertise. Without further research and development funding, the US is at risk of falling behind as this technology quickly develops.

⁶ Department of Energy

The CHAIRMAN. Thank you, Ms. Henly.

Mr. Narayanan, I had it right the first time.

[Laughter.]

Mr. NARAYANAN. Good morning. Thank you.

The CHAIRMAN. Good morning, Mr. Narayanan.

STATEMENT OF DR. ARVIND NARAYANAN, ASSOCIATE PROFESSOR OF COMPUTER SCIENCE, PRINCETON UNIVERSITY

Dr. NARAYANAN. Chairman Murkowski, Ranking Member Cantwell, members of the Committee, thank you for the opportunity to testify about blockchain technology and its implications.

I'm an Associate Professor at Princeton University. I'm a computer scientist. I've been researching cryptocurrency and blockchain technology since 2013. I'm the lead author of a textbook on this topic that's been used in over a hundred courses around the country and worldwide.

I'll address two topics today. I'll offer a view on what we can expect in terms of the energy consumption of certain blockchains, then I'll discuss potential applications of the blockchain technology in the energy industry.

I'd like to begin by highlighting an important distinction that's already been raised here today which is between public and private blockchains. Public blockchains are open for anyone to participate in. They were the foundation of cryptocurrencies, and the majority of public blockchains today are based on mining which involves the computation of a large number of mathematical calculations.

Private blockchains, on the other hand, are operated by a limited set of entities, such as a consortium of banks or a consortium of energy companies. They don't involve mining, they're not tied to cryptocurrencies and the applications in energy trading that we've heard about mostly involve private blockchains. This distinction is important when we talk about the energy consumption of cryptocurrency mining.

Mining today is carried out in large scale, commercial operations using purposed built computing devices that are specialized to the task of mining and nothing else. At present, the miners of bitcoin, the original blockchain-based cryptocurrency, are collectively calculating about 50 billion billion of these computations every second. That's a 20-digit number. This rate of calculation requires a large amount of power.

It's hard to estimate precisely. We've heard some estimates today. I've included my own estimate in my written testimony which is about five gigawatts for bitcoin mining alone today. Other blockchains also consume a substantial, but still lower amount of energy.

Now, as we've heard, mining-free blockchain technology is being developed. How will this affect the future of blockchain energy consumption? Let me offer a few points on this.

First, it's easy to design private blockchains that don't require mining, but it's proven much harder to get rid of mining in public blockchains that support cryptocurrencies. There are many technical challenges even if those are solved. The question remains as to whether all the existing mining-based cryptocurrencies will switch to a mining free model. My view is that this is unlikely.

So how will mining energy consumption evolve in the future? The main factor that governs the economics of mining is the exchange rate between cryptocurrencies and dollars. Roughly speaking, if the price of a cryptocurrency goes up, it will become more valuable to mine, more miners will enter the market and more energy will be used in mining it. If the exchange rate goes down, then less energy will be used.

So, what are the policy levers that can be used to influence mining? It's important to note that miners are very cost sensitive. That means that taxes and other policy incentives and disincentives could have a big impact in terms of where they locate their operations geographically.

Now let me turn to the implications of blockchains for the energy industry. Many exciting applications have been proposed: blockchains that underpin existing energy markets; new markets, such as the peer-to-peer trading of rooftop solar power; smart devices that adjust their operation based on dynamic price signals, et cetera. Blockchains are one possible technology platform among many for implementing these applications. Many of these applications inherently require the use of blockchain technology, and we should pick the best tool for the job on a case-by-case basis. Blockchain-based recordkeeping systems can be more efficient compared to paper-based records, but at the same time, compared to other types of electronic databases and platforms, blockchains are often less efficient.

Finally, let me turn to cybersecurity. Our electric grid and energy systems are becoming more computerized and more networked. That leads to new cybersecurity risks. If foreign adversaries are able to exploit digital vulnerabilities to penetrate these networks, that means they might be able to interfere with the grid's operation.

Now technology for improving the security and fault tolerance of computing systems has been developed for several decades. Cryptography is a key element of these defenses. For example, digital signatures help to ensure that a control command on the grid, for instance, was sent by an authorized person rather than an intruder. Other key cybersecurity technologies include things like consensus protocols and firewalls.

In some scenarios blockchains could augment the cybersecurity benefits of these classical technologies, and I've mentioned some examples in my written statement, but blockchain technology is not a necessary or core component of cybersecurity. It brings potential benefits, as well as new cybersecurity risks, and policymakers should view it as one tool among many.

Thank you again for the opportunity to testify and I look forward to your questions.

[The prepared statement of Dr. Narayanan follows:]

**Written Testimony of Arvind Narayanan
Associate Professor of Computer Science, Princeton University**

**United States Senate, Committee on Energy and Natural Resources
Hearing on Energy Efficiency of Blockchain and Similar Technologies**

August 21, 2018

Chairman Murkowski, Ranking Member Cantwell, and members of the Committee, I thank you for the opportunity to testify about blockchain technology and its implications for energy efficiency and cybersecurity.

My name is Arvind Narayanan, and I am an associate professor at Princeton University. I am a computer scientist, and my main research areas are information privacy and cybersecurity. I have been researching blockchain technology since 2013 and have authored numerous peer-reviewed publications in this field. I have taught courses on cryptocurrencies and blockchains since 2014. I am a lead author of a textbook on the topic that has been used in over a hundred courses around the country and worldwide.

In this testimony, I will address three topics. First, I will provide an overview of blockchain technology. Second, I will describe the energy consumption associated with certain blockchains. I will explain why these blockchains consume a large amount of energy for their upkeep, and offer an opinion on how we can expect this consumption to evolve. Third, I will discuss the potential applications of blockchain technology in the energy industry. I will focus on the potential to improve cybersecurity and efficiency, while also highlighting the limits of blockchain technology and alternatives that might achieve the same goals.

1. Overview of blockchain technology

The term blockchain is used to describe a loosely related collection of technologies. What they have in common is a sequence of records that is collectively maintained by a set of stakeholders and is designed to support the addition of records while resisting modification or deletion of existing records. This is achieved by a technological mechanism that protects the integrity of the data even if any minority of participants attempt to undermine it.

Hundreds of blockchains exist today. The vast majority follow one of two basic designs. In public blockchains, also called permissionless blockchains, anyone may become a maintainer of the records. In the most prominent public blockchains today, maintaining the records involves a computationally intensive process called mining.¹ It requires mathematical calculations by a large number of computers working in parallel. This is by deliberate design: the aim is to ensure that any adversary aiming to disrupt the integrity of the system must control and operate, at

¹ I discuss mining-free public blockchains in the next section.

least momentarily, as much computing power as the rest of the miners. In other words, it is precisely the computational difficulty of mining that makes public blockchains hard for adversaries to attack.

Miners are compensated for their effort in maintaining the integrity of the blockchain. To enable this, each public blockchain is paired with a virtual currency, also known as a cryptocurrency. Addition of records to the blockchain triggers an algorithm to issue new units of the cryptocurrency, which are paid out as revenue to the miners. The blockchain in turn supports the operation of the cryptocurrency by serving as an authoritative record of cryptocurrency transfers. Since the blockchain is resilient to modification, participants trust the veracity of this record and use it to execute trades and determine currency balances. In other words, the blockchain is the technology platform necessary for the operation of the cryptocurrency, and the cryptocurrency incentivizes miners to maintain the blockchain. Neither would exist in a useful form without the other.

In contrast to public blockchains, private blockchains (also called permissioned blockchains or consortium blockchains) work very differently. Maintenance of the records is limited to a pre-specified list of entities, such as a consortium of banks or a consortium of utilities. Since a majority of stakeholders are assumed to be trustworthy and there is no risk of an unknown adversary attempting to subvert the system, mining is not necessary. Nor is a cryptocurrency needed for the functioning of private blockchains, although many such blockchains do support cryptocurrencies or other digital tokens. Finally, private blockchains tend to have a high capacity or throughput, that is, they support a high rate of addition of new records, whereas public blockchains are inherently limited in this respect.

| | Public blockchains | Private blockchains |
|-------------------------------------|--|---|
| Who maintains the blockchain | Anyone may participate anonymously | Limited, known set of participants |
| Relationship to cryptocurrency | Must be paired with a cryptocurrency | Cryptocurrency not necessary, but sometimes supported |
| Energy efficiency | Most prevalent design requires energy-intensive mining | Does not involve mining |
| Capacity (records added per second) | Inherently limited | High capacity achievable |

Table: comparison of public and private blockchains

2. Implications of blockchains for energy efficiency

On today's most prominent public blockchains, mining involves the computation of a large number of mathematical calculations, called hashes, in parallel. For example, as of this writing, miners of Bitcoin, the original blockchain-based cryptocurrency, collectively compute about 50 billion billion hashes, or 50 billion gigahashes, every second.² Most mining is carried out in large-scale commercial operations using purpose-built computing devices specialized to the task of repeatedly computing these hashes — and nothing else. These devices are housed in warehouses dedicated to mining, usually called data centers. Substantial energy is required to operate the computing devices as well as to cool them to keep them within their operating temperature limits.

An accepted method for deriving an estimate of the energy consumption of mining is to assume that all miners use the most energy efficient mining device available on the market.³ Commercial devices are accompanied by published specifications listing the number of hashes that can be computed per second using the device, as well as the power consumption of the device in watts. It is then straightforward to calculate how much power is required to compute 50 billion billion hashes per second using the most energy efficient devices available. I performed such a calculation and obtained an estimate of around 5 gigawatts for Bitcoin mining alone today.⁴ This is slightly under 1% of world electricity consumption, or slightly more than the electricity consumption of the state of Ohio or that of the state of New York. Other public blockchains also consume a substantial, albeit much lower, amount of energy.

To understand how this number might change over time, economists use equilibrium models of mining.⁵ Miners produce a virtual commodity in a competitive market. Miners will enter this market if it is profitable to mine and drop out if it is not, driving the market toward zero profit. Further, mining is a zero-sum game: the total revenue that can be earned per time unit by mining a specific cryptocurrency is fixed. In the case of Bitcoin, it is roughly 12.5 bitcoins every 10 minutes.⁶ In mid-August 2018, the exchange rate is roughly USD 6,500 per bitcoin, making

² An estimate of the current rate of hash computation is available at: <https://www.blockchain.com/en/charts/hash-rate>

³ See Alex de Vries, *Bitcoin's Growing Energy Problem*, 2 *Joule* 801-805 (2018), [https://www.cell.com/joule/fulltext/S2542-4351\(18\)30177-6](https://www.cell.com/joule/fulltext/S2542-4351(18)30177-6); Arvind Narayanan et al., *Bitcoin and cryptocurrency technologies: a comprehensive introduction*, Princeton University Press (2016), <http://bitcoinbook.cs.princeton.edu/>.

⁴ The most energy efficient mining device known to be in widespread use is the Bitmain Antminer S9, which achieves an efficiency of 10 billion hash computations per Joule of energy, resulting in an estimate of 5 gigawatts for Bitcoin mining. Recent announcements of new devices have claimed higher mining efficiencies; if these are in widespread use, the true power consumption might be slightly lower than 5 gigawatts. On the other hand, some devices in use may be much less efficient, which would mean that the true power consumption might be higher. Further, accounting for the energy consumption of cooling of mining data centers would also increase the estimate.

⁵ de Vries, *supra* note 3.

⁶ To be more precise, Bitcoin miners earn revenue from two sources: from newly minted units of cryptocurrency, and from transaction fees that cryptocurrency users pay to miners. The rate of minting of Bitcoin halves every four years. It is next scheduled to halve in mid-2020 from 12.5 bitcoins to 6.25 bitcoins every 10 minutes. Transaction fees are determined by a market mechanism; they currently

the mining revenue worth roughly USD 80,000 per 10-minute period. However, the exchange rate tends to fluctuate substantially. Against these revenues, miners have costs including electricity, mining hardware, and other costs of operating mining data centers. Electricity costs are a substantial fraction of overall costs — a fraction that is relatively stable over time. For example, if we use a ballpark figure of 50% for this fraction, the equilibrium model suggests that Bitcoin miners currently collectively expend roughly USD 6 million per day worth of electricity.

To summarize, the main variable in the equation that governs the energy consumption of cryptocurrency mining at equilibrium is the exchange rate between the cryptocurrency and dollars. Other factors such the amount of cryptocurrency available to mine per unit time have minor impacts on energy consumption, but they cannot explain the orders-of-magnitude increase in mining energy consumption that we have witnessed over the last few years. Roughly speaking, if the price of a cryptocurrency goes up, more energy will be used in mining it; if it goes down, less energy will be used. Little else matters. In particular, the increasing energy efficiency of mining hardware has essentially no impact on energy consumption.

Several attempts have been made to design public blockchains that don't require mining. The security of these designs is not as well understood theoretically or as well tested practically as that of mining-based blockchains; it is an active area of research and development. The developers of Ethereum, the second largest blockchain by market capitalization, have announced a goal of switching Ethereum to a mining-free model. However, the developers and the community behind Bitcoin have strongly resisted major changes to its design. In my opinion, this is likely to continue. Since Bitcoin is by far the largest consumer of cryptocurrency mining energy, this makes it unlikely that the maturation of mining-free public blockchain technology will have a short-to-medium-term impact on overall energy consumption in the blockchain sector. However, the long-term impact is harder to predict.

The above analysis pertains to the cumulative, worldwide energy consumption of public blockchains. A question remains as to where miners will choose to locate their operations. Cost considerations tend to dominate the geographic distribution of mining activity: low electricity prices and cooler climates (which leads to lower data center cooling costs) are attractive for mining. Policy incentives and disincentives such as taxes can also play a significant role.

3. Implications of blockchain technology for the energy industry

Just as blockchains can be used to record transfers of cryptocurrencies, they can be used to record transfers of other assets that can be represented digitally, such as commodities and derivatives. Since a blockchain can record both transfers of assets and payments for those assets, it can serve as a platform for a digital market. Indeed, a blockchain-based market for Internet

remain a small component of mining revenue, at a fraction of a bitcoin per 10-minute interval, but it is believed that as the minting reward dwindles, transaction fees will gradually increase. Most other mining-based cryptocurrencies follow this overall reward structure, but with differing specifics. Regardless of these nuances, at a high level, the amount of mining revenue available to be earned is not directly affected by the number of miners who compete to earn it, or the computational power they bring to bear.

domain names has existed since 2011, and numerous other blockchain-based markets are in various stages of development and deployment.⁷

Proponents of such markets view it as a benefit that unlike traditional digital markets, no single entity acts as a gatekeeper with the power to determine who is allowed to trade and who isn't. Another purported benefit is that transaction data would be accessible by all market participants, enabling more efficient trading. Finally, blockchain-based markets may improve efficiency and decrease settlement time compared to one where trades are settled using paper records or otherwise require human intervention. On the other hand, they tend to be less efficient compared to centralized digital markets.

In practice, blockchain applications have often fallen short of claimed benefits. One recurring pattern is that the development new blockchains is difficult to accomplish without a degree of central coordination, and the technology developers inevitably possess a significant ability to control the resulting platform. In at least one instance, a blockchain-based platform for voting was controlled by a single company, arguably negating the putative benefits.⁸

In the energy sector, several initiatives exist for utilizing blockchain technology in the context of both wholesale and consumer markets.⁹ Such blockchains are typically run by consortia of utilities or energy companies, and are thus private blockchains. Some initiatives enable firms to trade bulk quantities of power, gas, other commodities, and options on those commodities. A blockchain-based market might be more attractive than a centralized trading platform if market participants are averse to a single company controlling the platform. Other initiatives enable customers to directly trade electricity with each other in a "peer to peer" fashion, for example, by buying and selling excess rooftop solar power. However, peer-to-peer trading still requires the cooperation of utilities who ultimately control the physical flow of electricity.

Another envisioned application combines energy trading with automated control of energy consumption. This requires granular electricity prices based on time and location. If such a market existed, a refrigerator, for example, might contain a software controller that monitors energy prices, ambient temperature, and other factors to make decisions about power consumption at any given moment. Such automated controllers are often termed smart contracts. In my opinion, smart contracts for controlling energy consumption and generation can be adopted largely independently of blockchain technology.

To summarize, blockchains have the potential to underpin various types of energy markets, both existing and new. Many of these applications are currently speculative and blockchain technology is only one route to realizing them, with potential benefits as well as drawbacks.

⁷ The market for Internet domain names is called Namecoin (<https://namecoin.org/>). Other markets include Augur (<https://www.augur.net/>), a prediction market, and proposals in the finance industry. See Nasdaq, *Building on the Blockchain Nasdaq's Vision of Innovation* (2016), <https://bit.ly/2BuXQuU>.

⁸ See David Gerard, *West Virginia and the Voatz "blockchain" voting system — scaling and security concerns* (2018), <https://bit.ly/2BnXoJR>

⁹ See David Livingston et al., *Applying Blockchain Technology to Electric Power Systems*, Council on Foreign Relations Report (2018), <https://bit.ly/2vWmfN>

Blockchains and the cybersecurity of the grid

As the nation's electric grid and energy systems become more digital, cybersecurity risks arise: adversaries who exploit digital vulnerabilities to penetrate networks might be able to interfere with the grid's operation even without physical access to critical infrastructure.¹⁰ Like any digital system, securing the computing systems that supply our energy comes down to the protection of their confidentiality, integrity, and availability.

Confidentiality means keeping sensitive information from falling into the hands of unauthorized parties. Integrity means preventing unauthorized modifications to data and authenticating its origin. Availability means ensuring the smooth operation of computing systems and the accessibility of information to authorized parties when needed.

Today's information security best practices incorporate cryptography as a key vehicle for achieving these goals. Encryption, when properly implemented, aids greatly in ensuring confidentiality. Similarly, digital signatures and message authentication codes are vital tools for achieving data integrity. For example, an attacker who infiltrates a part of the network might be able to spoof a signal from a sensor or an intelligent electronic device, resulting in the issuance of rogue control commands, such as tripping circuit breakers. A design that aims to mitigate such attacks would require control commands to be accompanied by message authentication codes. The receiving device or system would verify the code before executing the command. If the keys were stored securely, the attacker would not be able to spoof the code without a compromise of physical security.

While these techniques have similarities to the manner in which blockchains ensure data integrity, blockchain technology is not necessary for achieving most of the cybersecurity benefits of cryptography in energy systems. That said, blockchain technology has the potential to provide additional cybersecurity benefits.

First, blockchains offer an alternative route to data integrity that provides an authoritative record of the date and time of transactions and other messages. Second, blockchains can enable rapid detection of (and recovery from) breaches. If all actions taken in a system were required to be recorded on a blockchain, it would provide a comprehensive audit trail that would aid intrusion investigation and forensics. Finally, blockchains can help improve the availability and fault-tolerance of computing systems, although alternative technologies exist.¹¹

¹⁰ See Rebecca Smith, *Russian Hackers Reach U.S. Utility Control Rooms, Homeland Security Officials Say*, Wall Street Journal (2018), <https://on.wsj.com/2mCMAM4>

¹¹ Fault-tolerant computing technology has a pedigree of several decades. For an early survey, see Michael Barborak et al., *The consensus problem in fault-tolerant computing*, 25 ACM Computing Surveys (1993), <https://dl.acm.org/citation.cfm?id=152612>. This technology is widely deployed in the Internet industry for building online services. E.g. Laura Nolan, *Managing Critical State: Distributed Consensus for Reliability*, in Niall Murphy et al., *Site Reliability Engineering: How Google Runs Production Systems*, O'Reilly Media (2016), <https://landing.google.com/sre/book/chapters/managing-critical-state.html>

On the other hand, blockchain technology might also introduce new cybersecurity risks. Let me highlight one. Participants in a blockchain network tend to adopt the same or similar software platforms. This so-called monoculture means that a vulnerability in one part of the network is a vulnerability in all of them, leading to the potential for cascading, rather than localized, failures.

To summarize, blockchain technology brings potential benefits as well as risks to the cybersecurity of energy systems. It is not essential for achieving the foundational components of digital security, and policy makers should view it as one of several possible technical tools for addressing energy cybersecurity.

Thank you again for the opportunity to address blockchain technology and its implications for energy efficiency and cybersecurity at today's hearing. I look forward to your questions.

The CHAIRMAN. Thank you, Doctor.
Dr. Kahn, welcome.

**STATEMENT OF DR. ROBERT E. KAHN, PRESIDENT AND CEO,
CORPORATION FOR NATIONAL RESEARCH INITIATIVES**

Dr. KAHN. Thank you, Chairman Murkowski, Ranking Member Cantwell and members of the Committee, I appreciate your invitation to testify before you today at this hearing.

I'd like to summarize a few points for you from my written testimony today.

When invited to testify at this hearing, I was specifically asked to focus my remarks on the more general topic of digital objects (DO). As Senator Cantwell just mentioned, bitcoin is a specific example of a distributed ledger technology and, in my view, distributed ledger technology and blockchains are specific examples of the more general topic of digital objects. So that's where my focus will be.

What I want to address also is what I call the digital object architecture which we've been developing at my organization, Corporation for National Research Initiatives, to manage digital information structured as digital objects. This architecture was created, initially, with U.S. Government support, is non-proprietary, it's in the public domain and it has been implemented by many parties over many years.

The architecture is a logical extension of the internet with a focus on simplifying the task of managing information in digital form in the internet or other computational environments. So I think this is pretty important. It offers users great flexibility in determining how to structure their digital information and how to manage it with a degree of cybersecurity protection previously unavailable.

The initial internet protocols we developed enabled networks and computer facilities to work together, interdependent of what the components actually were. That's why the internet keeps working today, even though the underlying technologies have scaled by a factor of, perhaps, ten million over the years. And the most essential aspects of those internet protocols remain unchanged, even though other aspects have evolved quite a bit over time. This same basic design approach for evolution and scaling have been taken in the development of the digital object architecture so that it will continue to work with the new and yet-to-be-developed technologies of the future.

Simply put, a digital object is a sequence of bits or a set of such sequences with a unique, persistent identifier and which incorporates a work or other information in which a party has rights or interests or in which there is value. This is certainly relevant in the energy space as well.

The DO architecture enables one to structure this digital information in a way though that it's self-describing with its own integrated metadata so that if a digital object were to show up on your computer, the software on that machine would know how to interpret the arriving bits. And digital objects can also be linked together as has been the case with blockchains.

As far as trust in the digital object is concerned, it has to come primarily from the use of strong cryptography. If you trust the cryptography, that should be sufficient for many reasons, but often both belt and suspenders are used, perhaps just because one doesn't fully trust the cryptography.

The digital object architecture enables each digital object to be separately encrypted and enables users to interact directly with the objects through the protocol for secure operation. Its utility is potentially quite large in both normal as well as abnormal situations.

In particular, a troubling situation would exist, for example, if the energy grid were compromised and no one in a position of responsibility knew anything about it. This might be an area where an implementation of the DO architecture could help to reliably detect such intrusions, either before they happen or afterwards.

A user seeking information needs to be able to securely and accurately identify the information of interest. They need to rely on the strong cryptography it uses, perhaps for authentication, perhaps for encryption. And then, it has to trust that the system provider can defend against the systemic attacks that may be instigated, perhaps even, surreptitiously.

On this later point, I've included with my testimony a paper that I wrote entitled, "The Role of Architecture in Internet Defense." It describes an alternative approach to the never-ending task of defending against threats in the form of harmful bit patterns.

We don't defend our borders by looking for photons and electrons, specific patterns. Yet, that's what we do on the internet today. But because of a technique known as data typing, digital objects can be structured to enable harmful inputs to be flagged ahead of time with a degree of granularity not previously available.

On the issue of value, I've also included with my testimony, a paper entitled, "Representing Value as Digital Objects," with a focus on being able to transfer such objects and to do so with anonymity while enabling the object to retain its value. This, of course, is the essence of cryptocurrencies.

Finally, I'd like to comment on how one may reasonably expect to bring about social change as well as technological change when the value of the new approach is not yet widely understood or demonstrated in the industry. This was the challenge we had with the internet.

Fundamentally, one needs to identify an area that requires assistance for which a new and novel approach seems to make sense and, if possible, find one or more early adopters to apply that approach without the need to require them or need any commercial provider to make substantial changes to their existing technology and/or services. Sometimes, only small changes are needed, maybe even no changes are needed if you can augment those existing capabilities to demonstrate the new approach. Eventually, if the new approach has enough added value, industry will likely adopt it and then integrate it by themselves. This was the approach taken in deploying the early internet. This is also how progress can be achieved in advancing and protecting our energy infrastructure, in my view, while at the same time, enhancing our ability to manage the infrastructure and better understanding what is happening with it.

I would be pleased to share with you more detailed information on aspects of the digital object architecture or its implementation if you think it may further assist the Committee in its deliberations.

In closing, I appreciate the opportunity to testify and I'd be happy to answer any questions you may have.

Thank you.

[The prepared statement of Dr. Kahn follows:]

Testimony of Dr. Robert E. Kahn

To the U.S. Senate Committee on Energy and Natural Resources:
August 21, 2018 Hearing on Energy Efficiency of Blockchain and other related Technologies

Chairman Murkowski, Ranking member Cantwell and other members of the Committee, thank you for the invitation to testify here today. My name is Robert Kahn and I am President & CEO of Corporation for National Research Initiatives (CNRI), a non-for-profit organization in Reston, Virginia, which I set up upon leaving U.S. Government service to provide a leadership role in the private sector, working with industry and academia, to foster the development of National Information Infrastructure and related pilot projects that demonstrate how to improve our national capabilities. Key to this approach is the use of computational facilities and high speed digital communication networks, which, as you know, are all critically dependent on energy to operate. We have witnessed significant infrastructural advances over the years, many of which CNRI has been directly involved with; but, unfortunately, as a nation, we are still a long way from fully realizing those goals.

I started CNRI after having spent 13 years at the Defense Advanced Research Projects Agency (DARPA), where for much of that time, as Director of the Information Processing Techniques Office, we funded a significant fraction of the IT research in the country. Among other things, DARPA funded the first packet switched computer network, called ARPANET. I was responsible for its overall system design; and I led the design and/or development of several other such networks based on the use of satellites and ground radio. At DARPA, I then started a project to link together these very different types of packet networks into what ultimately became more widely known as the Internet.

I would like to relate my experience with those activities to those that you and others are currently dealing with in securing our nation's critical infrastructure, in particular the Energy Grid. I am well aware that human threats are perhaps the greatest danger one may encounter here, and that significant efforts have been made to address issues concerning the hardening of our critical infrastructure. Also, there are inevitably technological vulnerabilities discovered in the industrial control systems that operate the infrastructure; and software patching is used to regularly update those systems. Unfortunately, we can only fix what we know to be a problem, so this is no guarantee against further technical threats to these systems; and patching itself can introduce challenges even if one takes steps to manage the supply chain effectively. I also assume that embedded threats can still be implanted in these system, can be activated at any time in the future, and that we may not be aware of it until it's too late.

Even if a solution were at hand to ensure that this kind of problem could never occur, there is the substantial and fundamental problem of getting industry to buy into any solution that entails major reengineering of their existing systems. This is as much a community buy-in and coordination issue, as it is a technical or even a security issue. We had a similar challenge facing us in creating the Internet, where it was not practical to cause every existing network to change to provide a new and unproven capability for internetworking. Instead, we set out to make use of existing capabilities and to work around the existing systems (i.e., with very minimal change) via the use of gateways (now called routers) and new protocols that the research community

experimented with in their computers. It is not possible to keep up in real-time with all the ways in which our systems can be compromised, but we can detect if changes have been made to the software that operates them, whether those changes pose a threat or not. While the analogy is not exact here, I believe the kind of workaround strategy we used in creating the Internet is implementable in the Energy Grid with only a small amount of help from industry, and (importantly) without requiring significant reworking of their existing industrial control systems. Over time, however, I would hope that industry would integrate those changes as well, if it sees the need or merit in so doing.

I do not hold myself out as an expert on energy systems or, for that matter, energy related issues. However, I do have a PhD in Electrical Engineering and took courses in Power Engineering along the way. As a scientist and technologist, I am familiar with some of the critical issues that may arise in the design and operation of real systems, including energy systems.

The subject today basically concerns managing information in digital form, whether or not it concerns the control of an energy system, cryptocurrencies (which is where the notion of blockchain is most prevalent today), or other types of application. As I am sure you have heard from others, a blockchain doesn't itself secure any system, but rather is intended to provide a trusted record of events recorded in digital form in what are called blocks. There are other ways to do this. The notion of blocks goes back more than fifty years and, ultimately, the trust in any digital information will depend on the trust one places in 1) the ability to securely and accurately identify the information of interest, 2) the strong cryptography it uses, and 3) the ability to defend against attacks that may be instigated surreptitiously (before or after the fact). I would now like to turn my attention to the many ways of managing such information and many ways of structuring digital information using cryptography to develop trust.

An important architecture for managing digital information in the Internet, which I call the Digital Object Architecture, derived from earlier work by CNRI on mobile programs, and has been under development going back to the 1980s, much of it with DARPA support. The architecture is not proprietary and is widely used today in many applications. It is a logical extension of the Internet whose purpose is to simplify and make more efficient access to digital information. The basis of this architecture is the "digital object" which is a sequence of bits (or even a set of such sequences) with an associated unique persistent identifier, and which incorporates a work or other information in which a party has rights or interests, or in which there is value. Digital Objects are self-describing and enable interoperability based on the use of embedded data types. When this technology was first introduced to an industry group during the 1990s, there was general agreement that this was an agreeable method for organizing, identifying, authenticating and otherwise managing information in digital form, and structured as containers, cryptolopes, packages, or, more generally, digital objects.

Each digital object has an associated unique persistent identifier that is resolvable to "state information" about the object. The operation of the resolution system by an organization is a deployment choice that may be performed in a restricted environment, or it may be more widely distributed, as is the Internet. The resolution system will accept the unique identifier and return information about the location(s) where the digital object may be accessed, how to authenticate

it, public keys if needed, and more. This architecture may also be implemented more widely to provide for defense of the Internet by managing its information flows with increased granularity.

Blockchain technology represents a specific way of structuring digital objects. In my view, a blockchain is itself a digital object, and every block within a blockchain could be considered a digital object as well. Thus, a blockchain is, in reality, a digital object that consists of other linked digital objects. In the Digital Object Architecture, digital objects are managed by repositories, which are themselves digital objects; and a repository provides network based services that enable objects to be stored, processed, accessed and otherwise managed. Every organization that provides repository capabilities, to itself or to others, will likely want to decide how to manage its repository services. Whether to deploy one repository or many, provide mirrored operations or not, and perhaps even how to link their objects and possibly provide links to other such objects.

How does one acquire trust in a digital object? Ultimately, in the digital world, the strongest protection is in the cryptography that is used. As computers get more powerful, we may need to re-encrypt data that was once thought to have strong enough encryption, but we should have adequate warning (as in a decade or more) to get prepared. Questions that can be raised here are 1) was the information accurate before it was originally encrypted, 2) what if you don't have access to the decryption keys, 3) what if the information was structured in a proprietary data format that requires proprietary software to manifest the underlying information, and 4) what kind of computational environment (which may be antiquated) is needed to run the possibly antiquated proprietary software?

Blockchain technology is said to provide such trust, not because it uses strong cryptography, but because every block is cryptographically linked to another block, which (in turn) is cryptographically linked to yet another block and so forth. Multiple distributed systems and different organizations are typically involved, such that a change to any one block, or a subset of the blocks, could easily be determined. This approach requires many systems, much storage and the ability to maintain these records over suitably long time frames.

Is such an approach necessary to develop trust? Probably not. Are there other equally effective ways to generate trust? Almost surely? Are there better ways to develop trust? This is not entirely a technical question, or even a factual matter, but rather a question about comfort levels or perhaps beliefs. I would not argue that blockchain technology has no role here, since it is really one particular way to implement digital objects. But I would certainly urge that serious consideration be given to all the other ways in which one might protect, secure and hopefully trust that the Energy Grid is safe from corrupted operation.

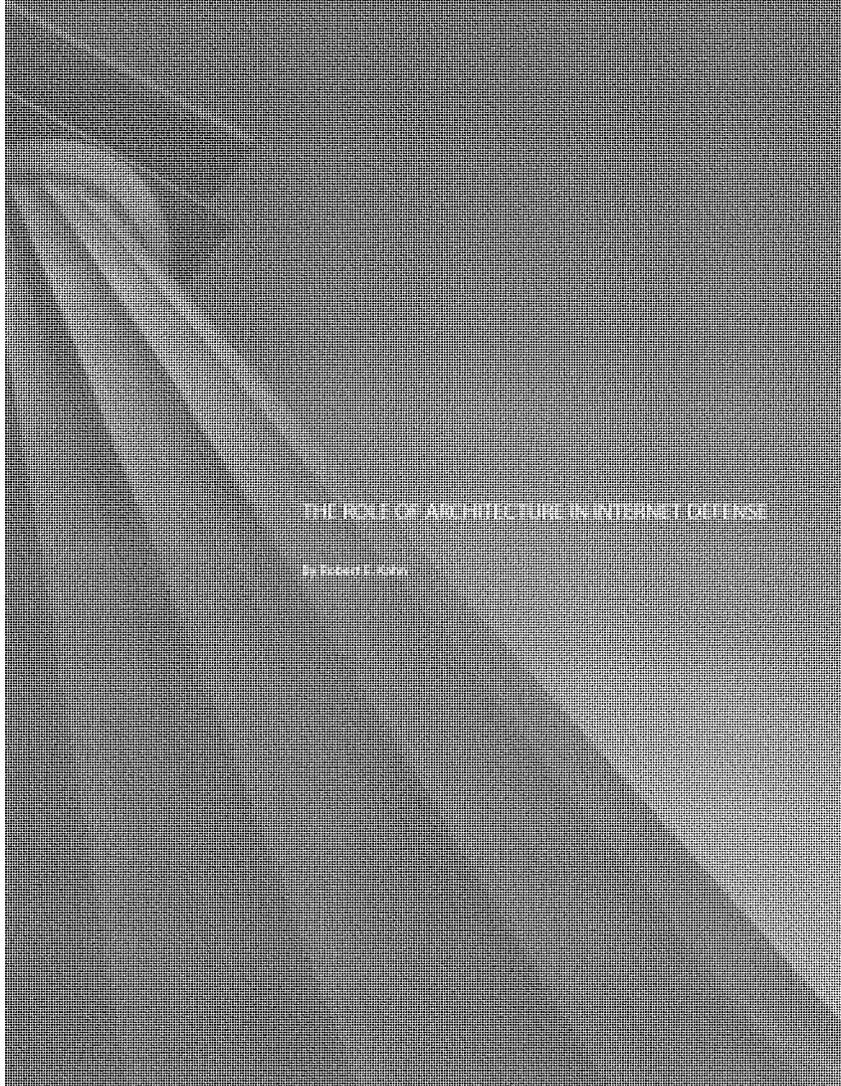
Ultimate trust in a digital object, no matter how you obtain it, would be based on the application of strong cryptography, whether just for authentication or to hide the contents. One size fits all is unlikely to be what is required for all applications in either the short or the long term. And the overall efficiency of the choices made will be an important part of the decision-making process.

To elaborate somewhat on the points I made today, I am including three attachments to my testimony, namely a paper I wrote on "The Role of Architecture in Internet Defense", a paper I

wrote with Patrice Lyons on “Representing Value as Digital Objects”, and, finally, a slide presentation I gave in March 2018 at an Asia-Pacific Blockchain Conference in Melbourne, Australia, entitled “Trusting Digital Entities”. The last two slides of that presentation also contain a number of other references you may find of interest.

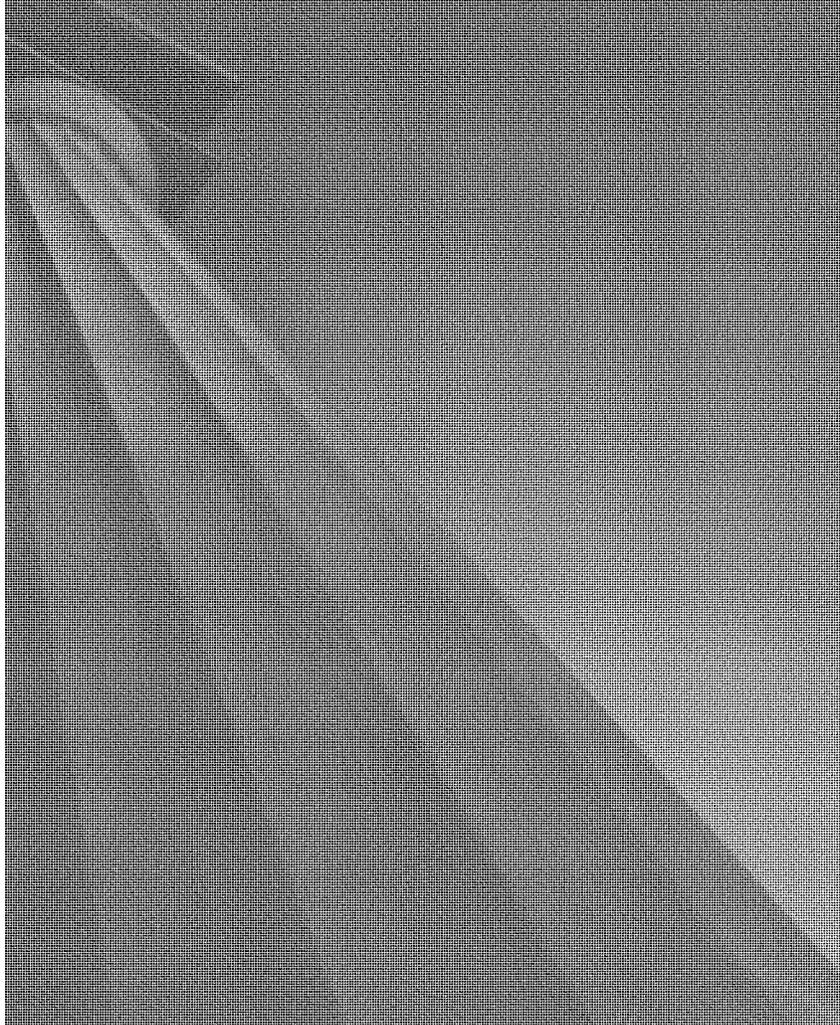
I would be pleased to share with you more detailed information on aspects of the Digital Object Architecture or its implementation if you think it may further assist the Committee in your deliberations.

In closing, I appreciate the opportunity to testify today and would be happy to address any questions you may have.



JUNE 2011

America's Cyber Debate
Security and Privacy in the Information Age



THE ROLE OF ARCHITECTURE IN INTERNET DEFENSE

By Robert E. Kahn

Since it was first introduced in the early 1970s, the Internet has met the growing needs of an ever-widening community of users with great benefits to individuals, organizations, governments and their associated disciplines. Yet, along with that growth and evolution has come an increasing downside, namely traffic that intrudes and may disrupt productive uses of the Internet. Worse yet, concerns exist that such unwanted and unwarranted intrusions may cause more extensive damage in the future. Managers of information systems and resources attempt to find ways to ensure that access controls are not breached, or that intrusions or disruptions have little likelihood of success: There are no guarantees, however, that a resourceful adversary will not find ways to subvert existing techniques to their own benefit. Since cyber insecurity is likely to persist, a rethinking of the architecture of the Internet, and how it might evolve to become more secure, is warranted.

This chapter explores the interplay between Internet architecture and the ability of users, network operators and application service providers to adequately defend against threats posed by others on the Internet. It introduces the digital object (DO) architecture and suggests a way of integrating certain defined functionality into the Internet based on the use of digital objects. This approach is compatible with existing Internet capabilities and has the potential to substantially improve our ability to detect and deal with intentional hostile actions. It would also deal with actions that are simply accidental or naively misguided, but which may have serious consequences.

Today's Internet subsumes a wide range of networks, devices and other computational facilities, as well as diverse services, processes and applications. In order to protect against real and potential threats, technical capabilities are required to understand what is transpiring within the Internet and its various constituent components, and to take steps to deal with emergent

situations that may require action. For example, most laptop users have little or no idea what is transpiring on their computers, and no effective way to find out in real time. They may only know that something is not working properly, or that the machine is running more slowly than usual. At present, the Internet landscape is sufficiently complex that the myriad exchanges of bits over the Internet cannot easily be differentiated by intent or function. Certain architectural changes to the Internet, which primarily affect the way the Internet is used, can help in mitigating these situations. Specifically, the DO architecture can help remediate this situation.¹

There are no guarantees that future threats, which require reconsideration of various architectural and design choices in the future, will not materialize; nor does use of the DO architecture guarantee that those who ignore or do not otherwise choose to take advantage of new architectural approaches will necessarily be harmed by that choice. At present, the Internet environment is tilted in favor of those with adverse motives, while the rest of the community must be on constant vigil to defend against harmful interference. However, over time, architectural changes become more pervasive. The assertion of this chapter is that the playing field will become more level in a way that provides architectural advantages for the defense of the Internet.

In the DO architecture, all system interactions involve the exchange of structured information in the form of digital objects, each of which has a unique identifier that can be resolved by a resolution system to state information about the object. Information, structured as a digital object, can be accessed and used by resources on the Internet based on its identifier, and is subject to any stated access controls or permissions associated with such objects. Even user commands, when invoked, can be converted into digital objects before being sent. This enables interoperability of the systems that embrace the protocol.

Digital Objects

A digital object consists of a data structure that is flexible, scalable and extensible. The data structure has a unique persistent identifier and may be used to create one of the following:

- A set of typeable attributes that describe the object and of which at least one is named the object's identifier, which is mandatory;
- A set of named "data elements" that hold potentially large bytes, references, hierarchical paths to one or more data files;
- A set of typeable attributes that are part of the data elements;

The elements of a digital object consist of "typeable items" that software at the destination and other locations can interpret for further processing. A protocol, known as the DO protocol, is responsible for managing the interactions between systems, services and other resources. This protocol enables actions to be taken based on the use of identifiers. The actions to be taken, and the targets of those actions, are specified by identifiers, which relate to digital objects that provide the actions. A resolution system to the target information. This system also enables operations of resolution by identifiers, and these operations are known as, with each identifier and resource also has at least one unique identifier linked, a state that uses resolvable identifiers depending on the particular use the individual is making of the resource. The system is whether they are regarding their endpoint or acting on an individual.

While many, if not most, interactions on the Internet are likely to be reasonable and legitimate, intrusions or hostile actions need to be flagged. Action must be taken to prevent damage, or other steps must be taken to quickly isolate matters. Even with the more structured view of the Internet provided by the DO architecture, the task is extremely challenging. Without such a view, the task is close to daunting, and would likely require semantic

interpretation of unstructured interactions, even if decrypted on the user's machine, that may be beyond the state of the art.

In the future, if arbitrary information arrives, the type of information will need to be understood from the structure of the information itself to enable further processing. Further, the environment into which the information arrives or is ultimately processed will require some degree of structuring, such as the structuring provided by the DO architecture, to determine with more specificity how best to deal with the information. In some cases, manual intervention may still be called for. In many other cases, however, automated processing may be possible based on interpretation of the structure of the actual information. For example, a medical reading sent by a remote wireless device might be understood from the structured information itself and placed in the user's medical record. Likewise, a remote financial transaction may be received and inserted automatically into a record of the user's daily transactions. Information collected in real time from remote sensors and appropriately identified can also be managed according to general rules and procedures adopted for such types of sensor information.

Overview of the Existing Internet Architecture

The existing Internet architecture was designed to enable the interconnection of multiple networks, devices and other computational facilities. Each potentially had a different design and performance, such that computers on different networks could communicate seamlessly and reliably with each other without having to know the location of the facilities, the intervening networks or how to actually route the information. More specifically, it enabled information in the form of packets of digital information to be communicated between computers without the need to first establish communication pathways between the computers.

As a result, the Internet has become a standard means of communication worldwide, not only for

traditional computer facilities, but also increasingly for digital representations of voice, video and sensor data managed by computers.³

At present, the Internet environment is tilted in favor of those with adverse motives, while the rest of the community must be on constant vigil to defend against harmful interference.

The Internet's creators based the existing architecture on two relatively simple notions. One was connecting networks with routers, which forward received packets by a process in which the routers act as relays with each step hopefully moving the packet closer to the eventual destination. The destination is specified by a globally unique identification known as an Internet protocol (IP) address that distinguishes the destination machine from all other destination machines on the Internet. The routers interpret the IP address to determine how best to route the packet. The process of communicating packets does not require the user to specify how to route the packets, which combination of networks to use, or even where the destination machine is located. Indeed, except for certain control information (such as the IP address) the contents of the packet may be encrypted. A dynamic routing protocol is used to adapt to changes in the underlying network components, such that if the packet can be routed to the eventual destination, it can be delivered in a timely fashion.

The second notion was the use of a host protocol, originally known only as the Transmission

Control Protocol (TCP), to enable the components to intercommunicate. TCP was later separated into two parts, one of which is IP, and the remaining part remained TCP. At the destination computer, TCP checks the validity of the arriving packets, discards duplicates that may have been generated along the way, reconfigures the data as appropriate and takes the necessary next steps in furthering the processing of the packets at the destination. In 1995, to clarify what the Internet actually was, the U.S. Federal Networking Council provided a definition of the Internet as a global information system that enables information resources of all kinds to intercommunicate by use of certain defined protocols (including IP) or their logical follow-ons and extensions.⁴

We note here that the overall objective of today's Internet is to ensure that global connectivity is achieved with low latency and reliable communication. While attacks on the network components of the Internet are possible, the Internet is far from completely defensible. Operators can take many types of precautions to ensure that traffic originating from users on their networks – and transit traffic from other networks – cannot directly cause actions within their networks (adverse or otherwise) other than to forward packets to their intended destination. However, although network operators can play a central role in helping to understand what is happening within their networks when adverse actions are reported or detected elsewhere, much of the concern still centers on vulnerabilities of the application service providers, their users and the underlying information systems they employ.

Vulnerabilities in Today's Internet

Various characteristics of the existing Internet make it especially vulnerable to harmful interference. One is the lack of overt security, which makes communications vulnerable to interference. Second is lack of identity management, which makes verification less secure than perhaps may be desired or necessary. Password protection is often used, but public key

At present, all communications are treated basically with equal significance, thus making it difficult to differentiate between those that are known and acceptable, versus those that are unknown and possibly undesirable.

systems offer greater protection assuming the private keys are not communicated over the Internet. Passwords, which are communicated, may travel in the clear or be included in email messages (or perhaps accessible files), and can be used by anyone to access a password controlled system if they know the account name. Third is freedom of communication without prior arrangement that can include desirable or essential communication; however, this also enables undesirable communications, which may range from simply annoying to potentially harmful. There is a role for anonymous and non-pre-arranged communication in the Internet. But, at present, all communications are treated basically with equal significance, thus making it difficult to differentiate between those that are known and acceptable, versus those that are unknown and possibly undesirable. The key to addressing this issue lies with architectural changes in how information is managed in the Internet, including, in particular, in the devices and other computational facilities that provide the application services.

Much has been done to protect the various networks that comprise the communications portion of the Internet, and serious ongoing efforts exist to build ever more robust and reliable computational facilities. But, for the most part, the most severe

vulnerabilities in today's Internet exist in those applications – in operating systems and in other resources – that cannot adequately defend themselves. The extent of the threat possibility is still unfolding, but the earliest examples of intrusive action are by now well known. For example, spam is unwanted email that consumes communications capacity and can overwhelm user systems. But spam is increasingly being filtered out with the help of commercially available software designed to distinguish between spam and non-spam communications. Generally speaking, these software packages are not perfect, but they do reduce the nuisance significantly. Since most spammers rely on the dissemination of lots of similar traffic relatively indiscriminately, certain charging schemes could mitigate the spam traffic. However, most spam is not intended to cause damage, and some unwanted advertising might actually be of interest to some. In most cases, however, it represents an intrusion upon an unwilling recipient.

Other actions can actually cause damage in some form. Intrusions that penetrate user systems can collect private information, can harm or degrade the operation of the user's system and in extreme cases can render it unusable. These harmful actions are usually achieved by exploiting vulnerabilities in the operating system or in one or more applications that run on the machine. These actions result from incoming traffic generated by usually unknown sources that may have immediate effect, or may be the result of implants which arrived over the Internet much earlier. Indeed, one of the loopholes that many users are unaware of is that such intrusive software and implants may result from devices such as memory sticks that transmit them when inserted into the user's machine. Any individual whose memory stick has been compromised can (in principle) compromise any system to which it comes into contact. If you change the word "compromise" to "infect," the analogy with epidemiology becomes clear.

Finally, every network capability can be compromised by what are known as distributed denial-of-service attacks. These generally require coordinated actions by lots of machines on the Internet; and certain known types of attack can be mitigated or denied by the network operators who detect or are otherwise made aware of them. The first line of defense here must be the network operators.

How Best to Deal with These Vulnerabilities?

What can be done to deal with this situation going forward? Three assertions are made in this chapter, each of which is discussed further below. First, the DO architecture will help to achieve increased visibility and awareness into the possibility of actions that threaten systems that are part of the Internet. Second, a greater use of identity-based transactions on the Internet will ensure that – with the user's concurrence – the parties and perhaps devices and systems/resources involved in the transactions can be determined from the transactions, while still supporting privacy and allowing anonymous operations, if desired. Third, the use of an identifier-based mode of interaction with Internet resources may help to circumscribe the kinds of actions that can be taken and thus help to clarify the landscape whereby intrusions may occur. None of these steps, by themselves, will prevent clever individuals from seeking workarounds; but the architectural constraints can help to make the commission of unwanted actions more visible and harder to accomplish.

INCREASING VISIBILITY AND AWARENESS

When we drive a car, we have a general idea of what the car is and what is normal and abnormal behavior. We can determine if a tire is flat, or a headlight is out by direct inspection. By other clues we know that gas is required to power the engine and can sense when the tank may be empty, and can see the tank level from the gauges on the dashboard. In general, we have a degree of visibility into the current operation of our car. Similar statements can be made for many other things we come into contact with and depend on. No such statement

can be made about the computational facilities on which we depend or, for that matter, about the Internet itself.

Internet operators may know quite a bit about their networks and other computational facilities from information accessible in their control centers, and they are in a position to readily respond to many types of outages and disruptions. In general, they tend to have visibility into their networks and are aware of their current state and what may go wrong. While there will always be new situations they have not encountered before and situations in which they have no idea what is happening, their forensic staffs will undoubtedly be engaged to deal with these situations quickly. No such thing can be said if the situation is such that significant parts of the Internet are compromised. Remedial action by one network operator may only solve a piece of a more complex problem. While a global means of responding to a widespread threat is needed, this is largely a matter for policymakers from multiple nations to address in a political arena.

Users are generally in the worst position to respond to attacks and would have to rely on Internet defenses provided by others or contained in the software they use. Users typically rely on their computational facilities to carry out well-known tasks, and are usually much less knowledgeable than technical staff working for the organizations providing Internet services. For example, there is no serious equivalent of a user dashboard that portrays for the user the most important aspects of its computer in such a way that the user will know when something unwanted has happened, or makes it possible for the user to take action to repair the problem. Turning a machine off and then back on does nothing to deal with an implanted and potentially harmful virus, for example. Virus checking programs can help to prevent such unwanted intrusions, but, with today's operating systems and applications, clever perpetrators will easily find ways around commercial virus checkers and even hide the presence of harmful actors on a user's machine from subsequent detection.

210 |

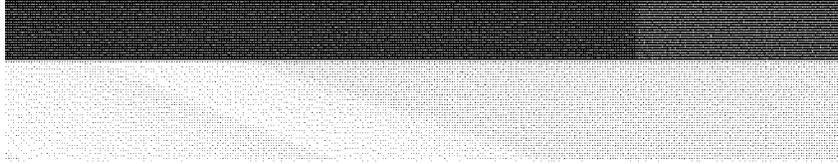
Users should be able to inspect their computers with as much facility as they can inspect their cars. What might they like to know? Perhaps some would like to visualize the "actual" memory map of their computer to know what is stored in the various parts of memory - "actual" meaning what is really there, rather than what a program may be fooled to think is there. In addition, a user might like to know when traffic that makes it into or out of his or her machine is notable for some reason. A user might like to know about information flow that is unauthorized and to locate (and remove) programs that may be extracting information and shipping it elsewhere without permission or authorization. Further, users may want to access audit trails that provide information about how the unauthorized program was put on their machines, along with certain information that may already be available such as the time it was created on the machine.

With the DO architecture, a basis would be in place for better understanding what is transpiring within the Internet, thus yielding greater visibility into and awareness of potential threats. In this mode of operation, all operations are explicit and, with authorization, can be logged and diagnosed. In addition, the same can be done for entire sessions consisting of many transactions in series. Programs and users will have a smaller set of well defined primitives to invoke in their instrumentation; and presentations of results can be more succinctly prepared along with more detailed semantic interpretation.

While much of this area is still likely to be the subject of research and development for many years, some aspects can be addressed immediately. It remains to be seen, however, just how much information the average user will need or want in order to be a more informed Internet user in the future.

IDENTITY-BASED OPERATIONS

Critical information about users and their intended actions on the Internet today is largely unavailable



Users are generally in the worst position to respond to attacks and would have to rely on Internet defenses provided by others or contained in the software they use.

from or not visible from the information communicated. Further, such information may be encrypted and, thus, the intent would be purposely hidden while the information is in transit. The communications are from one machine with an IP address to another and otherwise consist of a flow of undifferentiated packets. Authorized users who wish to make use of remote machines are usually required to log into the remote machine and supply a password of some kind. Some systems allow anonymous usage (e.g. most search engines), but take steps (usually by severely limiting the number of possible actions) to ensure that users cannot harm their systems.

Let us postulate that every user has the ability to obtain one or more unique identifiers from one of potentially many bodies, each of which is known and trusted to authenticate assertions in digital form about individuals, including the mapping between such assertions and their unique identifiers. Efforts are underway in several quarters to formalize this mapping process, but such formal processes may not be required in many customary cases. The most convenient way to handle this is via individual actions involving parties that know and trust each other. For example, if a patient has an identifier he is comfortable providing to his doctor, the doctor can rely on that identifier for the purpose of providing information to that patient, since the patient would have authorized use of

that identifier in the first place. If the identifier has associated with it a public/private key pair, and if the public key is accessible by use of the identifier, then a public key authentication can be invoked at any point the doctor or the doctor's information management system wishes to validate the patient. Similarly, if the patient contracts with a company to manage his or her health records, that company would have the obligation to make the connection between user and identifier.

An assertion about an individual that has a unique identifier acquired in connection with a desired task, process or service can be used to authenticate the user to a resource on the Internet. This provides a uniform way of validating the assertions. A similar process can be used to authenticate assertions about services, physical objects, organizations and other entities. When the service is remote, and the user learns of its identity from a third party, the user may elect to trust the third party (although this is not without its potential pitfalls) or to rely on bodies that maintain trusted information about such services.

However, users that do not wish to use their identifiers, or do not have identifiers, may still use Internet resources that permit such anonymous access. However, taking the route of anonymity may still allow services to be controlled in some situations where such control is deemed important or necessary. The main concern here is the provision of bogus identifiers by trust authorities or other entities. Using the term bogus does not mean that the identifiers are invalid, although that may be the case, but rather that the mapping of the identifier to assertions about the individual is not accurate or perhaps simply not known. These cases represent a kind of anonymity, but identifiers known to be linked to specific individuals may be unimportant in many cases, such as where payments are properly made or where accurate checking of identity is not critical. If problems were to arise here, one will know which identifiers were involved and perhaps who issued them in the first

place. Some regulation of the issuance of identifiers and the coupling of them to key pairs will be important, as is regulation of other trusted entities in society (such as banks).

Once a means of obtaining identifiers for individuals and organizations becomes routine, similar steps can be taken for Internet resources of all kinds. Systems and services can be given identifiers and users can validate them as easily as they can validate the users. Although accurate audits of information requested and disseminated can be enabled in this fashion, it also has the downside of enabling unauthorized accounts of such activity. In a free society, the balance of privacy versus security comes squarely into play here and requires careful examination from both regulatory and political perspectives.

Assuming all Internet information systems and other resources (including users, networks and devices, as well as the actual information or services being provided) have associated unique persistent identifiers, how would the operation of the Internet actually function in this context? How would informational resources be accessed in this manner? And why would it matter for Internet defense?

Circumscribing the Operations

If the main vulnerability of today's information systems comes from the operating systems and the applications that make use of them, an important first question is whether either or both of them can be avoided or if it is possible to otherwise constrain the vulnerabilities in some fashion. For some applications, the answer is clearly no, since they are essential to providing the desired user functionality. Most applications currently depend on underlying operating systems for many tasks such as storing files, scheduling multiple tasks and handling security and network functions.

Vulnerabilities in the operating system pose direct threats to the application, yet many operating system functions will still be required. If some of the operating system functions are not really needed,

however, perhaps that software can be simplified and made less vulnerable to attack.

Most of today's workstations, desktop and laptop computers are installed with a suite of application software, including office-related software for document preparation, spreadsheets and more. Downloads from trusted vendors are the norm, but subject to the vagaries of the user's system. Access to remote sites, such as those on the Web, are typically enabled via a Web browser, where each website complies with standard Web protocols and vulnerabilities in the browser protocols can have repercussions for users of the websites visited.

Reliance on structured information in the form of digital objects is another way to circumscribe the operations, since one knows both the nature of the operations to be performed and the targets of those operations. Digital objects, whether embodying what is traditionally viewed as "content" or actions to be taken on that content (perhaps in the form of executable code for which trust mechanisms can be invoked) can easily be incorporated within the DO architecture to enable a scalable and evolvable system going forward.

The largest growth in computational facilities has recently been with wireless devices, such as smartphones and tablets, where the devices may not be intended for use as general purpose computing platforms; and user desired functions that are not already installed on these devices are enabled by obtaining vetted computer programs (applications or "apps") usually written by others. Such apps can provide services of their own, or enable access to other resources on the Internet. Users typically activate these apps by touching the screen on their wireless device or taking an equivalent action. These apps can be customized by their providers to give a unique experience either using the device or in connection with a remote service or interaction. Thus, suppliers of such apps are usually not constrained by the technology to

any single set of application protocols or means of presentation, but those made available with the user's device are often the most convenient to use. By this measure, the Web, along with the Web browser, is but one very pervasive app.

Apps, in general, may not require many services typically provided by an operating system. In this chapter, it is assumed that the operating system may be viewed as a mini-version of a combined traditional operating system with a high-level programming language, which we call "MyOp" for short. MyOp is assumed to provide a well known programming language execution environment, network access, maintenance of address books and/or mailing lists, the ability to select and schedule resources for execution and the ability to execute public/private key encryption and decryption. It is assumed that usual file and folder operations are replaced by use of a special purpose app that provides repository functions and uses either internal storage (if necessary), external storage (if available) and possibly both under certain conditions. Synchronization functions are not discussed here, but these could be embedded in MyOp or combined in the repository app.

MyOp is assumed not to be programmable by third party computer programs, and since apps cannot directly interact with other apps except by communicating with them via information structured as digital objects, this should limit the vulnerability from external threats to manifest themselves through unknown installed "hooks." It remains to be seen whether it will be possible to inhibit apps from permitting the execution of third party digital objects that are executable programs. If not, the use of specialized sentinel programs called "bastion objects" that cordon off the range of operations of such apps may be required. If a user can be aware of all the downloaded apps on his device, he can be made aware if an unwanted app were somehow to arrive. In any event, since he would have taken no action to cause it to be downloaded (of which he was aware), either his

system would detect it to be unwanted and take appropriate action or, somehow, his system would have had to be fooled into making such a request (or getting his system to think such a request had been made). All this is to explain how the discourse of dealing with threats and defense against such threats would shift from a wide unknown range of possibilities to a situation in which various types of attack scenarios can be better described and thus dealt with both before, during and after the fact.

No other actions are allowed by any app relative to MyOp, and further no app is permitted to interact with any other app except by passing identified information, referred to here as digital objects. So, temporary or permanent storage of digital objects takes place via the internal repository app or by passing the information to an external repository. Digital objects are constructed by the repository app, or by APIs (application programming interfaces) that make use of it, according to a meta-level standard and parsable structure understandable by apps throughout the Internet; a unique persistent identifier is also associated with each such digital object. Thus, all arriving and departing information is in the form of digital objects, and internally generated information that does not leave the local computational environment is also stored as one or more digital objects.

Information in the form of digital objects flowing over the component networks of the Internet can thus be individually identified along with all incoming and outgoing information from any device or other computational facility. Although there is no requirement that any part of this information, including its associated identifier, be made visible in the network, users may wish to make the identifier part of a given digital object visible for any of several reasons. One is that the provenance of the information can be made available when the information becomes available. Another is that users can require that references to responses from their systems include the identifier of each digital object being responded to for cross-correlation or validation on receipt. Coupled with

timestamps and use of public key encryption, this approach can also be used to validate individual steps in a series of transactions or other operations taking place during a single session.

Large server farms will have very different needs than an individual user's computational devices, but their level of expertise can be expected to be much higher as well. No matter what the level of expertise, however, if such server farms require more sophisticated operating systems and related services to support distributed computing (sometimes referred to as "cloud computing") within and among the servers in the farm, care will have to be taken to identify, isolate and hopefully remove latent system vulnerabilities. Internet-based server farms, particularly if they store large amounts of data, provide specific targets for potential attackers. Thus, a combination of local storage and remote storage might provide a reliable approach in the event of sabotage or denial-of-service. Normally, one might rely on remote storage for day-to-day operations and only use the then-current local storage choice in those cases for which the remote storage is unavailable. If the remote storage is disabled or destroyed, or cannot otherwise be brought back up for days, weeks or months (or longer), a user can temporarily resort to the user's local storage capability.

It is assumed these server farms can be operated both reliably and securely. However, users may wish to store their digital objects in encrypted form, with the keys kept separate from the remote storage site. In this case, operations with the remote storage site will likely be of the warehousing variety with entire digital objects being passed back and forth. When encryption is not required or is not invoked, operations with the remote storage can be more fine-grained, and specific elements of the digital object may be accessed directly or after performing one or more remote operations without the need to retrieve the entire object. Recent developments have shown that remote interactions with encrypted objects are also possible in certain

cases, but this aspect is not explored further in this chapter. In cases of very large objects, which would consume bandwidth and take time to transport, the ability to access directly specified parts of the object would have obvious appeal.

In each of these cases, the potential number of digital objects can be quite large and users cannot, and indeed will not, be able to remember their identifiers, even if they can recall attributes of the digital objects to which they were assigned. Software known as registries serves the purpose of allowing users to register such objects, presumably automatically in most cases and manually (if desired) in others. These registries can be installed as separate apps on the user devices, or provided by server farms over the Internet. In both cases, the registry metadata will be produced either manually by the user or automatically at the time the original digital object is created. Indeed, the user should be able to annotate such metadata and have it apply to the metadata pertaining to a specified range of digital objects.

If a user's device is lost, he may lose the apps that were available on it, but some vendor implementations should permit the user to access such programs over the Internet at no additional cost and inhibit the operation of that app on the lost device. At a minimum, this capability would seem to require each such computational device to have its own unique identifier, and perhaps be able to hear about such loss via MyOp; however, other means of disabling such apps are also possible.

In this model, the role of IP addresses would remain unchanged, along with the role of routers and networks that interpret them. In addition, those components would have the added advantage of using the digital object identifiers to meet stated objectives as well. The DO architecture can thus be integrated into the existing Internet as well as working in other communication systems. To clarify this point, in a proposed modification to the 1995 Federal Networking Council definition,

the Corporation for National Research Initiatives (CNRI) recommended adding the words "or integrated with" to the section that talked about applications layered on the underlying protocols.⁵

In an architectural environment where all accesses to systems, services and other resources are managed using identifiers for each such resource, and all information is structured in the form of digital objects, the task of Internet defense is altered in several fundamental ways. When operations in the Internet can be made more structured, one no longer has to be on the lookout for bit patterns whose purpose and intent cannot easily be determined. If, as a result, most actions consist of a more limited set of types of basic operations (which the author refers to as "meta-level operations" to reflect the fact that they indirectly reference the actions to be taken and their targets), it may be possible to develop protective steps that are more effective. This is definitely not the case today. If the digital object architecture were integrated within the Internet, its operations and targets would be separately identifiable so that, from these identifiers, the digital objects that were involved could be determined from the metadata, and the users could (if they choose) retain all the associated digital objects for later analysis (if desired). Many other properties of the communication could also be acquired, such as timing data for each digital object (e.g., creation, dispatch and arrival) should that be of interest. This is particularly important in connection with emerging Internet capabilities that relate information about "things" to other information in the Internet.

A user who is well aware of what is happening on his device will ordinarily be in a position to take manual action if necessary. First, he has to be paying attention, which may not always be the case. Second, an attack may have significant negative impact within seconds, or even microseconds. Thus, the ability of a system to respond in kind would seem to be essential. Efforts to develop cognitive systems that understand their environment, their own capabilities and modes of behavior, and threats to their

operation have been undertaken in the past; but the task has remained daunting by virtue of the many degrees of freedom posed by the general problem. In other words, there are just too many things to have to know about, look for and react to. With the digital object architecture, the number of possibilities is greatly reduced and, thus, the likelihood of success is potentially much higher. An environment where threats could be internalized within a system, and where the system can defend itself with mobile programs specifically tasked and authorized to take actions against fast moving attacks, would provide an immediate benefit to the user by defusing the attack in real time. It could also serve to provide data for a post-mortem report on the attack.

As a matter of policy, it would be useful if users can work with the involved carriers or other relevant service providers when such problems arise to determine what happened. This can be helpful in determining what networks, proxy servers or other related infrastructure or resources may have been compromised, and how best to thwart any such ongoing incidents. This would potentially have the effect of enabling legitimate backpressure or other corrective action wherever required in the Internet.

Conclusion

The digital object architecture would impact the nature of many Internet activities by making them more explicit and, thus, potentially more defensible against attack. It would help to support an informed discourse about implementation of effective Internet defense strategies that are difficult to achieve today. The continuing transition to the DO architecture is an incremental process that may take years to complete. In the meantime, considerable progress could be achieved (especially for users) in understanding what is transpiring on the Internet (including on their machines and devices), and working with Internet service providers to ensure that undesirable events can be more easily diagnosed and prevented, or at least detected and hopefully defused before they cause substantial damage.

ENDNOTES

1. Peter J. Denning and Robert E. Kahn, "The Long Quest for Universal Information Access," *Communication of the ACM*, Vol. 53, Issue 12 (December 2010): 34-36, <http://dx.doi.org/10.1145/1859204.1859218>. See also Corporation for National Research Initiatives, "A Brief Summary of the Digital Object Architecture" (1 June 2010), <http://hdl.handle.net/4265337/5041>.
2. Sean Reilly, "Digital Object Protocol Specification," Corporation for National Research Initiatives (12 November 2009), http://dorepository.org/documentation/Protocol_Specification.pdf.
3. See Robert E. Kahn and Vinton G. Cerf, "What is the Internet (And What Makes It Work)," Corporation for National Research Initiatives (December 1999), http://www.cnir.reston.va.us/what_is_Internet.html; and Robert E. Kahn, "The Architectural Evolution of the Internet" (17 November 2010), <http://hdl.handle.net/4265337/5044>.
4. U.S. National Coordination Office for Networking and Information Technology Research and Development, "FNC Resolution: Definition of Internet" (24 October 1995), http://www.nicrd.gov/fnc/Internet_res.html.
5. Patrice A. Lyons, "The End-End Principle and the Definition of Internet," Corporation for National Research Initiatives (10 November 2004), <http://www.wgig.org/docs/CNRNovember.pdf>.

REPRESENTING VALUE AS DIGITAL OBJECTS: A DISCUSSION OF TRANSFERABILITY AND ANONYMITY*

ROBERT E. KAHN & PATRICE A. LYONS**

This article discusses the use of “digital objects” to represent “value” in the network environment. Deeds of trust, mortgages, bills of lading and digital cash can all be represented as digital objects. The notion of “transferable records” structured as digital objects is introduced, along with references to its application in real financial situations. Even in a formal information system, anonymity reflects the desire of a holder of value to remain incognito, except as he or she wishes to be made known. The use of unique, persistent identifiers and a resolution mechanism to fashion such a capability for anonymity and transferability is presented.

I. BACKGROUND

A basic element in commerce is the representation of “value” by a writing, or more generally, a “data structure,” fixed in a tangible form such as paper. The use of such instruments is so ubiquitous that they are often taken for granted in daily life. A business will take delivery of a new computer, desk, photocopy machine or some other good and sign a

* An earlier version of this article was published in DLib Magazine (May 2001), at <http://www.dlib.org/dlib/may01/kahn/05kahn.html>.

** Dr. Robert E. Kahn is Chairman, CEO and President of the Corporation for National Research Initiatives (CNRI), which he founded in 1986 after a thirteen year term at the U.S. Defense Advanced Research Projects Agency (DARPA). Dr. Kahn conceived the idea of open-architecture networking. He is a co-inventor of the TCP/IP protocols and was responsible for originating DARPA’s Internet Program, which he led for the first three years.

Patrice A. Lyons serves as Senior Legal Counsel to CNRI. While serving as a legal officer in the Copyright Division of Unesco (Paris, France; 1971-76), she participated in the preparation of the Convention relating to the distribution of programme-carrying signals transmitted by space satellite; as a Senior Attorney in the Office of General Counsel of the U.S. Copyright Office, Library of Congress (1976-87), she was called upon to assist in the drafting of regulations to implement the cable compulsory licensing system adopted by the U.S. Congress in 1976, and played a lead role in the preparation of the Semiconductor Chip Protection Act of 1984. Ms. Lyons later served as a Partner in the communications law firm of Haley, Bader & Potts (1987-90), and is currently in practice in Washington, D.C. at Law Offices of Patrice Lyons, Chartered.

document acknowledging receipt without a second thought about the validity of the process being used. This is not a recent development. For example, data structures such as “bills of lading” were used in the thirteenth century.¹

A promise to carry loads of produce to a country fair centuries ago may differ from a promise to perform “operations” on material in digital form to produce a required informational result. Additionally, promises of centuries ago may also differ from a promise to deliver a digital object, embodying a literary or musical work. Even so, the instruments evidencing the contract of carriage, the right to possession of the goods, or the receipt by a customer of the product or service, have basic elements in common. The issue addressed in this paper is whether and how such elements may be appropriately represented in a way that frees the transaction from the need for a physical manifestation, while allowing for both anonymity and transferability.

Representing a transaction in the form of a digital object does not preclude the production of a corresponding physical artifact upon demand. However, whether such artifacts are in fact necessary at all would depend more on the perceived needs of the participants than on the validity and reliability of the underlying mechanisms that can produce it. Transferability is achieved if the data structure may be transferred with authenticity from the party in possession to another party using verifiable techniques. While transferability would require a third-party trusted system to facilitate the transaction, the third-party system would only serve as an intermediary in a technical sense, but would not need to know who the current holder of the object is or maintain any information about the transaction. Anonymity is achieved where the party currently deemed the “holder” of a data structure is not generally known, or cannot be known, without the consent of that party. With such a third-party system in place, each party to a transaction can demonstrate a legitimate claim to the data structure before and then after the transaction has taken place. If an adequate confirmation of legitimate possession after the transaction cannot be made, the second party would normally reject the transaction.

Although a tangible fixation of an object provides a relatively easy means of displaying the data structure representing the intangible “value” being provided, we consider here only the case where the need for such a physical artifact is no longer present. As discussed in a report prepared

1. See, e.g., PAUL HALSALL, *MEDIEVAL SOURCEBOOK: BILL OF LADING 1248* (1998), <http://www.fordham.edu/halsall/source/1248billoflading.html>; SPYROS M. POLEMIS, *THE HISTORY OF GREEK SHIPPING*, http://www.greece.org/poseidon/work/articles/polemis_one.html (last visited Oct. 1, 2006) (noting similar mechanisms employed in ancient Greek and Roman times); *RULES FOR ELECTRONIC BILLS OF LADING* (Comite Mar. Int'l [CMI]), <http://www.comitemaritime.org/emidocs/rulesebla.html> (last visited Oct. 1, 2006) (recent effort by the Comité Maritime International to develop Rules for Electronic Bills of Lading).

for the United Nations Commission on International Trade Law (UNCITRAL) Working Group on Electronic Commerce,² there have been many attempts over the last few years to replace traditional paper-based bills of lading by electronic messages, and more generally, what was termed the “dematerialization of documents of title,” particularly in the transportation industry.³ It was thought useful to expand such efforts beyond maritime bills of lading to encompass other modes of transportation, as well as issues involving “dematerialized securities.”

In the United States, efforts to develop alternatives to paper-based documents have given rise to the concept of a “transferable record.” Initially, this work was carried out under the umbrella of the National Conference of Commissioners on Uniform State Laws (“NCCUSL”). Section 16 of the Uniform Electronic Transactions Act (“UETA”) was approved and recommended for enactment by NCCUSL in all States in 1999, and sets forth the general parameters of the “transferable record.” In essence, this section provides for the creation of “a record created, generated, sent, communicated, received, or stored by electronic means,”⁴ i.e., an “electronic record” as defined for purposes of UETA, “which may be controlled by the holder, who in turn may obtain the benefits of holder in due course and good faith purchaser status.”⁵

A more restricted definition of a “transferable record” was enacted into law by the U.S. Congress.⁶ Title II, sec. 201(a) of what has become known as the ESIGN Act provides that the term “transferable record” is limited to specific types of “electronic records” such as loans secured by real property. As experience is gained in this area, and technical systems and processes are developed to support electronic equivalents of paper-based loan documents, steps may be taken to expand the scope of the law to encompass other representations of “value” in commerce.

The digital object architecture has been under development by Corporation for National Research Initiatives (“CNRI”) for a number of years and is currently being implemented in several commercial contexts. This architecture may be of relevance to the evolution of the notion of a transferable record for purposes of the ESIGN Act, as well as the ongoing discussions in the United Nations relating to the transfer of rights in

2. U.N. Comm’n on Int’l Trade Law [UNCITRAL], Working Group on Elec. Commerce, Note by the Secretariat, *Legal Aspects of Electronic Commerce* 2, U.N. Doc. A/CN.9/WG.IV/WP.93 (March 2001), available at <http://daccessdds.un.org/doc/UNDOC/LTD/V01/812/31/PDF/V0181231.pdf>.

3. U.N. Comm’n on Int’l Trade Law [UNCITRAL], *Report of the United Nations Commission on International Trade Law on its Thirty-Fourth Session*, ¶ 288, U.N. Doc. A/56/17 (June 25, 2001).

4. UNIF. ELEC. TRANSACTIONS ACT § 2(7) (1999).

5. *Id.* at § 16.

6. Electronic Signatures in Global and National Commerce Act, Pub. L. No. 106-229, 114 Stat. 464 (codified as amended in scattered sections of 15 U.S.C. and 47 U.S.C.).

tangible goods and other rights.

II. PHYSICAL ARTIFACTS

Many applications involving physical artifacts, such as health records fixed on paper, often raise the notion of an original or authentic copy. In fact, in many cases there may be multiple originals of the same document like a contract that is signed in duplicate originals. In other cases, only one original record may exist, as in bearer bonds or in deeds to real property. For some applications there is no requirement of anonymity. The holder of the original record may be known by any of several means. In other cases, the holder may be completely unknown unless and until he or she produces the physical artifact. This is the case for issued paper money such as a dollar bill. Although the issuer of the official record or document is generally known to the holder and to anyone else who is permitted to inspect it, there can, but need not be, any record of the actual holders in due course of the record over time. Furthermore, it is generally understood that physical artifacts such as paper or other material objects are not required to maintain certain official records. For example, the issuer of an official document may retain a computer record of the issuance. This might be known by any of several terms such as a book entry, or journal entry and the official record is kept by the issuer or a known designated agent of the issuer. The issuer may also maintain a record of the "chain of title" to the entry. Various registries maintain this kind of information, such as a typical Recorder of Deeds, although the actual deed may be retained by others. Still, the prevailing mode of operation is to issue paper for many, if not most, of these applications.

In each of the above cases where only computer records are used, there is usually a trusted party that maintains the records, as well as the linkages between each record and the party to whom the record is currently "attached." Absent the maintenance of accurate records by the trusted party, proof of ownership may be compromised, perhaps fatally. Even though an official computer-based record may be kept by a trusted party, normally the issuing party or its agent, a copy of the record may be available in digital form at other locations. In order for the record to be negotiable, the bearer may be required to provide the record in digital form, but the authenticity of the holder as well as the record can be separately validated if the appropriate records are available.

The discussion below focuses generally on the case where a record of linkages is not kept, and thus, no equivalent "chain of title" is maintained by the trusted party. It also assumes that a generalized record-keeping capability need not be in existence, but that a trusted means of authentication is available. The digital object architecture described generally below can play a key role in facilitating the authentication process.

III. DIGITAL OBJECTS AND THEIR IDENTIFIERS

The term “digital object” is used to denote an identifiable item of structured information in digital form within a network-based computer environment. Generally speaking, a digital object is a set of sequences of bits or elements, each of which constitutes structured data interpretable by a computational facility, at least one of the sequences denoting a unique, persistent identifier for that object. Information of virtually any kind that is represented in digital form may be structured as a digital object. The identifier of a digital object may be of any form, as long as it may unequivocally be de-referenced to the digital object. The Handle System[®] is an example of such an identifier system.⁷ Some known part of the identifier could contain a cryptographic hash or fingerprint of the identified object, which could be used to help to authenticate the object.

The Handle System being developed by CNRI, serves as a “resolution system” and would typically contain “resolution information” sufficient to resolve an identifier to the “location” of the computational facility containing the object. However, the resolution information, nominally state information about the digital object, may not necessarily be publicly available in its entirety. Indeed, portions of the state information may be available only to the party that is the current owner or “holder” of the object. The resolution system is also assumed to be secure from tampering. This is achieved through a combination of mechanisms including the use of public key infrastructure, backup procedures, and protected physical equipment. It need be no less secure than, for example, other parts of an on-line banking system.

The location, if designated in the state information, may be merely the service point for obtaining the digital object. In fact, there may be multiple locations that can produce the digital object, and for informational purposes, any of these will suffice. However, it is assumed that only one of these objects is the official version, and the rest merely replicas. This leads to an important consideration: given the ease by which information can be replicated by computer and on a network, how can the official version be distinguished from the other identical versions?

IV. TRANSFERABILITY OF DIGITAL OBJECTS

In this section, the focus is on the transfer of an authentic version of a record or document in the form of a digital object. We begin by considering how a given digital object accessible on the network can be authenticated as having the proper information from the original issuer and possibly contain additional chain of title information where appropriate. The

7. The Handle System, <http://www.handle.net> (last visited Sept. 17, 2006).

possibility of encrypting each digital object may indeed be desirable for all of or parts of a digital object, especially where classified information comes into play. However, this capability is not essential to the basic system in which it is only assumed that the digital object is signed by its issuer using a strong encryption mechanism such as the U.S. federal digital signature standard. The authenticity of the digital object can then be verified directly from the digital object and its signature, if the signature can be assured. The use of a trusted public key infrastructure is one, but not the only way to achieve this result.

The Handle System can store digital object signatures to be used for authentication, and even bind the signatures tightly to the identifiers. The digital object will generally contain other information that can be used to show authenticity, but this is not necessarily required. For example, the inclusion of a sequence number, date-time stamp and/or the length in bytes would inhibit attempts to tamper with even weak signatures, or strong signatures made weak over time with increased computer power.

The question of determining which of N authentic digital objects is the original is, in some sense, an epistemological question since there is no way for a computer to know where a party providing bits to it "obtained them." If all instances of a digital object are identical and since bits are themselves fundamentally incorporeal there is really no notion of original bits. For purposes of illustration, four transferability mechanisms are identified below. The first two are equivalent to physical artifacts embodying data structures. The third is a hybrid situation. Only the fourth will be discussed in any detail.

Mechanism one is a tamper-proof device provided by the original issuer that contains the original information. It is assumed that the issuer only issues one such device, that others cannot replicate the device without destroying some critical part of it, and that no means exist to change the original information (although it may be possible to incorporate additional signatures to reflect chain of title). The device thus assumes the role of paper and ink and, for most purposes, can be viewed as equivalent to paper and ink. One transfers the data structure by transferring the physical device. Mechanism two is like mechanism one, in that the above assumptions apply except that the internal information may be read out of the original device and into another device. Assuming a means by which there is no possibility for corrupting the information in the transfer process (e.g., the receiving device will reject corrupted information), this leads to the issue of whether the receiving or sending device can insure that only one such transfer can occur. There may be cases where, in fact multiple transfers might be appropriate, but this possibility is not addressed here. Mechanism three is like mechanism two, except that one of the devices is not tamper proof. This would have to be assumed if one of the devices were a general-purpose computer. The techniques for ad-

dressing mechanism three are essentially the same as those that would be used if all the devices were general-purpose computers; and so we go directly to the fourth case.

V. DISTINGUISHING ORIGINAL INFORMATION ON THE NET

Mechanism four assumes that the original information is structured as a digital object and stored in a general-purpose computer or other computational facility on the net. The notion of “holder” is tied to the notion of unambiguously designating the computational facility that purports to hold the original digital object. For example, a transferable record such as a deed of trust could be the original digital object held at a particular moment in time in such a computational facility (referred to in this paper as the “holder facility”). While recognizing that this is a logical construct, the holder facility may be deemed generally equivalent to the evidentiary role played by a physical object. The evidentiary showing could entail demonstrating how the system works. For example, the showing could identify the particular holder facility as the authorized holder at a particular moment in time and producing the relevant digital object using the system. The identifier uniquely identifies the data structure stored within the designated holder facility. For an individual to claim to be the holder in due course of an electronic record structured as a digital object, the holder facility must be able to present the record to the appropriate party or parties for inspection on demand. It is asserted that only the authorized holder of the original digital object will be able to cause the desired object to be produced by the holder facility (unless, of course, it was trusted for safekeeping with untrustworthy associates). For example, if the holder was untrustworthy, it could present the material to a third party and claim it was holding the digital object on behalf of someone other than the party who is the authorized possessor.

The holder facility must be known to the resolution system, or a means of determining the holder facility must be uniquely derivable from the resolution system. While information about the holder of a transferable record need not be made available to others, the actual holder facility containing the object may also not be known publicly. However, it is mandatory that each holder facility only provide the original digital object to the bearer or his agent and in a form that allows the authenticity of the information to be verified. This can be achieved without the resolution system knowing the identity of the holder. In this case, the agent of the bearer might be a trusted computer system or its operator. A compromise of this trusted system would be equivalent to a loss of say a bearer document. A compromise of the resolution system could also result in a loss of such a document, but the latter compromise must be addressed on a system-wide basis. The former compromise (of a specific

trusted system) would be the responsibility of the bearer that selected it.

Each digital object can be validated by use of its fingerprint or signature, which is maintained by the issuer or its agent. The issuer may also elect to retain a replica of the original object, or only certain archival information about it such as its digital signature, length, date-time stamp of original issue, and possibly other non-personal identification information, such as sequence numbers. A transferable record itself consists of the original digital object and its signature, possibly along with additional information such as chain of title information added each time the object is transferred to another party. Certain elements of the additional information would be necessary for some objects and not for others. For example, bearer bonds would not usually have chain of title information, nor would digital cash. At the time of transfer, an instance of the digital object would be formed in a new holder facility corresponding to the new holder and the system would require that a change in the state information indicating the then valid holder facility be entered into the resolution system.

The Handle System has all the attributes necessary to provide the functionality of a trusted third party system. Specifically, system responses may be "signed" by the system upon request and each signature may be authenticated by a built-in certificate authority, if desired. The built-in certificate authority may itself be certified on a system-wide basis, and the cryptographic strength of the certificate authority increases as its purview widens. For example, the system-wide authority has the longest and strongest key. Each entry into the Handle System requires the use of a private key known only to the owner or its authorized agent. Further, various cross-checks carried out regularly within the system are designed to detect anomalies with respect to replication and mirroring of data. The top level of the Handle System is known as the Global Handle Registry and consists of a number of servers and services managed by a single trusted authority.

Entries in the Handle System for a newly designated holder facility would be made by the authorized holder at the time of transfer; the identifier for the data structure need not change, but the corresponding information in the Handle System would be changed to indicate that the data structure is now accessible from the new holder facility. It is not required that the entire Handle System be trustworthy in order to implement this capability. It is only required that a subset of the system be trusted, namely, a subset separately cordoned off to manage objects of value in which transferability and/or anonymity are needed.

VI. DIGITAL OBJECTS SENT VIA E-MAIL AND/OR AGENTS

Digital objects structured as mobile programs or software "agents"

may serve as their own transport mechanism or be used to transport other digital objects with appropriate access procedures to effect the authorized disseminations. Existing mechanisms such as email may also be used for the same purpose. Specifically, both email and agents may be viewed as ways to move the separately identifiable information contained within them, but these would not be an integral part of the Handle System *per se*. While in transit, the information may or may not have any status of value until and unless it arrives at its proper destination and is validated. Alternatively, the use of identifiers, such as handles, can obviate the need for an actual data structure to be communicated as the data structure can be retrieved independently if the ability to access it at a remote holder facility is enabled. If desired, a synchronization mechanism, familiar in distributed data base technology, may then be invoked to insure the designated object is moved from one holder facility to another and that only one such facility is the newly designated one. The Handle System can also provide the equivalent of this function. At that point, an email reply could go back to the sender confirming the transaction. For audit purposes, the reply itself could be structured as a digital object with its own unique identifier.

The case of network-based agents is in many ways the more interesting and also more complex topic. In this case, the value represented by a digital object may be present entirely in a mobile context, with the object never stopping at any computational facility for more than a transitory period of time. Interactions involving value transactions may thus take place in arranged meetings and rendezvous situations. Validation of the agents as well as their contained data structures and/or identifiers would be necessary. This could be carried out using the same techniques as for any other type of digital object, whether stationary in a repository or in transit on the net.

This paper does not purport to fully describe, much less specify, an entire system for representing value. There are many other issues remaining to be worked out on the way toward creating a viable system for identifying value based on the notion of a digital object. A starting point down this road would be the development of a general “type framework” for transferable records. The capability for such a mechanism exists in the current implementation of the Handle System. The notion of typed data, inherent in a digital object, is deliberately intended to be an open and extensible attribute of the system. If the digital object architecture were introduced in various areas of commerce, it would be possible to agree on specific “types” that are meaningful for specific subjects or industries. There may be multiple types for representing “value,” such as a category called “bill of lading” or “deed of trust.” A data structure would be assigned a “type” for purposes of resolution of digital objects that are designated by an issuer as conforming to the particular type. Types may

also be defined dynamically and resolved by the resolution system. Once agreement is reached on the use of "types" in such a system, consideration may be given to identifying possible standard operations allowed to be performed on a given type. For example, where dealing with the type: "transfer of copyright ownership," there may be a permitted operation: deposit for recordation in the Copyright Office.

While various notions concerning "value" and "typed data" require additional study in the network environment, the basic underlying resolution system, already in operation in Internet commerce, may be used directly to resolve typed data and to manifest value. The flexibility of a system based on the notion of a digital object may serve to open new avenues of commerce in a networked environment and contribute efficiencies and cost savings to existing methods of doing business.

Trusting Digital Entities

by
Robert E. Kahn
Corporation for National Research Initiatives
Reston, VA

117

A presentation at the APAC Blockchain Conference
Melbourne, Australia
March 13, 2018

What is the Internet?

- **Original FNC Definition still applies:**
 - Global Information System that makes use of IP (its logical extensions or follow-ons), TCP (its logical extensions or follow-ons, and other IP compatible protocols), and which supports applications based on the above.
- **Overall Architecture is still intact despite increases in the underlying technology by factors of 1 – 10 Million (computation, communication and storage)**

Bindings to Technology vs. Information

- Arpanet – 16 bit addresses → wires
- Internet – 32 bit IP addresses → machines
- Web - URLs → <IP Address/filename>
- DO Architecture – DO Identifiers → DOs
 - DO Architecture describes a means of managing information over both short and long time frames in which Digital Objects are the basic structures.
 - Compatible with the current Internet and builds upon it.

Fundamental Properties of the Digital Object (DO) Architecture

120

- Logical Extension of the Internet
- Based on the same architectural ideas embedded in the Internet's architecture, and which have sustained its evolution, the three most important characteristics being:
 - **Open Architecture** (defined protocols & interfaces)
 - **Independence** from the underlying technology
 - **Minimized Complexity** for users
- The DO Architecture enables interoperability across heterogeneous information systems, whether in the Internet or not.
- It is a non-proprietary architecture and is publicly available.

Basics of the DO Architecture

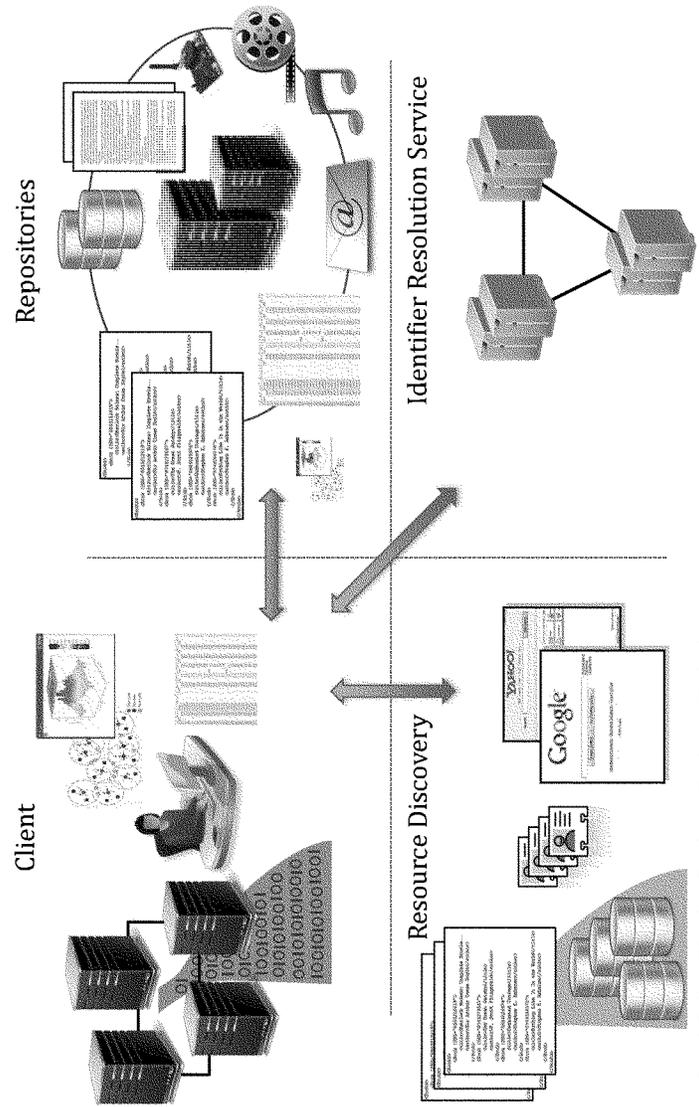
- **Digital Objects** are its basic structures (also known as Digital Entities). Each DO consists of information represented in digital form, and having an associated unique persistent identifier.
- The DO Architecture consists of three components:
 - a resolution component that resolves identifiers to “state information” about the desired information - - a resolution request yields a handle record;
 - Repositories that store DOs and enable access via their identifiers; and
 - Registries that store metadata about DOs and are used for searching.
- Resolvable “data types” are critical to understanding a DO by computer or otherwise.

ISO Effort to define the structure of

Types

- Not intended to define specific “types”
- But rather what a type specification should look like, where each type is represented as a separate DO with its own unique persistent identifier
- Every element of every DO consists of a pair of (type, value) entries.
- And every type is represented by its identifier.

Digital Object Architecture: Information Management on Networks



Search Engines, Metadata Databases, Catalogues, Registries, etc.

Critical Role of Identifiers in the DO Architecture

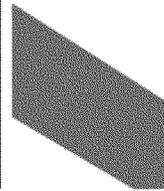
- Identifiers are used to designate users, system resources, networks, services and desired information of all kinds represented in digital form and structured as digital objects.
- The resolution system provides important real-time information to client software.
- Everything being identified has a public/private key pair; and the public key is accessible by resolving its identifier.
- This enables an integrated PK Infrastructure (PKI) that is essential for purposes of providing security and generating trust.

More Background on DO Architecture

- Started with the work of Bob Kahn and Vint Cerf at CNRI on mobile programs in the 1980s (i.e., Knowbots)
- Elaborated upon in the early 1990s in the Computer Science Technical Reports (CSTR) project.
- In 1997, the Cross-Industry Working Team (XIWT) supported the concept of digital objects and “stated operations” on digital objects, and noted the importance of chaining operations and managing value.
- The DO Architecture received the Digital Id World Award in 2003 for balancing innovation with reality.



*Cross-Industry
Working Team*



Managing Access to Digital Information: An Approach Based on Digital Objects and Stated Operations

126

*3Com
Alcatel Telecom
American Management Systems
Apple Computer
AT&T
BBN
Bell Atlantic*

May 1997



DOIP Protocol

- The Digital Object Interface Protocol (DOIP) is a simple, but powerful conceptual protocol for software applications (“clients”) to interact with “services” which could be either the digital objects or the information systems that manage those digital objects.
- The DOIP enables a user (or another DO) to interact with a DO based on the use of associated identifiers
 - Each action is represented by a DO; and the interface conveys the action’s identifier (ID1);
 - Each target of an action is also a DO; and the interface conveys that identifier as well (ID2);
 - The formal specification is written as a schema that is incorporated in a program typically run by a repository that serializes structured data.

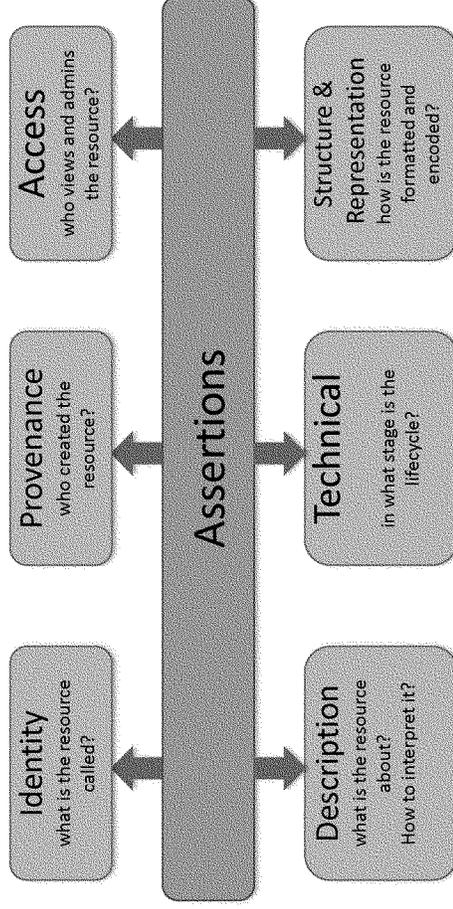
Framework for Discovery

ITU-T Recommendation X.1255

- Based largely on the Digital Object Architecture, ITU-T Recommendation X.1255: “Framework for discovery of identity management information” was approved in September 2013.
- Focused specifically on discovery and access to information in digital form, structured as digital objects, X.1255 is applicable to operational requirements for information management more generally.
- For purposes of X.1255, a digital object is defined as a digital entity; and the Recommendation describes a data model and interface protocol.
- Since the notion of a DO and a DE are nominally identical, from here on I will simply refer to both of them as Digital Entities (DEs).

What is Metadata

- People commonly define metadata as “data about data”
- A more complete definition:
 - Metadata is a set of (structured) assertions about an entity
 - Multiple parties may make those assertions
 - Veracity of those assertions is usually outside the scope of metadata
- Those assertions could be about



What is a Block?

- Blocks are not new!
- Historically, a block was viewed, essentially, as a sequence of bits, usually with a defined beginning and end.
- In the past, it may or may not have been uniquely identified other than by its arrival sequence in time.
- Blocks were also linked with other blocks
 - In the programming field, blocks were often linked or chained using pointers
 - In the communications field, they were usually linked in some time sequence and often involved encryption.
- Blocks were not usually managed separately from the application that invoked them- -but they could be.

Accessing Information about a Block

- This is the province of what is now called **metadata**.
- Part of the metadata may be self-contained within the block. This is sometimes referred to as key metadata.
- The amount of information that may be termed metadata about a block can be enormous and would normally be managed separately from the block.
- The use of blocks in information management was pioneered by CNRI in connection with mobile programs in the Internet.

General Observations about Blocks & Blockchains

- As previously mentioned, the context for the development of blockchain technology has been around for many years. Indeed, every block is an example of a Digital Object.
- A blockchain represents a particular way of structuring a Digital Object that comprises multiple DOs.
- DOs are stored in Repositories, which may be replicated (otherwise known as mirroring).
- Various mechanisms can be invoked to cross-check the multiple repository entries, if deemed necessary to augment trust.

Managing Mutable & Immutable DEs

- Blocks may be immutable; blockchains may not change, but are inherently mutable as they need to change when they are updated.
- Examples of blocks are transactions, contracts, bills of lading, digital cash; for example, see “Representing Value as Digital Objects: A Discussion of Transferability and Anonymity.”
- Immutable objects can be authenticated without reliance on external parties.
- Mutable objects rely on external mechanisms to validate.

Authenticating a DE

- If a DE has been signed by a user or a system resource with its private key, the DE will identify the signatory (or it will be conveyed separately by the access protocol); and the DE can then be validated from the signature.
- Parts of a DE can be signed or encrypted, if desired.
- If so, the usual approach is to treat that part as a separate DE in its own right and link to it from the other DE by including its separate identifier in the first DE.
- Alternately (and often in addition) the handle record obtained by a piece of client software for a given DE will contain the authentication information for that DE; and the handle record may be signed by the server from which it was obtained.

Authenticating an Immutable DE

- If a DE is known to be unchangeable in a given context (e.g., a contract, digital cash, or bond), then a simpler mechanism is available to authenticate the DE.
- Namely, the identifier can contain a powerful cryptographic hash of the DE (with an associated methodology to use the hash for authentication) so that the DE can be validated no matter how it is provided or obtained.
- In this case, the user need not rely on anything other than the strength of the encryption mechanism.
- And, a failure in any one case (perhaps due to a loss of a private key) will not compromise the rest of the system.
- Replication of the DE in multiple instances of a repository will increase the likelihood that a valid version of the DE can be accessed in the unlikely event of repository failures.

Trusting the Resolution Mechanism

- A key part of the Internet is the IP Addressing mechanism that is used to route packets from source to destination.
- Similarly, a key part of the DO Architecture, which is a logical extension of the Internet, is the identifier/resolution component.
- It is used to map identifiers for *digital entities* to useful state information about them: that system must be trustworthy as well.

Two Stage Resolution

- CNRI implemented a two stage resolution system in the early 1990s in which identifiers have the structure “**prefix/suffix**”. This implementation is in widespread use with more than a billion digital entities identified. The prefix, which is allotted to a specific party that wants to create resolvable identifiers, is unique to that party; and that party would start its identifiers with its prefix and add whatever suffix it wishes.
- Derived prefixes may be created by the party using a “dotted” convention. For example if prefix 35 is allotted, 35.1 or 35.HQ.1 may be derived from 35. The zero and one delimiter prefixes are retained in a distributed registry called the Global Handle Registry (or **GHR**). Multiple organizations around the world operate the GHR and coordinate with each other in maintaining its integrity.
- The actual identifier records - such as those corresponding to 35.1/abc - are retained in one or more local services they run, or contract to have run for them, and also managed by the party that created them. The system is inherently distributed. The local services can also be mirrored for reliability and security, as desired; and most organizations choose to do so.

What changes are in process?

- Fundamental changes took place in the Internet as the number of devices exceeded what were then a staggering number – like 100 Million.
- Today, it is envisioned that the number of devices in the IoT (or cyber/physical systems more generally) may come close to 100 Billion in the not too distant future.
- This will stress almost every aspect of the Internet - and especially those that involve information management.
- Many organizations are rallying behind the use of blockchains to provide trust, but this is but one of several alternatives; its use will provide its own challenges for managing information. Issues of interoperability as well as scalability, efficient performance and graceful degradation must be balanced against the need for architectural changes to provide enhanced defenses.

As the Internet Confronts Increased Complexity

- Mobile program technology may soon be needed in the context of implementations of the DO Architecture.
- Trust in the system of information management and the digital entities it manages are critically important especially when the DEs have value, as is the case with crypto-currencies.
- The need to protect rights, values and other interests that may be embodied in DEs, coupled with the sheer volume of information that will be available in digital for, requires a new paradigm for information management.
- The Digital Object Architecture can provide a sound basis for moving forward.

Some Background Reading

- Kahn, Robert E., Vinton G. Cerf, "An Open Architecture For a Digital Library System and a Plan For Its Development," The Digital Library Project Volume I: The World of Knowbots, (DRAFT) March 1988, <http://hdl.handle.net/4263537/2091>.
- Kahn, Robert E., Robert Wilensky, "A Framework for Distributed Digital Object Services," *International Journal on Digital Libraries*, (2006) 6(2): 115-123, https://www.doi.org/topics/2006_05_02_Kahn_Framework.pdf. (First published by the authors May 13, 1995, "A Framework for Distributed Digital Object Services", <http://hdl.handle.net/4263537/5001>).
- Managing Access to Digital Information, Cross-Industry Working Team, May 1997, <http://www.xiwt.org/documents/ManagAccess-1.pdf>.
- Denning, Peter J. and Robert E. Kahn. "The Profession of IT: the Long Quest for Universal Information Access". *Communications of the ACM*, December 2010, Vol. 53, No. 32, pp. 34-36, <http://doi.org/10.1145/1859204.1859218>.
- ITU Recommendation X.1255, "Framework for discovery of identity management information," was approved on September 4, 2013 (the work is based largely on CNRI's Digital Object Architecture; and Robert E. Kahn, CNRI's President, served as Editor), <http://www.itu.int/ITU-T/recommendations/rec.aspx?rec=11951&lang=en>
- Kahn, Robert E., "The Role of Architecture in Internet Defense", America's Cyber Future: Security and Prosperity in the Information Age, Center for a New American Security (CNAS), Volume II, Chapter XII, May 2011, http://www.cnri.reston.va.us/papers/CNAS_CyberSecurity_kahn.pdf.

Some Background Reading (Cont'd)

- Braswell, Jefferson, Lannom, Larry, Milne, Alistair, Northey, Jim, Paskin, Norman and Traub, Ken, "Response to the Financial Stability Board's Request for an Engineering Study on the Best Approach to Managing the Structure and Issuance of Legal Entity Identifiers (LEIs)," (2012), <http://doi.org/10.2139/ssrn.2197269>.
- Kahn, Robert E., and Patrice A. Lyons, "Representing Value as Digital Objects," *Journal on Telecommunications & High Technology Law*, Vol. 5, Issue 1 (2006), http://www.jthtl.org/content/articles/V5I1/JTHTLv5i1_KahnLyons.PDF.
*A patent application (US 20030233570 A1), titled **Authenticating and using digital objects**, and based in part on the ideas expressed in this article, was filed by CNRI. It specified that the technology may be applied in managing, inter alia, the issuance and authentication of financial instruments). The application was later abandoned when the claims were rejected by the U.S. PTO as covered by the now expired, CNRI Patent No. 6,135,646. **System for Uniquely and Persistently Identifying, Managing, and Tracking Digital Objects.***
- Lyons, Patrice A. and Robert E. Kahn, "The Handle System and its Application to RFID and the Internet of Things", in *RFIDs, Near-Field Communications and Mobile Payments: A Guide for Lawyers*, edited by Sarah Jane Hughes, Cyberspace Law Committee, 2013, <http://hdl.handle.net/4263537/5046>.
- *Definition of "block" as a "Digital Entity,"* Contribution of Corporation for National Research Initiatives (CNRI), DLT-I-048, ITU Focus Group on Application of Distributed Ledger Technology (Bern, Switzerland; 2018), http://www.cnri.reston.va.us/documents/DLT-I-048_CNRI_Contribution_FG_DLT.pdf

The CHAIRMAN. Dr. Kahn, thank you. I appreciate what you have shared with us and look forward to reading these various papers as well.

And to the other members of the panel, thank you. Very informative this morning.

Dr. Kahn, I am going to come back to you on the issue of trust, because I think that is something that is so integral to what we are talking about here.

But before I do that, I would like to ask several of you to touch further on the impact that we might anticipate, just from a consumer perspective, on electricity rates and the concern that some might have that, I am not, my family and I might not be ones that benefit from blockchain or bitcoin, and yet I am wondering am I, through my rates, going to be expected to pay for this infrastructure?

Can we have a little bit of discussion about, again, expected impacts on electricity rates? How we deal with consumers who are concerned about what they may consider to be paying or helping to effectively subsidize some of the costs that we build out of infrastructure?

Mr. Skare? Mr. Golden?

Mr. SKARE. Yeah, I think that you're hitting upon a fundamental aspect of how the power grid works. So anytime load increases, the only choice the utility has is either to generate more electricity or to import more electricity from its neighbors.

So when you're looking at a situation like this with the cryptocurrency mining is increasing the loads, while that would be at a very localized level in the distribution part of the power system, it still will lead to increasing the need for the generation.

Now whether the utility itself hits its limits of generation that it can provide, that determines typically and economically whether they should buy electricity from their neighbors.

I think if you take a look at the written testimony from the Chelan Public Utility District, they chronicle some of their interesting issues that they've had where they've declared some moratoriums on cryptocurrency mining and then look at the process of what is the impact on their grid and since they're a public utility, understanding what's the right way for them to address the issue to get to a policy that works. They thought they had it at one point and then they had to re-apply the moratorium when they found out their policy wasn't quite complete enough.

The CHAIRMAN. I think this is part of the complication here is understanding how you prepare for this short-term, mid-term, long-term. I mean, what is the long-term future here? Do the utilities build out for that or is this a shorter-term interval but you have an aggressive investment up front and then several years from now you might not necessarily need it?

I think these are some of the fears that I am hearing in terms of how do you address the demand, right now, but not knowing what this may look like in the years ahead, and in the very, very short-term.

Mr. Golden, do you want to or care to comment?

Mr. GOLDEN. Yes, Senator.

I think, like Paul mentioned, Mr. Skare mentioned, it's fundamentally, it's a supply and demand question, right? If obviously demand goes up then and supply doesn't keep up with it then you have a possibility of increased rates.

But I think it's important to note that utilities also have pretty robust planning processes in place to ensure that they can, sort of, understand where the grid is going, how much excess capacity in energy they have to serve new customers and work in partnership with new customers to determine, you know, how much load they're going to actually be generating for the grid. I think, so, fundamentally, I think there will be a chance for utilities to have that conversation.

And you mentioned building infrastructure at the beginning. I think many utilities will have, or power providers will have, the ability to have that discussion at the early outset of the load coming into their territory and talk about having them help pay for some of the costs associated with actually building out the infrastructure required. So I think making that smart decision and having those early discussions up front will help to alleviate some of the concerns when it comes to cost, some of the costs that may occur if the technology were to leave town.

The CHAIRMAN. Ms. Henly has mentioned that the United States is behind Asia, behind Europe in terms of just how we are approaching cryptocurrency.

There are some out there that would suggest oh, this is just the latest fad. It is a hype. It is going to be here today, gone tomorrow. It sounds like you all believe that this is very much a part of the future going forward. Is that a fair statement?

Shaking heads yes or no? I know Ms. Henly is a yes, for sure.

Mr. Kahn, you have been around for a long time observing this. Is this here to stay?

Dr. KAHN. As I said in my testimony and I strongly believe, this is one of the options that one ought to look at, just like when we did with the original internet there were all kinds of options from networking and computing. People should make their choices based on what's available.

This may be one that lasts. It may not be. I certainly would not argue either for or against it other than to say it's an option on the table. Figure out whether it works for you.

The CHAIRMAN. Very good.

Dr. KAHN. We'll see.

The CHAIRMAN. Very good.

Ms. HENLY. And just to add one more point on this.

I think it's definitely an option on the table, and I think what we are excited about is the potential. But I think right now it's so early that what, you know, we're seeing is that it's deserving of more research but as Mr. Kahn mentioned, it is, you know, in the grand scheme of investments in the electricity sector and cybersecurity, only one tool in the toolbox.

The CHAIRMAN. Very good.

Senator Cantwell.

Senator CANTWELL. Thank you, Madam Chair.

I really didn't think it was possible to thank you for a hearing in August—

[Laughter.]

—but thank you, thank you. Thank you for this hearing. I think the witnesses have done an excellent job of outlining the options here.

I would like to hone in on the security aspect, because I think that's one of the most interesting just because it is also one of our biggest challenges today.

Mr. Skare, you mentioned patch management which I am very intrigued with given the Equifax situation. My head just explodes when I think about the fact that they had an Apache patch but somebody in the organization just didn't apply it. And the notion that this technology could help us with the architecture on patching which is a lot more frequent than everybody thinks, right? How would that work?

Mr. SKARE. Well, one of the challenges with patches is you want to understand and be able to validate the provenance of a patch to make sure that the patch hasn't been altered from the time it was created until you're the person applying the patch.

So this is a way to provide a chain of custody, as it were, for that patch as it leaves the manufacturer until it gets to the asset owner who will be applying the patch. And I think that's one of the interesting things to help validate that no one has tampered with the patch is an important piece of this.

Senator CANTWELL. But would that also help us get patches implemented faster and more efficiently?

Mr. SKARE. No, this is a way of getting them implemented more securely.

Senator CANTWELL. Just to authenticate, you are saying. Just authenticated?

Mr. SKARE. Yes.

Senator CANTWELL. Interesting.

Which could cut down on the posers who are online posing as patches, right?

Mr. SKARE. Yes.

Senator CANTWELL. So, okay, I definitely think this is something to consider after the Equifax breach, just the amount of software that is going to be in our system, the amount that people are going to depend on and then the amount of updates and patches. Obviously, figuring that out, leaving that many people exposed just because a patch was not fixed at Equifax, is just mind blowing.

Ms. Henly, I think your testimony was quite helpful in the sense that it just reminded me of the 1980s when Microsoft said "a computer on every desk." Obviously we have come a long way since that motto from a company, but the notion that you are all discussing a digital ledger, you know, the computer as a digital ledger is really something, I think, for us. I think what we have to do is not overregulate here and make sure that we are continuing to invest in what those technology applications are, the level of efficiency that you could get from that, particularly on the energy side of peer-to-peer is very interesting.

What do you think we need to do to keep moving forward?

Ms. HENLY. I think that you make a great point.

There is a lot of promise for this technology. Some of the funding has already gone in to explore the cybersecurity benefits with PNNL and the DOE, I think it's a really good start.

I think that there are other programs that the U.S. Government can support—research, in particular—to answer some of these questions around energy use, but not only energy use, cybersecurity and not only cybersecurity, other applications of the technology in the energy sector. So what I would recommend is increased research and development, and coordinated research and development is a sign that the government and DOE is interested in the technology and wants to see the promise of it in this sector.

Senator CANTWELL. Mr. Kahn, have you heard about aviation applications for digital objects that work in blockchain?

Dr. KAHN. If they're applicable, I think they're applicable almost anywhere, so sure in the area of aviation. But you know also, autonomous vehicles on the ground, in the air, linkages between them, interoperability.

Some of the biggest challenges that we have in dealing with information systems is getting interoperability with other information systems. So we need a sound basis. The internet was all about getting interoperability between information systems.

Senator CANTWELL. Well—

Dr. KAHN. The computers.

Senator CANTWELL. I think one of the things we're interested in, obviously post 9/11, is making sure that people don't take over aircraft. One of the applications is to have this network be able to help with aviation if somebody is trying to hijack or take control of a plane, to have this kind of secure system that would have the plane land with this kind of architecture.

I think there are lots of applications, as you said, but I think there are some very specific ones that we should look into.

Dr. KAHN. And security is particularly important. I really think that by dealing with information at this large a level of granularity than we have before rather than just worrying about bits floating around the internet, we have all kinds of potential at our fingertips for doing a better job on managing security.

Senator CANTWELL. Thank you for saying that, very well put.

Thank you, Madam Chair.

The CHAIRMAN. Thank you, Senator Cantwell.

Senator Cassidy.

Senator CASSIDY. Thank you, Madam Chair.

Congratulations to you all in making something very technical something I can understand.

I am going to take this conversation a little bit afar from energy but remain with the blockchain technology.

I am interested in trade-based money laundering. Now trade-based money laundering, ideally, would be combated by having the people on both sides of the transaction. For example, I learned of a transaction in which the goods went from the United States to Guatemala but the invoice to Panama and then back to Guatemala and that interlude in Panama substantially changed the invoice so that they were able to misinvoice and, therefore, transfer dollars. Everybody with me so far?

I am gathering, and I will direct this to you, Mr. Narayanan, that blockchain, a public blockchain, could be maintained by the parties but with a central authority, I think Ms. Henly referred to, it would be a transparent blockchain. People could be looking at it, a central authority, to make sure it is not changing between Panama and Guatemala and you could trace this transaction throughout, ideally of course, to combat misinvoicing. Is that a correct assumption?

Dr. NARAYANAN. With public blockchains, there has been a tension between the transparency of the blockchain that is all of the data being out there for anybody to look at and trace as well as the anonymity or pseudonymity of the system which is that for participants to trade, using these blockchains they don't have to put their real name out there.

And so, in my—

Senator CASSIDY. Now let me stop you.

That may be current, but could you set up a system, as Mr. Kahn suggested, with preexisting rules that when it comes to international trade, yes, you would have to say that it is Rob Portman, Inc. sending a good to Guatemala and there would be some sort of bar code scan that uploads the manifest. But nonetheless, a central authority in each government could look and make sure that the invoice remained constant throughout.

Dr. NARAYANAN. Yes, Senator, you're absolutely correct. That system would be set up.

It would have to be accompanied by legislation and enforcement to make sure that people are using those regulated blockchains instead of ones that are harder to—

Senator CASSIDY. Now let me ask you again. I am directing this to you, but anyone can weigh in. How difficult would this be because I think Ms. Henly spoke, or one of you spoke, about what we are looking at now is beta chain, beta versions, but it continues to evolve. If we wish to put in that system now, could we put in that system now or no, the technology is still evolving?

Dr. NARAYANAN. If I may answer that?

I think technologically we're at a point where we can deploy those systems. I think—

Senator CASSIDY. Then let me ask, because obviously Guatemala would not have the resources of the United States but, is it possible to cloud base this so that we would, if you will, distribute, the U.S. could distribute this system to our trading partners and they could have authority? You would absolutely have to have authority, but nonetheless, they could participate in this function?

Dr. NARAYANAN. Senator, I think that's the challenging part which is not the technology side of things but instead how do we interface with our partners? How do we get everybody on the same table? I'm not the expert to speak about that, but what I will say is that the technology is not the hard part. It's all of these other things.

Senator CASSIDY. Ms. Henly?

Ms. HENLY. Just to add.

I think one of the great innovations of bitcoin was creating an incentive structure that aligned everybody's incentives around reinforcing the security of the system.

And so, while it's possible to create a structure and an application that you're describing that, you know, you'd have to essentially mandate or control or regulate every step of the way and ensure everybody is acting correctly.

But it's possible with a public blockchain network is to create a system of incentives so that everyone is incentivized to act and verify—

Senator CASSIDY. Let me ask because, theoretically, the incentive is one, to cut down on trade-based money laundering.

Ms. HENLY. Yup.

Senator CASSIDY. But also, to include tax revenue, for example, the government of Guatemala, this invoicing deprives them.

Ms. HENLY. True.

Senator CASSIDY. But can you prevent the enterprising programmer from being paid off and messing with it because this is the attraction? It seems like the folks in the U.S. or Panama would be able to look at some corrupt programmer, not to accuse programmers of corruption. But you see where I am going with this.

Ms. HENLY. Absolutely.

And this gets back to the protocol design of how do you set up a system where every party is bought in, sometimes literally, to the foundation of the blockchain where, you know, if the network is corrupted, everybody loses some, you know—

Senator CASSIDY. I thought I gathered from your testimony or your collective testimony that you can actually look and see if somebody corrupts.

Ms. HENLY. Yes.

Senator CASSIDY. And so it, kind of, flags.

I think you, Mr. Kahn, mentioned it flags.

Dr. KAHN. Yeah.

Senator CASSIDY. The code that has now been interpolated. Is that correct?

Ms. HENLY. Absolutely. The only question there is on endpoint security. So, how can you guarantee that the physical actions are being reflected in the digital ledger? And that's where you want to create a system of incentives so the physical actions, each person, is incentivized to contribute to the robustness of the network.

Senator CASSIDY. I am over time.

Thank you all, very stimulating.

The CHAIRMAN. Thank you, Senator Cassidy.

Senator Smith.

Senator SMITH. Thank you, Madam Chair.

And thanks to all of you, though I will only claim to have understood a fraction of what you described today. I appreciate the conversation very much and especially how the applications, how this can apply to everything from trade-based money laundering to cybersecurity and autonomous vehicles and so forth.

But I would like to just hone in particularly on an area that I am really interested in in Minnesota. It is very important in Minnesota which is the area of energy efficiency and renewables.

I would like you to just talk a little bit more. You know, in Minnesota we get about 25 percent of our energy from wind and solar, and I think while sometimes the challenges of incorporating that

kind of energy into the grid are overstated, they still are challenges around reliability and also storage.

Maybe starting with Ms. Henly, could you just talk a little bit about how this application might help us solve some of those problems around reliability and also storage? And then I am also really interested to hear from all of you about the kind of additional research that we need, particularly in that sector, that we ought to be putting our attention to.

Ms. HENLY. Absolutely, thank you for the question.

So, one of the ways in which, one of the reasons we're excited about blockchain is because there is an ever increasing, in the electricity sector, set of distributed energy resources that have the capacity to contribute productively to the grid, whether it's, you know, second-on-second demand response, whether it's long-term efficiency, that can contribute to balancing some of the intermittency of renewable energy.

But there are challenges in the electricity sector at the moment of leveraging those assets to productively contribute. And one of the things we're excited about with blockchain technology and that is very deserving of more research is the ability for blockchain given its distributed nature, given its potential to create low cost transactions with those devices, to coordinate, aggregate and leverage those devices to balance some of the intermittency of renewable energy. So we absolutely are excited about that. It is an early stage idea but deserving of more work.

Senator SMITH. Thank you.

I would be interested in what others on the panel think about this. Yes?

Mr. GOLDEN. So I would agree that distributed energy resources, when you think about solar or wind or others and even smaller, maybe home-based or maybe a hospital has a combined heat and power generation facility.

Today or in the past the way that it's been set up is you have centralized generation that's, sort of, a command and a control and how you command the generation to supply the demand that's out there on the grid.

With the advent of these distributed energy resources, they could very much help toward the grid, provide reliability, stability, resiliency, but it's hard to tap into. If you think of having, maybe, ten nodes today of generation and then expand it to three million or something, how does the one central, sort of, utility manage all that chaos?

So blockchain might be one of those technologies where you could use it the way that it works to basically help utilities and others manage the grid and be able to have those resources participate in a more meaningful way.

Senator SMITH. Go ahead, Doctor, please go ahead.

Dr. NARAYANAN. Thank you.

Let me address briefly the second part of your question which is how can we incentivize this kind of research and development.

As a researcher what I see is there is certainly a vigorous amount of research going on in the United States on these topics but perhaps what we could have more of is researchers from very different areas working together about the applications of one kind

of technology in a different sector, such as blockchain technology or other computing technologies in the energy sector. So perhaps funding that is strategically directed in route to incentivize these types of collaborations could be very fruitful.

Senator SMITH. What would be an example of that kind of collaboration that you are envisioning?

Dr. NARAYANAN. A collaboration between computer scientists and technologists and experts in the energy sector who know what are the most important problems that need solving. And so, when you bring those two types of expertise together that's when we can build actually useful technology solutions.

Senator SMITH. Okay.

Yes, Mr. Skare, my fellow Minnesotan.

Mr. SKARE. Yes.

I think another area of research that would be very beneficial here is expanding upon transactive control and trying to see if demand response can also apply to the cryptocurrency mining. That might be a way to provide financial incentives to the miners themselves to back off on their mining during a time of heavy load for the grid.

Senator SMITH. Yes, Dr. Kahn?

Dr. KAHN. So one of the areas it might be useful at a policy level might be to have people look into visibility strategies.

I mean, it's very clear to me that the main interest in the cryptocurrencies is going to ultimately depend upon the visibility of the currency flows. That's the only thing that's going to make governments really comfortable with what's going on. I think you can upscale that discussion to visibility and information flows more generally, and I think that was alluded to by the Senator's comment.

And so, I think one of the questions that could be raised is what is the right, appropriate way to develop public visibility into these kinds of flows when they are intended to be public. But there are going to be some that are non-public. And one of the best ways to do that, maybe in the energy world, there's a lot of information that needs to be visible, but not to the public. So, I think this whole area of the policies that should apply to visibility, what should be visibility and to whom.

Senator SMITH. Right.

Dr. KAHN. And what technologies should apply is really going to be important. Blockchain may or may not be part of that solution.

Senator SMITH. Thank you very much.

Thank you, Madam Chair.

The CHAIRMAN. Thank you, Senator Smith.

Senator Gardner.

Senator GARDNER. Thank you, Madam Chair.

Thanks to all of you for your time and testimony this morning.

Mr. Skare, I will start with you, if you don't mind.

In your testimony you talk about vulnerabilities to 51 percent of the nodes and the importance of endpoint security. How can we address issues such as security for Internet of Things, devices, that are going to be the backbone of, sort of, this peer-to-peer energy trading network?

Mr. SKARE. Thank you, Senator.

I think that that's a key question to be talking about for this whole space because what we're finding is that there's a lot of standards on how to build secure systems and run them from an operational point of view, but not just from a building point of view. So there's a lack of expertise and best practices as far as, how, if I was a vendor, when I was going to build products, how do I build them securely? So that way, we'll have less need for patches once they're operational. That's an area we've been doing research on at PNNL but in the bigger sense is missing in the industry right now.

Senator GARDNER. Senator Warner and I have legislation that tries to address this, not through a prescriptive mandate type of view or approach, but sort of a, if the U.S. Government is buying Internet of Things devices, can we use our purchasing power to influence industry standards?

Now we are buying billions of devices, billions of dollars' worth of devices. Can we set standards that when the U.S. Government purchases these things and patchable devices, no default or hard coded password from the factory that you have to have segmentation, I guess, and other provisions that would make sure that we have Internet of Things device securities? Is that the right approach, do you think, or are we barking up the right tree, so to speak?

Mr. SKARE. I think—I did some work with the Department of Energy on guidance on how to procure systems with, by adding cybersecurity requirements during the procurement process. I think that is a very valuable approach to help. It's one of the ways to influence the market as a major market participant.

Senator GARDNER. Thank you.

You also have received a DOE CESER funded project. Can you talk a little bit more about any other blockchain research going on at PNNL or the other labs that you are aware of?

Mr. SKARE. Yeah, so I think a number of the labs are putting internal investments into understanding how blockchain works and how it can be applied to problems in this space.

The current project that we have that's funded through CESER is really focusing on helping to flesh out and understand a number of use cases as well as, you know, both the pros and the cons of the issues.

And the example of trying to influence, you know, the voting nodes within a blockchain solution if you were to be able to gain control of 51 percent of those voting members. You could theoretically then alter the outcome of the blockchain. And so, it's just an example of one of the issues that we're looking into right now.

Senator GARDNER. Mr. Golden, your testimony talked about transactive energy. You talked about the possibility of two utility customers talking to each other, trying to sell power to each other.

Assuming the regulatory barriers to this were removed, it seems like a blockchain system could capture the accounting aspects and other aspects, but is that enough? Do you need more? Are there other things that we ought to be looking at? Would other systems or rules be required to ensure that the power transfer, it was safe for the buyer, the seller, the distribution system? How would that work? What needs to be done?

Mr. GOLDEN. Yeah, Senator, I think there's a ton of unanswered questions in that space. I mean this is very aspirational to have that, sort of, free market that allows for customers to buy and sell and then interact with the utility.

Lots of things could come into play about how the grid is actually operated—how to understand where that power is flowing and when it's flowing to make sure they maintain reliability and resiliency for the customers. So I think there's a tremendous amount of research that's required in this area to really move that from, you know, sort of an idea or hypothesis into actual, into an action item.

Senator GARDNER. Could you talk a little bit about the research that is being done right now, both in the private sector and the federal R&D, that is taking place and what needs to happen at the federal R&D level to further this research?

Mr. GOLDEN. Yeah, today, I mean, we've set up, sort of, a bench lab test of what does that technology look like if you were to set up several nodes and see the interaction between, sort of, the buyers and the sellers in the market. It's very small scale.

So I think additional funding to, sort of, support that research. It's one of those things where lots of times you're trying to figure out what's in the next couple years and this is maybe five to seven years out. You have to spend the money today to get to that, sort of, five- to seven-year future. And so I think that it's hard to get a huge amount of investment today, but it's something that's very much needed in order to get those technologies to move forward.

Senator GARDNER. The Secretary of Energy and I were out at NREL this past week where we talked about some of the grid work that they are doing in some of the facilities and opportunities that they have to do a lot of this research and development, the testing there with the power coming into the system from wind energy, gas turbine energy, solar energy, traditional energy, other traditional energy sources. I think that is the kind of approach that we need to continue to utilize, the expertise with the private sector and Federal Government researchers.

I commend you all for your work and thanks for trying to help us understand some of the great technologies that we are on the cusp of achieving.

Thanks.

The CHAIRMAN. Thank you, Senator Gardner.

Senator Heinrich.

Senator HEINRICH. So a couple of you have mentioned this potential for peer-to-peer trading using blockchain technology. How much do we know about how and whether that works well?

Mr. Skare, I think you have worked a little bit on the Brooklyn project. Can you just maybe elaborate on is this working? Is this the right architecture to facilitate that kind of market?

Mr. SKARE. I think that's a really great question because what we're trying to understand is, you know, where are all the best use cases for this technology. I think it's in the Brooklyn microgrid transactive control type of scenario. It's a good place to try it out.

I think additional use cases should look for other opportunities to do that early stage, kind of a demonstration of how each use case applies.

Senator HEINRICH. If this proves to be a good application for peer-to-peer trading, do any of you think that you could similarly, could you take blockchain and use it to facilitate participation of aggregated, behind the meter generation and storage assets to even go so far as potentially participate in bulk power markets?

Ms. HENLY. That is one application that we're quite excited about at Energy Web Foundation, and I know many of our affiliate companies are interested in specifically that type of aggregation.

To the point around is blockchain the solution or one of the technologies that might be useful? It's really one of the technologies. And in order to realize an application like that, you're going to have to not only invest in the blockchain technology itself but also other technologies that will connect with devices and be able to realize something along those lines. But definitely a key part of the puzzle.

Senator HEINRICH. Ms. Henly, you said that the U.S. is behind in dealing with this architecture. Can you elaborate a little bit and, specifically, can you comment on how Congress, PUCs, FERC, et cetera, ought to be crafting regulation with these protocols in mind?

Ms. HENLY. Absolutely.

So there is really valuable and excellent research happening in the U.S. I don't mean to suggest that there isn't any research happening. However, what we're seeing from the industry perspective is that most of the development, especially in the energy sector around blockchain, is happening in Europe and some of the key core developers are based in Berlin and a lot of the demonstration projects around energy applications are also in Europe.

I think there's a real opportunity for the U.S. to put together additional programs, funding sources, perhaps the DOE, perhaps more broadly on blockchain and related technologies that would support foundational and fundamental research also in the U.S. and, in addition, demonstration projects that could be showing the value of peer-to-peer and other applications.

Senator HEINRICH. So energy related blockchain transactions are going to need to be energy efficient, unlike the bitcoin example. They will need to be scalable. They will need to have reasonable transaction costs if they are going to be implemented.

Would any of you like to, sort of, talk about the different security protocol architectures that you mentioned, the proof of work, the proof of authority, proof of stake, even alternative architectures like tangle and what some of the positives and negatives of those different architectures are proving to be and how you think that is going to apply, specifically, to energy transactions?

Ms. HENLY. I'll jump in quickly, but then others should comment.

In the energy sector we realize this is an issue and, for example, one of the reasons Energy Web Foundation was created was to address this issue and to create a blockchain that did not use bitcoin's energy intensive mining practices.

And so, that's why we are launching our blockchain next year with proof of authority. Ethereum currently processes 1.3 million transactions per day on average and has announced, the Ethereum Foundation has announced, a move to proof of stake, also uses orders of magnitude less energy. We expect those to be the trans-

active foundations for applications in the energy sector and in other sectors.

Mr. GOLDEN. I would say that I think transaction speed is obviously very important. We think about instantaneously balancing supply and demand. On the grid, you can't have latency that's going to not really, sort of, make sense of what's happening in making real-time decisions on how to balance everything.

So, I would say that the proof of stake, like you mentioned, there's other protections that, sort of, are pointing toward being faster in their transaction speed.

Senator HEINRICH. Great.

Thank you, Madam Chair.

The CHAIRMAN. Thank you.

Senator Portman.

Senator PORTMAN. Thank you, Madam Chair. I appreciate the witnesses being here today to talk about blockchain and energy and also some of the cybersecurity threats.

As you may know in my home State of Ohio, blockchain has gained a little momentum recently. The Governor, John Kasich, recently signed legislation saying that blockchain was an electronic record like other electronic records and that, in turn, gives blockchain the same legal protections as other types of electronic records.

Cleveland, in particular, and other parts of Ohio are becoming a home to some of the innovations going on here and have big interest in the Cleveland business community in developing blockchain thinking it is going to be a big part of the future.

And what I've heard today is probably, maybe, who knows what the next great technology is, but it looks like it is very promising for a lot of applications.

In the National Defense Authorization bill last year and again in the Homeland Security—which is a law now—and the Homeland Security bill which is out of Committee this year, I worked with colleagues to put in measures to have the government do more studying about blockchain, its opportunities and challenges in particular. Again, this focuses on how do we get our hands around cybersecurity threats that are proposed.

On energy efficiency this Committee has done a lot of work in this area and passed legislation. Not all was passed by the Congress, some that has energy efficiency legislation that Senator Shaheen and I introduced, for instance, would be equivalent of taking 22 million cars off the road in terms of emissions and then you have blockchain. According to one of your studies on, particularly, the use of mining and some of the transactions.

Dr. Narayanan, in your testimony you cited your work in determining the amount of energy required for bitcoin mining alone is slightly more than the electricity consumption of my entire State of Ohio. Is that accurate?

Dr. NARAYANAN. That is my best estimate based on the data available at this point. There can be some uncertainty in that estimate, but the order of magnitude, I think, is very clear.

Senator PORTMAN. And the order of magnitude is enormous and, you know, concerning.

What ways do we have to verify the methods of blockchain transactions that are less energy intensive than others and how can we get our hands around this issue?

Dr. NARAYANAN. Senator, I think that is one of the key areas where we need more research. We've had some discussion today already about alternative mining for any blockchain technology.

So far, it's proven tricky to apply those to public blockchains which support cryptocurrencies. That doesn't mean it's impossible. We're just not quite there yet in my view in terms of the technology and more research, more funding, more encouragement for the development of alternative technologies can certainly help.

Senator PORTMAN. Yet there have been some discussions already in terms of limiting the amount of mining—and any of you jump in here and Ms. Henly, maybe you have a thought on this—is that a practical solution?

Ms. HENLY. I would say our strong recommendation is to invest in researching alternatives to energy intensive bitcoin mining practices. I think that Mr. Skare, excuse me, mentioned some other approaches to time of use pricing or demand response that could be offered to miners as an incentive to not mine at peak times.

Senator PORTMAN. Creating incentives?

Ms. HENLY. Yes, exactly.

Senator PORTMAN. My time is coming to an end so let me just get to the opportunities now because the challenges are clear. Again, it would be ironic if all this work on energy efficiency was countered by this great new technology that does consume even more energy because the idea is to use it to be more energy efficient, to, as my colleague from New Mexico, Mr. Heinrich said, make it scalable and practical and timely, to be able to take distributed energy and make it work more practically. And all those opportunities are out there.

What do you see as the best near-term applications of blockchain to improve energy efficiency?

Open it up, Mr. Skare, maybe you can give a thought on that.

Mr. SKARE. Well—

Senator PORTMAN. Short-term, now.

Mr. SKARE. Short-term, I think the key is that blockchain is a technology as well as a number of other parallel technologies that can achieve the same goals.

I think that energy efficiency can be tackled with or without blockchain and that the speed of getting energy efficiency gained is kind of a separate, separate orthogonal problem from blockchain technology.

I think the big opportunity right now for blockchain is the fact that there is not a standardized definition of blockchain so we can continue to work on the research on how to make it more energy efficient and private blockchains. And that's probably the single most, immediate thing.

And as we work toward cryptocurrency mining, figuring out ways to maintain the needs there in less energy intensive ways is also a—

Senator PORTMAN. How about tracking individual energy use? Is that a short-term application that could be helpful? Mr. Golden? Ms. Henly?

Ms. HENLY. To the tracking point?

Senator PORTMAN. Yes.

Ms. HENLY. The Energy Web Foundation is building a reference application that tracks renewable energy credits that we've launched in Alpha earlier this year.

To the point around energy efficiency, one of our affiliates, Elia, which is a TSO in Belgium, is currently building a blockchain application to run their demand response program. So there are present day applications for energy efficiency, for tracking of blockchain that can be supported by the technology at this current stage.

Senator PORTMAN. Mr. Golden?

Mr. GOLDEN. Senator, I think you'll continue to see, if you look at data centers and that first, sort of, came to the front, people were worried about the amount of energy usage. And those data centers continued to drive down with energy efficiency so that they can, sort of, arrive at a lower amount of energy used.

I think, if you look at blockchain mining and proof of work mining, as price drops, for instance, you might see them get a little bit tighter with the way that they want to spend their money. They're at \$19,000 which it doesn't make a whole lot of sense to look at energy efficiency plays because you're making so much money. Whereas, if that does drop you might start to look at how do you make this operation more efficient.

And then as larger plays come in, I think, there will be another opportunity for them to, sort of, consolidate and look at how they can make their operations more efficient maybe through the data centers.

Senator PORTMAN. So, we incentivize the business, not the market side of—my time has expired but any other thoughts you guys have, I would love to hear them as a matter of the record.

But I think it is, again, challenges with huge opportunity as well. Thank you, Madam Chair.

The CHAIRMAN. Thank you, Senator Portman.

Senator Cortez Masto.

Senator CORTEZ MASTO. Thank you and, again, thank you for this hearing. I think it is an exciting time for us. I thank all of you for the conversation and the information that you have provided.

I do believe we are at the dawn of a new age with the potential for blockchain technology. We cannot squander it. We have an opportunity right now to invest in the R&D for the very reasons that you all have identified.

I appreciate the conversation and will follow up with respect to what we can do at a federal level, but there is a role for the states to play as well. And we have talked a little bit about that.

I am proud that Nevada is one of the leading states in the nation encouraging the growth of the blockchain applications. For instance, in Nevada, our state legislature recently passed a bill incentivizing blockchain start-ups to locate their businesses in the state and that had the support of our governor as well, our current governor. It is an exciting time.

But let me talk a little bit about some of my concerns and how you can help us address this. We talked about the architecture, and we talked about the framework. And as we build out this frame-

work and we do the R&D, we talked about cybersecurity and security in general which, I think, is an important guardrail that we start looking at and putting in place as we build out this infrastructure instead of trying to layer it on after the fact which is much more difficult.

But one thing we have not talked about yet is the privacy component. So let me open it up to all of you. What privacy considerations should be examined for consumers using digital currency or other blockchain applications? Please.

Dr. NARAYANAN. If I may speak to that, Senator.

With cryptocurrencies consumers have, sort of, a privacy dilemma which is that they can interact with other individuals and businesses using cryptocurrencies without providing their real name.

However, all of the transactions that they make are permanently recorded on this public database in a way that they cannot take back. As we've heard, you cannot go back and modify the blockchain records later which means that if a consumer's identity at any point gets associated with any one of their transactions on the blockchain then it can further be linked to all of the transactions that they have made using a cryptocurrency.

This is a very new type of system. It's not like the credit card system. With credit cards we don't really have the danger that all of my credit card transactions ever are going to be publicly, you know, displayed with my name attached to it on the internet.

We don't quite know how to manage this privacy dilemma from a consumer perspective. A lot of research is happening, for example, there are newer blockchains which try to hide all information from public view. However, that raises other questions if those blockchains get adopted how can law enforcement, for example, still do their job. And so, I don't think we've quite figured out an answer yet which is both technologically sound, which is useable for consumers and also meets the legitimate needs that we have for law enforcement and investigative purposes. Definitely very much a work in progress.

Senator CORTEZ MASTO. Thank you.

Doctor?

Dr. KAHN. Yeah, I think that we didn't really have a chance to get into the technology of the digital object architecture, but as I stated earlier, blockchain is one example of a digital object, but just one.

And yet, I know the focus of this hearing is on blockchain and so, that's why most of the questions are about it. But I think if you look at this issue more broadly, for example, the question of managing information shows up all the time in the research data alliance.

This is the something the U.S., the European Union and Australia set up where many people who deal with very large data sets are worrying about how to curate it, share it, protect it, secure it and the like. And within that context the notion of blockchain almost never shows up because it is one choice and as they look at their needs, they don't see that as the critical, initial thing.

Within the digital object architecture every aspect of it can be self-identified. For example, every individual can have an identity,

including anonymous, that could be an identity. Every system can have an identity, every piece of information can have an identity. And that's how privacy is generated.

So, if it were a health record that you were looking for, part of that health record would say, the following identifier for people are appropriate or programmed. So whatever it took to access this information. Privacy is essentially, inherently, built in.

But by asking a question only about blockchain, it seems to me, that you omit all of the other potential applications that might be useful in society for which blockchain may not be the solution.

This is not to take a position for it or against it, as I've said before, but with the focus here only on blockchain, I think you need to understand that there's a broader universe of applications of which blockchain is exactly one option.

Senator CORTEZ MASTO. Okay. Thank you.

I notice my time is up. Thank you again for the conversation.

The CHAIRMAN. Thank you.

Senator Daines.

Senator DAINES. Chair Murkowski, Ranking Member Cantwell, thank you for holding this hearing today.

I came from a tech background, as some of you know. I was in the cloud computing business for 12 years before I came to Capitol Hill.

I have seen the progress of blockchain technology for cybersecurity and other industries. It has been fascinating to watch what's developed here. It is also interesting from an energy point of view.

For example, bitcoin mining requires an enormous amount of electricity, sometimes surpassing even traditional mining projects. In my home State of Montana, we have facilities in Bonner and Butte that collectively require about 80 megawatts of electricity. There are plans by developers of these facilities to increase the energy demand and news of other bitcoin operations planning to move to Montana. Why are they thinking of moving to Montana? Well, we have lower cost energy. One mining operation projects the next few years to expand to 100 megawatts, making it one of the largest energy consumers in the state.

Montana has cheaper electricity, we have a colder climate and we have less expensive real estate—not true everywhere in Montana, but in many places that is true.

This activity can create a net benefit to our economy. In fact, we could see some changes though in the near future. There are Colstrip units one and two in Montana that are planning closure. There are threats currently to units three and four.

We may see electrical prices go up and energy production in Montana go down. We are a net energy exporter today. If we lost Colstrip units one, two, three and four, we would become a net energy importer.

As the demand from bitcoin miners increases and the supply of cheap, reliable electricity from coal generation decreases, this could pose a threat to the expansion of bitcoin operations and an even greater threat to energy supply and prices for Montana as a whole.

Mr. Golden, you noted that bitcoin miners prefer to locate in areas that have low energy prices like Montana. They rely on continuous, steady streams of electricity like other data centers.

I remember when I was running and hosting operations for our company. We had data centers all over the world. But these mining operations usually run 24/7 with continuous demand. Areas with strong baseload power are, therefore, attractive locations, like Montana, with robust hydro as well as coal generation. In fact, I recently heard of a coal-fired plant in Australia which is reopening to provide power to a new bitcoin mine.

The question is, how do you see the increase in bitcoin mining and the reduction in traditional baseload generation resources, like the closure of Colstrip, affecting the grid in the cost of delivered energy, especially in places like Montana?

Mr. GOLDEN. Thank you, Senator.

So, I think, fundamentally, for the supply and demand piece, if your supply is dropping and your demand is increasing, then obviously you could see a change in the amount of the pricing structure for the service territory. So, I think, you know, obviously, that's one thing to consider.

But there's also the fact that utilities are consistently looking at what is happening on their service territory and making decisions on what technology or what generation to start up or shut down.

So, as they look at that, I think there will be a lot of decisions made in their planning processes if those coal units do go down, what would come in to replace that so they could continue to, sort of, provide their charter which is least cost energy for their customers. So, I think it'll take a lot.

We don't know where these bitcoin mines go. If they someday fold or they don't—if they're not there any longer and it's difficult to make those decisions, but you know, having those, sort of, planning tools in place and processes that utilities are very much used to, helps them make those decisions.

And I think having partnerships and frank discussions with many of these companies at the outset is a very important thing for the utilities to get engaged with.

And one of the things we've done at EPRI is we've started a thing called Utility Blockchain Interest Group which is simply to get as many of our members around the table to have these discussions in a frank, meaningful way so that maybe your utilities from Montana can talk to somebody who's had this problem in New York or maybe in Washington and some lessons learned could be passed.

Senator DAINES. Or Australia it sounds like.

Mr. GOLDEN. Or Australia, right.

Senator DAINES. Right.

They brought a coal-fired plant on to supply reliable, low cost energy for bitcoin mining. Along that line is this need, I guess, for agility and responsiveness.

In your written testimony you discuss the difficulties of predicting future energy consumption of bitcoin miners and the potential problems associated with this uncertainty. You were alluding to that.

As the value of bitcoin rises and falls, so does the incentive, perhaps to mine more bitcoins. This could lead to the investment in infrastructure that might be used for a few short months or could extend many years. How can communities and energy companies be more prepared as bitcoin miners move into their regions?

Mr. GOLDEN. I think, if these operations are, sort of, green field sites and they're starting to build brand new, I think it's a great time for utilities in the area to have discussions about what kind of infrastructure will be required to be put in place and help make sure that those companies are providing, sort of, their fair share for providing that infrastructure.

I mean, other opportunities that we've, sort of, researched and seen is that these companies will often look for areas that may be, for instance, a car wash that had a high electric load that could then be utilized, that that connection could then be utilized to, sort of, hook up to the utility. So in that case it's, sort of, a benefit where you're realizing infrastructure that maybe was defunct before that.

I think those two areas of looking for infrastructure that's already in place and built out. They could look for that and then also, having utilities have those discussions right up front with new customers on how to best set that infrastructure up.

Senator DAINES. Chair Murkowski, as Mr. Golden has alluded to, uncertainty is something that's difficult to manage but having more optionality by having a balanced energy portfolio, I think, is part of helping address uncertainty going forward. I hope we can learn from places like Australia, like Taiwan, like Germany that move too fast in one direction and kind of lost sight of having a diversified energy portfolio.

Thank you.

The CHAIRMAN. Senator King.

Senator KING. Thank you, Madam Chair.

This has been a very informative discussion. You have noticed that the Senators have come and gone which tells you that the real software challenge, forget bitcoin and blockchain, is scheduling Senate hearings in a rational way so that we do not have to be in two and three places at once.

[Laughter.]

That is a challenge the world has never been able to tackle. I will just mention that.

I am interested in the energy consumption issue. To what extent do these facilities which essentially are server farms, is that correct? Isn't that correct? That is what they are.

You are nodding. Could you say yes?

Ms. HENLY. Yes.

Senator KING. Nods do not get in the record.

[Laughter.]

Ms. HENLY. Yes.

Senator KING. To what extent are they dispatchable in the sense of being subject to peak load pricing, to load shedding and is this an opportunity to rationalize the utilization of the grid on time of day?

Yes, sir?

Mr. SKARE. Yeah, I think this, there is an opportunity to move forward there.

Today's installations do not have that built in ability to regulate based on either time of day rates or other incentives. There are a number of places in the United States that do have time of day

rates that can use it. And the opportunity is there to add that demand response or even load shed capability.

Senator KING. But as I understand it, this whole system is built upon incentives and this is another incentive that could be built in. Would that be feasible in terms of the mechanics of this particular business, that they could ramp up and come down according to the cost of the energy that they are consuming?

Mr. SKARE. Yeah, I mean, as you said earlier, these are basically like server farms, so you could really turn some of the mining, you know, computers off during times of difficult load and then run them when the load isn't as such.

Senator KING. Well, one of the realities of the grid is it is terribly inefficient in the sense that there is huge slack at night, for example. And to the extent we can shift things like charging electric cars or running server farms, we wouldn't have to build a lot of additional infrastructure because the infrastructure is there. The wires are there, the generating capacity is there and it is scaled back at night.

On this issue of individual sharing energy transactions, why does it take blockchain? Why can't the ISO do that now, the ISOs in the business of turning on plants and turning off plants according to demand? I don't quite understand why that can't happen under the current technology with additional software.

Ms. Henly?

Ms. HENLY. Well, I was just going to refer to Mr. Golden's comment previously about the number of devices where today we have maybe hundreds or thousands of devices that a central operator is trying to optimize.

In, you know, as increasingly even, you know, in the last few years, distributed energy resources, small devices, in households have started to expand. We're talking millions of devices and that requires a different architecture in order to coordinate them. But perhaps—

Mr. GOLDEN. I would agree.

I mean, also, you know, you think about the ISOs and they have, sort of, minimum barriers for entry. You might have to be one megawatt of generation in order to participate in the market. We're talking about a situation where we have kilowatts.

Senator KING. But there is no law that says that. I mean, they could alter their software in such a way to accommodate smaller transactions.

Mr. GOLDEN. Right. Then it becomes, I think, a question of the cost of doing business.

If you're having these micro transactions and having to pay a certain percentage for all of them, it almost doesn't seem, sort of, worth it at some point.

Senator KING. What is worrying me is the development of a whole different—we have already got a system for turning power off and on and monitoring the grid and determining when there is a need. If you build a whole new system on a blockchain basis, it is still going to have to integrate with the ISO at some point. Do you see what I am saying?

Mr. GOLDEN. I think you're completely—I think, like we've been sort of saying—

Senator KING. Completely right is good. Finish that.

[Laughter.]

Mr. GOLDEN. As a group I don't think that we know that blockchain is certainly the end-all, be-all and that's the right decision to go with that technology to solve that problem.

I think there's, obviously, many other ways to solve that problem and you know, obviously, those are the things we're examining.

Ms. HENLY. And the only thing I would add is that there is a current system in place that is quite effective for operating the current grid.

Senator KING. Right.

Ms. HENLY. But if, as you have just, you know, reflected, I think, quite well, is that the slack capacity on the system right now is enormous and to be able to create a system that is better optimized, more responsive, it may be necessary to add additional functionality where blockchain can aggregate small assets to hook into an ISO or to even, eventually and this would be, you know, a many year type of situation, to switch over to a more dynamic control system that is more distributed.

Senator KING. I have often likened the grid to a church that is built to accommodate Christmas and Easter and on a Sunday morning in February there are a lot of empty pews. We need to, I think, think about ways to more efficiently utilize this enormous investment that we have.

The best example to me is charging electric cars overnight which you could do without any additional infrastructure whatsoever. It would only be energy cost, no additional capital costs.

Well, thank you very much, Madam Chair. Thank you for scheduling this fascinating hearing. I appreciate it.

The CHAIRMAN. It has been very interesting. Thank you, Senator King.

Just a couple more questions. I would like to follow up.

Ms. Henly, you were asked, I don't know, maybe it was by Senator Cortez Masto, on the issue of just privacy from a broader perspective and recognizing that if using blockchain for energy transactions, you might not, perhaps you don't need the same level of anonymity, I guess. But a question would be, how do you ensure then if you have this anonymity that is, we are saying okay, this identification is not required. How do you ensure that you are not selling a product that might be subject to sanction?

Ms. HENLY. Thank you for the question, Senator.

I think that this is an area which deserves a lot of attention and interest and actually, engagement from regulators because, I think Dr. Narayanan stated it very well that privacy functionality is being developed by the industry, but it is not robust enough at this stage and is still in an early phase of development and needs to be developed, in particular, with specific use cases in mind.

What is it that regulators care about in terms of privacy? What is necessary to require? There is a productive collaboration there between technologists in continuing to work on the technology and regulators in, you know, setting specifications and requirements for the direction the technology and privacy, in particular, moves in.

The CHAIRMAN. Let me ask about the issue of regulation, because that is what we do up here is we advance legislation.

Is it too early for us to be discussing the potential for legislation in this space?

Mr. Golden, you mention in your testimony that regulatory barriers exist that restrict the transactions between the customer and the local utility. Should we expect some kind of a request to adopt federal regulations? You have some interstate issues connected to these transactions. Speak a little bit about your perspective on the need for regulation at this point.

Mr. GOLDEN. Yeah, EPRI doesn't generally, sort of, get into, jump into the regulatory, sort of, realm, but really, it's about looking at the technology itself and whether or not the technology is ready. I think there's a lot more work that has to be done as far as whether or not this technology is ready to, sort of, participate in a transactive energy way.

I think one area that maybe is also important for this transactive energy is the use of smart contracts. When we think about peers on a network selling electricity back and forth to one another there needs to be some sort of contract in place. And I don't know that there's any sort of recognition from a legal framework of whether or not those smart contracts are actually enforceable. So that's one thing.

And I think one of the Senators mentioned that in Ohio they had looked at, sort of, allowing for that contract to be something that is true and real and is recognized by the legal system.

So, I think those areas are areas for you to look at and then also just keep an eye on the technology as it develops like we will plan research and development.

The CHAIRMAN. Mr. Skare?

Mr. SKARE. I'd like to just add a little bit to that in that I think the area of concern besides personally identifiable information being made part of an immutable record that you also have to look at anything else that the transaction parties might be placing within that block so that you can put in, you know, illegal data into these blocks and then they stay there. So that's another area of concern that should be reviewed as part of that process.

The CHAIRMAN. Well then, that takes me back to an issue that I raised with you, Dr. Kahn, after your testimony and that was the issue of trust.

You explained that the focus of blockchain is not creating blocks of data but in creating data that can be trusted and trust is one of those words that we don't often hear when we are talking about cybersecurity.

I guess I would ask you to just speak a little bit further about what more can be done to ensure that we are actually achieving this trust. You mentioned that blockchain is just one option out there. But how does the trust that is created by this blockchain technology compare with other techniques that might be used in gaining that trust for the digital objects?

Dr. KAHN. Well, that's a wonderful question. Thank you for that. And it's not going to be very easy to answer it directly because you have to get into all the ins and outs of how you trust blockchain and how it works and the like.

I don't think we really have time for that right here and probably other people could address that equally well.

But let me just say that in developing trust in any informational system and I might point out that things are not new, as in the Internet of Things, because when we did the original work on computer networking the things we connected were big computers. So we're really still in the business of connecting things, big informational systems, little ones, whatever.

You need to know what information you want. That means you've got to have a way of describing it and it can't be semantics. It will be different in different languages.

That's why the importance of unique identifiers associated with digital objects. You can say, the information I want is in the object who has this identifier, wherever it is in the information universe. And then you need to have some way of going from that identifier to actually accessing that information. Now it might be you can't manifest it because it's all encrypted, but you've got to be able to get to it.

So let me assume that you can interpret it when it shows up. The next thing you want to know is can I trust that this information has not been tampered with. And that's, kind of, an interesting question that deals in authentication technologies of one sort or another.

Now, often times you think you need some other set of records to know whether it's accurate, but in many cases you're dealing with information that is immutable, that is it never changes. And if the information never changes, then you can actually create an identifier for that information, cryptographically, based on the content of the information so that from the identifier you can get the information and from the identifier you can validate whether that information was accurate or not. It doesn't require anything else. It's self-built in to the way kind of system would operate.

Now, if the information is changeable, you probably need to depend on another system. It could be a blockchain system. It could be many of the kinds of systems.

But there is one that's widely used throughout the world in the publication industry and options trading and in managing building activities. It's called the Handle system, well that's the trademark for it. But it's an identifier resolution system that's very powerful and it's been on the internet for almost 30 years, actually 25, 27 years now, widespread in the publications industry.

It's hard to find any scientific journal that doesn't rely on this to identify references in those journals.

So that's an important aspect of how you develop trust. You might develop it because the system itself is intrinsically, the trust level doesn't depend on anything else, anywhere else. It doesn't depend on other service systems or whatever.

The question is can you get the information? This issue also has shown up in other areas where you have to learn to trust the cryptography like if you could actually send information over another nation's satellite system, would you trust the information that you got? Well, if you could get the bits and you believe in the strength of the cryptography, maybe the answer should be yes or certainly yes in times of dire need.

There's also the question of how you can trust in a particular object that has value when it's the bitstream itself that purports to

have value, as in cryptocurrency, when you can actually transfer that, perhaps even anonymously from one party to another. So if I transfer a bitstream that's got \$100 of value to another party, who's got the \$100 of value if I kept a copy of that same bitstream?

You have to have a way of understanding which is the real bitstream even though they're both identical. And that can be done in a variety of ways, including through, if you read that paper I wrote on representing value with Patrice Lyons, you'll see that we actually go through that in detail and it's not a blockchain solution, but it is one that's another alternative.

And then finally, you need to know that communication world that you're dealing with can actually get information through reliably because if you can't get the information then all of the trust that you might have in it will never really materialize.

Cybersecurity is important in terms of maintaining the flows but also making sure that things can't be changed within the system. If I pull down an object which is, let's say, a piece of legislation, I'd like to know is this the real piece of legislation? Well, you can't tell that necessarily from something else, you might have to look at the object itself to tell. You might want to know is this the latest version? Has it been amended? You may want to get the prior version.

And so, all of those things go into building trust and that's one of the things that we've thought through very seriously in the digital object architecture, but it's like a set of building blocks. It's not a cookie cutter, one size fits all and every example that we have been involved with, we've built probably dozens of systems of various kinds, including some for managing options trading around the globe, managing the construction of buildings and smart cities and things like that.

Every one is different and there is no single solution that, you know, universally applies to everything. You have to look at the issue of trust and managing the information and protecting it and securing it individually in every case. There is no universal solution that's going to work for everything now and in the future.

The CHAIRMAN. Well, I thank you for that. You are right, it is not something that you can respond to in a couple minutes blurb. This is a much broader and exceptionally important and very, very timely.

Not a lot of trust that you have around the halls of Congress right now. Quite honestly, there is not a lot of trust that the American public has.

Just think about what people get on their news, they are all wondering is this real? You know, we are not trusting much of anything nowadays. So making sure that we have architecture that can be trusted, I think, will be exceptionally important moving forward. So I appreciate you outlining that.

Senator King, any final words?

Senator KING. No.

The CHAIRMAN. I appreciate what you have provided to the Committee this morning. It has been very interesting, very enlightening. I think we have all learned a lot.

This has been good to fill our heads with as we move forward, and I thank you for what you have contributed.

With that, the Committee stands adjourned.
[Whereupon, at 11:56 a.m. the hearing was adjourned.]

APPENDIX MATERIAL SUBMITTED

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Paul Skare

Questions from Chairman Lisa Murkowski

Question 1: Bitcoin miners are reportedly demanding significant amounts of electricity in certain regions. In fact, we received a statement for the record from Steve Wright, General Manager of the Chelan County Utility District in Washington State, describing the utility's experience with bitcoin mining and noting two important issues.

First, because bitcoin mining is highly portable and dependent on the price of bitcoins, the electricity loads from bitcoin mining can and do move away without notice. And second, loads can change their demand requests substantially. A request for 10 megawatts (MW) in ten locations can transform to 100 MW in one location, which can then change to 1 MW in 100 locations.

- Since these two characteristics of bitcoin mining suggest no long-term stability, how should utilities prepare for bitcoin mining?
 - This is a critical question, but one without a clear answer today. I believe, we should begin by ensuring that utilities serving as load serving entities are able to share their experiences and develop best practices for addressing the challenges bitcoin mining presents. Having a national conversation on what these impacts are and approaches various utilities have had in planning and responding to these impacts should be encouraged. From this conversation and development of best practices, policies *best suited for each particular business and technical situations*, can be applied. Specific policy approaches from the Chelan PUD and others that the broader community may find valuable include the use of mining moratoriums and licenses to mine with up front fees to cover impacts on the distribution system. In addition, load serving entities can share best practices on detecting unplanned mining operations and effectively and safely interceding in operations found to be in violation of existing policies.

- How can we encourage more stable outcomes in bitcoin mining?
 - Two technical actions could be beneficial here. First, policymakers and the electric grid service providers could engage directly with Bitcoin and encourage them to move away from a Proof of Work model to some other model (Ethereum is reportedly considering an alternate model, which would dramatically reduce the electricity used in mining). Second, bitcoin miners could be encouraged or required to implement a demand response capability that allows them to accept payment for either reducing or curtailing consumption at times of high load or stress.

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Paul Skare

Question 2: We have been hearing about blockchain in the energy industry more frequently, much in the same way that we heard a few years ago about how “cloud storage” of data was going to improve energy markets.

- What do you think are the reasons for the hype? Is it driven by sellers of computing solutions? Or by customers looking for better answers?
 - I think both marketing drivers and consumer desire for silver bullets in a time of increasingly visible cyber threats both play a role in the growing attention to blockchain technology. In the case of blockchain, as I addressed in further detail in my testimony, there is both promise for adding cybersecurity for some applications, there are also potential vulnerabilities and high costs for other approaches that need more research before widespread adoption. This is the driving impetus for the research project PNNL is performing for DOE CESER’s CEDS program on blockchain.
- Will blockchain as a concept simply fade away when the next “big thing” arrives to promise solutions for the energy industry?
- Or does blockchain technology offer something deeper that has real value?
 - Technology forecasting is a tenuous art, and long-term views are hard to validate. My best insight is that the techniques used in private blockchains will be an additional tool for utilities to provide secure Operational Technology for specific use cases. Because adding a network of nodes to validate transactions is a higher cost approach initially than the current, more centralized approaches utilities deploy, use cases involving financial transactions or where the consequences of cyber intrusions are otherwise extremely high will be the leading contenders for adoption. Blockchain technology to be used in Information Technology will likely continue to have appeal wherever a transfer of funds can be performed without middlemen.

Question 3: We have been talking about a series of issues that are not always that closely related. On the one hand, blockchain has cybersecurity impacts that may be helpful in securing the power grid. On the other hand, blockchain has potential for improving the efficiency of the organized wholesale markets.

- As this committee explores blockchain, can you tell us where we should focus our time and resources? Is it cybersecurity? Markets? Energy consumption? All of the above?
 - My recommendation is to continue exploring cyber security and energy consumption in tandem, while separating the terms ‘mining’ and ‘blockchain’. Consider blockchain as an additional tool that can be used for cyber security especially in use cases where it has the potential to add security to various grid transactions. Separately, I would urge the Committee to continue to work to better

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Paul Skare

understand Bitcoin mining as a source of greatly increased energy consumption. For the latter, there may well be insights to be gained from similar past experiences, as when utilities had to learn how to work with Aluminum Smelters and other non-conformant loads.

- Do you see a need for legislation? Or is it too early to know?
 - I believe it is too early to legislate specifically on either bitcoin mining impacts or the application of blockchain for cybersecurity. As we discussed at the hearing, PNNL is currently in the first year of a three year research project funded by the DOE's Cybersecurity for Energy Delivery Systems program to investigate the application of blockchain to various use cases on the grid to enhance cybersecurity. I look forward to briefing the Committee in the future on the results of this project, which may well provide relevant insights as you consider legislation. As the Committee heard from the whole panel at the hearing, I think it is important for the Committee to continue its important work promoting energy and cybersecurity research and development funding and encouraging public-private partnerships that can aid in the transition of new technologies to industrial applications.

Question 4: I want to explore how blockchain miners might fit into an energy market:

- In a blockchain system used for trading energy, will miners be used to verify a transaction? Or are miners even needed?
 - Mining is used with public blockchain applications that use a "Proof of Work" model, as described in my testimony and that of my fellow witnesses. For energy trading applications, *private* blockchain technologies could be used, which do *not* use "proof of work" for accessing the blockchain and thus not require mining operations.
- Will miner compensation be added to the consumers' bill?
 - My recommendation is to avoid technical approaches that use mining in applications for energy trading.
- If no miners are needed – such as in a proof-of-stake or proof-of-authority system – what is the benefit of blockchain over a centralized system?
 - Blockchain applications in this use case could be more highly automated, reduce the need for oversight, and remove the need for middlemen. All of these benefits could potentially reduce the administrative and transaction costs for consumers compared to a more centralized system. Further, a blockchain application in energy trading could increase the resilience of the energy trading system to disruption through the inherent redundancy of its distributed ledger approach. While these are exciting potential benefits, further research, including the ongoing work we are performing for DOE's Cybersecurity for Energy Delivery Systems

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Paul Skare

program, is needed to demonstrate these benefits can be realized and explore potential downsides of the blockchain approach.

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Thomas Golden

Questions from Chairman Lisa Murkowski

Question 1: Bitcoin miners are reportedly demanding significant amounts of electricity in certain regions. In fact, we received a statement for the record from Steve Wright, General Manager of the Chelan County Utility District in Washington State, describing the utility's experience with bitcoin mining and noting two important issues.

First, because bitcoin mining is highly portable and dependent on the price of bitcoins, the electricity loads from bitcoin mining can and do move away without notice. And second, loads can change their demand requests substantially. A request for 10 megawatts (MW) in ten locations can transform to 100 MW in one location, which can then change to 1 MW in 100 locations.

- Since these two characteristics of bitcoin mining suggest no long-term stability, how should utilities prepare for bitcoin mining?
 - Since utilities are obligated to “serve all willing customers” this is in part, a regulatory issue. Utilities should think about ways to incentivize miners to utilize existing resources, where appropriate, to avoid new infrastructure investment.
- How can we encourage more stable outcomes in bitcoin mining?
 - There may be market incentives that could more closely align load/generation. For example, market incentives might be structured such that mining operations could be curtailed to reduce load if required, as part of a demand response program.

Question 2: We have been hearing about blockchain in the energy industry more frequently, much in the same way that we heard a few years ago about how “cloud storage” of data was going to improve energy markets.

- What do you think are the reasons for the hype? Is it driven by sellers of computing solutions? Or by customers looking for better answers?
 - Blockchain provides three core characteristics: transparency, immutability, and security that offers a value proposition for some processes. It provides a disintermediation capability between buyers and sellers which reduces costs and speeds transactions. Blockchain, because of some of its design tradeoffs that emphasize these core characteristics, is not applicable for all business processes. It can provide increased trust or this disintermediation role, which may be an attribute.
- Will blockchain as a concept simply fade away when the next “big thing” arrives to promise solutions for the energy industry?

U.S. Senate Committee on Energy and Natnral Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Thomas Golden

- Or does blockchain technology offer something deeper that has real value?
 - Blockchain offers three core characteristics as noted above. If an improved technology comes along to displace it, then the new technology will be the “next big thing”. Often it can be difficult to recognize disruptive technologies when they emerge. The disruptors often underperform the incumbents in some way (with blockchain it is speed of transactions) yet provide a new value proposition that is attractive to new customers. When the underperforming characteristic is addressed as the technology improves, the new technology can be broadly disruptive.

Question 3: We have been talking about a series of issues that are not always that closely related. On the one hand, blockchain has cybersecurity impacts that may be helpful in securing the power grid. On the other hand, blockchain has potential for improving the efficiency of the organized wholesale markets.

- As this committee explores blockchain, can you tell us where we should focus our time and resources? Is it cybersecurity? Markets? Energy consumption? All of the above?

As noted above, disruptive technologies can have broad implications, requiring a comprehensive understanding. Regulatory constructs may need to be in place to foster peer-to-peer market trading. Although built from the “ground up” with security in mind, blockchain has not been immune from hacking or bad actors that ran scams on either exchanges or Initial Coin Offerings (ICOs).

- Do you see a need for legislation? Or is it too early to know?

As a non-profit research and development organization, EPRI refrains from policy and legislative recommendations.

Question 4: I want to explore how blockchain miners might fit into an energy market:

- In a blockchain system used for trading energy, will miners be used to verify a transaction? Or are miners even needed?
 - The mining operation (Proof-of-Work based) gets revenue from mining (tokens, e.g. Bitcoin, Ether), and transaction fees. Revenues must offset expenses for the mining function to continue. The lowest cost provider of the transactional functions will likely have an advantage over other architectures which suggests a Proof-of-Stake, Proof-of-Authority, tangle (which does not require “mining” [the PoW problem solving mechanism]) would be more attractive in the long-run over Proof-of-Work which is more energy intensive.

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Mr. Thomas Golden

- Will miner compensation be added to the consumers' bill?
 - State regulatory policy will be determined by the appropriate governmental entities.
- If no miners are needed – such as in a proof-of-stake or proof-of-authority system – what is the benefit of blockchain over a centralized system?
 - Immutability would be one benefit. Blockchains are “append-only” meaning that records can only be added. Previous records cannot be changed. This is the distinction between blockchains and older distributed databases that had full create-update-delete functions on prior records.
 - Personal Identifying Information (PII) is nominally better protected due to the nature of how keys (the secret account numbers) are distributed, controlled, and used.
 - The promise of smart contracts is that they can be executed automatically when the terms are met, e.g. Renewable Energy Certificate (REC) created, but while smart contracts existed in the past, they did not have the built in security inherent in blockchain (per the bullet above).

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Ms. Claire Henly

Questions from Chairman Lisa Murkowski

Question 1: We have been hearing about blockchain in the energy industry more frequently, much in the same way that we heard a few years ago about how “cloud storage” of data was going to improve energy markets.

- What do you think are the reasons for the hype? Is it driven by sellers of computing solutions? Or by customers looking for better answers?
- Will blockchain as a concept simply fade away when the next “big thing” arrives to promise solutions for the energy industry?
- Or does blockchain technology offer something deeper that has real value?

The hype is driven by two groups. On the one side the blockchain industry as a whole has raised hundreds of billions from its initial coin offerings and subsequent token valuation increases. This market has recently cooled off (which I believe is a good thing) but there is still a significant amount of capital in the space looking to go to work to solve market problems and create long term value. The energy market is a large and compelling target.

On the other side the energy sector and major energy companies are in the midst of significant transformation driven by digitalization, deregulation and distributed energy resources. Energy companies are looking for technical solutions and approaches to solve challenges presented by these changes as well as capture potential value.

While it is still early days for blockchain applications in the energy sector there have already been several promising demonstrations that give both the blockchain industry and the energy industry hope for the value of the technology to the energy sector:

1. Renewable energy certificate tracking and trading which addresses the high broker fees, opacity and high barrier to entry nature of certificate markets ([link](#))
2. Electric vehicle charging station sharing and streamlined payment which unlocks underutilized electric vehicle charging stations ([link](#))
3. Live trading between large energy companies in electricity and wholesale natural gas markets cutting costs of settling these markets ([link](#))
4. Local energy trading in Brooklyn addressing an increasing demand from residential customers to preferentially purchase electricity locally ([link](#))

Both the blockchain industry and forward-looking energy companies are invested in experimenting and establishing the value of blockchain in the energy sector. While it is too early to say if blockchain will live up to its current hype, existing pilots as well as pilots and commercial projects under development give a strong indication that there is real value and actual usage of the technology will not fade away but only continue to grow.

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Ms. Claire Henly

Question 2: I want to explore how blockchain miners might fit into an energy market:

- In a blockchain system used for trading energy, will miners be used to verify a transaction? Or are miners even needed?
- Will miner compensation be added to the consumers' bill?
- If no miners are needed – such as in a proof-of-stake or proof-of-authority system – what is the benefit of blockchain over a centralized system?

The majority of the energy-blockchain pilots and demonstration projects are running either on blockchains that do not require miners (such as Energy Web Chain) or on blockchains that currently use miners but will soon transition away from miners (such as Ethereum). We strongly believe that mining based blockchains should not be used to run energy sector blockchain applications in the long term because of the required electricity use and cost to the electricity sector and electricity sector consumers.

Proof-of-work blockchains create value because they allow a decentralized and trustless network of users to transact with each other and agree on a common record of those transactions without the need for an intermediary and with a high degree of network security. Proof-of-stake and proof-of-authority blockchains are designed to maintain the benefits of proof-of-work systems: a similarly high levels of network security and decentralization, while dramatically decreasing energy use. They do this, in the case of proof-of-stake, by replacing the cost of mining with a staked deposit and cost of capital, and, in the case of proof-of-authority, by using a large number of known and trusted entities to validate transactions, essentially ensuring if a transaction is incorrectly validated, the authority responsible can be pursued. In both cases, just as in proof-of-work, the validators would still be a large and diverse group and would need to be dominated by malicious actors in order to corrupt the overall integrity of the system. Early signs are very positive that these alternate approaches will be able to maintain (and even improve) the benefits of proof-of-work based systems.

Still, these approaches are at an earlier stage of technical development than proof-of-work blockchains. Further technical development and testing is required to ensure that proof-of-stake and proof-of-authority (and perhaps many other approaches) can maintain the benefits of proof-of-work blockchains without the energy costs.

**United States Senate, Committee on Energy and Natural Resources
Hearing on Energy Efficiency of Blockchain and Similar Technologies**

Responses to Questions for the Record

**Arvind Narayanan
Associate Professor of Computer Science, Princeton University**

Chairman Murkowski, Senator Manchin, and members of the committee, thank you for the opportunity to respond to questions for the record. I would also like to take this opportunity to correct an error in an estimate that I provided in my written testimony. I have done so in an Appendix to this document. I sincerely regret my error.

Questions from Chairman Lisa Murkowski

Question 1: You have testified that a blockchain “is a sequence of records that is collectively maintained by a set of stakeholders and is designed to support the addition of records while resisting modification or deletion of existing records.”

- At what point does the length of a blockchain provide diminishing returns for security and verification purposes?
- Would setting a cut-off point for the number of blocks on a chain (i.e., dropping blocks on the back end as new blocks are created) allow for any energy efficiencies – and less energy needed for mining purposes – for proof of work or data storage requirements?

Response. There are broadly two sources of inefficiency in blockchains. The first is the need for participants to track transactions, verify their validity, and perform other maintenance activities. It is indeed true that as a blockchain grows, keeping track of all transactions leads to increasing data storage requirements. Most blockchain designs do include the ability for participants to drop or “prune” old records to decrease their storage costs. This is an area of ongoing research and innovation.

However, pruning does not address the second source of inefficiency, namely the energy needed for mining purposes. The large computational requirements of mining arise not because of the need to track or verify transactions. Rather, it involves a separate type of calculation that performs no useful function except to deter adversaries from attempting to disrupt the blockchain’s operation. Thus, it is precisely the computational inefficiency of mining that allows it to fulfill the function of securing the blockchain. The difficulty of the mining calculations does not vary based on the length of the blockchain.

Fortunately, mining is used only in public blockchains that support cryptocurrencies. The envisioned applications of blockchain technology in the energy industry involve private

blockchains which have no need for mining. This is because a majority of stakeholders are assumed to be trustworthy and there is no risk of an unknown adversary attempting to subvert the system.

Question 2: Blockchain is often promoted as a means to achieve anonymity in market transactions. But much of the trading in energy markets has no need for anonymity. For example, two state-owned utilities may not need or want to keep their transactions anonymous.

· In a market where anonymity is less important, does that reduce the value of blockchain techniques?

Response. Indeed, there is no need for anonymity in energy trading and other applications in the energy industry. While many blockchains, especially public blockchains, enable a degree of anonymity, that is not an inherent property of blockchains. Blockchain-based energy trading can be combined with strong identity assurance of market participants.

Questions from Senator Joe Manchin III

Question 1: In effort to reduce overhead costs, developers look to areas that offer cheap electricity and cool climates. However, some communities across the country are dealing with the impacts of increased electricity prices due to the demands of the data centers. According to your written testimony, “the increasing energy efficiency of mining hardware has essentially no impact on energy consumption.” Can you please discuss the efficiency of these deployed mining machines, and what local leaders should be aware of for those communities where they are being considered by cryptocurrency developers?

Response. Inefficiency is deliberately designed into mining-based blockchains. When more efficient mining hardware is developed, mining becomes more profitable, which incentivizes more computing power to be dedicated to mining. Blockchain protocols are designed to react to such an increase in mining capacity by proportionally increasing the computational difficulty of mining. In this way, the effect of the increasing efficiency of mining hardware gets nullified by the increasing difficulty of mining, and thus there will be little or no net impact on mining energy consumption.

Cryptocurrency miners are different from many other types of industrial electricity consumers because their revenues are dependent on cryptocurrency exchange rates, which tend to be highly volatile. Demand for mining tends to increase or decrease relatively quickly in response to cryptocurrency exchange rate fluctuations. This unpredictability might create difficulty for communities. One potential way for local leaders to manage this unpredictability is to establish (possibly informal) communication channels with the cryptocurrency mining community, which would give leaders advance knowledge of trends in mining demand.

Question 2: What role, if any, can FERC have to provide oversight of blockchain technologies in the energy industry?

Response. Blockchain technology is relatively new and there are uncertainties in a number of areas, such as the legal status of transactional records stored on blockchains. Another concern is the “endpoint” cybersecurity risk, that is, the risk of compromise of the devices that store participants’ private keys and protect their digital assets stored on the blockchain. There is a need for regulatory oversight in promulgating requirements for record-keeping systems to be considered authoritative, as well as in setting cybersecurity standards for the computing systems used in the grid.

Appendix: correction of an error in written testimony

In my written testimony, I stated:

An accepted method for deriving an estimate of the energy consumption of mining is to assume that all miners use the most energy efficient mining device available on the market.¹ Commercial devices are accompanied by published specifications listing the number of hashes that can be computed per second using the device, as well as the power consumption of the device in watts. It is then straightforward to calculate how much power is required to compute 50 billion hashes per second using the most energy efficient devices available. I performed such a calculation and obtained an estimate of around 5 gigawatts for Bitcoin mining alone today.² This is slightly under 1% of world electricity consumption, or slightly more than the electricity consumption of the state of Ohio or that of the state of New York. Other public blockchains also consume a substantial, albeit much lower, amount of energy.

The 5 gigawatt estimate of Bitcoin mining energy consumption, including the footnotes, remains correct to the best of my knowledge. However, the comparisons to the electricity generation / consumption of the world and U.S. states are incorrect. An electricity consumption of 5 gigawatts translates to 44 terawatt hours per year. According to the International Energy Agency, world electricity generation in 2017 was 25,570 terawatt hours,³ which puts Bitcoin mining energy consumption at around 0.2% of this figure, rather than “slightly under 1%” as I stated previously. I sincerely regret the error.

¹ See Alex de Vries, *Bitcoin's Growing Energy Problem*, 2 *Joule* 801-805 (2018), [https://www.cell.com/joule/fulltext/S2542-4351\(18\)30177-6](https://www.cell.com/joule/fulltext/S2542-4351(18)30177-6); Arvind Narayanan et al., *Bitcoin and cryptocurrency technologies: a comprehensive introduction*, Princeton University Press (2016), <http://bitcoinbook.cs.princeton.edu/>.

² The most energy efficient mining device known to be in widespread use is the Bitmain Antminer S9, which achieves an efficiency of 10 billion hash computations per Joule of energy, resulting in an estimate of 5 gigawatts for Bitcoin mining. Recent announcements of new devices have claimed higher mining efficiencies; if these are in widespread use, the true power consumption might be slightly lower than 5 gigawatts. On the other hand, some devices in use may be much less efficient, which would mean that the true power consumption might be higher. Further, accounting for the energy consumption of cooling of mining data centers would also increase the estimate.

³ International Energy Agency, *Global Energy & CO2 Status Report 2017* (March 2018).

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Dr. Robert E. Kahn

Responses from Robert E. Kahn to
Questions from Chairman Lisa Murkowski

Question 1: In your written testimony, you testify that this nation has an opportunity to dramatically reduce the cybersecurity problem in the power grid. However, a fundamental roadblock is the existing energy infrastructure that is already in place. In particular, you state:

“We had a similar challenge facing us in creating the Internet, where it was not practical to cause every existing network to change ... I believe the kind of workaround strategy we used in creating the Internet is implementable in the Energy Grid with only a small amount of help from industry, and (importantly) without requiring significant reworking of their existing industrial control systems.”

- Can you elaborate on how you think that the energy industry can be persuaded to invest in new ways of reducing its cyber vulnerabilities?

Answer: My testimony addressed a key aspect of the energy grid, namely the way in which we monitor and control it. Except as indicated below, wholesale reimplementation of the existing technology seems quite impractical, and so I advocate an incremental approach in which new capabilities to detect, inhibit or respond to threats is demonstrated mainly by external monitoring and/or control mechanisms that work with the existing systems with only minimal change or intrusion to them, if any. The number of possible approaches here may potentially be very large, thus I don't focus on any particular threat scenario here. Eventually, if the value of such new capabilities can be demonstrated (including both functional value and affordability), then industry will likely be incentivized to invest in making the necessary changes on a timescale they can manage.

Should the vulnerabilities become too great for the status quo, however, a less conservative approach may be required. This can occur if the energy grid is actually compromised, or can be seen to be compromised, or if a known threat is deemed so imminent as to require an immediate response. At that point, the situation changes dramatically; and all necessary forces would need to be marshalled.

- And what is your confidence that the grid itself can be fundamentally improved?

Answer: I restrict my remarks here to the monitoring and control of the grid, rather than those aspects that involve the means of energy generation and transmission for which others will have much more expertise.

I assume that external communications are required to monitor and/or control each of the major sites in the energy grid. Otherwise, the threat profile is greatly reduced and largely dependent on the internal actions of people and the associated personnel management techniques. Situations have been reported where even disconnected systems were compromised, but this has to have

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Dr. Robert E. Kahn

happened from within and due (most likely) to violation of internal rules and procedures. It is fundamentally difficult to protect against certain insider attacks, but one can make it increasingly difficult for one individual to carry it out without other insider help. On the other hand, it seems appropriate to examine the extent to which internal procedures are sufficient to protect the grid; and, it would be important to know when systems are compromised due to internal as well as external actions, and then how to respond.

As for external threats due to cyber-related vulnerabilities, I am confident that fundamental improvements can be achieved through combination of 1) architectural changes (with integrated security) that more tightly control the flow of information to and from the system, 2) sophisticated rapid evaluation processes that make the current system status more visible, 3) design changes to seal off conventional maintenance channels that are subject to compromise, and 4) new internal and external maintenance approaches that take advantage of the new architectural changes and integrated security.

Finally, while one can demonstrate such approaches ahead of time, changes in the field can only take place with the active participation of those sectors of industry that are involved in the actual operation of the energy grid.

Question 2: You described digital objects as fundamental to an understanding of both blockchain and cybersecurity.

- What is the advantage of looking at blockchain from the viewpoint of digital objects?

Answer: Much has been written about blockchain being a new conceptual development, but I do not see it that way. Rather, I view it as a particular way of structuring a digital object, which happens to involve linkages with other digital objects. My objective in pointing this out is to put the blockchain notion into perspective by providing context, and to enable a wider discussion about the properties of blockchain technology compared to those that can be achieved with the Digital Object Architecture (DOArch), more generally.

- Can you point to any success you've had in creating systems that use the digital object architecture?
- Are any digital objects being used today? How and where?

Answer: The two parts of this question are answered together below.

CNRI has created reference software implementations of the three components of the DOArch, namely the identifier/resolution system, DO Repository and the DO Registry; and the last two components have been combined into a single package called Cordra. The CNRI software is publicly available at the following CNRI sites: <http://www.handle.net> and <http://www.cordra.org>.

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Dr. Robert E. Kahn

There have been quite a number of effective applications of the DOArch. I will describe a few of them here:

- a. During the 1990s, the library community and the publishing industry both saw the value of the DOArch; and today, a large number of university libraries use the identifier/resolution component of the DOArch to designate their digital library materials. In addition, the Library of Congress was perhaps the earliest adopter of this component to identify certain collections; and this usage continues to the present.
- b. The publishing industry adopted the identifier/resolution component of the DOArch in the 1990s; and it has been in widespread use for almost two decades in the science, technology and medical communities. A joint initiative of three trade associations in the publishing industry (International Publishers Association, International Association of Scientific, technical and Medical Publishers; and Association of American Publishers) created the DOI System in 1997; and they later established the International DOI Foundation to develop and manage the System that is now an ISO Standard. For description of DOI System, see *DOI Handbook*, <http://www.doi.org/hb.html>
- c. The entertainment industry applied the DOArch technology to the management of important entertainment industry assets, in particular movies, television programs and related activities. Under the leadership of the Entertainment Identifier Registry Association, they collaborated on creating and making use of a registry system for management of their entertainment assets structured as digital objects. For overview of Entertainment Identifier Registry (EIDR) as of May 2018, see http://eidr.org/documents/2018-05-08_EIDR_Overview_FINAL.pdf
- d. The Association for National Numbering Agencies (ANNA) commissioned a system to oversee the trading of derivatives. The resulting system is based on the DOArch and enables real-time distributed tracking of such trades, represented as digital objects, made by securities dealers around the world. A demonstration version was first released in 2016 (<https://www.anna-web.org/anna-presents-demo-version-derivatives-service-bureau>); and a production version was later released in 2017.
- e. The British Standard Institute (BSI) joined with the Royal Institute of British Architects (RIBA) to commission a pilot system that can record and make available to authorized parties relevant information about building construction in their country. The resulting system is based on the DOArch and will enable such information to be available well into the future, without the need to recreate the information as technology changes.
- f. The European scientific community has implemented the identifier/resolution component of the DOArch for persistently identifying scientific datasets and related

U.S. Senate Committee on Energy and Natural Resources
August 21, 2018 Hearing: *The Energy Efficiency of Blockchain and Similar Technologies*
and the Cybersecurity Possibilities of Such Technologies for Energy Industry Applications
Questions for the Record Submitted to Dr. Robert E. Kahn

scientific resources; and efforts are underway to make use of the other components of the DOArch to enable scientific workflows and cross-discipline activities. The European Identifier Consortium (ePIC) is centrally involved in these efforts (<https://www.pidconsortium.eu>).

The DOArch is a logical extension of the Internet to deal explicitly with information management and a few principal objectives. One, to manage digital information in a way that is independent of the underlying technology, much as the Internet achieved for communications. Second, to support persistent access to digital objects without the need to expend additional effort in recreating existing digital information (including its metadata) to make use of new underlying technologies as they emerge. Third, to enable interoperability across heterogeneous information systems. And fourth, to integrate security into the architecture rather than adding it as an overlay after the fact.

Industry did not embrace the Internet at the outset, and so the early efforts to build the Internet capabilities involved applications level developments. For example, the TCP/IP protocol was first implemented as a user program; and what would be called routers today were assembled by research groups from component hardware and software made available as a result of research efforts. Eventually, industry recognized the value of adopting the Internet technology and actively participating in its evolution. A similar set of developments is expected here as the value of the Digital Object Architecture is more widely understood, along with the recognition that it is available in the public domain.

