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 the1990s, 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.