Faculty Research Working Papers Series

Ensuring (and Insuring?) Critical Information Infrastructure Protection

Kenneth Neil Cukier, Viktor Mayer-Schönberger, and Lewis M. Branscomb

October 2005

RWP05-055

Ensuring (and Insuring?)
Critical Information Infrastructure Protection
Report of the 2005 Rueschlikon Conference on Information Policy[1]

Kenneth Neil Cukier / Viktor Mayer-Schönberger / Lewis M. Branscomb

Introduction:

To understand the way critical information infrastructure is vulnerable to failure, consider what happened on September 11, 2001:

After the second plane crashed into the South Tower at 9:02 am, telephone calls increased up to ten times the normal traffic volume -- so much congestion that only a handful could get through. Major news Web sites -- CNN, the BBC, The New York Times and others -- were so clogged with traffic they became temporarily unreachable. By 9:39 am many radio stations in the city went dark (most broadcasters had transmitters on the towers). When the first tower collapsed at 10:05 am, and then the second at 10:28 am, they destroyed a vast amount of telecom infrastructure in the vicinity, complicating communications even more.

To be sure, in many instances the systems proved resilient. For instance, network technicians struggling to repair systems coordinated their activities using mobile text messages since their cell phones couldn't handle calls. And as many noted afterwards, the Internet worked when the phone system didn't. Indeed, at 9.54 pm the Federal Emergency Management Agency alerted all stations to prepare in case primary communications methods failed -- and did this, ironically, by email.

But here is the nub: as bad as all this sounds, the actual event did not do too much damage to the information infrastructure -- yet subsequent problems with other networks began to cause havoc. For instance, a fire at a building on the periphery of the World Trade Center knocked out a power station upon which telecoms equipment elsewhere depended. A falling beam from an unstable building in the vicinity crashed into an operator's central switching office, damaging the machines. By late evening, systems that had survived went down simply because they overheated. And telecom services were disrupted when backup generators ran out of fuel because trucks carrying new provisions were blocked from entering lower Manhattan.

In short, the incident highlights both the vulnerability and resilience of information infrastructure -- and importantly, its interdependence with other infrastructures. For instance, the communications network is dependent on the electrical grid; the back-up generators are dependent on the roadway network. And of course, it bears noting, that the target of the attack in New York was not communications infrastructure at all, but two office buildings. What might have been the consequences if critical information infrastructure had been targeted as well?

---

[1] For more information on the conference series, please visit http://www.rueschlikon-conference.org

The success of the information society can be seen in the way it is ubiquitous and taken for granted. However, critical information infrastructure (CII) is more abundant, and also more fragile, than we often admit.

In this context, Professors Lewis M. Branscomb and Viktor Mayer-Schönberger of Harvard University's John F. Kennedy School of Government convened the fifth annual Conference on Information Law and Policy for the Information Economy together with Swiss Re, from June 16-18, 2005, at the Swiss Re Center for Global Dialogue in Rueschlikon, Switzerland. Over 30 experts from industry, government and academia discussed the nature of the problem, obstacles to addressing it, and possible solutions. As is Rueschlikon tradition, the discussion was on a strict not-for-attribution basis, to encourage frank dialogue. Following the conference, a report is produced to document the discussion, which in recent years has been made publicly available as a way to contribute to the broad technology-policy community.

The informal consensus was that industry lacks an incentive to fully address the matter, but that government regulations would probably be ineffective since technology moves too fast and it would distort the emphasis from real security to regulatory compliance. Instead, it may be best to create a market for CII protection -- and participants believed that the insurance industry could play an important role. Insurance is a means to transfer risk, aggregate liability, and create a market for uncertainty. Applied to CII, this could create incentives for companies to invest in protecting infrastructure, while minimizing the impact of potential failures when they occur.

However, a key obstacle is that the insurance industry, in order to operate, requires detailed knowledge of the risks, which is currently impossible to attain because infrastructure owners are loath to share data about their vulnerabilities. This is partly due to antitrust concerns, partly for fear of potential liability. Thus, there may be a role for government to provide benign assistance, by acting as an observer for both the insurance industry and the infrastructure players to meet and exchange information in a way that does not leave them legally vulnerable. This would allow the insurance industry to create products for coverage, and thereby create a market for CII protection.

Importantly, it offers the beneficial side-effect of better preparedness, as industry has a clear incentive to take reasonable action (as happens with fire insurance and fire protection rules). This falls in line with the historic role played by the insurance industry to create incentives for good practices, from healthcare to auto safety. Moreover, the insurance industry, through a market-based, risk-analysis approach, could foster better planning and response to disasters by prioritizing and valuing risk.

Moreover, an idea was advanced that to help create such a market, governments could indemnify the insurance industry for catastrophic CII failures, akin to the United States's Terrorism Risk Insurance Act established after 9/11 (which unless renewed by Congress, will expire at the end of 2005). It provides for the federal government to reimburse insurance firms against insured losses up to $100 billion a year that arise from terrorism. Applied to CII, a similar sort of indemnification would remove the biggest stumbling block -- the consequences to the insurance industry itself of catastrophic CII failure -- and allow the marketplace to concentrate on the more manageable, and more common, aspects of CII

protection. Indeed, a full assessment of CII vulnerabilities is impossible to attain in part because terrorism risks have elements of uncertainty -- of motives, intent, and the capability of the perpetrators -- that is uncommon in other sources of high risk.

As a start, the Rueschlikon participants agreed that it is imperative to bring the insurance industry, infrastructure operators and governments to the table immediately to discuss these matters together, and more importantly, to act.

In that spirit, this report is designed as an overview of the Rueschlikon discussion as well as a roadmap for further work to take place. It comprises five sections. The first considers what is meant by critical information infrastructure (and from here on uses the abbreviation CII) and its vulnerabilities. The second section examines the economics of protecting CII and the paucity of information about risks. Section three identifies examples of good crisis management that can serve as models for CII protection. The fourth section notes the possible role of different stakeholders. The fifth section considers potential solutions (and specifically examines how the insurance industry can act).

The British writer Samuel Johnson once quipped: "When a man knows he is to be hanged in a fortnight, it concentrates his mind wonderfully." However, the problem with CII protection is that without a real incentive to address it -- and with a plethora of commercial pressures to be complacent -- the matter is ignored until the noose is knotted, and by then it is too late. As such, the group shared the uncomfortable but pragmatic ethos of "thinking the unthinkable," in the famous phrase of Hermann Kahn, which was originally applied to considering the effects of nuclear war. It is with this sense of purpose that participants examined the vulnerabilities of CII and, more optimistically, the ways to protect it.

I. Identifying the Problems

> *"The sky is not falling, but it is raining pretty hard."*
> -- Rueschlikon participant on continuing CII problems

A. Defining CII

There are three problems to understanding critical information infrastructure: What is meant by critical? What constitutes "information" systems? And what is infrastructure? This is not to be pedantic, but a quick set of definitions are needed in order to understand the issue. A good starting point is to concentrate on the services that the systems provide.

For example, the telecom network is important because it drives other critical infrastructures, from the power grid and gas distribution, to civil aviation and emergency services. In the modern banking system, money is simply data that swishes around a network -- and thus might be considered a component of CII, too. The US Marsh Commission in 1997 identified seven critical infrastructures (of which "information and communications" was but one); today, the official list notes 13, and is growing.

Among the ganglion of interconnecting networks that comprise CII, there is some commonality in the way that it is treated. Security experts identify four elements of CII

protection: prevention, detection, response and recovery. Each step builds on the previous one. Moreover, all entail tradeoffs, and in a context of uncertainty, manageable risks and affordability. In some instances, these tradeoffs do not represent shortcomings, but versatile approaches towards CII protection. "We may not be able to defend, but we may be able to detect and respond," explained one participant.

A high threshold is needed when one attempts to identify something as truly critical. For instance, many important systems are self-repairing or self-healing, such as the way that Internet traffic routes around damage, or how motor traffic still continues to flow even if traffic lights go out. Thus, even in failure, some operations can still be sustained. What is needed is a far more difficult level of judgment and calculation. "We can live with the loss of services and many killed. In evaluating security risks, we have to consider national security more than simply the loss of life -- it sounds cold, but it shouldn't," explained one participant.

As for what is "infrastructure," it is a system. It may be comprised of many different components and networks in a single company, or how many firm's networks interoperate. It entails both tangible objects like computers and fiber-optic cables, as well as intangible things like software and data. Meanwhile, the issue of what is "information infrastructure" can be murky. Looking at computers and communication networks are only a part of the picture since they are the bedrock of all other infrastructures. For example, so-called SCADA communications (for "Supervisory Control and Data Acquisition") drives many of the key services on which the public depends, from the energy grid and water supply to aviation. It is important to narrow the issue down to those systems that deal primarily with data-processing and information-transfer within and among firms.

Understanding how to protect CII becomes problematic because the definition of CII means different things to industry (which tends to own and operate the infrastructure) and government (which relies on it, often regulates it, and at times must protect or restore it). One participant explained the difference in approaches to what constitutes "critical" this way:

> *** Private sector**: "A consequence to the firm of failure, of sufficient severity as to warrant purchasing insurance or investing the equivalent cost for substantially reducing the risk of failure."
>
> *** Public sector**: "A consequence to society of failure, justifying the deployment of emergency measures and resources, and the enactment of legislation to reduce vulnerability."

A useful approach to understanding CII, then, is to appreciate how important its failure is not just to a single actor (such as a company or a service), but the consequences system-wide and its impact on other critical infrastructures. "It is not about protecting the infrastructure, but the things they serve -- it is a tool," explained one participant. These working definitions let us better recognize some of the risks concerning CII.

B. Understanding the Risks

CII protection is made more difficult due to the inherent features that make it so useful. It is decentralized, interconnected, interdependent and controlled by multiple actors (mainly private) and incorporating diverse types of technologies. The effects can be severe, even if the outages last a short duration. It is almost axiomatic of CII that localized physical disruptions affect systems much farther away. Failure in software can affect hardware. Likewise, cyber problems have direct, physical-world consequences. Indeed, the Internet can be used as a "force multiplier" to amplify the effects of a traditional attack, either by spreading misinformation or disrupting the network so that there is a lack of information, among other things.

Failures can come in a variety of ways and may be due to a myriad of causes. In most cases, they are either intentional (vandals, criminals or terrorists), accidental (natural disasters or human foul-ups), or due to poor decisions (be it in engineering, management or regulation). With the current focus on terrorism, it bears recalling that the majority of high-consequence events come from natural disasters or human error and bad design. When people are responsible, the problems most frequently come from insiders.

Meanwhile, problems cascade across the network, and the speed at which this happens has quickened. For instance, in 1999 the Melissa virus took three days to spread across the Internet; in 2001, Code Red took minutes. When a blackout blanketed the Northeast of the United States in 2002, it was due to a relatively minor power surge in one part of the grid that gave operators a mere 42 seconds to respond before bringing down the entire system. "That is the time frame we're looking at now," explained one participant.

There are a number of inherent paradoxes regarding CII protection. For example, the entanglement of different networks that comprise CII is a source of complexity and vulnerability -- though it can also be the source of redundancy and protection, too. Moreover, as businesses strive for efficiency in their operations, they eliminate the slack in the system that creates resiliency, thus making themselves even more vulnerable and less able to recover from failures.

What is more, redundancy does not always remedy problems: it may simply make problems recursive, such as in the case (possibly apocryphal) of the aircraft whose second engine failed after the first one had died, because both were serviced -- improperly -- by the same ground crew. Indeed, back-up systems do not help much if they simply replicate the technology of the primary system, since it may replicate a fault, too. To this, many participants expressed a deep concern about the homogeneity of IT, such as for PC semiconductors and software, and Internet routers. These "technical monocultures" represent key sources of vulnerability, and it was recommended that a diversity of products be encouraged.

Meanwhile, the Internet's much-vaunted decentralization, though helpful, is not the answer, either. This is because, contrary to myth, the topology of the Internet is not totally decentralized; rather, it looks more like a classic "centralized" hub-and-spoke system -- there are users at the edges, and content at the center. Moreover, not all nodes on the Internet are equal, and some handle far more traffic than others, called "super-nodes."

This requires us to rethink our conception of the architecture of our networks relative to CII protection. For example, we built the Internet precisely to overcome one set of vulnerabilities to the telecom systems yet unwittingly introduced others. The ARPANET was hailed for being decentralized and redundant -- but all the hardware was built by a single company, BBN, which created risks. The situation is little changed today, with the omnipresence of Microsoft and Cisco across PCs and the Internet.

The tradeoffs of CII are multifold. What is robust (that is, difficult to break) is often not resilient (flexible in times of crisis). What is most reliable from a technical perspective may not be the most efficient from an engineering standpoint. And in terms of protection, what is "optimal" is by definition not "best." Thus, as businesses strive for optimal solutions, are they doing enough or placing themselves in jeopardy? The good news is that the idea of super-nodes underscores the issue that not all elements of information infrastructure are truly critical -- and what is vital is to identify the elements that are, and to secure those.

C. Operational Realities

Technology is always improving, and as it does, new vulnerabilities emerge. In 2003, the Computer Emergency Response Team (CERT) at Carnegie Mellon University identified around 137,000 separate intrusions and attacks under a voluntary reporting system. It was a major jump from 82,000 in 2002 and only 53,000 the year prior. That they cause disruptions, even if temporary, suggests there are technical problems to addressing cyber-security. Moreover, thousands of these incidents are considered to be "test attacks" whereby the perpetrator seeks to identify vulnerabilities that can be exploited in subsequent attacks -- something that makes hacking different from other forms of attacks. And for all the doomsday data, the self-reporting approach was considered by CERT researchers to undercount the extent of the problem -- capturing only about 1 in 100 security incidents -- so that the annual study was discontinued. CERT, meanwhile, continues in its role as the coordination center for American IT security, in partnership with the Department of Homeland Security.

Of course, the technical means to protect against CII failures are developing. But it can never do so fast enough or perfect enough. Among Rueschlikon attendees, some technologists stated that the tools already largely exist to protect CII from nearly all the potential failures, but they are not deployed -- which may point to an economic failing rather than a technological one. However, this view was not shared by others; the majority of participants believed the shortcoming of technology to protect CII will always exist. Indeed, it must be considered the starting point for CII protection discussions, not the objective. What technology may do, however, is increase the comfort level to a point where CII is reasonably protected.

In this context, concentrating on the weaknesses in commercial software and hardware products is counter-productive. While more effort must take place to ensure that the systems are better designed, imposing regulation or tort liability on the manufacturers would penalize the maker rather than the miscreant who exploited the vulnerability, as well as squelch the innovation that the products bring users.

Rather, the technological environment itself may provide lessons on how to reconcile the different interests of the public and private sector in CII protection. One participant explained this by identifying the way in which CII networks are "modular," and the benefits this brings. Consider the topologies of networks. The most efficient is a hierarchical network: there is no redundancy, and a high degree of specialization to maximize economies of scale. Though the most efficient, it is also the most vulnerable. Alternatively, the other type of network is a nonhierarchical, random network, which is more robust. One gives up efficiency for robustness against failure.

In these network topologies lies the efficiency-resiliency tradeoff. Generally speaking, it is in the interest of public-sector networks that there is robustness, and in the interest of the private-sector networks that there is efficiency. These seem to pull in opposite directions. But in the real world, networks are not strictly hierarchical or random, but a blend of the two. (In fact, they are "scale-free" networks, with hubs or super-nodes that emerge, and are highly clustered.)

What is most striking is that there is a certain modularity in how these networks emerge, without it actually having been designed as such. What modularity provides is not efficiency but adaptability. It accommodates complexity and change. It is also open to innovation. In the business world, modularity is something desirable. For instance, computers are very modular systems, and this allows for interchangeable parts and a diversity of components. Modular networks are robust against random failure, yet vulnerable to targeted attacks on the hubs. Yet the tradeoff between what is in the private and the public interest is not as severe as one might think, since networks are not really just efficient or just resilient, but a combination of the two qualities. Clearly, we can engineer modular networks to embody the values of CII protection that we want to most promote.

More broadly, just as the owners of infrastructure and the threats are diverse, so too should be the solutions to address them. As a first step, it is important to protect systems from the inevitable stupid mistakes -- and small investments would eliminate some of the most common vulnerabilities. It should be done first, but not lull people into a false sense of security that the issue of CII is being addressed. "It is necessary, but not sufficient" said one participant.

The bigger issues, however, are far more difficult to treat because they entail tradeoffs among scarce resources. As one participant put it: "Do you build a dyke -- or do you let it flood over?" That is, does one pay for prevention, or the clean-up costs after the damage? It entails a cost-benefit analysis that raises economic questions, something considered in the next section.

II. The Economics of CII

> *"Critical things are things we have to protect, even though there is no amount of money that we can spend to really protect them perfectly."*
>
> -- Rueschlikon participant on the dilemma of CII security

A. Whose Costs, and Who Pays?

The amount spent on information-technology security worldwide comes to around $100 billion annually, and is growing between 5% and 10%. Companies generally spend around 5% of their IT budget on security, according to the research firm IDC; 40% of IT managers rank it their top priority. Will more money make the systems more secure? How can we be sure that we are not actually overspending on security relative to the risks? This, considering that the efforts simply reduce the probability of failure, but can not eliminate it. Moreover, when there is an outage, the business disruption is usually temporary and only harms the affected firm -- it does not mean a loss economy-wide.

"The idea still persists that the problem is that we're not working hard enough. But a plausible argument can be made that there is not much more we can do to improve security very much -- and that we do the right thing by ignoring the problem," offered one participant. Said another: "All things are resource-limited -- money, brains, political will. This idea of protection 'at all costs' is just hot air. A lot of money is to be made by scare-mongering." Thus, should we protect CII? Perhaps it is more sensible to accept occasional outages?

The answer, of course, is "no"; the very nature of "critical" information infrastructure is such that we are obliged to secure it as optimally as possible given the costs and benefits. Yet to make a decision on how to do this, it is necessary to briefly examine the economics of CII protection. The issue is complicated because risks are difficult to estimate. Markets rely on information, but in this area, there is very little data about the probability of failures and their financial costs.

That makes it difficult for firms to reach decisions on how to handle the matter, and for a "market" to form to address it. On the contrary, there is a perverse commercial incentive for firms to internalize the risks and costs. Competitive pressures increase the reluctance of companies to invest in avoiding the consequences of CII failures. Coupled with this is a big free-rider problem: if most firms invest to secure their systems, some players will capture the benefits without paying, and thus no one wants to pay.

Taken together, the situation suggests a market failure for CII protection. Understanding why will help when it comes to identifying potential solutions. Thankfully, basic economic theory goes a long way towards explaining both the problems and possible answers.

Simply put, CII protection falls into a dimension of economics and regulation that are typified by two sorts of scenarios, known as "joint-care" and "alternative-care" cases. In alternative-care cases, it is efficient for only one of the multiple parties to use precaution, generally the party who would pay the least to remedy the problem. The result benefits everyone, but it happens by placing a big burden on just one party. For example, to prevent sparks from railroads starting fires, either the farmer plants his crops far away from the tracks, or the trains have to slow down considerably. Of course, neither option is good for either party.

On the other hand, in the joint-care case, it is efficient for multiple parties each to use some precaution -- for instance, railroads can pay a little to install spark-guards, and farmers plant a little further back. Both parties thus share the cost in a manageable way, and the risk is greatly reduced. As it happens, joint-care cases are far more common than alternative-care cases because there are diminishing marginal returns to precaution and as a back-up to human error. Moreover, joint-care cases are more common because the most efficient party to take the precaution (say, a bad driver) might be immune to effective sanctions (i.e. probably wouldn't have car insurance were it not required).

Unsurprisingly, it is large-scale, joint-care cases that pose the most challenging problems to a society, from highway safety to environmental protection to treating communicable diseases. Protecting CII clearly fits into this category. The situation poses a classic collective-action problem, made more complicated by the large number of entities, their diversity of interests, ever-evolving security technologies and the ever-increasing number of vulnerabilities and threats. It raises the thorny questions of who can effectively use precaution; what are the costs; how much precaution and what type should be the responsibility of each player, and how to enforce this?

What is most important is that market forces alone do not solve large-scale, joint-care problems. Economic actors normally seek to maximize their private interests -- and investing in more precaution comes out of one's own pocket while the benefits are shared. This creates a pernicious sort of market equilibrium, whereby everyone hopes everyone else is using precaution, but in reality no one is, because no one is incentivized.

Classic market mechanisms that could lead to CII protection are not working. For instance, the firm's interest to protect its property is not sufficient, since there is a financial incentive to ignore the risks -- complacency is, in economic terms, a "revealed preference"; there is no return on investment for CII protection. The threat of legal liability does not exist since minor failures can be covered up, and there are no precedents of major financial penalties against CII providers for failing to secure their systems. Indeed, courts may find this sort of "negligence" is in reality an "efficient error," in that it would be excessively costly to avoid these risks. Meanwhile, contract and association regulation is not working since there is no incentive for this to happen.

Finally, the area where the insurance industry traditionally plays a role -- in its ability to encourage good practices -- is not brought to bear, since there is no "market" for CII protection. This is largely due to the lack of information in the area.

B. Information Problems and Quantifying Risks

Markets rely on information. But in the area of CII, very little information exists. The consensus among Rueschlikon participants was that this single factor constituted the central obstacle preventing better CII protection generally, and the development of market solutions to address it, specifically.

The paucity of information is twofold: first, about the risks to CII, and secondly, about the costs of failure. It is fueled by incentives of individual firms and entire industries to ignore

the matter. As previously noted, firms have a rational tendency to internalize the costs of CII security, so the situation is not adequately examined and in many firms -- with notable exceptions -- data is not collected. There is an obvious resistance to disclose security incidents due to the damage to one's reputation. Moreover, some firms treat infrastructure protection as a matter of competitive advantage, and prefer the private gain to the system-wide welfare. Clearly, companies would be loath to share information with others should it fall into the hands of rivals. Transparency also offers an information asset to bad actors who might use it for the purpose of harm rather than protection (albeit this is easily overcome through third-party intermediaries who could anonymize the information and provide it only to eligible entities).

As a result, the extent of the risks and costs of CII protection is unknown. Furthermore, the classic models of measuring uncertainty do not so easily apply to CII protection. CII is inherently fragile: there is a major asymmetry between risks and consequences. In other areas, the increase in certain behaviors leads to the probability of measurable danger; for example, driving faster increases the chances that a wreck will be worse. But with CII, even very small vulnerabilities can extract huge consequences; for example, a small hole in browser software can be used to breach the entire network, or a distributed denial-of-service attack can happen, where many individual computers can be harnessed to take down a major site -- just as Gulliver awoke to find himself tied down by a plethora of tiny Lilliputians.

One major element of uncertainty in assigning risk is the fact that thousands of attacks on the Internet, of unknown motive, are in fact effective -- but although there have been some attempts at massive cyber-war type attacks aimed at bringing down a major portion of CII, none have so far succeeded. The Internet's root server system, for example, is under continual attack, and although one particular attack in the autumn of 2002 caused a temporary disruption to a server, the system itself was unscathed. This points to the fact that the uncertainty surrounding CII makes creating a market for insurance very difficult. Indeed, CII vulnerability has the unique feature that the attacker hides behind a remote and probably unknown border, beyond the reach of law enforcement agencies.

The axiom "If you can't measure it, you can't manage it," applies. Even attempts to quantify the risks have failed, as in the case of CERT, whose voluntary disclosure system led to such inadequate data that instead of reforming the study, the study itself was discontinued. "We tend to over-exaggerate the problem -- using phrases like 'a cyber nuclear-attack' -- when we don't know how serious or likely a problem it is. That is counter-intuitive. It scares away the private sector from helping," explained one attendee.

The situation is compounded because markets solve problems through competition and diversification -- different assets and a variety of approaches can diversify risks. But in many cases, infrastructure exists in monopoly situations (because fixed costs are high and marginal costs are low). So one regularly sees underinvestment in areas like infrastructure protection, and may need government regulation to require it.

Ultimately, the situation creates a vicious cycle: because there is no information, there is no market; meanwhile, the absence of a market results in no information being generated. The question, therefore, is how to break this cycle -- how to form a "marketplace" for CII protection?

C. Making a Market

Markets cannot act if they lack the information necessary to set signals. The conundrum is how to create an incentive for firms to first generate, and then share, information about their CII risks and costs, so that market solutions can form. Companies may not be willing to disclose security incidents unless legally required to do so. And then, disclose what, to whom, and how frequently? What is required is a way to generate information in a manner that is in the self-interests of firms to contribute, and with major disincentives to remain mute.

The benefit of establishing a market for CII risks is that it can aggregate cost and risk information over a large number and diversity of actors. Moreover, a well-functioning market can divvy up responsibilities by identifying who is the most appropriate body to pay for certain precautions. Market approaches are often reasonably self-enforcing and flexible with respect to technological change.

One way to achieve this informational output is through law. As it currently stands, legislation requiring CII security or disclosure of security incidents (data breaches, infrastructure failures, etc.) are either narrow, embryonic, ill-defined or ineffective. In the United States, a number of laws require some form of data-security, audits and adherence to best practices. However, they narrowly apply to publicly-traded companies (for accounting compliance and corporate governance under Sarbanes-Oxley); to hospitals and healthcare providers (under rules covering patient privacy); and to financial services firms. The small number of enforcement cases does not suggest that all is well in Camelot, but that the rules are probably not being applied as rigorously as intended.

Moreover, US policy was revised in 2005 so that telecom outages are no longer required to be disclosed to the Federal Communications Commission as a matter of public record. The feeling was that in the post-9/11 world, the information could be used by terrorist, though critics point out that it is far more likely a way for telecom operators to hide poor performance.

The one success has been a pioneering California law in 2004 that obliges companies to notify customers if their personal data is compromised -- the reason why over the past 12 months a large number of security breaches have come to light in the US. A host of federal legislation along the same lines is being considered in Congress, which could mark a huge first step towards the generation of information about the risks and costs of CII protection.

In Europe, the 1995 data-protection directive requires companies to have detailed knowledge of how they handle and share personal information, including about the security of the IT systems in place. The law imposes stiff penalties on companies that violate the rules. However, there are three major shortcomings. First, the directive focuses on the data itself, but says little about the infrastructure in its own right. Second, as in the US, there has been almost no enforcement (in fact, few companies comply with the letter of the law). Finally, there is no requirement for public disclosure of breaches. As a result, there is little information available to the market related to CII.

In order to achieve the degree of transparency that the market requires to assess risk, there are a number of things industry -- working independently or alongside government -- can do. First, associations of network operators can form and agree on performance standards, and firms can certify their own compliance with those standards. Furthermore, they can agree to undergo regular security audits against those benchmarks. What is more, the results of those audits can be made available in anonymous form to a neutral intermediary, who would share it with eligible parties. This would facilitate the information-exchange that is needed, which, in turn, could be used to create a "market" for CII protection.

However, markets operate not only with positive incentives but negative ones, as well. As such, the approach could be coupled with fines and other penalties for firms that fail to live up to a specific standard of security. Moreover, the adherence to a set of best-practices suggests that if they were not followed, firms could be held liable for negligence. Taken together, firms would have both carrot and stick -- and a market could form.

The idea of tying CII protection to market forces seems simplistic or naïve. Yet in many other instances of "large-scale, joint-care" situations, it has been shown to work. In the words of one attendee: "There is a reluctance to believe that the market is a solution to the problem. In the 1970s, in the environmental movement, there was the strong belief that there was a fundamental inconsistency between environmental values and marketplace techniques -- but now they are the biggest supporters of markets. The only way to solve this is through the market process. There is a knowledge-problem over how to assign responsibility to parties in society. Government is one way, but markets are better."

However, some felt that this reliance on market mechanisms was only one dimension of the cure, and that it should not foreclose other activities, such as classic (but hopefully enlightened) government regulation. For instance, one person expressed skepticism that relying on the market to solve CII security would work, since it seemed to fall too neatly into the modern ideological mantra that markets solve all problems. "Doctors right before the great plague were asked what they could do for it, and replied with the belief that religion is the ultimate solution to everything," he explained. "That was the 14th century answer," he concluded.

III. Elements of Good Crises Management

> *"In 1914, we were caught totally unprepared. In 1940, we were totally prepared -- for the First World War."*
>
> -- Rueschlikon participant, quoting a European military officer.

The paradox of critical information infrastructure is that no matter how much security and protection is brought to bear, vulnerabilities and risks of failure remain. The customary ways in which these kinds of problems are addressed -- through technology, markets or regulation -- are not effective. The technology is insufficient, the marketplace doesn't operate, and governments have been reluctant to act. However, from an institutional framework, a

number of early-stage initiatives involving both the private and public sectors have taken place. A brief examination of them is useful, since they may serve as models for how CII protection can be established on a formal, institutional basis.

There was an important recognition among Rueschlikon participants that CII protection requires a system-wide approach. On one hand, this is because failures are rarely ever isolated, but inherently affect other infrastructures. Thus, any solution must take into account all CII players. On the other hand, the system-wide approach is needed for a more subtle reason: the way in which a network can be harnessed by an attacker, to itself be the attacker.

This idea is best highlighted by the anthrax scare in postal networks after 9/11. Though the problem would appear to be Internet-like, in that it seems to be an end-to-end concern -- one bad sender at the edge of the network, linked to one vulnerable recipient also at the edge -- this view profoundly misunderstands the character and magnitude of the threat. Instead, the anthrax incident represents a system-wide problem, because a single letter containing anthrax actually contaminates all other mail as it passes through postal machinery. The crisis was neatly summed up by one participant, in his insistence that it required a change of thinking from: "My God! My system can be attacked!" to "My God! My attackers are attacking me with my system!"

Preparing for the unexpected -- an oxymoron, of course -- is necessary. Yet many of the very policies in place are adapted to certain sorts of situations that may no longer apply. For instance, civil aviation procedures for hostage crises prior to 9/11 were: "get the plane on the ground and negotiate." Obviously, it was rendered an obsolete and useless approach in a new age of terror, where planes themselves became weapons.

As such, an institutional look at different approaches to protect network infrastructure is in order. This section considers two examples of cooperation (among national postal operators, and American telecom carriers), and concludes on broader lessons for CII protection.

A. The Paris Initiative of Postal Operators

In the weeks following the terrorist attack of September 11, in the autumn of 2001, there were numerous incidents where letters containing anthrax were sent through the US postal system. It led to the closure of some Congressional offices, the evacuation of some business offices and widespread public panic; remarkably, only five deaths were reported. In the midst of the crisis, the question of whether service should be suspended was seriously considered -- which would have created an unimaginable backlog of the nation's 700 million pieces of mail the US Postal Service handles daily.

There were even many copycat cases of false alarms around the world, such as white powder landing on the lap of the *Wall Street Journal's* Malaysia correspondent as she opened her mail (innocuous talcum powder, it turned out). Or the hapless Bostonian who screamed upon opening his newspaper when white power tumbled out. Emergency responders raced to the Financial District, where under questioning he admitted that he had enjoyed breakfast at Dukin' Donuts, and indulged his taste for donuts dusted with powdered sugar. More than

simply frightening the recipients, these incidents led to huge disruptions. In some cases, decontamination teams were forced to collect the material and re-sanitize the environment, and the police had to launch investigations. Public fear, along with misinformation, spread quickly.

Six months later, in April 2002, experts from Ecole Polytechnique in Paris, in conjunction with La Poste in France, began a national "debriefing process" to identify the key lessons of the incident so as to better prepare for future crises. It quickly became apparent that considering the issues at a national level was insufficient, since the postal system is inherently international. Preparations for a global dialogue were therefore begun in June 2002, via a pre-existing organization, PostEurope, a trade association comprising at the time of 42 national mail operators across Europe.

In November 2002, one year after the anthrax incident, nearly 30 national postal operators met in Paris to share experiences, note operational capabilities and identify crisis-management lessons. In addition to European operators, the US Postal Service also participated, as did two major mail organizations, the Universal Postal Union and the European Postal Regulation Committee.

One tangible outcome was the agreement to create a new communication structure among all European postal operators, so they can respond within the first 24 hours of a crisis with a joint strategy. The network went live on January 15, 2003 -- and on that day had its first formal emergency alert. The US Postal Service issued an advisory of a suspected anthrax contamination in the post office responsible for the US Federal Reserve. This made other operators aware of the incident and assess the scope of the risks as it affected them.

Overall, there is much to praise about the initiative. Broad international discussions were held, and concluded with concrete action. However, there are a number of noteworthy shortcomings, as well. The meeting took one year to organize; only 30 operators participated (not even all of PostEurope's membership); and it is limited to national operators in an era of postal liberalization and the rise of private mail firms. Moreover, a communications structure is only a first step towards adequate cooperation. Unless discussions continue regarding how to actually respond to crises, and include private postal operators and ones from all nations, the industry may be lulled into a false sense of security. It will have surely "fought the last war" -- in the words of the aphorism by the French general that started this section -- at the expense of preparing for new, unforeseen threats.

B. Telecom Security Coordination in the US

The attack on the World Trade Center on September 11, 2001 severely damaged communications infrastructure, as the introduction to this report explained. However, the far more remarkable aspect of that day and its aftermath regarding telecommunications is the speed with which the systems recovered. For instance, many recall that New York's stock exchanges remained closed for four trading days after the attack; few realize that this was largely to calm market jitters and a bow to fairness for some traders who lacked reliable service, not for technical reasons -- much of the critical communications infrastructure was operational within two days following the attack.

Consider the magnitude of the recovery: after the attack, around 200,000 telecom circuits were down, 3.6 million data lines went out, 10 cell sites were destroyed and service was off for 40,000 businesses and 20,000 residential customers. The result? Emergency responders (police, firefighters and ambulances) received priority service and their communications never experienced interruption; the roughly 1.5 million data lines and 2 million telecom circuits that comprise the stock exchanges were operational within days. Business service was restored quickly, and eventually, residential service fully back within three months, not a matter of years. Where businesses prepare their response and recovery to a single crisis, the attack in New York and Washington, DC, represented multiple emergencies, explained one participant. "We had to respond to the devastation and be on high alert, not knowing if further attacks were coming."

There are three main reasons why the nation's telecom infrastructure was able to respond: first, decades of cooperation among industry and government; second, preparation for crisis situations; and thirdly, experience with emergency-response on a company-wide, as well as industry-wide, basis. A brief look at all three areas follows.

Many organizations comprise America's National Communications System (NCS), which was created by President John F. Kennedy in 1962 after the Cuban missile crisis, when communications problems among the US, USSR, NATO and other countries complicated the crisis. (Europe, in contrast, does not have such elaborate telecom-security institutions, which is examined in the next part of this section.) Today, three key organizations coordinate telecom security in the US.

As a starting point, the National Security Telecommunications Advisory Committee (NSTAC) was created in 1982 as a way to ensure that the president and a handful senior officials retain communications even under crises situations. It comprises around 30 chief executives of firms mostly in communications and information technology, and acts as an advisory council on telecom and cyber-security policy.

Additionally, the Network Reliability and Interoperability Council (NRIC) identifies best practices (but not standards) based on collaboration among industry, alongside the government and public. It was formed in 2002, out of the Federal Communications Commission's Network Reliability Council, which itself was established in 1992 to ensure disaster recovery strategies. There is a high degree of compliance, and the government sets broad goals but participates as an observer and listener.

Lastly, the National Coordinating Center's Telecommunications Infrastructure Information Sharing and Analysis Center (Telecom-ISAC) is made up of around 30 operators and trade groups to exchange information among industry and between industry and the government. It manages real-time coordination among telecom operators in crisis situations. ISACs in general were formed for numerous industries in response to a 1998 presidential directive on protecting critical infrastructures. The Telecom-ISAC's "watch and analysis operation" started running full-time, 24 hours a day and seven days a week, in September 2001.

The positive aspect of these groups is that issues are discussed, information shared and preparations for emergency responses can be examined. (For instance, the major telecom

operator servicing New York had a plan for disaster recovery in the city and drills had happened; on 9/11, those plans were activated.) However, it is unclear whether having many different groups dealing with different facets of the issue is optimal, or if a degree of centralization would improve preparedness and coordination.

The second reason for the quick recovery of communications after 9/11, in addition to organizational factors, was previous preparation for crisis situations, specifically, the experience of preparing for the Year 2000 date changeover. Prior to the turn of the millennium, there was a serious fear that many older computer systems that recorded the year in two-digits (i.e. "99" for 1999) would confuse the year 2000 with 1900 and malfunction. In reality, no such crises occurred.

But it turns out that the experience was extremely valuable for unanticipated network security purposes: it forced firms to perform an inventory of their IT systems and assess their state -- something many had never done before. There was a safe-harbor for information sharing among companies. Even the public markets exerted a watchful eye -- the Securities and Exchange Commission required regular company disclosures on the state of Y2K compliance. And, most importantly, there was a non-negotiable deadline for all this to be completed: December 31, 1999, of course!

The third reason that accounts for the recovery of communications after 9/11 are previous emergency situations. These served as valuable training for the far larger crisis that 9/11 presented. Most notable was Hurricane Isabel, which tore across the East Coast of the United States in September 2003, resulting in over $3 billion in damages. It forced telecom operators to treat the crisis with a "cross-sector response" since not just communications, but other infrastructures like power, gas and road transport were affected as well.

To be sure, there are dissimilarities with other crises. For instance, in this case there was advance notice of the impending disaster due to weather reports. That allowed telecom operators to implement pre-existing plans for marshalling resources. Still, the coordination within companies, industries and across sectors provided useful lessons for disaster recovery, and gave individuals experience they would use on 9/11.

Taken together, these examples provide both good and bad news for the prospects of CII protection. Beneficially, it shows what preparation and practice can accomplish when the unexpected strikes. However, it may ultimately offer false optimism by obscuring the need for cross-sector coordination. As one participant stated: "The telecom world has a good way to talk to one another. The problem is within structures and across structures."

C. Other Institutional Lessons

In addition to the experiences of postal operators and the American telecom industry, there are a number of other specific recommendations for CII protection. These concern organizing groups, prioritizing problems and the pace of addressing them, organizational structures, and the ability of organizations to change.

The first is the role of a convening agent, to bring businesses together within and across industries. The central problem in both preparedness and response is the presence or lack of collaboration, and the prioritization of efforts (and, of course, the key issue is one of trust). The Paris Initiative does this for national postal operators, and America's National Communications System does this for telecoms. In Europe, the coordination is just starting to be developed. Although telecom-security institutions exist at national levels (and at different levels of sophistication), only recently has it been established across the European Union, now boasting 25 member states.

To forge cooperation is the task of the European Network and Information Security Agency (ENISA). It was formed in the past few years and is poised to begin operations. In the summer of 2005, it prepared to grow from five initial officials to a full staff of 40, and open a permanent office in Greece. Once operating, it is intended to act as a clearing house for information (like the US CERT, mentioned in section two), identify best practices and encourage information sharing among groups regarding cyber-security and CII protection.

The ENISA initiative marks an important first step. It is a tangible example of a broad, international CII protection program. Where the US has a more developed telecom security system, this is partly due to the luxury of being one nation, rather than a community of twenty-five nations, to make coordination easier. Furthermore, where the US is characterized by multiple, overlapping groups, ENISA is endowed as the only trans-national EU organization addressing the matter. However, the proof will be in the results. Considering the delays and diplomatic obstacles that ENISA initially faced, the program raises the concerns whether governments are the best-placed entities to lead CII security.

Another lesson is the need for companies to be in "crisis mode," in the words of one participant, so that CII protection is treated seriously in calmer times and preparations can take place. During emergencies, there is incredible cooperation among firms, but that dedication to address the issues dissipates as soon as the emergency is over, making deeper CII security cooperation difficult to achieve. "When the adrenaline goes down, people go back into their boxes. It takes months to set up a meeting, and then half the people show up," explained the participant. "The day-to-day management becomes more important than long-term planning," the person concluded.

Regarding institutional structures, one idea that gained support was for groups to form based on an informal, narrowly-defined remit and then disband once the goal was achieved. It was referred to as "speed dating," by one participant; explained another: "They should be short-term love affairs, not long-term marriages." This approach may let problems be addressed faster and better than having to turn to larger bureaucratic structures that handle a wider mission, and where the very continuity of the group's existence lets them defer decisions. Another benefit is that it avoids institutional elitism, whereby the organization itself becomes more important than the issue it was designed to address. Avoiding formal institutional structures may prevent this, by putting the emphasis on collegiality instead of command-and-control, collaboration rather than czarism, and pragmatic solutions, not rule books.

On a final note, it is useful to remember that institutions are capable of moving fast if the problem is severe enough. This was the chief lesson of a crisis that hit the Coca-Cola Company in June 1999, the biggest crisis in its over 100 year history. After some 250 people,

mainly children, throughout Belgium and France became ill after drinking Coke, the company was slow to respond. It denied there was a problem, even after Belgium banned the drink and officials from other countries grew very concerned -- as did consumers. The stock price tumbled from $78 to around $40.

Yet, once Coke finally comprehended the gravity of the emergency (bacterial contamination that was toxic at a bottling plant), the response was fast. "The company, overnight, changed how it addressed the situation," explained one participant. It recalled about 30 million cans and bottles; it communicated openly with the public; the CEO was forced to resign.

Indeed, it is the lesson suggested by Dr. Johnson's aphorism cited in the introduction to this report, that nothing concentrates the mind more than the threat of being hanged. What is relevant in such incidents is that the hapless fellow never actually goes to the gallows -- his mind thus concentrated, he devises a solution. Ingenuity lets one avoid an otherwise unpleasant fate.

IV. The Role of Different Stakeholders

> *"We don't have any silver bullets. We have a hail of bronze bullets -- or, maybe they're rubber…."*
>
> -- Rueschlikon participant on the difficulty of finding solutions

So far, we have examined what constitutes critical information infrastructure and its vulnerabilities. We identified the economic factors responsible for why the issue is inadequately addressed. Also, we considered institutional frameworks for network security, and its lessons for CII protection. Now, we shift focus. Where the first three sections of this report concerned what is happening and what has come before, the final two sections look towards the future. In this section, we turn to the different players relevant for CII protection. What are their strengths and weaknesses to treat the issue -- and what are the best roles they can play? The final section takes an even more pragmatic bend, by noting specific steps proposed by Rueschlikon participants.

Three main groups are called to play a role for CII protection: industry (both the IT and telecoms sector that furnish the gear, and the firms that actually operate CII); government and the insurance industry. A point about the latter: though it may seem intellectually contrived that the discussion would concentrate on insurance -- as if it were an attempt by attendees to contort the theme to the interests of our host -- this was absolutely not the case. Rather, there was a growing realization over the three-day dialogue that the central obstacle to adequate CII protection was the lack of a market-based response, and that insurance will be an essential feature of any solution.

A. Business Community

An irony of CII protection is that while industry is on the frontlines as the owner of most infrastructure, and must deal with the consequences when systems fail, firms lack an incentive to take adequate precautions or adopt recovery procedures. Companies even resist

cooperating industry-wide. There is no incentive to share information, let alone collect it in-house (in fact, knowledge of one's vulnerabilities could potentially leave a firm liable if it failed to address them and something went wrong). Business partners and the public, not to mention the government, are largely left in the dark about the security of corporate IT systems, including CII systems.

To understand the challenges of instilling CII protection within an industry sector, consider the case of US energy companies. In the 1990s, the US government convened a meeting of the power industry to encourage better security practices. The officials made what one might see as an appealing offer to improve protection and lower costs. "We will collect your data, sanitize it, anonymize it, aggregate it and give it back to you, so you have information on attacks, to help the entire sector," went the offer. The response was "quick and immediate," said one Rueschlikon participant, "No, and hell no!" As the person explained, the feeling in industry was: "What I know about survivability under attacks may be my competitive advantage, and you want me to give it to you?!" He concluded: "After 9/11, they deleted the 'hell' but otherwise the answer is the same."

Ultimately, there is little commercial incentive for firms to act individually and jointly. Indeed, in many cases it is not even in a company's interest to turn to law enforcement when it encounters security problems, since it entails such things as impounding computers and taking time away from business activities to help the investigation. Many firms would rather absorb the cost of security failure for all but the large-scale attacks, breaches and IT failures. Indeed, mechanisms for collaboration cannot be developed because economic incentives are not there. Some firms fear sharing would expose them to liability or cause the market to react negatively.

Changing this will take time, and require new approaches. Businesses must change the culture within their companies, both at the board level and among the managers with day-to-day responsibility. "There are not technical issues standing in our way -- there are human issues standing in our way," noted one Rueschlikon participant. The role of the business community, as a first step, is to collect the data about the security of their IT systems for their own self-appraisals. This will give firms a clearer picture of what is happening in-house, in order for them to make rational decision.

But it is only a first step. The next responsibility is information-sharing industry-wide, and more importantly, across industry sectors. "Businesses need the mandate to talk to one another. This is a role for government, which can be the convener for those activities," said a Rueschlikon participant from the private sector. Indeed, it must be a two-way street. "We need to share information with government -- but the government needs to share information with us so that we can be ready (to protect ourselves)," the participant added.

Yet rather than just convene a meeting -- as the US did with the energy sector in 1990, to no avail -- viable reasons for deep cooperation must exist that play to the interests of industry. Those benefits can be two-fold. First, it can provide valuable data that helps firms manage CII protection. Second, it must lower their costs of CII protection. Today, these cost are subsumed into general business activities and not made apparent, for the economic reasons discussed in section two. It is vital to breakout this information as a cost in its own right. One way to do that is by governmental mandate of disclosure of IT security assessments and

breaches; another way to do it is through a marketplace intermediary, the insurance industry. A look at both institutions thus follows.

B. Government

In Marxist times, there were well-worn mantras about "capitalist pigs" and "the exploited proletariat." In modern times, the knee-jerk clichés have simply been replaced with "industry self-regulation" and "the logic of the marketplace." Is government really that bad? Are commercial approaches really the only answer? Rueschlikon attendees never speak with one voice. However, though there was generally an ideological preference for industry solutions rather than government regulation, there was also an appreciation that government action is both necessary and useful. The question, as always, is: in what way?

The drawbacks of government involvement in CII protection are notable. Often, government lags behind the private sector in understanding the threats and the state of technology to address them. Government is slow to respond or adopt to new situations. It tends to politicize issues rather than remain focused on the substance. And governments usually place the emphasis on the tools they know best, top-down regulation, which may not be the most effective approach. Moreover, the response by industry sometimes shifts the focus away from the purpose of the regulations to mere compliance with them. As one participant put it: "government can't do it in a good way -- it is always regulating previous behavior." Said another: "Governments cannot solve information problems," and ends up with rules that are intrusive and overly comprehensive.

That is not to say there is no role for government -- on the contrary, Rueschlikon attendees identified numerous areas were action is needed. As mentioned in the previous discussion of industry's role, government can serve as the convener to bring parties to the table. It can compel -- either through persuasion or regulation -- the sort of information sharing that many believe is needed. Government can fund long-term research into IT security (and many participants chided the US government for dropping its support of long-term, basic research in this area in favor of short-term homeland security projects).

Moreover, government can use purchasing criteria to create a market for products that conform to certain specifications, like security standards. For instance, the US Dept. of Homeland Security is trying to detail the functional requirements for communications devices that first-responders will use, as a way to avoid having thousands of different municipalities devising their own requirements. If the initiative comes to pass, it will make it more efficient for the IT industry to create products.) Furthermore, the market for these products enables the business community to adopt them as well, with the same high security standards.

More importantly, government can create an environment that facilitates information sharing and coordinated action. First, it can create a "safe harbor" so that companies are free to share data without fearing antitrust action or legal liability from the information, which would chill disclosure. Additionally, government can endorse sound practices that may provide immunity to liability exposure in the wake of externally-caused CII failures. "Why does a crisis bring change?," asked one participant. "It is because there are no other

alternatives but to act, so the risk of antitrust and trade concerns are suspended -- it requires collaboration."

There is one important caveat to these cooperative activities. It must be done clearly for the purposes intended and in the public interest. Rueschlikon attendees noted that there is a risk that these measures are used for little more than protectionism. Businesses have a long history of using security as a commercial weapon against competitors. For instance, AT&T for years prohibited third parties from connecting their own devices to the network on the grounds that they harmed the system -- including, in 1968, banning DARPA-funded researchers from connecting primitive versions of modems, routers and computers that would one day be the Internet. If activities to secure CII were distorted into protectionist policies, it would be a gross abuse of the public's trust, and the responsibility of both government and industry.

C. Insurance Industry

"Could there be a fire-brigade for the Internet that intervenes for problems, but not necessarily state owned?," asked one Rueschlikon participant. That is the question. And in choosing the metaphor, the person conjured the insurance industry's lineage as far back as its origins in the 1700s. The financiers who met at Lloyd's Coffee House in London in the 1680s may have had their eye initially on merchant ships that risked plunder by pirates or perishing by Poseidon's powers. But soon, their commercial acumen extended to other risks. By the early 1700s, some of the first fire brigades in Britain where established by insurance companies, as a way to protect customers from loss -- and thereby, of course, protect the insurance firms themselves.

Insurance companies do not simply cover against losses. They do this by a number of ancillary activities that turn out to be equally vital. First, they collect information about risks and consequences. Next, they identify practices that would minimize the frequency of incidents and their costs. Then, they create a market for these things, which encourages good practices to reduce risk. It makes the uncertainties inherent in life more manageable. As one participant neatly put it: "Insurance companies do not like risk, even though they are in the business of it. They found ways to aggregate individual risk, and with the law of large numbers, found ways to mitigate it."

There are qualifications to this rosy assessment, of course. In the words of one attendee: "The insurance industry is not a social police force." Indeed, the industry will only act if it can make a profit. Insurance, as a sector, can widely accumulate knowledge and identify best practices, and reflect this in the price structure of its products. It can even require policy holders to buy assessments and advisory services in order to qualify for coverage. But all this is provided the risks are both measurable and manageable. If CII represents intolerable risks that, if failures were to occur, no amount of money would have been spared to prevent, then it is beyond the ability of the insurance industry to address these alone.

Meanwhile, not being able to accurately judge the risks, in order to establish a price on them, means that it is more difficult to offer coverage. This, in turn, means that only a limited market can be created to deal with that risk, through an ambiguous approach towards

differential pricing of insurance policies, premiums and deductibles. Thus, the chief obstacle is the lack of information.

The role of insurance in CII protection then can be to establish a market through information. This would require establishing a sort of "underwriters laboratory" for CII risks. To do that, it would be imperative to overcome the asymmetry of information between insurance firms and potential policy holders, in terms of both risks and solutions to offset those risks. What is more, involving the insurance sector may provide not only a market-based risk-analysis, but economic incentives for preparedness, which would likely foster cooperation and collaborative arrangements among CII operators themselves.

Yet turning to the insurance industry is clearly not the final word; there are still obstacles to overcome. For example, because information systems are interdependent, no one insurance company can solve the problem alone -- joint action is required by the insurance firms and their re-insurance partners, who serve as the veritable "central banks" of the industry. Moreover, potential losses might be too big to justify a decision by them to cover the risk, and so they might not act. As an industry, insurance firms themselves would need a "safe harbor" from antitrust action in order to cooperate to understand the risks. And the state may be called on to serve as the "insurer of the last resort," as happens in the airline industry.

But it all relies on information, which for the moment does not exist. In the words of one participant: "The insurance industry will only get involved when there is a sense of fear, mixed with the scent of opportunity."

V. What Is to Be Done

> *"I came here thinking it was a technology problem, and I am relieved to learn that it is really an information problem. We have risks and uncertainties and we know a lot about them, but we don't have a market mechanism, since markets need some information to function. The fact that there is no insurance is a symptom of that."*
>
> -- Rueschlikon participant on attaining CII protection

Engineering and insurance share in common this: both are predicated on the belief that paying a known cost upfront is better than unknown costs in the future if things don't work out well. Fittingly, there is a natural compliment in bringing to bear the activities of insurance on the technology of critical information infrastructure. This section is meant as a practical roadmap on how this might be accomplished, based on the informal dialogue among Rueschlikon attendees. Just as the conference represents a series of cooperative, brainstorming conversations among experts from wide ranges of backgrounds, so too do the suggestions mark initial ideas, not a fully-formed plan of action.

There were two general areas of consensus. The first was that a market-based solution would likely be most efficient, and that the insurance industry has a vital role to play. The second area of agreement was that government action was indispensably needed to help the

insurance industry to act -- yet must be done carefully not to disrupt the private sector but to help it. The previous sections of this report explain how these conclusions were reached. This section lays out concrete ways in which industry and government can join forces to enable better CII protection, by looking at each group individually.

A. Possible Steps by Industry

In order for the insurance industry to act, and a market-based approach to be established for CII protection, a number of subsidiary activities need to take place. Specifically, the commercial and political will must exist; an institutional framework needs to be created; benchmarks must be set; information must be collected and shared, and market-based enforcement needs to happen, so that financial interests (as opposed to regulations) compel good practices.

Partnership

The first step is identifying a "convening agent" to bring parties together from CII operators, the insurance industry and government. Most likely, a research organization can be the convener. If one business tried to do it, this would raise suspicions of capture by a single interest; if a number of firms tried, this would give rise to antitrust concerns. Insurance firms, however, can act as the "lead responder" to forging such a partnership. This sort of meeting can set broad goals and identify steps to take to achieve them, particularly in regards to benchmarks and information sharing. Government should support and encourage this activity, as well as monitor the process.

Institution

The institutional framework needed to address CII protection does not yet exist, and it is unclear if existing forums on related themes provide the appropriate institutional design to serve this role. More likely, a new institutional arrangement is needed (being very conscious of the drawbacks to such organizations, as identified in section three). The group should be non-bureaucratic, comprised of specialists and with a focused remit. Importantly, the framework needs to bring together not only different firms and different sectors, but also do this on an international basis. Beneficially, the private sector tends to forge these types of alliances better than governments, where political considerations outside the issue itself sometimes come into play. Moreover, the structure needs to balance openness to public scrutiny with confidentiality to the firms and the data itself, which could be misused if made public.

Benchmarks

Certain standards, benchmarks and best practices need to be defined, against which firms can be measured. The private sector has a tradition of drafting such quasi-voluntary industry standards that emerge as de facto requirements, such as for corporate financial statements in America with the Generally Accepted Accounting Principles (GAAP). In this case, what may be needed (in the words of one participant) is "GASP," that is, Generally Accepted Security Practices.

The GAAP is a good model. It is maintained by the Financial Accounting Standards Board, which is a private organization comprised of representatives from the accounting industry. GAAP is enshrined by the private sector, not law (although the SEC requires it for financial

reporting); indeed, many non-US companies take pains to tally their accounts under GAAP in order to attract investment.

As it pertains to CII, a number of pre-existing standards for IT security exist, such as the ISO 17799. As such, benchmarks, standards and best practices for CII protection would not be meant to supplant them but incorporate them into a broader set of recommendations. The era of GASP may be upon us.

Audits

The new institution, with its formal benchmarks, can require companies to perform regular self-audits. This need not be overly burdensome. Companies perform audits for numerous other things anyway, from accounting to sector-specific guidelines. Commercial vendors can make auditing easier. For example, global IT consultancies sell services for firms to assess their adherence to privacy legislation; many software products exist to help companies ensure they comply with Sarbanes-Oxley rules for financial oversight.

By making the process a self-audit, the procedure is similar to regulations for self-certification, such as Federal Communications Commission rules for wireless devices, where closer scrutiny only happens if problems emerge. Moreover, because the self-audit does not take place to address a legal requirement, it may avoid the "compliance trap" discussed earlier, whereby firms act to prove adherence to the law regardless of whether their actions truly treat the problem. By setting specific targets but remaining agnostic about how the audits and compliance take place, it lets the private sector devise competitive products to achieve the goal.

Compliance

The issue of self-audits raises a number of questions about compliance. For instance, where does the information go; how is it used; and what enforcement exists to ensure the audits take place? At the first level, the data can be used internally so companies have a better sense of the security of their infrastructure, and can take steps to address potential problems. At the second level, the data can be shared with an institution that acts as a trusted intermediary, which would anonymize the data and aggregate it, and then use it to generate risk profiles to be used by both the insurance industry as well as the operators of CII. This organization can also enforce that audits occur, with the sanction of dropping a firm that fails to perform one satisfactorily (with the consequences that the marketplace is alerted that something is amiss).

At the third level, the auditing and compliance activities can serve another function, by improving the security practices of partners and suppliers to the CII operators. This trickle-down effect exists in many industries, such as retail (where, notoriously, Wal-Mart forces its suppliers to adopt certain technologies in order to do business with it), the automotive industry and the financial services sector. Most CII is at large companies -- and the concept of "supernodes" suggests that even then, the number of critical pieces of infrastructure is limited -- so addressing the major players as a starting point seems reasonable, and permitting the trickle-down to occur over time. This also represents a way to use commercial activities as the basis of better IT security and CII protection, relying on the market rather than classical government regulations.

Information Sharing

The biggest obstacle identified by Rueschlikon participants was devising a way to compel CII operators to share data so that the insurance industry could craft policies for those same CII firms -- a classic chicken-and-egg problem. One way to require information-sharing is through government regulation, but it is not the only way. Another is by mandating it happen through industry practices. The new institution's benchmarks can require that firms share the data they collect about the security of their infrastructures. A company that chose not to participate would thereby raise questions about it among other firms in the industry, its partners, customers, and its investors if it is publicly-traded (just as a public company does when it delays issuing an earnings statement).

Here, the marketplace is imposing the requirement to cooperate, not government. Moreover, the insurance industry can require such information-sharing to happen in order to let firms qualify for coverage, which would give them a strong incentive to do so, particularly if coverage were a normal practice among firms in the industry. Insurance firms traditionally require policy-holders to adhere to certain practices. Moreover, it has experience serving as the aggregation point for data. Together, a marketplace carrot and stick can exist to move firms to share information.

Disclosure

There is a tradeoff between openness and confidentiality. A circumscribed amount of transparency is clearly beneficial in the area of CII protection, but if a firm feels that sharing data would leave it open to liability from government, swarms of tort lawyers, or retribution from the stock market, it would have a chilling effect on disclosure. At the same time, exposing security problems is important in order for the marketplace to act as a force to correct the problem. At Rueschlikon, the question was raised, but not answered, on what constituted the right amount of openness, and whether public disclosure of major CII security incidents was useful or actually created new vulnerabilities.

It was acknowledged that industry self-regulatory bodies could mandate things like disclosure, akin to how the National Association of Securities Dealers regulates securities professionals, which the private group was delegated to do by the SEC. Such disclosure would create a form of informational transparency that benefits the insurance industry as well as the larger market: a company with unsound practices would be exposed and thus have a commercial incentive to perform better. The SEC's Y2K disclosure requirement compelled firms to take steps to upgrade their IT systems and may have averted potential disasters (which, perversely, may have led many people to believe it all unnecessary). The case of the California law obliging firms that suffer data breaches to notify the individuals whose personal information was compromised is instructive. It alone has forced the entire financial and database industry to take the issue of data-security far more seriously than ever before. Sunlight, as the old adage goes, is the best disinfectant.

If major CII security incidents are reported, there remains the question whether this will leave firms open to legal liability -- and if this is indeed a good or bad thing. Were the stock market to penalize a firm, its credit rating falls and insurance premiums rise, it all might be considered positive, as capitalism's corrective. In the case of many industries, it is easy to make the case that public accountability is beneficial. But CII, as this report has shown, is usually a case apart. Such definitive determinations are less easy to make. It is the nature of

CII that special treatment be granted to it; it is critical, after all. For instance, one can imagine a tradeoff whereby firms enjoy legal immunity in return for disclosure, or non-public "escrowed" disclosure in order to maintain confidentiality for national security purposes. The importance of CII to society may require that new approaches be tried.

Eligibility for Coverage

The insurance industry has many tools to encourage certain behaviors. It uses the push-and-pull of premiums and deductibles to elide the financial interests of policy-holders to their actions. Firms then take certain steps not only to lower their costs, but more importantly from a social-welfare point of view, to decreases their risks. This, of course, helps the insurance industry, too, as a business. Yet there is another tool: requiring eligibility for coverage. Insurance firms can mandate that only firms that agree to perform certain actions can qualify for a policy.

One example is insurance for kidnapping and ransoms, which emerged in the 1970s and 1980s in dangerous places. The insurance industry took on a social problem characterized by high stakes (with potentially tragic outcomes), and an unpredictable level of risk. It addressed the matter by treating coverage as part of a broader service. Firms were unable to purchase coverage without first buying risk assessment, advisory services and individual protection. With these services on the front-end, the insurance industry was able to act on the back-end. The result had a bifurcated effect: insurance seemed to increase the risk that one was kidnapped, but reduce the chances that one was killed. As a society that prizes life, this produced the best possible outcome out of many bad alternatives.

As it pertains to CII protection, insurance may have a similar effect -- that is, not so much preventing attacks, but decreasing the chances of failure if one occurs. Again, it represents a form of marketplace enforcement, telling industry that insurance will not let firms shift their risk unless they adhere to certain practices and take steps to maintain records.

Corporate Governance

The final area where industry ought to act is in making CII protection a board-level matter, concomitant with general corporate governance activities. Boards should see that a "chief security officer" is appointed. It requires support from senior managers, and the person should come from those ranks. Too often, the position is pushed down to relatively low-level IT professionals, partially to create a scapegoat if things go wrong. However, CII protection is not a technical but mainstream business matter.

Moreover, the role may at times clash with other senior executives. For instance, chief information officers are largely responsible for squeezing their IT systems to attain operational efficiencies and new sources of revenue, which may be at odds with the requirements for suitable CII protection (which constitutes a cost only). Yet, just as chief financial officers must play a dual role in their organizations -- keeping the firm's finances in order so it can operate, as well as ensuring the integrity of those accounts to the firm's investors -- so too should the CSO act to maintain CII operations, as well as be a check against improper practices.

B. Possible Steps by Government

Markets rely on governments in order to function, and the area of CII protection is no exception. There are many places where the activities of government are required in order for industry to address the issue.

Encouraging Collaboration

Government can act as the convener to initiate dialogue among stakeholders, or support other parties who do this. As discussed above, governments need to provide protection against antitrust action so firms can talk about these matters with their competitors. Likewise, government can act as an observer to the dialogue. More elaborately, government may want to provide some sort of limited immunity to companies who share data, so that the information is not used against them if they adhere to industry best practices. Government can sponsor or help channel investment in long-term research and development for CII protection technologies, so that industry and government alike can tap into the fruits of the R&D. Lastly, government can use a combination of carrot and stick, agreeing to defer to industry approaches but letting it be known that unless the issue is adequately addressed, it will step in.

Immunity from Liability

Government can encourage the insurance industry to act through a more sophisticated policy action. It could provide the first-line and last-line of immunity from liability to the insurance industry for CII failures that arise from acts of terrorism. This could apply to relatively small-scale disruptions and catastrophic large-scale ones, to enable insurance companies to concentrate on the middle of the market for risks, where most vulnerabilities lie, and offer competitive products.

A model for how this might work is the Terrorism Risk Insurance Act (TRIA), which was enacted in the US after 9/11 (and unless Congress renews it, will expire at the end of 2005). Under the law, the US commits to reimburse insurance companies 90% of their insured losses above $5 million per claim per event due to terrorism, up to $100 billion annually. This provided a financially viable way for insurance companies to create policies for terrorism insurance, while it assessed the new risk profile following 9/11. Elsewhere, government aid is commonplace in times of crisis. In the case of hurricanes and other natural disasters, emergency funds are often brought to bear to help communities recover.

Insurance can only act if there are risks it feels are knowable and to some degree controllable; it is reluctant to take on issues that entail high levels of unpredictability -- which is why, for instance, most policies exempt coverage in times of war. The insurance industry is traditionally uncomfortable covering risks that arise from politics. By removing the modest-scale terrorist activity from the risk equation, as well as major failures that if they occurred would bankrupt the insurance sector, the industry can compete in the middle tier. For example, the September 11 terrorist activity resulted in $32 billion of insured losses, including property, life and liability claims.

In the case of CII protection, like the debate over TRIA, the question it raises is whether government, by serving as the ultimate provider of risk insurance, might not distort the market by creating less incentive for policy-holding firms to take preventative security

measures than if insurance companies shouldered all the losses. TRIA benefits go to any firm that suffers catastrophically, regardless of whether they have taken steps to mitigate their vulnerability. Surely the government could adapt legislation so that it pushes toward a market for better cyber-security, while also proving indemnification. An additional concern is whether the law might discourage private insurance firms from the market, which they would otherwise enter more vigorously. What is certain is that while the insurance industry considers tailoring products for a new area of coverage, it will entail a partnership between the public and private sector, and innovative approaches may be called for.

International Law-Enforcement Coordination

Major insurance and reinsurance companies are based in Group of Seven (G7) countries, that already share a substantial degree of political and commercial cooperation. Furthermore, the majority of CII are based there, too. Together, the G7 as an institutional entity can be another dimension to forge international cooperation in regards to CII protection. This collaboration already happens for cyber-security. The G7's Lyon Working Group brings together law enforcement officials to discuss issues and cooperate on initiatives. This, in turn, has laid the groundwork for cooperation with other countries outside the G7.

The benefits of the G7 is that it lacks a standing secretariat; it is organized based on strong political commitment at the presidential and ministerial levels, but its main activity occurs on a regular basis by experts at the working-level of national administrations. In regards to CII, it may be a forum to bring countries together in a way that lets it focus on the topic outside of bureaucratic politics, yet also adheres to the norms of intergovernmental diplomacy.

Conclusion

At the World Economic Forum in Davos, Switzerland in January 2005, around 700 leaders from business, government, academia and the media were asked to identify the six most pressing problems the world faced. The protection of critical information infrastructure didn't make it onto their list (which included poverty, climate change and education). This might well be a good thing, since around three-quarters of the participants in the "Global Town Hall" said that none of the vital problems are likely to be solved under today's governance structures or existing leadership, be it government or corporate. Similarly, CII security has no controlling governance and no recognized leader -- which might also represent a good thing, in that it may offer a chance to solve the problem.

Protecting critical information infrastructure is similar to predicting an earthquake. In geology, we know where a quake will strike, and at what magnitude, but what we do not know is when. Likewise, the insurance industry has elaborate models to predict the frequency of a risk occurring, and the economic costs when it does -- but not when it will. It is managing this uncertainty that will test our fortitude as businesspeople, technologists, policy-makers and scholars of CII in the days and years ahead.

What is certain is that a new institutional mechanism is needed, rooted in the private sector, where the majority of CII lies. Yet it must be cross-sector and international in scope. This in itself makes addressing the matter hard. Yet the difficulty is increased because of the

challenge of compiling data in this environment, in order to understand the risks and how to deal with them. This information is a pre-requisite for the insurance industry to offer coverage and establish a market for CII protection. Yet it requires new forms of cooperation. The role of government is important as well, to monitor the process, support the creation of such a market, and act if problems escalate yet the concerns are inadequately addressed.

The consensus among attendees was that CII security is too important to be left to technology, companies or government. Rather, all parties must play a role. In order to be effective to guard against threats known and unknown, today and in the future, a market-based approach seems the wisest course, facilitated by the insurance industry. This aggregates and transfers risk, as well as compels best practices.

As the awareness develops for market-based solutions, one thing missing is a roadmap for advancing the issue. The Rueschlikon Conference on Information Law and Policy for the Information Economy now marks its fifth year. Previously, the dialogues have made a significant contribution to technology-policy by identifying nascent trends and establishing the conceptual framework for businesspeople and policy makers. This year, it is hoped that the event goes a step further, in serving as the basis of concrete activities.

Amid the vulnerabilities to CII, the opportunity exists to address it -- and the challenge must be taken up. It will require the work of numerous stakeholders. Yet the feeling among participants was measured optimism. Provided the right institutional framework for cooperation among private parties can be established, the issue has a good chance of being treated. Or, as the conference co-chairman Lewis Branscomb quipped in closing the event (to gently break from Rueschlikon rules forbidding attribution): "Global economic activity might be the force that drives this, even if global harmony does not!"