Systemic Cyber Risk: Linking Measurement, Market Incentives, and Policy

Susan Athey, Neil Gandal, and Tyler Moore *

1 Introduction

All organizations have an incentive to reduce the risk that the key personal, financial, or business secret information held by the firm will be breached. Employees often receive reminders about the risk of "phishing" attacks or other approaches used by hackers to break into the organization. However, cyber hackers do not just attempt to break into one organization at a time; instead, they often seek a more systemic cybersecurity weakness that applies to many organizations. *Product-level systemic cyber risk* occurs when there is (i) concentration among technology suppliers and (ii) the products of those suppliers exhibit significant cyber vulnerability. *Firm-level systemic cyber risk* occurs when dominant firms in a sector exhibit significant cyber vulnerability.

There is another type of cyber risk that comes from interconnected systems; it is typically referred to as interdependent cyber risk. See Kunreuther and Heal (2003). Such risk comes from contagion via connected networks. This type of cyber risk is similar to the contagion risk associated with finance, where one failed financial institution can lead to a cascade of failures, and public health where the contagion effect works through the spread of disease. While such risks are sometimes present in cybersecurity, we focus on a simpler and much more pervasive form of systemic risk in which correlated failures can arise, like when a single piece of widely-used vulnerable software results in many firms being attacked at the same time.

As traditionally offline activities like medical care and transportation increasingly adopt digital technologies, the risk of systemic events escalates. The most severe cyber incidents can lead to losses amounting to billions of dollars and also have the potential to result in physical harm or loss of life, and trigger political turmoil. Moreover, past hackers used artificial intelligence software as an aid to human efforts, but in the last year or so, AI tools have advanced to the point where they can act as autonomous agents, searching for the same weaknesses across many enterprises. This likely increases systemic cyber risk considerably. Hence, autonomous AI agents make the economic analysis in our paper critical for policy makers moving forward.

We show how cyber vulnerability can be measured proxy measures such as numbers of security flaws in products or fewer observed security controls deployed by firms. We can then build on entity-level vulnerability to define systemic cyber risk. While the risk of catastrophic failures cannot be completely eliminated, we will point out opportunities to create incentives and institutions to mitigate and reduce systemic cyber risk.

This paper will apply economic principles to the rising danger of systemic cyber risks. The paper is organized around several central economic lessons and uses the case studies, new data, and policy discussion to develop and support them. The central economics lessons are:

1. Cybersecurity investment is below the social optimum because buyers have limited ability to assess the security quality of software and hardware ex ante, switching costs are substantial, and software providers often face limited liability for downstream harms. Additionally, addressing systemic risk is made more difficult by a public goods problem: the consequences of a breach applies across organizations. No single actor has an incentive to internalize the costs of a systemic

*Susan Athey is Professor of Economics, Graduate School of Business, Stanford University, Stanford, California. Neil Gandal is Emeritus Professor of Economics, Tel Aviv University, Tel Aviv, Israel, and Professor of Cyber Studies, The University of Tulsa, Tulsa, Oklahoma. Tyler Moore is Professor of Cyber Security, The University of Tulsa, Tulsa, Oklahoma. Gandal is the corresponding author at neil-gandal@utulsa.edu.

breach, and the result will be an inefficiently low level of investment by institutions to reduce these risks.
2. Information technology product markets are concentrated for various reasons: network effects, natural market power associated with high fixed costs and very low marginal costs in software, and regulatory or other barriers to entry. So any security flaw is therefore likely to create widespread impact.
3. Policy responses are most effective when they target the underlying market failures directly rather than focusing exclusively on compliance requirements.

1.1 Background and Overall Trends

It is helpful to provide some background on the level of cybersecurity risk. There is no authoritative list of total cyber attacks, nor of the subcategory that would be classified as systemic. However, an array of reports suggest that concerns about cybersecurity are large and growing for private firms, nonprofits, and government. Recent examples include *Global Cybersecurity Outlook 2026* from the World Economic Forum, (World Economic Forum (2026)), the *Cyber Security Report 2026* from Check Point Research (Check Point Research (2026)), and the *Cost of a Data Breach 2026* report from IBM (IBM (2026)).

We do have good data on other measures associated with systemic cyber risk: the number of "ransomware" attacks per year, the number of Common Vulnerabilities and Exposures (CVEs), and the number of Known Exploited Vulnerabilities (KEVs).

Ransomware has become the dominant cybercrime threat to organizations, and it is growing fast. Ransomware is a form of malicious software that encrypts an organization's files or systems, rendering them inaccessible until the victim pays a ransom demand. In recent years, attackers have started threatening to publicly leak stolen data if the ransom is not paid. Attackers often breach a target through unpatched vulnerabilities in information technology products. Ransomware attacks have the potential to be systemic because information technology product markets are very concentrated and the same vulnerability may be exploited to attack multiple victims.

The U.S. government's National Vulnerability Database repository hosted by the National Institute of Standards and Technology (NIST) tracks software vulnerabilities and distributes them publicly (National Institute of Standards and Technology (NIST) (2026)). Vulnerabilities are standardized using the Common Vulnerabilities and Exposures (CVE) system. Each disclosed vulnerability is reviewed and assigned a unique CVE identifier once verified. The NVD enriches each CVE record with severity assessments, affected products, remediation guidance, and references. CVE count is a decent proxy for the security of a software product, as insecure software will invariably have more vulnerabilities in it. However, there are downsides to utilizing CVE count directly, because not all vulnerabilities are actively exploited by attackers.

To that end, CISA started publishing the Known Exploited Vulnerabilities (KEV) Catalog in 2021 (Cybersecurity and Infrasructure Security Agency (CISA), 2026). The KEV catalog focuses exclusively on the subset of vulnerabilities published in the NVD for which there is confirmed evidence of active exploitation. Hence, products with KEVs can be interpreted as especially risky.

Table 1 reports trends in the number of ransomware attacks, CVEs, and KEVs. As the Table shows, the number of ransomware attacks has skyrocketed in recent years, quadrupling from 2021 to 2025. The Table also shows that the number of CVEs published in 2025 are approximately 2.4 times the number published in 2021. While the number of new KEVs per year has been fairly stable since 2023, nearly 1,500 different vulnerabilities have been exploited since 2021, when the CISA began collecting

this data.

Table 1: Ransomware Attacks, CVEs, and KEVs per year over time (2019–2025). Note that the CISA began to report KEV data in 2021. Sources for the data: For KEVs, www.cisa.gov, for Ransomware Mousqueton (2026), and for CVEs National Institute of Standards and Technology (NIST) (2026).

Year	Ransomware Attacks	CVEs	KEVs
2021	1,816	20,177	311
2022	3,168	25,000	555
2023	5,412	28,848	187
2024	6,058	39,925	186
2025	7,985	48,154	245

2 Economics and Systemic Cyber Risks: The SolarWinds Case Study

We begin our analysis with a case study of a cybersecurity event at the SolarWinds company in 2020, which brought these issues to the forefront.

In 2020, the Orion platform of enterprise software supplier SolarWinds was hacked (Tidy, 2020). In all, over 18,000 organizations in 36 countries downloaded the malicious code, including many governments and large enterprises with seemingly sophisticated cybersecurity programs in place — including the US Departments of Treasury and Commerce and a majority of Fortune 500 companies. This software, which previously few people had even heard of, was deployed by many of the world's largest corporations and governments to monitor their most critical information technology assets. The vulnerability granted unauthorized access to the most sensitive systems at thousands of high-value targets in an instant. This attack created a huge shock to the security ecosystem. Many of the world's largest corporations and governments were at the time using Solarwinds software to monitor their critical information technology assets.

Attackers gained unauthorized access to SolarWinds *build platform/system*, used to compile and package the Orion software. In this case, the build system was the critical point of failure that enabled the attacker to inject malicious code into the compiled SolarWinds Orion software. When users installed the compromised software, the malicious code was executed on their systems, giving the attacker unauthorized access to their systems. The attacker was then able to steal sensitive data from their systems, such as credentials, intellectual property, and customer information.

Afterwards, other significant vulnerabilities in the Solarwinds software and operational practices came to light. Perhaps most egregiously, it was found that they had used a default password of "solarwinds123" for configuring their systems, which the CEO blamed on an intern Fung and Sands (2021). But there had been earlier warnings about cybersecurity at Solarwinds. In the year before the breach, SolarWinds was ranked by Intangic (a firm that estimates cyber risk, often for insurance purposes) as a high-risk company because, relative to its competitors, in part because SolarWinds had relatively high spending on mergers and acquisitions but relatively low R&D spending (as reported by Nolan and Fixler (2021)).

In spite of weak security practices, the SolarWinds Orion platform operated for many years as a leading provider in a niche area of corporate information technology software. How did this come to be? First, like many technology products, network effects tend to produce dominant products within market segments. There are benefits to scale in enterprise software, just as we have seen in operating

systems (say, Microsoft Windows for desktop computing and the iOS and Android operating systems for mobile phones). Second, the enterprises that purchase software lack the expertise to ascertain the security quality of the products they are purchasing. This information asymmetry enabled Solarwinds to invest fewer resources to securing their products, as their customers could not judge whether they were more secure than their competitors. Third, the breach generated huge negative externalities in that thousands of Solarwinds customers were harmed by the breach much more than the vendor itself.

One might expect that the market would provide an incentive for Solarwinds to invest more in security, because customers would respond to a hack by switching to competitors. In fact, we see only limited movement away from Solarwinds after the attack. We identified three major competitors to SolarWinds Orion product using Gartner Peer Insights Gartner (2024): Compuware Dynatrace, Zoho ManageEngine, and Zabbix. Because the attack was not put into place (via updates) until early 2020, and was not discovered for many months, it makes sense to consider the 2020 data on product installations to be the "before" period, while 2022 is the "after" period.

The market share of enterprises with each of these products in 2020 and 2022 is shown in Figure 1 (left). Within two years of the SolarWinds hack, several of its competitors did experience some growth, while the installed base of SolarWinds Orion did decline somewhat. Nevertheless, as Figure 1 (left) shows, the market reaction was limited and the SolarWinds product was still used by a large number of enterprises several years after the attack.

A plausible rationale for the limited change is that in software, the switching costs to move from one enterprise software to another are usually high. Solarwinds Orion is typically deeply embedded in the information technology environment of an enterprise; changing to a competitor is a substantial undertaking that introduces new risks that the executives in charge of information systems may not want to take on.

Another possible explanation for why more firms did not switch could be that SolarWinds improved its security practices after the incident. The evidence does not support this explanation. As discussed above, the National Vulnerability Database (NVD) posts known vulnerabilities in software products. Although SolarWinds patched the Orion system after the attack and released new versions, remarkably (and unfortunately) SolarWinds had nearly as many vulnerabilities reported in 2023 as they did in 2020 and 2021, as shown on the right-hand side of Figure 1.

It is also instructive to consider what policy remedies were available and how they fared. First, firms like Solarwinds can disclaim any liability for defects in their software products, which naturally dulls any incentive to proactively seek out and fix defects. Nonetheless, the Securities and Exchange Commission (SEC) requires publicly traded companies to disclose cybersecurity practices and risks. In 2023, the SEC charged Solarwinds and its Chief Information Security Officer with fraud and internal control failures. The SEC claimed that in filings up until December 2020, SolarWinds downplayed its cybersecurity risks (and misled investors) by disclosing only generic and hypothetical risks, even though SolarWinds internal documents assessed that its product was not very secure, which was consistent with the external assessment of FDD and Inganic (US Securities and Exchange Commission, 2023). However, the SEC dropped the charges in 2025, which unfortunately sends a signal to firms that the SEC cybersecurity reporting rules are not very effective (US Securities and Exchange Commission, 2025).

Given the absence of regulatory remedies, one might expect that the private market for cyber insurance could help to manage the risks. In fact, insurers also bore only a small fraction of the overall damages (about 1 percent by some estimates). Policies were written in such a way that the incident was either excluded from coverage or they did not trigger a claim.

This case study highlights the systemic risk created by market concentration, the incidence of the

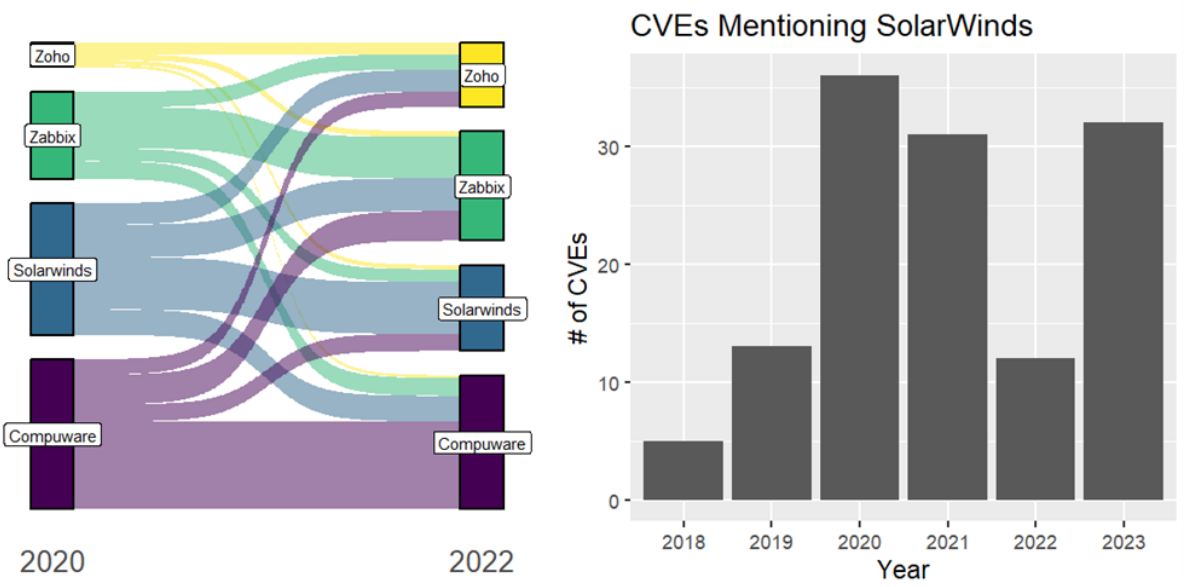


Figure 1: Solarwinds Before and After the Hack. Sankey diagram (left) shows the migration of firms between vendors between 2020 and 2022. (The vertical bars represent market shares.) The bar plot (right) shows the total CVEs added to the NIST database each year which mentioned SolarWinds in the description. Source: Authors' calculations based on databases we discuss below.

costs of the attack across software customers, vendors, and insurance, and the apparent failure of market incentives to improve security in the years after the consequences of the incident became apparent. The case further illustrates that market incentives do not seem to have mitigated the systemic cyber risk created by market concentration.

3 Measuring Systemic Cyber Risk

Perhaps surprisingly, there is no universally accepted definition for systemic cyber risk (Forscey et al., 2022).

We use the term "cyber vulnerability" to refer to the vulnerability of a single entity to cyber attacks. Cyber vulnerability can manifest through security flaws in software or underinvestment in defenses by organizations.

We can then build on entity-level vulnerability to define systemic cyber risk, broken into two categories: *Product-level systemic cyber risk* occurs when there is (i) concentration among technology suppliers and (ii) the products of those suppliers exhibit significant cyber vulnerability. *Firm-level systemic cyber risk* occurs when dominant firms in a sector exhibit significant cyber vulnerability.

In this section, we seek to measure cyber vulnerability both at the product-level and firm-level. Notably, the Cybersecurity and Infrastructure Security Agency (CISA), an agency within the US Department of Homeland Security, has called for measuring such sources of systemic cyber risk: "[S]imply put, what gets measured can be managed. CISA's goal to reduce systemic cyber risk is centered around finding concentrated sources of risk that, if mitigated, not only provide the organizations cost benefits for heightened risk management but also manage critical risks to the Nation's security and economic security" (Cybersecurity and Infrastructure Security Agency, 2024).

Many firms and sectors within the global economy depend on a common set of critical technology

products and software services, thereby consolidating risk across a small number of potential failure points. This convergence has undoubtedly enhanced efficiency, expanded capabilities, and increased flexibility for many organizations. However, the widespread reliance on a limited set of specialized products and services has simultaneously given rise to many single points of system vulnerability. Such concentration of risk amplifies the potential fallout from security breaches, especially since attackers have strong incentives to specifically attack such products and services.

In this section, we develop a methodology for measuring systemic cyber risk using available data. We do this by combining two types of data sets we have assembled: (i) installations of information technology/security products at the firm/site level and (ii) vulnerability data at the product level. The dataset on information technology security dataset provides both detailed data on firms that sell cyber security products and on firms that adopt cyber security products. We have data from the Spiceworks Company Intelligence Database (CIDB) Spiceworks Ziff Davis (2024) on worldwide firms spanning three years (2018, 2020, 2022). The data are on 11,000 installed information technology products at the entity/site/class/subclass/manufacturer/model level. Each of the observations includes detailed information about the entity/site including name of firm and site, location and industry (by Standard Industrial Classification code). Each observation has a single product that is installed at that entity/site.

The takeaway for economics researchers from this economic framework is that it is possible to measure systemic cyber risk throughout the economy, and to draw lessons that can be used to guide policy. In the following subsections, we show how this can be done at both the product-level and the firm level.

3.1 Measuring Product-level Systemic Cyber Risk: The Case of Cisco

Cisco is the market leader in routers, with more than 50 percent of this market. Routers form the backbone of the Internet; they are used to direct traffic across internet service providers, as well as within enterprises. Table 2 shows the market shares of the leading suppliers in routers in 2018, 2020, and 2022 in the United States, and shows the Herfindahl-Hirschman measure of industry concentration for each year.

We built a data set of (1) Common Vulnerabilities and Exposures per vendor in the router category and (2) the installed base of routers per vendor, usign the CIDB data. Including Cisco, the correlation between Common Vulnerabilities and Exposures and installed base is 0.89, which is a *very* high positive correlation. Since Cisco has such a large installed base (and many router vulnerabilities), we also calculated the correlation with Cisco excluded, which remains quite significant at 0.46. These correlations suggest that there is systemic cyber risk in routers, because there is both high cyber risk (in terms of many vulnerabilities in the leading products) and high systemic risk (in terms of product concentration).

Attacks via Cisco routers have caused widespread systemic damage. The "Salt Typhoon" cyber event, widely attributed to China, has been described as the "worst telecom hack" in US history. Though the attack likely began in 2022, it was only discovered at the end of 2024. Cisco routers were specifically targeted and "Salt Typhoon" attackers gained access to telecoms through Cisco devices. According to US officials, the objective was to provide Chinese operatives with sustained access to US telecommunications networks by compromising devices such as routers and switches. A report by Cisco Talos, Cisco's threat intelligence unit, noted that in one instance, the "Salt Typhoon" attackers likely gained access to Cisco devices by exploiting a seven-year-old critical vulnerability in Cisco's IOS XE operating system (as reported by Kapko (2025a)). Additionally, Recorded Future — a subsidiary of Mastercard that provides global threat assessments — wrote that five telecom networks were

Year	Vendor Market Share (Routers)								HHI
2018	Cisco	55%	Mikrotik	21%	SMC	12%	Dell	3%	3,667
2020	Cisco	51%	F5	8%	Mikrotik	7%	Asus	4%	2,830
2022	Cisco	51%	F5	13%	Juniper	6%	Aruba	4%	3,257

Table 2: Market Shares of the Leading Suppliers in Routers: 2018, 2020, 2022. The Herfindahl–Hirschman Index (HHI), a standard measure of industry concentration, is defined as the sum of the squares of the market shares. According to Department of Justice guidelines U.S. Department of Justice (2024), an HHI index less than 1,000 usually means a very competitive market; the 1,000 to 1,800 range typically means a moderately concentrated market; and above 1,800 is usually considered to be concentrated.

hit by Salt Typhoon via two other vulnerabilities in Cisco's IOS XE at the end of 2024 . They write that the attackers targeted over 1,000 Cisco routers globally during late 2024 and early 2025, with a primary focus on telecommunications networks (See Kapko (2025b)) and Recorded Future Insikt Group (2025)).

3.2 Firm-level systemic cyber risk

We can also measure systemic cyber risk at the firm level. A promising start is the five categories of security controls from the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF): Identify, Protect, Defend, Respond, and Recover National Institute of Standards and Technology (2024). A brief description of the five categories of controls can be found at the NIST website at https://www.nist.gov/cyberframework/getting-started/online-learning/five-functions. We can assign security products by function to one of more of these five categories. We can then examine the number of categories covered by firms.

Three of the 24 product series categorized by the CIDB are specific to security: Security Devices, Information Technology Security Information, and Network Security Monitoring. These three product series include a total of 507 security products. Data from the CIDB enables estimation of firm-level cybersecurity posture based on the number of security controls utilized by firms.

Using 2022 data, Table 3 illustrates the type of data that we can employ to measure cyber risk at the product and firm level by category.The first column is the twelve CIDB security categories. The second column is our mapping of the CIDB security categories to the five NIST Cybersecurity Framework functional categories. Thus, we mapped all 507 security products to the five NIST Cybersecurity Framework categories. The third column of Table 3 is the number of products in each category. Column four is the percentage of firms that have installed a product in the category.

Table 3 also reports a count of Common Vulnerabilities and Exposures (CVEs) and a list of Known Exploited Vulnerabilities (KEVs) by product category from the US Cybersecurity and Infrastructure Security Agency. Column five is the total number of CVEs in products in the category, while column six is the total number of KEVs in products in the category

3.2.1 Firm Coverage of NIST CSF Categories

From this data, we can calculate the percentage of firm sites that cover each of the five NIST CSF categories. (Many firms have multiple sites.) Of the firms observed to employ at least one security control, just 39 percent of sites have controls covering two categories, while 21 percent cover three of the categories. Just 10 percent of the sites cover four of the categories, while only 2 percent of firms have controls covering all five categories. This is surprising, given that the National Institute of

Category	NIST CSF Category	# Products	% Firms	CVEs	KEVs
IT Asset Management	Identify	8	3.7%	49	0
Advanced Threat Protection	Protect	18	11.6%	125	0
Data Loss Prevention	Protect	4	5.0%	14	0
Firewall Software	Protect	35	26.9%	1,237	28
Identity Access Management Software	Protect	63	45.8%	615	3
VPN	Protect	17	30.5%	858	15
Anti-virus	Detect	138	17.6%	803	95
Email Security	Detect	7	46.7%	72	0
Endpoint Security	Detect	42	22.2%	1,467	24
SIEM	Respond	85	26.6%	711	2
Disaster Recovery	Recover	27	11.6%	39	0
Other Security	Multiple	33	23.4%	232	0

Table 3: Security product categorization and installation metrics. Note: NIST CSF refers to the Cybersecurity Framework from the National Institute of Standards and Technology. SIEM refers to "security information and event management" products. CVEs are Common Vulnerabilities and Exposures and KEVs are Known Exploited Vulnerabilities. Data Source: CIDB (2022). Table Source: Hajizada et al. (2024).

Standards and Technology recommends that all organizations adopt a baseline set of controls in all five categories. Even if we restrict attention to large firms with large sites with more than 500 employees, only 21 percent of the sites cover four of the NIST categories, while 6 percent have controls in all five categories. It is likely that some firms have greater coverage than reported here, but the overall picture remains troubling.

Following this discussion, we can now expand a bit on our earlier working definition of firm-level systemic cyber risk: Large firms that employ many employees and serve many markets are a source of systemic cyber risk if they do not have security products in at least three of the five of the NIST CSF categories.

3.2.2 Sector Coverage of NIST CSF categories

We can aggregate from the firm level to the sector level by using the North American Industry Classification System (NAICS). This system is used by US government statistical agencies to classify firms by sectors. There are 96 three-digit sub-sectors in the North American Industry Classification System (NAICS) structure. This delineation is known as the NAICS3 level.

Table 4 presents the 10 lowest-performing sectors in terms of the share of firms implementing controls in at least three of the Cybersecurity Framework categories at the NAICS3 level. In the Table, the first column is the number of firms in the NAIC3 sector, while column two is the percent of firms in the sector that employ products in least three of the five NIST categories. Column three is the average number of NIST categories covered by firms in that sector. Finally, column four is the name of the sector.

Overall, only 29 percent of firms in NAICS3 sectors implement controls in three or more of the five categories. Especially concerning are Truck Transportation and Educational Services. In these sectors, less than 12 percent of the firms had products at least three of the five Cybersecurity Framework functional categories. While Table 4 shows that some other NAICS3 sectors have even lower levels of implementation, we focus on these two because of their importance: Transportation and Manufacturing

NAICS3	≥ 3 NIST (%)	Avg # NIST	Sector Name
238	6.5%	1.32	Specialty Trade Contractors
624	7.5%	1.37	Social Assistance
712	7.5%	1.36	Museums, Historical Sites, and Similar Institutions
321	9.2%	1.43	Wood Product Manufacturing
337	9.3%	1.44	Furniture and Related Product Manufacturing
713	10%	1.40	Amusement, Gambling, and Recreation Industries
484	11%	1.52	Truck Transportation
332	11.1%	1.46	Fabricated Metal Product Manufacturing
925	11.9%	1.51	Administration of Housing Programs, Urban Planning, and Community Development
611	12.0%	1.46	Educational Services

Table 4: Security control adoption level for the 10 worst-performing Sectors (as categorized by NAICS3) Note: NAICS3 refers to the North American Standard Industry Classification at the three-digit level. NIST refers to the National Institute for Standards and Technology. sectors. Based on the CIDB data set for 2022. The label $\geq 3NIST(\%)$ is the percentage of firms in the sector employing products in at least three NIST categories. The heading "Avg # NIST" is the average number of NIST categories covered by firms in that sector. (Source: Hajizada et al. (2024).)

are included in the 16 critical infrastructure sectors, and Education is extremely important as well.[1] In these two NAICS3 sectors, column three of the Table shows that the average number of Cybersecurity Framework controls employed is only roughly 1.5 out of the five Cybersecurity Framework categories. These results show that we can identify sectors with high cyber risks. The next step is to check these sectors for firms with concentrated market share, which would indicate elevated systemic cyber risk.

3.2.3 A Lagging Industry Sector

We take one closer look at an industry sector that has especially lagged in cybersecurity investment: municipalities, which are critical infrastructure services. (The CIDB data includes municipalities.) Municipalities typically provide Water and Wastewater services.

Like many commercial enterprises, governments have increasingly employed digital technologies to supply services. While this undoubtedly brings benefits, there are costs as well, including increased exposure to cyber attacks. Whereas federal and state governments are reasonably well equipped to protect against cyber threats, cities and towns are not. Cybersecurity and Infrastructure Security Agency and Federal Bureau of Investigation and National Security Agency (2022)

Several recent reports suggest that ransomware attacks against municipal governments have increased significantly in recent years. Multi-State Information Sharing and Analysis Center's (MS-ISAC's) 2022 survey, found that ransomware incidents against municipalities had increased by 51 percent Fox-Sowell (2024). MS-ISAC's 2025 report also found that there also was a significant increase in cyber attacks on critical municipal infrastructure in 2024 MS-ISAC (2025).

In July 2026, sophisticated hackers attacked many US municipal water systems in seven states.[2] In

[1]The sixteen "Critical Infrastructure Sectors" are defined in the National Infrastructure Protection Plan (NIPP) U.S. Department of Homeland Security (2013). They are as follows: Chemical Sector, Commercial Facilities Sector, Communications Sector, Critical Manufacturing Sector, Dams Sector, Defense Industrial Base Sector, Emergency Services Sector, Energy Sector, Financial Services Sector, Food and Agriculture Sector, Government Services and Facilities Sector, Healthcare and Public Health Sector, Information Technology Sector, Nuclear Reactors, Materials, and Waste Sector, Transportation Systems Sector, and Water and Wastewater Sector.

[2]U.S. and state officials believe that the attacks were by Iranian hackers.

Group	All governments		Municipality		State		Federal	
	#	%	#	%	#	%	#	%
Identify	969	2%	236	0.3%	498	12%	235	12%
Protect	10,368	17%	7,080	13%	2,372	57%	916	48%
Detect	9,877	17%	7,095	13%	2,016	49%	766	40%
Respond	3,749	6%	1,432	3%	1,677	41%	640	34%
Recover	1,822	3%	622	1%	764	18%	436	23%
$\geq$ 1 NIST category	15,934	27%	12,236	23%	2,655	64%	1,043	55%
$\geq$ 2 NIST categories	5,557	9%	2,770	5%	2,043	49%	744	39%
$\geq$ 3 NIST categories	3,096	5%	951	2%	1,540	37%	605	32%
$\geq$ 4 NIST categories	1,535	3%	371	1%	763	18%	401	21%
5 NIST categories	663	1%	137	0.2%	326	8%	200	10%

Table 5: Governmental coverage of the Cybersecurity Framework Categories of the National Institute of Standards and Technology Note: Based on the CIDB data set for 2022.

several cases, hackers managed to change passwords and network configurations. This meant that the municipalities could not access the automated control systems that manage pumps, valves, and other infrastructure. Most of the more than 50,000 water systems in the US serve less than 10,000 people. Many depend on decades-old infrastructure that were not engineered for a networked environment where ransomware and state-sponsored attacks pose serious problems. (See Volz and Londoño (2026) and Easterly (2026)). It is likely that the systemic cyber risk associated with the use of common infrastructure enabled so many simultaneous attacks.

Fully 80 percent of US municipalities reported having fewer than five employees dedicated to cybersecurity. In order to cope, many municipalities are switching to third-party providers including "managed service providers" to address growing information technology and cybersecurity needs. Several of these managed services providers serve many municipalities. While these services can be helpful for municipal cybersecurity, they also creates a common point of vulnerability across customers.

To illustrate how valuable the CIDB data can be for measuring municipal security control adoption, we examined the security technologies for all US government organizations gathered in 2022. Table 5 cuts the data in two ways. First, it counts the number of NIST Cybersecurity Framework categories covered by each governmental unit. Second, it compares the overall results for local, state and federal governments. We can see that municipalities badly lag state and federal governments. Of the municipalities in the CIDB database, 23 percent have at least one of the five categories, but it falls off quickly after that. Just 5 percent of municipalities have controls covering two categories, while 2 percent cover three of the categories. Just 1 percent of the municipalities cover four of the NIST categories, while only 0.2 percent of municipalities have controls covering all five categories. While some of the poor performance can be attributed to gaps in CIDB coverage, these findings underscore how far behind municipalities are in terms of cybersecurity technology adoption.

3.3 Summary: Measuring Systemic Risk

We showed that cyber risk can be measured, both at the product level and the firm level – and measurement is a first step in managing such risk.

Systemic cyber risk arises in settings with concentration. The literature on innovation and market structure in economics identifies a tradeoff: concentration enables scale economies and allows firms to profit from their innovation, but oligopolistic market structures can lead to low innovation due to

lack of competitive pressure. The empirical evidence presented here is not sufficient to demonstrate a causal link. But it is consistent with a belief that, when it comes to cyber risk, concentration is not creating sufficient investment incentives to alleviate systemic risk or overcome free riding, policy failure, oligopoly behavior, and the lack of customer responsiveness to poor quality.

4 Policy Options for Mitigating Systemic Cyber Risk

Having demonstrated the widespread presence of systemic cyber risk and its economic drivers, we now turn to discuss policy interventions that could counteract the risks. We begin with traditional regulatory approaches and illustrate their limitations in this context. We then turn to non-traditional mechanisms that can help promote cybersecurity public goods and address information issues driving systemic cyber risk. Finally, we consider the potential and pitfalls of insurers in mitigating systemic cyber risks.

Throughout the discussion, we point to the economic forces that are critical in guiding policy. We also present several policy recommendations that we believe will improve cybersecurity, streamline regulation and increase competitiveness in government procurement which will reduce market concentration, and help mitigate systemic cyber risk and reduce its costs.

4.1 Promoting Secure Practices

When faced with problematic behaviors, such as widespread insecurity of software or technology products, policymakers traditionally respond in one of two ways: by enacting regulations in advance intended to limit such behavior or by imposing liability after undesirable events have occurred. Given US rules about tort liability, a policy of unlimited or joint liability for software makers could easily make developing new software, even software intended to fix cyber problems, into an investment that no firm would wish to undertake. It seems especially difficult to impose liability for a product that seemed safe when issued and tested, but later a hacker found a point of entry.

Cybersecurity safety regulation has fared a little better. Policymakers have proposed instituting baseline security requirements for operators in critical infrastructure industries. This approach has recently gained traction in the European Union via the Cyber Resilience Act, which broadly requires that manufacturers follow cybersecurity rules in the planning, design, development and maintenance of their products — sometimes including assessment by a third party. In July 2026, the Commission published practical guidance for compliance with these rules.

One danger with EU-style detailed regulation — especially in the context of the fast-moving array of technologies potentially affected by cybersecurity concerns — is that time, energy and resources can be wasted on negotiating such rule and accompanying compliance procedures without understanding whether there are corresponding benefits. Such a focus crowds out sensible "best practice policies" - and creates fixed costs that act as a barrier to enter for small and medium-size firms. For example, Weishäupl et al. (2018) find that cybersecurity investment decisions are frequently driven by compliance obligations, with learning about investment effectiveness occurring in an ad hoc and largely reactive manne. Related evidence suggests that firms may also rely on peer behavior when making security decisions: a study of Finnish information security managers documents herding behavior in cybersecurity investment, with firms following a "let's follow others" compliance strategy when selecting security controls Shao et al. (2020).

4.1.1 Case Study: from Compliance Obligations to Best Practices

Until very recently, the US government was using a "compliance obligations" approach to cybersecurity for defense contractors: all U.S defense contractors were supposed to fully implement the more than 100 security requirements included in NIST Special Publication 800-171 by November 2026. This is often referred to as the Cybersecurity Maturity Model Certifications (CMMCs). These rules were to apply whenever a system owned, or operated by or for, a contractor processes, stores, or transmits "controlled unclassified information," a regulatory category created by the Executive Order 13556 that refers to nonclassified information where the government still wishes to safeguard and limit or prevent dissemination.[3]

No empirical justification was given for why these more than 100 requirements were chosen. The implementation was required for all firms. This policy imposed large costs, but benefits were not obvious, and innovation incentives were near-invisible. Unsurprisingly, adoption had been minimal, even for large firms that could absorb the costs.

During a "Promoting Competition in AI" conference at Stanford University in 2024, former US National Security Advisor and Secretary of State Condoleeza Rice was asked "how you think the US government is doing today with respect to being able to leverage the commercial sector to acquire the technologies that are clearly going to be needed for national security purposes…". She replied: "I think we're doing very poorly. Very poorly. And it's largely because the procurement and acquisition processes in the defense department are baroque to say the least."[4] She went on to note that it is virtually impossible for small companies with great ideas or products to"penetrate the Pentagon".

In this way, cyber security issues create barriers to entry and affect competition and especially in the defense procurement sector, but also across other sectors. This is because a number of departments, including NASA, the Department of Energy, and the General Services Administration (which acts as a centralized buyer for the federal government) all had various regulatory changes in process to roll out mandatory implementation of NIST 800-171 standard. Indeed, this standard was to apply for all US federal government contractors, because the standard underlies the protection of US "controlled unclassified information."

Thus, rather than leveraging US defense contracting to take advantage of the ecosystem of innovative small and medium-sized technologies to improve cybersecurity, these firms instead would have been shut out, and many potential important cybersecurity innovations that could help to reduce cyber risk would never have come into the market. This important economic cost is typically overlooked.

It is important to emphasize that NIST 800-171 was not a reasonable step forward in cybersecurity for any firm, including large ones. Although large firms could afford the cost of implementation, the low adoption rate shows that they did not believe that the cybersecurity benefits exceeded costs.

Fortunately, the US government recently pivoted to a best practices approach. In July 2026, the DoD/DoW suspended implementation of these more than 100 security requirements. The DoD/DoW announced that they will propose new regulations that lower barriers for small and medium firms, and replace bureaucratic compliance with "scalable, resilient cybersecurity measures". This is a best practices approach. See https://dodcio.defense.gov/CMMC/.

One way to implement a "best practices" approach to scalable, resilient cybersecurity measures is to use explicit cost-benefit methodology. Barnir et al. (2025) showed that the number of requirements

[3]For additional detail, see the Defense Federal Acquisition Regulation Supplement (DFARS) 252.204-7012 Safeguarding Covered Defense Information and Cyber Incident Reporting. It is available at https://www.acq.osd.mil/dpap/dars/dfars/html/current/252204.htm#252.204-7012.

[4]For a description of the conferences, see the Stanford website at https://siepr.stanford.edu/events/promoting-competition-ai. The quotation from Rice is taken from a transcript of the conference.

included in NIST Special Publication 800-171 could be reduced by more than 50 percent without sacrificing much value, where value is the difference between benefits and costs. This "best practice" methodology facilitates opportunities for streamlined enforcement, focusing effort on the most important risks, particularly on those that affect systemic cyber risks. Such an approach will benefit all innovative firms. It will also likely reduce market concentration by enabling entry and competition.

This example shows that economics matters, especially for innovation incentives. Publication 800-171 implicitly assumed that technology evolves exogenously and that regulation does not have any tools at its disposal to shape either the rate or direction of technological change.

4.2 Imperfect and Incomplete Information

Imperfect and incomplete information generate suboptimal and ineffective cybersecurity. We now discuss four types of information issues that cause problems.

First, there is incomplete information on losses experienced by third parties. An organization knows when they are attacked, but they do not always know when their peers are. Even when they do know an attack has occurred, the magnitude of the resulting loss is usually unknown. This means organizations do not know the cost associated with cyber attacks until they are directly experienced, which in turn makes it harder for firms to justify proactive investment in cyber defenses before attacks happen. Such incomplete information can even contribute to negative externalities in that firms underinvest in cybersecurity defenses, which in aggregate can create systemic firm-level cyber risk.

Second, there is imperfect information on the effectiveness of security controls. Organizations struggle to determine which technologies are more effective at countering threats. Controls are selected based on measurable factors including cost, ease of use, integration into existing infrastructure. Determining which products do a better job at stopping future attacks is very hard to measure, so the choice is not made on this basis.

Third, there is imperfect information on the frequency of attacks and the relative likelihood of different threats being realized. This makes it hard for organizations to use more sophisticated risk management approaches that can take attack likelihood fully into account.

Fourth, there is imperfect information on the attackers themselves. It is often unknown from where attacks originate, what the attack's motivation is, and what the attacker's capabilities are.

What should governments do to address imperfect and incomplete information problems? The most promising solution is to create missing "information" markets and take steps to help shape these markets. Such information sharing markets could include threat intelligence, detailed loss variation and compromise frequency reports, evaluations of security measures, and consistent incident reporting. Such information would typically constitute a public good, the details of how such goods could be disseminated will be discussed in the next section.

However, we need to be careful that implementation does not diminish incentives to conduct R&D to find new security threats. If we required sharing of all such information, the incentives to conduct R&D would be reduced since proprietary information is a source of competitive advantage. But nevertheless certain information could be in the public domain.

Information sharing may also yield a pro-competitive effect. Hajizada et al (2026) found that cyber attacks are also associated with changes in vendor portfolios. Compared to non-attacked firms, attacked firms were more likely to shift to vendors with larger market presence and broader product scope. While these large vendors' products may be higher quality, it is also possible that the switch to them is due to incomplete and imperfect information about vendor product quality.

4.3 Promoting Cybersecurity Public Goods via Data Sharing

Cybersecurity has a public good aspect. Many of the steps to promote cybersecurity — once they are discovered and produced! — can be shared broadly at near-zero marginal cost. The problem arises when entities can benefit from this cyber-risk prevention without contributing, leading to a market failure where security is under-provided at a systemic level. Firms may under-invest in cyber-security, expecting others (competitors, government, or even information technology service providers) to cover the costs of securing the overall infrastructure, especially critical infrastructure operators. A mixture of institutional innovation and public rules can help to limit the extent of underinvestment in these public goods.

As one example, consider the role of cybersecurity data on vulnerabilities, attack strategies, and defenses. Once made public such data can be shared and used by many without diminishing its availability to others (that is, it is "non-rivalrous"). Just as public institutions provide detailed weather reporting for farmers (and other industries), public institutions can provide detailed data on how, for example, cyber-criminals are profiting from their incursions, and ensure regulation is aligned with reducing the potential profits from activities like fraud, identity theft, and ransomware. Data sharing across companies about fraud, suspicious accounts, and other yellow or red flags would be extremely beneficial. Cyber insurance companies could play an important role in information sharing here. (We discuss insurance below.)

While there is broad agreement that data sharing on cyber risk issues would be beneficial, organizations will be concerned about not making public data about customer lists or other business secrets. Data sharing could be limited so that the sharing only happens between firms, as already happens with sector-specific Information Security and Analysis Centers (ISACs). Another option is to require disclosure to government regulators but not the public, as has been adopted in EU regulations. Some proposals suggest that mandated disclosure of cybersecurity incidents might be paired with some shield against liability for the damage caused. The challenge here is to build a regulatory framework for data sharing; for example, it would provide specific guidance about how third party companies can combine and share data across firms for cybersecurity prevention purposes. Governments could make a substantial contribution in defining rules for such a framework.

Recent innovations in artificial intelligence makes it cheaper for good actors to create software tools that identify and fix vulnerabilities or test updates before releases. A variety of digital public goods could be procured by governments and provided to industry actors. Novel procurement approaches for this software could be used, including prizes, contests, or "paying for outcomes" whereby vendors could be rewarded based on outcomes such as a low number of economically damaging cybersecurity incidents.

A third approach to addressing the public good aspect of cybersecurity might be carried out by "bug bounty" platforms that utilize crowd-sourcing to find potential software vulnerabilities; for an examination of how these platforms work, Zrahia et al. (2024) offer a useful starting point. Bug bounty platforms are two-sided markets that host bug bounty programs and researchers and enable security researchers (sometimes called "ethical hackers") to be rewarded for finding software vulnerabilities. The programs are structured as tournaments where companies pay monetary rewards for unique vulnerabilities found.The two most prominent bug bounty platforms are Hacker1 and Bugcrowd. The popularity of bug bounty platforms is increasing, and some governmental agencies have begun to use them. For example, in 2021, the Cybersecurity and Infrastructure Security Agency launched a federal civilian enterprise-wide "crowdsourced" vulnerability disclosure policy platform (CISA (2021)). Nevertheless, these platforms have so far remained relatively small. While public disclosure of vulnerabilities is the default recommended policy for public programs at Bugcrowd and Hacker1, it is not required.

Systematic cyber risk could be lowered by ensuring that nearly all vulnerabilities discovered on these platforms are (at least eventually) publicly disclosed. As above, advances in artificial intelligence may fundamentally alter the behavior of such markets, as the cost of discovering new vulnerabilities is lowered.

4.4 A Private Insurance Approach to Systemic Cyber Risk

The private cybersecurity insurance industry offers another set of institutions that might reduce systemic cyber risk. Total US premiums for cybersecurity risk were $9 billion in 2024; premiums in the rest of the world were another $6 billion (National Association of Insurance Commissioners (2025)). However, despite the likelihood that cyber risk is rising, the total paid in cybersecurity insurance premiums has been largely flat in the last few years. One likely reason is that insurers quite rationally want to avoid systemic risk, because that can lead to correlated losses. They have largely succeeded in writing cybersecurity insurance policies to avoid covering systemic risks, which of course makes such insurance coverage less valuable to purchasers.

Nonetheless, insurance has the potential to play a critical role in measuring and managing systemic cyber risk, because insurance firms can gain data and expertise from evaluating and responding to cyber events across firms. Moreover, insurance firms have incentives to engage in research and development of knowledge, processes and software that helps assess cyber risk.

The key to getting social benefits from cybersecurity insurance is to ensure that it drives efficiency through scale economies (and the improved capabilities they bring) coupled with better aligned incentives, where insurance firms have substantial profit and reputation at stake in identifying and remediating problems. For example, insurance firms can use analysis from firms like Intangic that rate cybersecurity posture, and they can assess vulnerability to ransomware attacks by measuring the time it would take for a firm to recover compromised databases from backup (where longer times or uncertainty about capabilities mean firms are more vulnerable to ransomware demands). Firms with better demonstrated recovery capabilities are lower risk for ransomware payouts, and thus can be offered lower premiums. Insurance companies can use their experience to design more effective cybersecurity audits, and they can share learnings from prior incidents (such as unexpected challenges in recovering from backup) with insured firms.

The challenge is to propose regulatory or policy approaches that will further approaches that increase the incentvies for cyber insurance providers to contribute more directly to improving the risk profile of the firms they insure, either through incentives or the provision or curation of specialized services that are efficient and effective at reducing risk. Along these lines, a White House (2024) report from the White House on the "The National Cybersecurity Strategy Implementation Plan" acknowledges the imperative (in Strategic Objective 3.6) to deepen our comprehension of how potential catastrophic cyber incidents might impact the cyber insurance marketplace (White House, 2024).

Any proposed measures to bolster private cyber risk insurance must take into account the effects on the reinsurance market as well. Systemic cyber attacks can generate losses that no single insurer can absorb alone. Reinsurance makes it possible for insurers to underwrite large and complex cyber risks; without it, the insurance industry's ability to respond to catastrophic cybersecurity events would be seriously compromised. Because of the prevalence of systemic cyber risk, more than 50 percent of global cybersecurity insurance premiums are passed along to a reinsurance company; for property and casualty premiums, the equivalent figure would be about 20 percent Additionally, the cyber reinsurance market is highly concentrated: the five largest players underwrite roughly 62 percent of gross written premiums, while the top ten collectively control 87 percent of the market. Should any major reinsurer

curtail or exit its cyber capacity, the consequences for primary insurers — and for the market as a whole — could be severe (National Association of Insurance Commissioners (2025) and Willard (2024)).

5 Conclusion

Systemic cybersecurity risks are large and likely rising, while countermeasures against those risks seem inadequate. From an economic perspective, the underlying dynamics seem clear enough: cybersecurity efforts have a substantial public good component, and individual firms (whether firms at risk of hacking or firms providing insurance against hacking) lack sufficient incentive to take the necessary steps. Instead, many of them free-ride on the efforts of others.

In our own work in this area, we have been surprised that there is not more mapping of the specific types of market failures associated with systemic cyber risk to policies that would re-align incentives. It is perhaps unsurprising that non-economists have not focused on the incentive issues; in the case of economists, perhaps this shortcoming arises because the topic may appear too tactical for publishable economics research.

Our policy recommendations to reduce systemic cyber risk can be summarized as follows. The extent of concentration, competition, and interconnectedness affect the level of systemic cybersecurity risk and also the incentives of firms to address that risk. Incentives are shaped by market forces, institutional design, and other policies. We have pointed to areas where incentive alignment can be improved via policies, regulations, and private institutions such as insurance. But this essay is only a stepping board toward the further research and concrete actions that are needed to address systemic cyber risk.

6 Acknowledgments

Gandal and Moore gratefully acknowledge support from the US National Science Foundation (NSF) Award No. 2147505 and Award No. 2452738 and the US Israel Binational Science Foundation (BSF) Award No. 2016622 and Award No. 2021711. We are especially grateful to Tim Taylor (Managing Editor of the Journal of Economic Perspectives (JEP)) for help in rewriting the paper as well as comments and suggestions that substantially improved the paper. We are also grateful to Heidi Williams (Editor of the JEP), and Jonathan Parker & and Jeffrey Kling (Coeditors of the JEP) for important comments and suggestions that significantly improved the paper. We also thank Andrew Morin, Michael Riordan, Kostadis Roussos, and Vincent Scott for insightful discussions and suggestions.

References

Barnir, N., Gandal, N., Moore, T., and Scott, V. (2025). A cost-benefit approach to optimizing security control adoption. *forthcoming, Journal of Information and Computer Security, also available as SSRN Working Paper 4691404*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4691404.

Check Point Research (2026). Cyber security report 2026. https://www.checkpoint.com/lp/security-report-2026-gsem.

CISA (2021). Cisa announces new vulnerability disclosure policy (vdp) platform. https://www.cisa.gov/blog/2021/07/29/cisa-announces-new-vulnerability-disclosure-policy-vdp-platform.

Cybersecurity and Infrasructure Security Agency (CISA) (2026). Known exploited vulnerabilities catalog. https://www.cisa.gov/known-exploited-vulnerabilities-catalog.

Cybersecurity and Infrastructure Security Agency (2024). Systemic Cyber Risk Reduction. https://www.cisa.gov/resources-tools/programs/systemic-cyber-risk-reduction.

Cybersecurity and Infrastructure Security Agency and Federal Bureau of Investigation and National Security Agency (2022). Russian state-sponsored and criminal cyber threats to critical infrastructure. Technical report, Federal Bureau of Investigation. Accessed: 2026-01-07.

Easterly, J. (2026). Small towns shouldn't have to defend America's water supply from Iran. *New York Times*. https://www.nytimes.com/2026/08/02/opinion/iran-water-hack-trump.html.

Forscey, D., Bateman, J., Beecroft, N., and Woods, B. (2022). Systemic cyber risk: A primer. https://carnegieendowment.org/2022/03/07/systemic-cyber-risk-primer-pub-86531.

Fox-Sowell, S. (2024). Cyberattacks on state and local governments rose in 2023, says CIS report. *StateScoop*. https://statescoop.com/ransomware-malware-cyberattacks-cis-report-2024/.

Fung, B. and Sands, G. (2021). Former solarwinds ceo blames intern for 'solarwinds123' password leak. *CNN*. https://www.cnn.com/2021/02/26/politics/solarwinds123-password-intern.

Gartner (2024). Explore Enterprise Software Categories | Gartner Peer Insights.

Hajizada, A., Adams, M., and Moore, T. (2024). Construction and analysis of a large-scale firm-level cybersecurity posture dataset. In Carter, M., Fadel, K. J., Meservy, T. O., Armstrong, D. J., Deokar, A., and Jensen, M. L., editors, *30th Americas Conference on Information Systems: Elevating Life through Digital Social Entrepreneurship, AMCIS 2024, Salt Lake City, UT, USA, August 15-17, 2024*. Association for Information Systems.

IBM (2026). Cost of a data breach 2026. https://www.ibm.com/reports/data-breach.

Kapko, M. (2025a). Salt typhoon gained initial access to telecoms through cisco devices. https://cyberscoop.com/cisco-talos-salt-typhoon-initial-access/.

Kapko, M. (2025b). Salt typhoon remains active, hits more telecom networks via cisco routers. https://cyberscoop.com/salt-typhoon-china-ongoing-telecom-attack-spree/.

Kunreuther, H. and Heal, G. (2003). Interdependent Security. *Journal of Risk and Uncertainty*, 26(2):231–249.

Mousqueton, J. (2026). Ransomware.live. https://ransomware.live.

MS-ISAC (2025). Strengthening critical infrastructure: State, local, tribal, & territorial progress & priorities volume 1: Brief report. https://learn.cisecurity.org/strengthening-critical-infrastructure-sltt-progress-priorities-brief-report-vol1.

National Association of Insurance Commissioners (2025). Report on the cybersecurity insurance market. https://content.naic.org/sites/default/files/inline-files/2025_Cybersecurity_Insurance%20Report.pdf.

National Institute of Standards and Technology (2024). The NIST Cybersecurity Framework (CSF) 2.0. Technical Report NIST CSWP 29, National Institute of Standards and Technology, Gaithersburg, MD.

National Institute of Standards and Technology (NIST) (2026). National vulnerability database (nvd). https://nvd.nist.gov/.

Nolan, C. and Fixler, A. (2021). The economic costs of cyber risk. https://www.fdd.org/analysis/2021/06/28/the-economic-costs-of-cyber-risk/.

Recorded Future Insikt Group (2025). Redmike (salt typhoon) exploits vulnerable cisco devices of global telecommunications providers. https://www.recordedfuture.com/research/redmike-salt-typhoon-exploits-vulnerable-devices.

Shao, X., Siponen, M., and Liu, F. (2020). Shall we follow? Impact of reputation concern on information security managers' investment decisions. *Computers & Security*, 97:101961.

Spiceworks Ziff Davis (2024). Data Intelligence. https://swzd.com/products/data/.

Tidy, J. (2020). SolarWinds: Why the Sunburst hack is so serious. *BBC News*. https://www.bbc.com/news/technology-55321643.

U.S. Department of Homeland Security (2013). National infrastructure protection plan. https://www.cisa.gov/sites/default/files/publications/national-infrastructure-protection-plan-2013-508.pdf.

U.S. Department of Justice (2024). Herfindahl-Hirschman index. https://www.justice.gov/atr/herfindahl-hirschman-index.

US Securities and Exchange Commission (2023). Press release: SEC charges SolarWinds and chief information security officer with fraud, internal control failures. https://www.sec.gov/newsroom/press-releases/2023-227.

US Securities and Exchange Commission (2025). Litigation release no. 26423—SEC v. SolarWinds corp. and timothy g. brown. No. 1:23-cv-09518-PAE (S.D.N.Y. filed Oct. 30, 2023).

Volz, D. and Londoño, E. (2026). Scope of hacks on U.S. water supply widens as evidence points to Iran. *New York Times*. https://www.nytimes.com/2026/08/01/us/politics/iran-cyberattack-water-systems.html.

Weishäupl, E., Yasasin, E., and Schryen, G. (2018). Information security investments: An exploratory multiple case study on decision-making, evaluation and learning. *Computers & Security*, 77:807–823.

White House (2024). National Cybersecurity Strategy Implementation Plan, Version 2. https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/05/National-Cybersecurity-Strategy-Implementation-Plan-Version-2.pdf.

Willard, J. (2024). Ceded reinsurance, an ever smaller proportion of total premium over the last five years: Howden Re. https://www.reinsurancene.ws/ceded-reinsurance-an-ever-smaller-proportion-of-total-premium-over-the-last-five-years-howden-re/.

World Economic Forum (2026). Global cybersecurity outlook 2026. https://www.weforum.org/publications/global-cybersecurity-outlook-2026/.

Zrahia, A., Gandal, N., Markovich, S., and Riordan, M. H. (2024). The simple economics of an external shock to a bug bounty platform. *Journal of Cybersecurity*, 10(1):10–1093.