4 STEPS TO BECOMING Cyber Risk Savvy

How To Be a Smart Customer of Cyber Insurance

BY SHAYNE KAVANAGH, ROB ROQUE AND TERI TAKAI

Cyberattacks are a clear and present danger for all organizations, but local governments are particularly vulnerable. A 2020 study showed that local governments are more likely to be the targets of a ransomware attack than any other kind of organization and that 44% of ransomware attacks targeted local governments in 2020, a portion similar to 2019.[1] The trend does not seem to be abating: 2021 saw a nine-fold increase in ransomware attacks on government organizations between 2020 and 2021.[2]

Local governments are attractive targets for cybercriminals for a few reasons.[3] First, local governments are "soft targets." This means that networks are typically not very secure. For example, smaller local governments may not have dedicated IT staff, much less dedicated cybersecurity staff. On top of that, local governments often operate many disparate services, which creates a lot of "surface area" for an attack. In other words, an attacker could gain access to a city government's network through information systems in public works, community development, or any other department. Second, local governments maintain sensitive data like tax records, voter information, citizen and employee health-related data, and employee social security information. They also provide essential services that can't be interrupted. A soft target with sensitive information and essential services is the proverbial "low-hanging fruit" for the cybercriminal. A third and, perhaps, surprising reason is the public profile of local governments, which refers to transparency requirements, open data sets, public-facing internet-enabled transactions, and more. This public profile means hackers have an advantage in calculating an effective strategy to penetrate a local government's defenses. This compares to private firms that have a greater ability to conceal their activities from the public and, therefore, cybercriminals.

Cyberattacks are expensive. Cities like Atlanta and Baltimore have made headlines with the extreme cost of a cyberattack. These cities are reported to have incurred over $15 million each, including data recovery costs and the cost of downtime and lost revenue.[5] The risks are not limited to large governments. In 2019, the City of Stuart, Florida, (population 16,000) was hit with a ransomware attack and a demand for $300,000. The city elected not to pay and had to incur about 2,000 hours of staff time to manage the recovery and work-arounds and spent a significant sum on replacing/upgrading hardware and software.[6] Further, a study of the costs of cybercrime across industries showed that there was barely any relationship between the size of the victim organization and the size

Local governments are more likely to be the targets of a ransomware attack than any other kind of organization.

WHAT IS RANSOMWARE?

Ransomware is "a type of malicious attack where attackers encrypt an organization's data and demand payment to restore access."[4] Organizations fall prey to these types of cyberattacks by clicking on a malicious web link in an unsuspecting email (phishing) or visiting an unprotected website and unknowingly downloading and activating malware.

of the loss.[8] In other words, a smaller organization does not necessarily translate into lower potential losses from cybercrime.

The potential extreme consequences of a cyberattack have caused many local governments to turn to cyber insurance. Given the potential losses from an attack, transferring the risk of an attack to the insurance market could be an attractive proposition. However, cyber insurance is a relatively new type of insurance instrument compared to traditional insurances, like property and liability insurance. Also, the cost of a policy or the availability can change dramatically in a short time. In fact, as of this writing, many governments have experienced rapidly increasing premium costs. This article will help local governments approach cyber insurance in a risk-savvy manner and make smart decisions about how to invest in protection against cybercrime.

As a first step, let's understand three fundamental issues with cyber insurance that an informed consumer must be aware of.

First, insurance is remedial, whereas controls (cybersecurity measures) can be preventative. For example, training on safe computing practices can make it less likely that an employee clicks on a malicious web link in an email, thereby avoiding an attack that could have otherwise succeeded.

Prevention is generally preferable to remediation. Cyberattacks can have consequences beyond what insurance can cover. For example, the City of Stuart found that even if it had been able to use insurance to pay the ransom, the files that would be "restored" by the cybercriminal would go to one folder, with all new names and no file extensions! Insurance is not an "undo button" for a cyberattack. There are also indirect effects of a cyberattack that are best avoided, such as the hit to the reputation of a local government. Reputation is not an inconsequential intangible. A loss of public faith in government has consequences. A perceived vulnerability to cybercrime also could have consequences for bond ratings.[8] *This means that local governments must be savvy in choosing when to invest limited resources in better cybersecurity controls versus investing in cyber insurance.*

Second, commercial insurance, by design, is a "bad bet" for the insured, on average. If it weren't, insurance companies would go broke. This is why governments can sometimes reduce costs by self-insuring. This does not mean local governments should never buy commercial insurance. Commercial insurance is great for protecting against catastrophic losses that government isn't capable of absorbing. *This means local governments must be savvy in determining when to accept the risk (self-insure) and when to transfer risk to commercial insurers.*

Third, the market for cyber insurance continues to change and evolve with the level of threat posed to governments by cybercrime. The cyber insurance market is relatively underdeveloped, and fewer actuarial models exist compared to other kinds of insurance markets—which have been around for decades and maybe centuries. Hence, the market

for cyber insurance is evolving rapidly as insurance sellers and buyers come to understand the nature of the peril better and the financial implications of insuring it. As of this writing, the market for cyber insurance is tightening up, with policies becoming unaffordable or unavailable for local governments that don't have adequate controls to prevent cyberattacks. *This means local governments must be savvy about recognizing the evolving nature of the cyber insurance market and not assume that today's coverages will be available at comparable prices in the future.*

With these issues in mind, how should a local government approach cyber insurance? The rest of this article will take you through a step-by-step procedure for considering the costs versus the benefits of cyber insurance.

Risk Mitigation vs. Risk Transfer, or Cybersecurity Controls vs. Cyber Insurance

We will start from the premise that local government has limited resources, so a dollar invested in cyber insurance is a dollar not invested in controls. The advantage of controls is that they can be preventative; they can stop the attack from doing damage in the first place. A software patching strategy leaves fewer vulnerabilities for cybercriminals to exploit. Controls can also reduce the potential damage from an attack if an attack succeeds. For example, high-quality data backups make it easier to recover lost data.

Insurance is always remedial; it cleans up the damage after it has happened. The advantage of insurance is that it can provide some relief from catastrophic losses, where it is impractical to develop sufficient controls. Hence, there is a trade-off to consider. How can this trade-off be analyzed? We will present a four-step process:*

Step 1—Know the basics of your cybersecurity situation

Step 2—Quantify your risk

Step 3—Examine the potential of insurance

Step 4—Periodically reassess

STEP 1

Know the basics of your cybersecurity situation

Some local governments will have a good handle on their existing cybersecurity situation, but others may not. There are three questions to ask as part of Step 1:

What are the most important assets you need to protect? Technology assets with sensitive data or that administer mission-critical functions are the most important. These may include social security numbers, credit card information, bank account information, any kind of health data that might be protected by law (e.g., the U.S. Health Insurance Portability and Accountability Act), and criminal justice data. Examples of critical systems might include enterprise resource planning (ERP), tax revenue systems, or public health or public safety systems.

What threats are most important? Today, ransomware attacks are the most prevalent threat. Other possible threats include denial of service attacks, leaks of sensitive data, or cyber sabotage of various forms. Ransomware attacks will likely continue to be the top threat because there is a clear financial incentive for the perpetrator. It is worth noting that these threats can combine. For example, a ransomware attack could lead to data leaks.

What is the state of your controls? State and local governments have been challenged with finding resources to keep up with cyber threats. Important controls include multifactor authentication, firewalls, encrypted data storage, encrypted data backups, incident response planning, training staff to avoid phishing attacks, software patching, and endpoint detection response.** In a 2021 survey,[9] respondents indicated that spending on cybersecurity focused on software, hardware, backup, monitoring, and training. Incident response was listed as a lower priority. Only 57% of responses indicated that cybersecurity training was done annually for all employees. The focus areas for business continuity in the face of a cybersecurity attack were data backups and recovery, operational business plans, and ensuring manual work-arounds in case of an outage.

There are comprehensive frameworks for addressing cybersecurity risks, like CIS Top 18 (perhaps the most accessible for local government), COBIT, NIST, and ISO. These are valuable for organizations with the sophistication to use them. However, even a basic assessment of whether you have the controls we described here, or not, can be useful for Step 1. At the end of Step 1, many local governments will find that they have

CAN YOU ELIMINATE RISKS?

One strategy in risk management is to eliminate risks by eliminating risky activities. In the world of cyber insurance, an opportunity might be to reduce the amount of sensitive data that government collects and stores. You might ask if collecting and storing certain types of sensitive data is necessary and worth the exposure it brings.

* The four steps of this process are based on the "Cyber Loop" method described in: "Protecting Today. Safeguarding Tomorrow. The Cyber Loop: Managing Cyber Risk Requires a Circular Strategy," published by Aon in 2019. https://www.aon.com/cyber-solutions/wp-content/uploads/Aons-Cyber-Solutions_The_Cyber_Loop.pdf.

** If you are not familiar with the controls in this sentence, please see the Appendix.

opportunities to invest more in cyber controls. In particular, multifactor authentication, firewalls, patching, and training employees on safe computing practices are potentially valuable controls and may represent a wise investment in cyber risk prevention.

STEP 2

Quantify your risk

It will be difficult, if not impossible, to make a savvy decision about the trade-offs between investing in controls and purchasing insurance without quantifying the risks. "Risk" can be defined as the chance of the occurrence of a loss, disaster, or other undesirable event multiplied by the magnitude of the loss. This definition implies that risk is a quantifiable property.*

People have attempted qualitative risk analyses in the form of a "risk matrix" or "heat maps," where risks are classified along a scale such as "low," "medium," and "high," and color coded according to severity. However, research has shown that this kind of analysis can lead to worse decisions![10] A reason for this is an "illusion of communication," where decision-makers falsely believe that everyone who is part of the decision has a similar understanding of the risk.[11] The problem is that categories like "low," "medium," and "high" are vague and invite different interpretations by different people. Imagine one of your colleagues is an inveterate sports gambler and another has never so much as purchased a lottery ticket. These two people probably have very different definitions of "low" risk. However, if risk is quantified, like "we believe there is a 10% chance of a ransomware attack costing us more than $100,000 in the next year," there is less room for interpretation.

Another reason that risk matrices can be counterproductive is they act as an "analysis placebo,"[12] where decision-makers think they understand the risk because they have subjectively characterized the risk as "high," "low," etc. But, because the risk matrix is not based on hard data about the chance of loss and the potential magnitude of loss, decision-makers are overconfident about how well they understand the risk.

Typical Risk Matrix Often Leads to Worse Decisions About Risk

Likelihood \ Impact	Negligible	Minor	Moderate	Significant	Severe
Very Likely	Low Med	Medium	Med Hi	High	High
Likely	Low	Low Med	Medium	Med Hi	High
Possible	Low	Low Med	Medium	Med Hi	Med Hi
Unlikely	Low	Low Med	Low Med	Medium	Med Hi
Very Unlikely	Low	Low	Low Med	Medium	Medium

Though risk matrices are easy to create, easy to understand, and inexpensive, if they lead to lower quality decisions, they aren't a good deal. The alternative to a subjective risk matrix is to quantify risks.

In this article, we will not get deep into the details of how to quantify risk. The details of quantifying risk are best left to professional risk analysts. Instead, we will show concepts that will help you think about cyber risks in a way that is consistent with a quantified approach and that will help you ask the right questions of the risk professionals who are versed in the details of risk quantification. If you would like to dig deeper into risk quantification, here are three sources for further detail:

- GFOA has built a sample Excel ransomware risk model that uses the same methods to quantify risks that insurance companies use but is built using the open Probability Management standard.[13] This model is not a substitute for professional risk analysis and is intended only as an educational tool for ransomware risk. It will provide you with a basic understanding of how the risks of a cyberattack could be quantified. It is not intended to provide a comprehensive analysis of your cybersecurity risk. The content of Step 2 in this article will be largely based on the sample model but will not cover all of the details in the model. You can get access to the model at gfoa.org/cyber-insurance.
- Finally, GFOA has found that some insurance companies are taking steps to provide clients with richer quantification of risk. They believe that more informed customers will be better long-term customers. Understanding the concepts in this article will help you ask insurance providers for the right information and make the best use of the information.

*Loss also includes things that are sometimes thought of as "intangible," like community trust, reputation, etc. These losses are also measurable, though not as easily as some other losses. For more on this subject, see: Hubbard, D. (2014). *How to measure anything: Finding the value of intangibles in business.* Wiley.

Before we start our discussion of quantifying risks, we'd first like to acknowledge that quantifying risks is often not the normal course of business for local governments. As such, it is natural that there might be some skepticism about the potential for quantifying risks. We'd like to present three common objections to quantification posed by the skeptic and our response:[14]

Objection 1: Quantifying risk is more appropriate for insurance industry analysis and is unlikely to be appreciated by local governments looking for practical advice.

Answer: It is common for us to underestimate the capabilities of other people relative to our own.[15] GFOA has presented quantified risk information to many elected officials and government staff and has yet to find one who could not at least grasp the essential point. As for practicality, given that subjective methods (like a risk matrix) often lead to worse decisions, we would suggest that it is the subjective methods that don't work in practice.

Objection 2: The cyber insurance market is volatile, so decisions based on a quantitative model will be wrong.

Answer: Insurance companies have been making decisions based on quantitative methods as early as the 17th century. This does not mean that every decision an insurance company has ever made is perfect. But it is understood within the insurance industry that it would be foolish to attempt to compete without quantitative methods.[16] The next objection is also relevant to this issue.

Objection 3: Within cybersecurity, there are too many complexities changing too quickly to make a reasonably accurate assessment.

Answer: One way or the other, a government has to decide on how to invest in commercial insurance, self-insurance, and controls for cybersecurity. A government can either take a wild guess and hope for the best or take a more rigorous approach. No quantitative model will be perfect, but a model can still be useful. To be useful, a model does not have to be perfect; it just needs to outperform the alternative, which is a subjective judgment. Because a quantitative model forces rigor and transparency in how you think about a question, there is a chance that even an imperfect model will outperform subjective judgment.[17]

With the common objections to quantifying risk addressed, the first step in quantifying your risk is to get data on how likely a loss from cybercrime is and how big that loss might be. First, we must recognize that definitive data is going to be very difficult, if not impossible, to come by. But remember, a model does not have to be perfect; it just needs to outperform the alternative (e.g., guesswork). That said, let's start with the chance of a successful ransomware attack, defined as multiple computers infected *and* files are successfully encrypted. This means the local government is *not* able to stop the attack once the computers were infected. Our off-the-record conversation with a local government risk pool found that their pool members experienced roughly a 5% to 10% chance of a successful ransomware attack for a pool member in a given year. Moving on to damages from a successful attack, according to the NetDiligence Cyber Claims Study: 2021 Report, the five-year average total incident cost averaged $267,000, but with a median of $98,000.[18] This tells us that average is pulled upwards by a small number of catastrophic losses. The data showed 10% of incidents cost more than $638,000, and some cost much more: millions of dollars. Total incident response includes costs like forensics, business interruption, recovery, and paying the ransom (if one is paid). Finally, we should recall that cyberrisk is an evolving threat, so these figures could change year to year, perhaps significantly.

Next is to visualize this data to understand the implications of your baseline level of risk. There are many ways data could be visualized, but we'll use what is known as a "loss exceedance curve" (LEC). An LEC presents risk in the way that

REAL-LIFE EXPERIENCES | RENEWING CYBER INSURANCE IN 2022 FOR LOCAL GOVERNMENTS

On the GFOA member forum, we asked people to share their experiences with renewing cyber insurance for 2022. The two quotes below capture the experience of people who replied.

"The renewal quote has nearly doubled, and the retention amounts, particularly for ransomware incidents, have increased substantially, to the point where an individual government would face significant (think potentially seven figure) out-of-pocket exposure to a cyber event before any insurance coverage would kick in."

"We had cyber insurance until this past year. Upon renewal, the insurance provider needed a brand new questionnaire with far more significant requirements, multifactor authentication, as well as a proven and regular phishing training program and other quite significant requirements. As a result, we have been declined for this year and are working to see if we can get back onside with the requirement."

insurance companies think about it and is commonly used in different industries to depict risk. An LEC can be constructed for specific applications (e.g., ERP), departments (e.g., police), risks (e.g., ransomware), or any other relevant perspective. Exhibit 1 shows an LEC for a successful ransomware attack. The vertical axis shows the chance of a given loss (or greater) occurring, and the horizontal axis shows the loss. For instance, there is about a 40% chance of losing at least $160,000 because an attack was successful. This is because the blue line passed through the 40% mark at about $160,000. The blue line skims along the bottom of the graph for some distance, which indicates a small chance of catastrophic losses.

However, the damages from a successful attack must be considered against the chance an attack will succeed in the first place. Exhibit 2 shows an LEC with the chance that a successful attack will occur factored in. You can see that the blue line that intersects the vertical axis has a much lower chance in Exhibit 2. This is because a successful attack is not a high-probability event.

The blue lines in Exhibits 1 and 2 show what is known as "inherent risk." This is your baseline level of risk, reflecting the controls you have in place now. The analysis can show how the curve would change if you invested in additional controls. For example, perhaps you could invest in better data backup to reduce the damage from a successful attack—and in better training for employees to guard against phishing attacks to reduce the chance of a successful attack. Exhibit 3 shows what a 10% reduction in the chance of a successful attack and a 30% reduction in potential damages would look like via the orange line. You can see that the orange line intersects the vertical axis at a lower point, which means you've lowered your chance of experiencing damages. There is also a substantial gap between the orange and blue lines all along the curves. This gap represents the lower potential damages from the mitigations.

The orange line in Exhibit 3 is also known as "residual risk." This is the remaining exposure that would be left after making optional investments in additional controls. In the sample risk model, you determine the size and type of the investment, and you could explore different options for investing in controls. Making additional investments in controls shifts the curve downward, which means the risk profile becomes more favorable. There are two caveats to consider here, though. First, controls can fail, be poorly implemented, or otherwise not live up to expectations. Hence, a good control strategy is diversified so that you are not dependent on any single control. Second, residual risk can't reach zero. Not only is this a theoretical impossibility, as long as the government uses information technology, but it is also a practical impossibility, given the limited resources available for cybersecurity. Hence, risk savvy is a matter of identifying the point where you are willing to make additional security investments, where you will rely on insurance, and where you will absorb risk.

The quantification of your baseline (or inherent risk) and of the potential to invest in controls (or residual risk) accomplishes two goals:

First, it helps you evaluate the value of investing in additional controls. For example, local governments may find there is a strong case to invest in new controls such as training on safe computing practices for staff, multifactor authentication, virtual private networks, and data encryption and backup services. In particular, this kind of analysis can show the value of training. Decision-makers can see the reduction in risk that training provides. Research suggests that local governments have substantial opportunities to improve their controls. One study showed that local government was among the least effective sectors in stopping a ransomware attack before data could be encrypted. This same study showed two sectors most successful in stopping attacks (media, leisure, entertainment, and distribution/transport) were about 60% more successful than local government.[19]

If your controls are already strong, the analysis might highlight the limited benefit available from additional investment. For example, if you had to spend $1 million on new controls for an average reduction in your damages of $100,000, you might reasonably question if that is a good investment! The GFOA sample risk model for ransomware walks you through some return on investment calculations for controls.

The second goal that quantification accomplishes is to set the stage for making a wise decision about investing in controls versus insurance. We'll take this up in more detail in Step 3.

EXHIBIT 1 | LOSS EXCEEDANCE CURVE FOR A SUCCESSFUL RANSOMWARE ATTACK

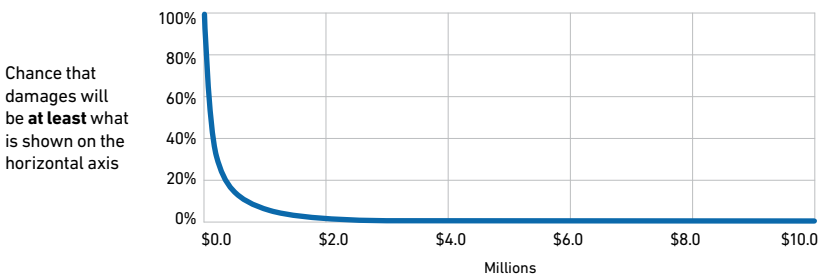

STEP 3

Examine the potential of insurance

First, "self-insurance" should not be overlooked. Local governments often set up self-insurance for all types of risks. There is no reason that self-insurance couldn't work for cyber risk as well. Self-insurance might be especially important in a tight market for commercial insurance for two reasons: First, to reduce the cost of a commercial policy to an affordable amount, governments might be forced to accept a higher retention amount* on the policy. A retention amount is a form of self-insurance. Second, if a policy is unobtainable, self-insurance might be the only option left past the point where investment in additional controls ceases to be practical.

EXHIBIT 2 | LOSS EXCEEDANCE CURVE, GIVEN THE CHANCE OF ONE OR MORE SUCCESSFUL ATTACKS IN A YEAR

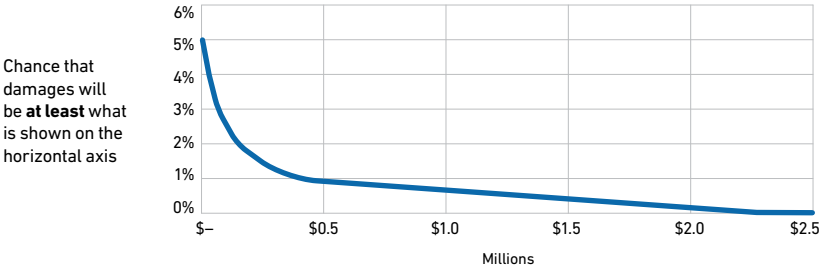

EXHIBIT 3 | LOSS EXCEEDANCE CURVE WITH THE IMPACT OF NEW CONTROLS ADDED

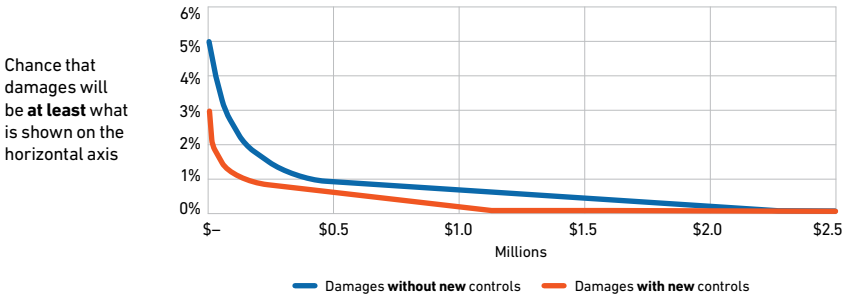

EXHIBIT 4 | LOSS EXCEEDANCE CURVE WITH AN AMOUNT AVAILABLE FOR SELF-INSURANCE

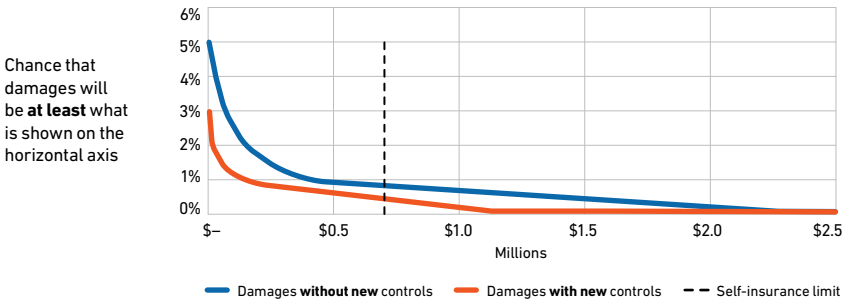

For these reasons, Step 3 should include an analysis of self-insurance capacity. This is a matter of determining the amount of risk you are willing to absorb via self-insurance. Exhibit 4 adds to our LECs from Step 2 by including the amount a government is willing to put aside for self-insurance—$700,000 in this case. This could be derived from the number of liquid resources a government has available to respond to unplanned emergencies (e.g., reserves). You could then determine the chance that you will exceed this amount and compare that chance to your appetite for risk. We have indicated the chances in Exhibit 4, and the GFOA sample model shows the chances for any self-insurance amount you enter. Would you be comfortable with an 8% chance (or one in twelve years) that self-insurance would be inadequate for the losses you experience in a year or, put another way, a 92% chance that self-insurance would be adequate? If not, you might need to consider commercial insurance if further self-insurance is impractical.

Self-insurance is often most valuable at a point where: A) investing in more controls loses cost-effectiveness, and B) commercial insurance can be made more affordable by accepting a higher retention. Knowing the amount available for self-insurance is a good place to start in considering the role of commercial insurance.

Commercial insurance is most useful at the far end of the loss exceedance curve. There is some unavoidable risk in operating a modern local government. For example, a local government could reduce a lot of cybercrime risk by severing all of its connections to the internet, but that would present an unacceptable cost in lost operational efficiency. This means that the risk of extreme losses is unavoidable. The far end of the loss exceedance curve is where the potential losses are too high to absorb via self-insurance.

* Retention is the total amount of a loss the insurance policyholder must pay out of pocket. It includes the deductible but is not limited to the deductible. For example, insurance policies have limits on how much will be paid out. The government is retaining the risk that the cost of a cyber incident could be more than the policy limit.

Commercial cyber insurance can, *theoretically*, cover a variety of different losses. Exhibit 5 provides an overview of the different coverages that *could* be available.[20] Make note of our use of conditionals like "theoretically" and "could." Market conditions will determine if an insurance company is willing to sell you any of these policies. In fact, GFOA spoke with one large reinsurer that was refusing to underwrite cyber policies.

That said, as of this writing, many insurance providers are willing to sell policies. Even so, they may place limits on the policy (i.e., boundaries on what is covered and what is not). Smart customers will understand these limits and their implications. Let's examine the limitations that appear in cyber insurance policies in the next sections.

Underwriting

Underwriting is the process insurers use to determine the risks of insuring your government. The underwriting process has intensified in recent years. Many insurance companies are using specialized cyber risk consultants to help them assess risk more accurately. Underwriters are increasingly looking for the insured to have key security features as a prerequisite for a policy. Such features might include multifactor authentication, incident response planning, encrypted data storage, patching cadence, and endpoint detection response.* Governments with inadequate internal security might have trouble getting a policy or might face increased costs. Earlier in the article, we quoted a GFOA member who could not secure a policy due to more intensive underwriting. Another member reported facing a doubling of premiums unless they implemented multifactor identification.

Key questions to ask: How can you make the best impression on your underwriters to convince them you are a good risk? Do you have cost-effective opportunities to improve your security controls?

* See the Appendix for definitions of these controls.

EXHIBIT 5 | CYBER COVERAGE OVERVIEW

Operational Risks

Network Business Interruption—Covers lost net income caused by a network security failure, as well as an associated extra expense.

System Failure—Expands coverage trigger for business interruption beyond computer network security failure to include system failure.

Dependent Business Interruption/Dependent System Failure—Coverage for lost income caused by a network security failure of a business on which the insured is dependent, as well as an associated extra expense.

Cyber Extortion—Coverage for expenses incurred in the investigation of a threat and any extortion payments made to prevent or resolve the threat.

Digital Asset Restoration—Coverage for costs incurred to restore, recollect, or recreate intangible, nonphysical assets (software or data) that are corrupted, destroyed, or deleted due to a network security failure.

Privacy and Network Security Risk

Privacy Liability—Coverage for defense costs and damages suffered by others for failure to protect personally identifiable or confidential third-party information.

Security Liability—Coverage for defense costs and damages suffered by others resulting from a failure of computer security, including liability caused by theft or disclosure of confidential information, unauthorized access, unauthorized use, denial of service attack, or transmission of a computer virus.

Privacy Regulatory Fines and Penalties—Liability coverage for defense costs for proceedings brought by a governmental agency in connection with a failure to protect private information and/or a failure of network security pursuant to applicable laws or regulations.

Media Liability—Coverage for defense costs and damages suffered by others for content-based injuries such as libel, slander, defamation, copyright infringement, trademark infringement, or invasion of privacy.

PCI Fines and Penalties—Coverage for a monetary assessment from a payment card association (e.g., MasterCard, Visa, American Express) or bank processing payment card transactions (i.e., an "acquiring bank") in connection with an insured's noncompliance with PCI Security Standards.

Breach Event Expenses—Reimbursement coverage costs to respond to a data privacy or security incident. Covered expenses include certain computer forensic expenses, legal expenses, costs for a public relations firm and related advertising to restore your reputation, consumer notification, call centers, and consumer credit monitoring services.

Cybercrime Insurance Coverage

Social Engineering Coverage—Coverage for direct financial loss as a result of fraudulent instructions provided by a third party that is intended to mislead an insured through the misrepresentation of a material fact.

Funds Transfer Fraud—Coverage for direct financial loss as a result of fraudulent instructions provided to a financial institution that authorize the transfer of the insured's funds by a third party impersonating an insured.

Computer Fraud—Coverage for direct financial loss sustained resulting from the unauthorized seizure of funds from their computer network by a rogue employee or malicious third party.

Miscellaneous Cyber Insurance Coverages

Reputational Income Loss—Coverage for lost net income caused by bad publicity resulting from a security event.
Bricking Coverage—Coverage to replace hardware rendered inoperable due to a security breach.

Claims Avoidance Coverage—Coverage for expenses incurred as a result of the insured's reasonable investigation of a potentially covered claim.

Reward Payment Coverage—Coverage of payment for information that leads to the conviction of any individual committing or attempting to commit an illegal act relating to a security event.

Betterment Coverage—Coverage for expenses incurred to update, restore, or improve computer systems to a level beyond that which existed before a security event.

Overview of Coverages Provided Courtesy of Aon

Payout Limits and Sublimits

A policy limit is a maximum amount a policy will pay out. Sublimits are a traditional part of insurance policies. Sublimits are a limit on the reimbursable loss for a particular type of risk that is less than the total limit on the entire policy. The savvy customer will review all policy language and make note of any sublimits. Sometimes sublimits are clear on the declarations page of the policy; but other times you will need to review the policy definitions and endorsements to find sublimits. Sublimits are important because you may find that you have less coverage for a particular type of risk than the limit on the entire policy might have led you to believe. Common sublimits include:

- **Ransomware**—Limiting the total coverage available for a ransomware attack versus the total limit for all cybercrimes.
- **System failure**—Limiting the coverage for a cascading system failure, where a failure in one system leads to failures in other integrated systems. For example, staff may not be reimbursed for personal devices that were damaged as a result of connecting to an infected network at work.
- **Bricking**—Limiting the reimbursement for replacing hardware that is rendered unusable by a cyberattack. For example, the company may cover "hard bricks," where the device is made inoperable, but not a "soft" bricked device, where part of the device may be operable or repaired.

Key questions to ask: What are the sublimits in your policy? Do these sublimits change your understanding of the level of coverage you have? What implications does that have for your investment in cyber insurance (commercial or self-insurance) versus cyber controls?

Retentions

Retentions are another traditional part of an insurance policy. Retention is the risk that is retained by the insured or, put another way, the amount of damages the insured will have to pay out of pocket outside of what is covered by insurance. Lower retentions are not necessarily better because a policy with lower retention will cost more. Higher retention could be a way to reduce the cost of the policy if the government can self-insure for the larger retention. Note that "retention" commonly refers to policy deductibles but does encompass other retailed risks. For example, if a policy has a low limit, then the risk that an incident will cost more than the limit would also be retained by the government.* Another issue is "single highest retention." Exhibit 6 shows a hypothetical cyber insurance plan with five policies. An attack happens and triggers three of the policies. The single highest retention looks across all three policies and selects the highest retention (deductible) as the amount the insured pays. The multiple retention policy sums up all of the retentions. Though a multiple retention policy would likely be less expensive, a single retention policy might be preferable because it will be easier for the insured to estimate the retention cost of a given attack. In Exhibit 6, the insured need only look across the retentions of all five policies, find the minimum and the maximum, and that is the range of possible retentions for a given attack under a single retention policy. For a multiple retention policy, the range is the smallest single retention to the sum of all the retentions in the policy. The latter would be more difficult to plan for.

Key questions to ask: What balance between retention (self-insurance) and policy price (commercial insurance) is best for you? Is your policy single highest retention?

EXHIBIT 6 | SINGLE RETENTION VS. MULTIPLE RETENTIONS

	Retention
Security liability	$500,000
Regulatory liability	$500,000
PCI (payment card industry)	$500,000
Breach response	$750,000
Business interruption	$500,000

↓

Attack happens and triggers these policies:
PCI (payment card industry)
Breach response
Business interruption

If you had **Single Highest Retention**:	
PCI (payment card industry)	$500,000
Breach response	$750,000
Business interruption	$500,000
You pay **max** of the group above	**$750,000**

If you had **Multiple Retentions**:	
PCI (payment card industry)	$500,000
Breach response	$750,000
Business interruption	$500,000
You pay **sum** of the group above	**$1,750,000**

*Other examples of retained risk include 100% self-insurance strategies that are apart from a commercial policy or risks that a government chooses not to insure at all (self or commercial).

Panel Requirements

Next, it is critically important to know the requirements to secure assistance from a preapproved cybersecurity contractor in the event of a breach. The list of preapproved contractors is known as a "panel" and may include all aspects of support for a breach (e.g., legal counsel, technology support, etc.). This is similar to personal automobile insurance, where, in the event of an accident, the insurer requires the insured to use an approved auto repair shop. With cyber insurance, if you use a cybersecurity contractor that is not on the insurer's list of approved providers (known as going "off panel"), then you may lose all coverage for a breach response. Some insurers provide options on which contractors you may use (or may require you to use a single contractor), but it is important to know the requirements and your options at the start of the policy period. Finally, we should note that panel requirements are not necessarily a bad thing. For example, paying a ransom might require paying in cryptocurrency. An experienced consultant, like would be found "on panel," will be able to secure the right kind of cryptocurrency faster than a government.

Key questions to ask: What obligations do you have to use a specified security contractor to help respond to a breach? What options do you have to select between contractors?

Exclusions

Insurance exclusions are policy provisions that waive coverage for certain risks or loss events. Smart customers understand the exclusions; otherwise, their policy may not provide coverage for risks that the customer assumed would be covered. According to one cybersecurity expert we spoke with, in almost every cyber insurance event they've been involved with, the insured did not take the time to understand the exclusions, to their great detriment. An example of a common exclusion for cyber policies is civil and legal liabilities for breach of personal data.

Ransomware is the leading type of cyber threat for local governments, so let's examine some of the germane exclusions to that type of policy. In the previous section, we discussed what are known as "panel requirements," or the requirement that the insured use only preapproved cybersecurity contractors to respond to a breach. Going "off panel" is a form of exclusion, where using an unapproved contractor could result in an exclusion from coverage.

An evolving issue is federal government policy on responding to ransomware attacks.* From the perspective of an individual organization that is the victim of an attack, the logical response is for the insurer and the customer to figure out if it is better to pay the ransom or pay the cost to recover the affected systems without access to the ransomed data. However, this presents a collective action problem: When any single victim pays a ransom, it encourages cybercriminals to launch more attacks. Therefore, federal law enforcement officially discourages ransom payments and has outlawed payments in some cases.[21] An insurance policy would exclude coverage for ransom payments when making the payment would violate federal policy. State governments could join the federal government in creating an exclusion. In April 2022, North Carolina became the first state in the U.S. to prohibit state agencies and local governments from paying ransoms.[22] Another sticky area is exclusions of "acts of war." Damages from acts of war are excluded from many types of insurance policies, not just cyber. The reason is that an act of war would presumably result in widespread destruction, and an insurance company could not afford to cover large losses occurring to many customers simultaneously. In some cases, cyberattacks are perpetrated by common cybercriminals, but many cybersecurity experts consider state-sponsored cyberattacks to be a significant risk. For example, it is thought that North Korea sponsors ransomware attacks to raise money for the North Korean regime. As of this writing, Russian cyberattacks are thought to be a risk from the war in Ukraine. If a state-sponsored cyberattack is considered to be an "act of war," it might be excluded. That said, it is often difficult to attribute a cyberattack to a particular attacker, much less to determine if the attacker is state sponsored. Nevertheless, a customer should recognize that state-sponsored cyberattacks are a real threat and policy exclusions could complicate receiving coverage from state-sponsored attacks.

Lastly, an exclusion with broader implications than just ransomware is if the insurance policy lists specific types of hardware, data, or other IT assets that are excluded from the policy. For example, the cost to replace "bricked" computers (i.e., computers that are rendered useless by a cyberattack) may be excluded from a policy. Similarly, cyber insurance policies generally don't cover bodily injury and property damage. Examples might include modems and connectivity devices for internet-enabled physical assets or damage to water or sewer control systems from a cyber sabotage attack. Also note that property damage insurance policies may have broad cyber exclusions, thus leaving the insured with no coverage under any policy.

Key questions to ask: What exclusions does the policy contain? What are the exclusions specific to ransomware attacks? What are the exclusions for particular types of IT assets you might own?

Definitions

The customer should be familiar with key provisions within the definitions of the policy. First, we'll reiterate the importance of understanding the requirements to

Insurance is not intended to protect against "average" conditions; it is intended to protect against extreme conditions. Therefore, the cost-benefit analysis must examine the value of insurance at the extremes.

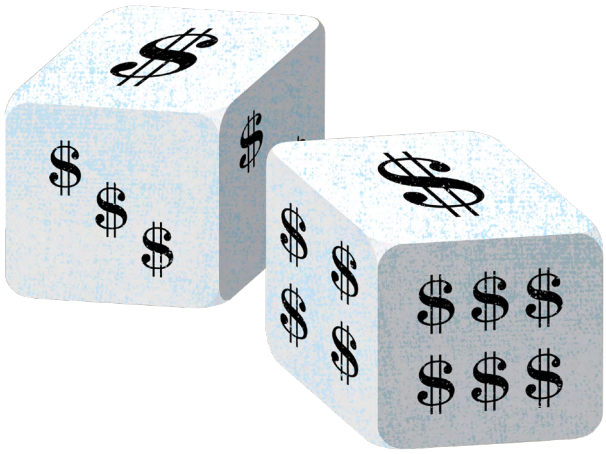

use certain preapproved cybersecurity contractors in the event of a breach (the panel requirements).

Next, be aware of provisions on when the insurance provider must be notified of claims and how that relates to your knowledge of when an insurable event has happened. For example, if malicious software infiltrated your network two months ago but you just find out about it today, then you can only report it today. Know how your policy would cover that situation. If your policy went into effect last week, would you be covered?

Be aware of definitions around internal security control standards that you are required to maintain as a condition of the policy. Earlier, we described how the underwriting process has intensified. Insurance companies are expecting customers to have more robust internal security in place as a prerequisite for the policy. Unsurprisingly, the insurance company will also expect the insured to maintain those standards over the life of the policy. The definitions are important for making sure the customer understands and can meet the requirements. For example, does the policy require that the customer remain current with software updates? Many local governments are not in the habit of applying new patches immediately because of the risk that a patch breaks important operational functions of the software—or because an update might disrupt the integration of software that needs to remain compatible. Hence, it would be important to know what "remain current with updates" means. Gaps in maintaining the standards also arise from renegotiating software contracts, changes in key personnel, and broken equipment.

Key questions to ask: Do you understand important definitions in your policy, such as requirements to use specified cybersecurity contractors when responding to a breach, notice of claims, and security standards?

The bottom line from the limits we just reviewed is that if there is a cyberattack that your policy addresses, there is a nontrivial chance that you might not recover as much from your insurance policy as you might have expected if you didn't understand the limits. The savvy customer understands this risk and weighs it when deciding where and when to invest in insurance versus controls.

We'll cover one other potential pitfall of insurance purchasing that is primarily a function of the customer's psychology and purchasing behavior. That pitfall is buying a policy that is overly focused on a narrowly defined risk. Generally, the more focused a policy is on a specific risk, the less beneficial it is for the insured. This is because the insured is insuring against a lower probability event (the narrowly defined risk) rather than the higher probability event (the broadly defined risk).

Insurance customers can fall into this trap due to what is called "recency bias." This means that recent events cause us to overestimate how likely a similar event is to happen in the future. A good example is flood insurance. Right after a flood happens, more people obtain flood insurance. Years later, many of those people have let the coverage lapse, though the underlying risk of flooding is the same. In the cyber world, a local government might experience a certain kind of cyberattack and then buy a policy to cover similar types of attacks in the future. The government would realize lower premiums by buying a specific policy. However, the likelihood that the local government will experience some kind of cyberattack in the future is greater than the likelihood that the government will experience an attack similar to the one it experienced in the past. Also, in a rapidly changing market, a narrow breadth of coverage

could be dangerous because the nature of the threats is rapidly evolving. This means local governments should make sure that **past** first-hand experiences with cyberattacks or stories from peer governments aren't being overweighted in the design and selection of insurance policies to protect against **future** and **evolving** risks. If the cost of a broader policy is prohibitive, it might be wise to consider if the money is better spent on preventative controls.

Key questions to ask: Are past (and painful) experiences with cyberattacks clouding your judgment in preparing for future risks? Is your cyber insurance policy too narrow and not providing adequate coverage for evolving threats? If a broader policy is cost-prohibitive, might you be better off investing in preventative cybersecurity?

The end of Step 3 is to ask: What prices/offers can you get from different providers? The market is rapidly changing. Working with a good insurance broker is important. Be sure to compare providers on:

- **Claims payment history:** Do customers get the coverage they thought they were buying? As we saw, limitations can cause customers to recover less than they thought they bargained for.
- **Pre-breach offerings:** Can the insurer offer useful advice for strengthening your preventative posture?
- **Flexibility on vendor utilization:** Does the insurer offer a reasonable number of options on contractors to support you in the event of a breach?
- **Experience in the public sector:** Does the insurer understand the risks that characterize the public sector?

Also, as with most if not all forms of insurance, there may be significant benefits available from pooling risk with other local governments, either as part of self-insurance pools or joint purchasing of commercial insurance. For example, the Municipal Excess Liability (MEL) Joint Insurance Fund of New Jersey has provided its members with cyber insurance coverage since 2013.[23]

The quantification of risk we discussed earlier can be extended to include commercial insurance policies. The cost-benefit of the policy could be weighed based on the retention and the limit of the policy. An important nuance, though, is that "on average," an insurance policy will be a financial loser for the insured; otherwise, insurance companies would go out of business. However, insurance is not intended to protect against "average" conditions; it is intended to protect against extreme conditions. Therefore, the cost-benefit analysis must examine the value of insurance at the extremes. The GFOA sample ransomware risk model walks you through how you could quantitatively analyze the value of insurance under extreme conditions.

ADDITIONAL CYBERSECURITY RESOURCES

- CIS 18 Critical Security controls: **cisecurity.org/controls**
- Cyber Resilience and Financial Organizations: A Capacity-building Tool Box: **carnegieendowment.org/specialprojects/fincyber/guides**
- FS-ISAC Cybersecurity Resources: **fsisac.com/resources**
- CIS Critical Security Controls: **cisecurity.org/controls/v8**

STEP 4

Periodically reassess

Because cybersecurity threats are constantly evolving, a government's posture toward those threats must evolve as well. A reassessment is critical after a cybersecurity event but should be done regularly even if no events have occurred to give you a better chance of your good fortune continuing. The Step 4 reassessment can ask many of the same questions we asked in Step 1. The objective is to find out if there are new vulnerabilities perhaps due to:

- Evolving methods of attack used by cybercriminals.
- Changed or new technologies, operations, etc., that increase or change the attack "surface area" presented by the local government to cybercriminals.

The reassessment can also look for opportunities to improve controls as new technologies and methods become available. There may be an opportunity to improve the local government's preventative security posture and reduce reliance on insurance.

Conclusion

Cybercrime is an evolving threat to local government. Savvy risk management requires making smart use of strategies to manage that risk, including reducing risk by implementing cybersecurity controls, absorbing risk with self-insurance, and transferring risk to the insurance market by purchasing a commercial insurance policy. Local governments can accomplish this by:

1. Knowing the basics of your cybersecurity situation.
2. Quantifying your risk.
3. Examining the potential for insurance.
4. Periodically reassessing your situation

Shayne Kavanagh *is the senior manager of research for GFOA's Research and Consulting Center.* ***Rob Roque*** *is the technology services manager in GFOA's Research and Consulting Center.* ***Teri Takai*** *is the executive director of the Center for Digital Government.*

[1] Coker, J. (2020). Local government organizations most frequently targeted by ransomware. *Info Security Magazine*. The article cites a study by Barracuda Networks. https://www.infosecurity-magazine.com/news/local-government-targeted

[2] 2022 SonicWall cyber threat report. *SonicWall*. https://www.sonicwall.com/2022-cyber-threat-report

[3] Information in this paragraph is based on an article written by: Dr. Oren Eytan, CEO of Odi-x. Eytan, O. (June 22, 2021) Municipal cyberattacks: A new threat or persistent risk? Forbes.com. Information is also from personal correspondence between Dr. Eytan and the author of this article. https://www.forbes.com/sites/forbestechcouncil/2021/06/22/municipal-cyberattacks-a-new-threat-or-persistent-risk/?sh=673e79ee3ffb

[4] Definition from the National Institute of Standards and Technology (NIST).

[5] Duncan, I. (2019). Baltimore estimates cost of ransomware attack at $18.2 million as government begins to restore email accounts. *The Baltimore Sun*. https://www.baltimoresun.com/maryland/baltimore-city/bs-md-ci-ransomware-email-20190529-story.html

Deere, S. (2018). Confidential report: Atlanta's cyber attack could cost taxpayers $17 million. *The Atlanta Journal-Constitution*. https://www.ajc.com/news/confidential-report-atlanta-cyber-attack-could-hit-million/GAIjmndAF3EQdVWIMcXS0K

[6] Information shared with GFOA directly by the City of Stuart.

[7] NetDiligence Cyber Claims Study: 2021 Report. (2021). *NetDiligence*.

[8] Rising insurance costs add to US public finance cyber pressures (2021). *Fitch Wire*.

[9] Survey conducted by Center for Digital Government (2021).

[10] Hubbard, D. W. (2020). *The failure of risk management: Why it's broken and how to fix it* (2nd ed.). Wiley.

[11] Budescu, D. V., Broomell, S. & Por, H. (2009). Improving communication of uncertainty in the reports of the intergovernmental panel on climate change. *Psychological Science*, 20(3): 299–308.

[12] Hubbard, D. W. & Seiersen, R. (2016). *How to measure anything in cybersecurity risk*. Wiley.

[13] The methods we are referring to are Monte Carlo analysis and computer simulation. Insurance companies will vary in the specifics of how these methods are applied. https://www.probabilitymanagement.org

[14] The authors would like to acknowledge the assistance of the attendees of the ProbabilityManagement.org March 2022 conference for their assistance with these answers.

[15] This is known as "overplacement bias," which is a subset of the well-documented psychological phenomenon of "overconfidence bias."

[16] Discussion of insurance industry taken from: Hubbard, D. W. (2009). *The failure of risk management: Why it's broken and how to fix it*. John Wiley and Sons.

[17] There is no shortage of research that shows quantitative models regularly outperform human judgment. In the context of government finance, see: Kavanagh, S. & Williams, D. (2017). *Informed decision-making through forecasting*. Government Finance Officers Association. This book also discussed relevant research from other fields.

[18] Note that the average figures we cite were compiled by NetDiligence from observations that provided them with more complete data, so it is slightly higher than the overall average across all of their observations. See: NetDiligence Cyber Claims Study: 2021 Report (2021). NetDiligence.

[19] "The State of Ransomware 2021." A white paper published by Sophos (April 2021).

[20] Definitions of coverages provided by Aon.

[21] "Updated advisory on potential sanctions risks for facilitating ransomware payments." An advisory letter from the U.S. Department of the Treasury (September 21, 2021). The letter was associated with the U.S. Department of the Treasury's Office of Foreign Assets Control's Sanctions Compliance and Evaluation Division.

[22] North Carolina becomes first state to prohibit public entities from paying ransoms (May 2, 2022). *National Law Review*.

[23] MEL Cyber Risk Management Program (2nd ed.) (March 8, 2021). https://njmel.org/wp-content/uploads/2021/03/MEL-Cyber-Risk-Management-Program-v2.pdf

Appendix | Definitions of Key Cybersecurity Controls

Multifactor Authentication (MFA)—A multilayered approach to security where a second step of authentication is required to complete a transaction. An example of MFA is entering a username and password to log into email as a first step. But a second step of receiving a code to your registered cell phone is required and entered to access the email. The user must enter the code; otherwise, the user cannot access email even if the user entered the correct username and password. MFA is used to prove the user is legitimate.

Incident Response Planning—Development of a plan based on a risk portfolio of potential cyber events. Each type of risk is typically assigned a priority, and a mitigation strategy is developed for each incident. Service level agreements may be applied to outsourced technology services as part of the incident response plan.

Patching Cadence—An established frequency for applying patches or fixes to software and other applicable technologies. The cadence should be considered from two perspectives. A vendor may establish a cadence for normal fixes and bugs. In these cases, the customer (the second perspective) takes into consideration the cadence to apply the fixes. For example, the vendor may release a patch each month. The customer may choose to apply the patches each quarter to ease testing efforts. The strategy should be defined and included in a risk mitigation strategy.

Endpoint Detection and Response (EDR)—A process of monitoring endpoints of technologies (e.g., devices, nodes) for suspicious activities and, in most cases, removing the risk automatically. Antivirus software can be considered a simple EDR tool since it is designed to actively monitor a device and remove issues that fit within certain risk categories. Advanced EDRs are constantly learning, analyzing, and communicating to develop mitigation strategies to respond to evolving cyber threats.

Firewall—A combination of devices and software that separate internal networks from external networks (i.e., the internet). Firewalls are configured to guard against intrusions and other unauthorized network traffic via device ports and other hardware, software, or telecommunication means.

Training—Almost all cyber incidents start at a system's weakest link—the user. Users should be trained on how to identify suspicious emails, conduct good password practices, use multifactor authentication, and other general safe computing practices.

Data Backups—This is the practice of backing up enterprise data in a safe means according to a backup methodology. The approach is typically based on a risk mitigation strategy that defines the type of data to be protected, the frequency of the backups, the physical requirements to back up the data, and the disaster recovery response requirements.

Encrypted Storage—Storing data in a format that cannot be read without a key and code interpreter. If this data is stolen, it cannot be read by the criminal and is useless—unless the criminal has access to the key and code interpreter. Encrypted storage is not commonly used since it can slow down computing resources during the encryption and reading process. It can also be expensive to encrypt the data and store it. Some organizations will select the type of data that is encrypted to avoid latency and minimize costs.