AIRA
AMERICAN IMMUNIZATION REGISTRY ASSOCIATION

Cybersecurity Tips

2026 Cost-Effective Strategies & Resources for IIS

August 28, 2026

Contents

Introduction

The cybersecurity tips and resources below highlight cost-effective strategies that can help immunization information systems (IIS) increase cybersecurity knowledge, adopt effective security practices, and better protect IIS and the data that they contain. Many of these recommendations leverage free government services, open-source tools, and existing technologies. Together, they provide practical, cost-effective opportunities to strengthen cybersecurity practices.

IIS staff should review these tips, along with the accompanying *2026 Security Guidance for IIS* document, with IIS IT implementation teams, including jurisdictional IT, information security teams, and/or IIS vendors to determine what may already be in place and what additional measures may be pursued.

Tip 1. Join Threat-Sharing Programs (InfraGard, CISA) – Situational Awareness

Recommendation for IIS: Ensure at least one IIS team member is an InfraGard health care and public health infrastructure liaison member and establish processes for sharing InfraGard knowledge as applicable. This single action provides a continuous, curated stream of threat intelligence without ongoing subscription costs.

Description: The Federal Bureau of Investigation's InfraGard[1] program is a formal public-private partnership established to create an official, trusted liaison relationship between the federal government and private-sector organizations that operate within the nation's critical infrastructure, including health care and public health.

Homeland Security has officially designated **critical infrastructure sectors** including power generation, waterways and water supply, food and agriculture, hospitals, health care IT, and financial systems. Each sector has its own InfraGard program. The health care/health care IT sector is directly applicable to IIS.

Individuals may apply to become an InfraGard member. Once approved, members gain access to a curated stream of threat intelligence. Alerts cover ransomware campaigns, emerging threat actor groups, nation-state activity targeting health care, and sector-specific cybersecurity vulnerabilities. Most alerts can be shared with the broader IIS community, although some are classified or protected by non-disclosure agreements (NDAs) and may be shared only within the InfraGard network.

Example: The IIS manager, an InfraGard member and liaison, receives an alert indicating that a bad actor is actively exploiting Microsoft Exchange servers used by health care organizations in the United States. Because IIS staff use agency Microsoft Exchange email accounts, the IIS manager coordinates with the IIS IT team and the state IT department to

[1] See https://infragard.fbi.gov/.

confirm patches have been applied and reminds staff to be vigilant for phishing attempts targeting IIS staff credentials.

Cost: Although enrollment, most events, and available resources are free, organizations should account for staff time needed to enroll, review cybersecurity alerts and threat intelligence, and share relevant information with appropriate stakeholders.

Tip 2. Participate in Information Sharing and Analysis Centers (ISACs) – Situational Awareness

Recommendation for IIS: IIS programs and their technology partners should consider membership in the Multi-State ISAC and/or Healthcare ISAC. If your organization is already a member, follow up to ensure applicable information and resources are shared with the IIS team in a timely manner.

Description: Information Sharing and Analysis Centers (ISACs) are sector-specific, largely virtual organizations that serve as structured clearinghouses for peer-to-peer cybersecurity threat intelligence. ISACs were formed to help fix critical cybersecurity asymmetry. Malicious actors were effectively collaborating and sharing attack intelligence with each other, while defenders were often isolated and unable to learn from each other's incident experience. These centers support sharing indicators of compromise (IOCs), specific technical evidence of an attack (such as malicious IP addresses, file hashes, phishing email patterns, or malware signatures) in near real time with all verified members.

Relevant ISACs for IIS include the Multi-State ISAC (MS-ISAC) for state and local governments (https://www.cisecurity.org/ms-isac) and the Healthcare ISAC (H-ISAC) for health-care-adjacent organizations, including hospitals, health systems, payers, and health IT organizations (https://health-isac.org/h-isac-membership).

Examples:

- Hospital A in Chicago begins experiencing a novel ransomware attack on Monday morning. Through H-ISAC, Hospital A can immediately post: "We are being attacked by actor X. Here are the file hashes, C2 server addresses, and phishing lure characteristics we've observed." Every other H-ISAC member can then immediately scan their own environments for those indicators and preventively apply mitigations or patches. This transforms a single organization's painful experience into community-wide protection in minutes or hours rather than weeks.
- Through an ISAC notification, the IIS manager learns that another jurisdiction experienced a ransomware attack after a vendor's remote support account was compromised. The IIS manager works with the IIS IT team and the IIS vendor to review all vendor accounts, confirm multifactor authentication (MFA) is enabled, and verify that unnecessary remote access has been removed.

Cost: Annual membership fees vary based on the organization's annual revenue (H-ISAC) or annual operating budget (MS-ISAC). For the H-ISAC, eligible nonprofit organizations may qualify for a discounted membership, with annual costs estimated at approximately $1,500 to $2,400 per year.

Tip 3. Train and Certify an Internal Cybersecurity Staff Member – IIS Internal Skills Development

Recommendation for IIS: Support an IIS team member in obtaining cybersecurity training and/or certification and ensure processes are in place to share the knowledge gained and resources within the team.

Recommended Certificates: Each cybersecurity certificate for IIS staff should be relevant to IIS technologies. Some certification categories are:

1. Vendor-independent certificates for the technology "stack" being used by the IIS (such as the networking hardware, operating system, databases, application software, etc.).
2. Vendor-specific certificates for the IIS technology stack.
3. Durable cyber certificates that have been available for more than 20 years and have been kept current. Examples include:
    a. CompTIA – https://www.comptia.org/en-us/certifications/#cyber
    b. ISC2 – https://www.isc2.org/certifications
    c. SANS Institute – https://www.sans.org/cyber-security-courses
4. Several cybersecurity certificate programs have a career track from beginner through advanced. This certificate progression could become part of the IIS staff member's multiyear career plan. As noted above, examples include CompTIA, ISC2, and SANS Institute.
5. Secure Controls Framework offers certifications for both organizations and individuals. For more information, visit https://securecontrolsframework.com/scf-certified.

Description: Identify an IIS staff member who has demonstrated interest and aptitude in technology and cybersecurity. Support their progressive professional development through cybersecurity certifications. The key principle is that each certification must map directly to a concrete job objective. After completing a certificate, the staff member's role should include actively applying the new knowledge within the IIS, such as auditing a specific system, implementing new controls, documenting a procedure, etc.

Over time, this training investment compounds. The staff member builds organization-specific expertise, institutional knowledge that an external consultant cannot provide, and the capability to perform ongoing security functions year after year without recurring large consulting fees. As this expertise matures, the impact can be expanded by cross-training

additional IIS staff so that critical cybersecurity knowledge is shared across the team, reducing reliance on a single individual and strengthening operational continuity.

Example: An IIS analyst completes security training and learns the importance of periodic access reviews. During a review of IIS user roles, the analyst discovers that several former employee accounts remain active. The analyst works with the IIS IT team to remove access and works with the IIS manager to refine policies and procedures for user account management.

Cost: In many cases, cybersecurity training and certifications represent a cost-effective investment, enabling internal staff to perform certain security functions that might otherwise require third-party consulting support.

Tip 4. Use the CISA Free Vulnerability Scanning Service – Vulnerability Assessment

Recommendation for IIS: All IIS (hosted by a health agency, self-hosted, or vendor-hosed), should request enrollment in the CISA scanning service. This will act as an important, and free, cross-check to make sure the IIS is not exposing certain vulnerabilities publicly. CISA cyber hygiene scans access only IP addresses, ports, and services that external bad actors also are scanning and attempting to exploit. The application process is simple, and the value is significant.

Many cloud providers now use multiple IP addresses for both inbound and outbound connections. This makes CISA scanning and firewall IP address allowlisting difficult. To overcome this challenge, the IIS can have its cloud configured to use an API proxy gateway so that all inbound and outbound connections go through a small list of static IP addresses.

Description: CISA, the Cybersecurity and Infrastructure Security Agency, operating under the Department of Homeland Security, offers a free, continuous, external vulnerability scanning service. This service is similar to what commercial third-party vendors charge between $5,000 and $15,000 per month to provide.

CISA initially restricted this scanning service to federal agencies. Over time, it was extended to state and local government agencies, then to health care entities, and most recently to most private companies with a static IP address. How it works:

1. As a prerequisite to enrollment, IIS staff should coordinate with the IT teams at their vendor or agency to avoid duplicate reports, avoid false alerts (CISA scans could trigger false alerts indicating the IIS is under attack), and keep the IIS software hosting provider from blocking the CISA scanning service. A best practice is for a member of the IIS staff to directly register for and obtain the CISA reports to avoid delays in reporting and to allow the IIS team to cross-check their CISA reports with those being provided by others.

An IIS manager may find the following sample language helpful when discussing enrollment with the IIS IT team or vendor: The CISA Cyber Hygiene service scans only the publicly accessible services that are already visible to external users and attackers. The service is provided at no cost and helps identify known vulnerabilities affecting internet-facing systems. Enrolling the IIS demonstrates due diligence by providing an independent assessment of our public-facing services without requiring additional security tools or licensing costs. We believe this is a practical, cost-effective step to help strengthen the security of the IIS environment.

2. An authorized person at the IIS completes a straightforward application process with CISA. Enrollment is initiated at the following landing page: https://www.cisa.gov/cyber-hygiene-services.
3. Within about two weeks, CISA begins scanning the IIS's externally facing infrastructure, from the outside, simulating what an external unauthenticated attacker would see. Scans are performed against only public IIS static IP addresses supplied by the IIS. Scans attempt to find vulnerable services (like remote access tools or database servers) exposed to the internet.
4. CISA delivers a detailed vulnerability report identifying specific weaknesses (e.g., "Your firewall model X running firmware version Y has a known critical vulnerability: CVE-XXXX-XXXX"). Reports are delivered via encrypted attachments to an email.
5. CISA continues scanning on an adaptive cadence until each vulnerability is remediated. Currently, systems with significant vulnerabilities are scanned every 12 hours, systems with moderate vulnerabilities are scanned every 2 to 6 days, and systems with no identified vulnerabilities are scanned every 7 days.

CISA's scanning is external; it scans from the internet looking in. This pairs perfectly with an internal scanner (see Tip 11. Use Open-Source Vulnerability Scanners), which can scan from within the network looking for vulnerabilities. Together, these two types of tools provide internal and external vulnerability scanning with no licensing costs.

Example: The IIS manager partners with the IIS IT team to enroll the public-facing web user interface in CISA's vulnerability scanning service. A scan identifies an outdated Transport Layer Security (TLS) configuration on the IIS web server. The IIS IT team and the IIS vendor remediate the issue, improving the security posture of the IIS environment without requiring the purchase of additional security tools.

Cost: Staff time is required to enroll in the service, review findings, and address identified issues.

See also: Tip 11, which discusses additional vulnerability scanning tools.

Tip 5. Implement Network Segmentation – Enforce Security Boundaries

Recommendation for IIS: Request a network segment dedicated to production IIS system deployment. Consider the following sample language in communicating with your IIS vendor or IIS IT and/or security team: "Our IIS PHI system is a special, high-priority, regulated environment. We need to ensure it lives on its own isolated network segment that cannot be reached from our general office network, email systems, or any other enterprise systems, even internally, without extra authentication steps such as a bastion host."

Description: Network segmentation is the practice of dividing an organization's computer network into multiple isolated sub-networks (segments). Each segment is restricted from communicating with other segments without explicit authorization. In a properly segmented network, a compromise of one segment does not automatically grant access to other segments.

Many small organizations operate with a single, "flat" network where all computers, servers, printers, and devices share one network. In this model, if any one device is compromised, an attacker can use that foothold to probe and attack every other device on the same network. Protected health information (PHI)-containing servers, email servers, HR databases, and office desktops are all equally reachable from a compromised machine.

Segmented architecture for IIS: The PHI-containing data and applications should live on their own dedicated network segment with strict rules:

- No direct communication with the general office staff network
- No direct communication with HR, payroll, or email systems
- No direct access from any general-purpose enterprise device
- All access must go through explicitly approved, authenticated pathways (see Tip 6. Use Bastion Hosts)

Attack scenario comparison:

- *Without segmentation:* A staff member receives a malicious email. They open it, or it executes without a click via a zero-day vulnerability (see Tip 7). Their machine is compromised. Malware propagates across the flat network and within hours reaches PHI servers. Patient data is exfiltrated or encrypted with ransomware.
- *With segmentation:* The same attack occurs. The staff member's machine is compromised. Malware spreads across the general office segment but hits a barrier. The PHI segment is invisible and unreachable from the office segment. PHI remains fully intact and uncompromised.

Example: The interoperability analyst works with the IIS IT team to ensure that the HL7 interface engine receiving VXU messages from providers resides in a separate network

segment from the IIS database. If the interface engine is compromised, attackers cannot directly access patient immunization records.

Cost: Staff time is required to configure and implement network configurations. Network segmentation is a built-in capability of virtually all modern firewalls, managed switches, and network infrastructure equipment. It requires configuration, not additional hardware or software licensing.

Tip 6. Use Bastion Hosts – Reduce Attack Surface Area

Recommendation for IIS: Follow up with your IIS IT and/or security teams to inquire about using bastion hosts (or other strategies, see Tip 5. Implement Network Segmentation).

Description: A bastion host (also called a jump host or jump box) is a hardened,[2] purpose-built intermediate computer that serves as the single, controlled access point between a commonly attacked network (like the general office network or the public internet) and a secured, sensitive network (such as systems containing PHI). The name comes from a fortification term; a bastion is the part of a castle or fortress wall most exposed to attack, deliberately reinforced to be the strongest point.

Access to bastion hosts is limited to users with elevated privileges who perform select administrative functions, such as managing the database. The bastion host provides an additional layer of security by serving as a controlled access point between the general office network and the PHI network. **Without** a bastion host, the IIS user with elevated privileges directly logs into the privileged IIS environment. **With** a bastion host:

1. The IIS user with elevated privileges logs into the bastion host using strong authentication (password and MFA, certificate, etc.) from their regular workstation.
2. The bastion host verifies the manager's identity and authorization.
3. From within the bastion host session, the manager can now log in to the PHI environment.
4. The session is logged, auditable, and time limited.

The result is that an attacker seeking to reach PHI must now successfully breach three separate layers: (1) the general office network or internet perimeter, (2) the hardened[3] bastion host, and (3) the PHI network segment itself. Each layer represents a significant independent barrier.

Example: The IIS implements a bastion host to protect its PHI. An attacker gains access to the IIS general office staff and developer staff's computers. However, the attacker cannot

[2] "Hardening" is the process of taking a desktop computer, smart phone, server, etc. and making it more resistant to attack by removing unneeded services, adding extra firewall rules, implementing policy changes, etc.
[3] See footnote 2.

find or access the PHI-containing IIS system database because the attacker cannot access the bastion host, which requires MFA.

Cost: Staff time is required to deploy and configure a bastion host. Minimal hosting expenses may apply, depending on implementation. A bastion host can be implemented as a small virtual machine in a cloud environment for approximately $5/month. Alternatively, an older physical computer (e.g., one that is too slow for everyday use) can be repurposed as a dedicated bastion host with no additional hardware costs.

Tip 7. Implement Protections Against Zero-Day Vulnerabilities and AI-Driven Mass Attacks – Implement Resilience to New Unknown Attacks

Recommendation for IIS: Implementing protections for zero-day and AI-based attacks can be partially mitigated by implementing the strategies described in this document, including Tip 5. Implement Network Segmentation or Tip 6. Use Bastion Hosts and other controls listed in *2026 Security Guidance for IIS.*

Description: A zero-day vulnerability is a security flaw in software or hardware that is known to attackers before it is known to the vendor or the security community. The term "zero-day" refers to the fact that on the day the vendor becomes aware of the vulnerability it has had zero days to prepare a patch or mitigation.

Zero-days are uniquely dangerous**.** With a zero-day, no patch is available to fix the vulnerability. The available responses include shutting down the affected systems, absorbing the attack while trying to minimize damage, or relying on architectural defenses that do not depend on vendor patches (network segmentation, bastion hosts, etc.).

For example, attackers discover a critical flaw in Microsoft's web server before Microsoft's own security team is aware of it. The attackers share this knowledge within their network and begin systematically exploiting web server installations across thousands of organizations. When Microsoft or the industry finally discovers the flaw and announces it on "day zero," there is no patch available yet, only an advisory warning. Organizations are left exposed until a patch can be developed, tested, and distributed.

Historically, zero-days were discovered manually by skilled security researchers, which limited their volume. However, AI is dramatically changing this equation. AI security tools enabling systematic identification of zero-days at scale are now available to both defenders and attackers.

The practical implication for IIS is that reactive, patch-based security strategies will become increasingly insufficient. Proactive architectural defenses, particularly network

segmentation, layered security, zero trust,[4] bastion hosts, pinholes, application allowlisting, and access controls, become the primary line of defense because they often remain highly effective even before a vendor patch for a zero-day exists.

Example: The IIS manager partners with their vendor or IIS IT/security team who will segment the IIS environment and restrict administrative access through a bastion host (see Tip 6). When attackers use AI-powered tools to rapidly identify and exploit a previously unknown (zero-day) vulnerability affecting the IIS provider web-user interface, the network segmentation and bastion host controls help contain the attack, reduce administrative access, and protect sensitive IIS data while the IIS vendor develops and deploys a security patch.

Cost: Implementation requires staff time, and costs vary depending on the selected controls. However, many recommendations can be implemented by configuring existing technologies rather than acquiring new ones.

Tip 8. Deploy Honeypots – Detect Internal Attackers

Recommendation for IIS: Follow up with your IIS IT and/or security teams to inquire about the use of honeypots to detect internal and external threats.

Description: A honeypot is a deliberately designed decoy system, a computer, server, database, or service that is intentionally made to appear valuable and accessible to an attacker. A honeypot contains no real production data and is instrumented to detect and record any unauthorized access or interaction.

Honeypots work as early-warning systems. In a normal environment, an attacker can probe and move laterally for days, weeks, months, or sometimes even years before triggering a security alert, because their activity may blend in with legitimate traffic. A honeypot changes this because the honeypot has no legitimate general users and no legitimate general business purpose, so *any* interaction with it is suspicious. The moment an attacker or a malicious insider "touches" the honeypot, an alert is triggered immediately.

Specific use cases:

- *External intrusion detection:* Place a honeypot server in the general office network segment. If an external attacker breaches the perimeter and begins scanning the

[4] The NIST Glossary defines zero trust architecture as "A security model, a set of system design principles, and a coordinated cybersecurity and system management strategy based on an acknowledgement that threats exist both inside and outside traditional network boundaries. The zero trust security model eliminates implicit trust in any one element, component, node, or service and instead requires continuous verification of the operational picture via real-time information from multiple sources to determine access and other system responses." Please also see the NIST zero trust architecture conceptual design at: https://www.nist.gov/publications/zero-trust-architecture.

internal network, it will likely discover and probe the honeypot, triggering an alert before the attacker reaches PHI-containing systems.

- *Insider threat detection:* Honeypots are one of the few security controls that are effective against malicious insiders. A disgruntled employee with legitimate network access who begins snooping beyond their authorized scope is likely to stumble upon a honeypot. The alert fires immediately, something that traditional perimeter defenses cannot detect since the employee is already "inside."
- *Attacker intelligence gathering:* Honeypots can be configured to observe and log attacker behavior, tools, and techniques, providing valuable intelligence about emerging threats.

A free, open-source, comprehensive honeypot package is provided by Deutsche Telekom, called "T-Pot." It includes a framework for running multiple honeypots, a real-time dashboard, the ability to search and filter alerts, and more. More information is available at: https://github.com/telekom-security/tpotce.

Example: The IIS IT team deploys a decoy server named "IIS-Archive-DB" in a monitored environment. When an attacker attempts to access the system using stolen credentials, alerts are generated and investigated before any production IIS resources are affected. The IIS manager is notified through the incident response process.[5]

Cost: Staff time is required to deploy, monitor, and maintain a honeypot. Many honeypot solutions are available at little or no cost. A honeypot can be implemented as a small virtual machine in a cloud environment for approximately $5 per month. Alternatively, an older physical computer (e.g., one that is too slow for everyday use) can be repurposed as a dedicated honeypot with no additional hardware costs.

Tip 9. Conduct Penetration Testing with Internal Staff – Cost-Effective Self-Assessment

Recommendation for IIS: Train an internal or an agency IT staff member, such as a senior technical analyst, in penetration testing through online penetration testing training programs.[6] This person would have basic IT skills such as networking and basic Python or shell scripting, along with a curiosity for exploring new technologies. Once trained, the staff member conducts regular structured attack simulations against the IIS's own test environment. While penetration testing, or "pen testing," should be conducted by all IIS IT teams as a matter of due care, IIS staff-driven penetration testing can be uniquely valuable. For example, IIS staff-driven penetration testing can include systems not tested by the IIS

[5] One of the authors of this paper, Eric Heflin, wishes to disclose that he has a commercial interest in honeypot technologies.

[6] A free, reputable resource designed to teach new penetration testers can be found here: https://docs.rapid7.com/metasploit/metasploitable-2/.

vendor, such as operating system remote access or adjacent virtual machines in the same hosting environment or other connected systems like the single sign-on system itself.

Description: Penetration testing is the practice of intentionally attacking your own systems, mimicking the techniques of real attackers to discover vulnerabilities before malicious actors do. External pen testing firms and Red Teams[7] can charge tens of thousands of dollars for a single engagement. A cost-effective alternative is to enroll an internal staff member who has expressed interest in and aptitude for cybersecurity in online pen testing training programs. Once trained, that person's position description can be updated to formally focus a fraction of their time on pen testing. Then that staff member can research IIS potential vulnerabilities and conduct structured, pre-approved, attack simulations against the IIS's own test environment.

IIS penetration testing should be done in close coordination and with written approval (scope, allowable tests, dates allowed, etc.) of the system owner of any system being tested.

What internal pen testing looks like in practice:

- Simulating a ransomware actor attempting to access the PHI network from the general office segment
- Attempting brute-force password attacks, using credentials obtained from dark web sites, password stuffing attacks,[8] or creating phishing or spear phishing credential-stealing attacks
- Testing whether the bastion host can be bypassed
- Attempting to exploit known vulnerabilities that the vulnerability scanner reported as present
- Escalation of privileges from a role with fewer access rights to a role with more access rights
- Testing resistance to social engineering by sending simulated phishing emails to staff or contacting the IIS patient support team and trying to bypass authentication

These exercises reveal real, organization-specific vulnerabilities in a controlled, safe environment before a real attacker discovers them. The internal pen tester also builds substantial familiarity with the IIS's specific architecture, which external consultants spending a few days on-site cannot replicate.

---

[7] NIST defines a Red Team as "A group of people authorized and organized to emulate a potential adversary's attack or exploitation capabilities against an enterprise's security posture. The Red Team's objective is to improve enterprise cybersecurity by demonstrating the impacts of successful attacks and by demonstrating what works for the defenders (i.e., the Blue Team) in an operational environment. Also known as Cyber Red Team."

[8] "The use of distinct passwords is important to avoid 'password stuffing' attacks in which an attacker uses a compromised password from one site to access the user's account on other sites." NIST https://pages.nist.gov/800-63-4/sp800-63b.html#passwordusability.

Examples:

- Periodically, the help desk specialist II, with full oversight and written pre-approval, attempts to access systems using the credentials of a recently off-boarded employee to confirm that person's access has been removed. Any anomalies are escalated to security staff.
- Each time a new release of the IIS software is installed, the help desk specialist II researches any recent vulnerabilities (Common Vulnerabilities and Exposures (CVEs)) published against the IIS software and attempts to exploit each published vulnerability to confirm the vulnerability has been fixed.

Cost: Costs vary widely. Online training can cost anywhere from approximately $50 per month for self-paced generic training programs to several thousand dollars for recognized industry certifications.

Tip 10. Prioritize Protected Health Information (PHI) – Protect Critical Assets First

Recommendation for IIS: Implement a PHI-first strategy.

1. Create a dedicated, isolated network segment exclusively for PHI-containing databases and systems (as described in Tip 5. Implement Network Segmentation).
2. Apply as many security controls as possible, up to the full set of available SCF and NIST controls, to that one smaller, well-defined environment first.
3. Because the PHI segment has no email, no office productivity applications, no general-purpose browsing, and no connections to general enterprise systems, it can be smaller in scope to secure than the broader enterprise (a reduced attack surface).
4. Once PHI is well secured, expand the security program progressively outward to the general office environment, HR systems, email infrastructure, and so on over time.

Description: Prioritizing PHI is perhaps the most strategically important tip for IIS with limited resources. A comprehensive cybersecurity controls framework may contain many hundreds of individual controls spanning the entire organization. Attempting to implement all of them simultaneously across every system, network, and device is overwhelming, expensive, and likely to result in partial implementation everywhere rather than strong implementation anywhere.

Reducing the scope of initial cybersecurity controls deployment to just the PHI-containing systems can deliver the highest return on investment (ROI); PHI is simultaneously the most sensitive data an IIS holds and the primary target of health-care-sector attackers (ransomware operators know PHI extortion often creates maximum pressure to pay the ransom). Securing PHI first means the most catastrophic breach scenario, PHI exfiltration or encryption, is addressed before expanding to lower-priority systems.

Example: The IIS manager identifies the IIS database, disaster recovery backups, web user interface, IIS administrator accounts, and IZ Gateway connection as the highest-priority assets. Security improvements focus on these systems before addressing lower-risk applications.

Cost: Staff time is required to assess and validate security controls, document evidence, and implement priority improvement. As a planning estimate, an experienced professional with Secure Controls Framework knowledge may be able to **validate** approximately one control per hour.

Tip 11. Use Open-Source Vulnerability Scanners and Tools – IIS Vulnerability Management Program Inputs

Recommendation for IIS: Work with your IIS IT and/or security teams to ensure routine internal and external vulnerability scanning is in place, leveraging open-source tools as applicable:

- OpenVAS scans from *inside* and *outside* the network, identifying internal vulnerabilities, misconfigured services, and unpatched software.
- CISA's free scanning service scans from *outside,* identifying what an external attacker can see and exploit (see Tip 4. Use the CISA Free Vulnerability Scanning Service).
- Together, these two free tools provide the same coverage that organizations typically pay $5,000–$15,000/month to obtain.

Description: A vulnerability scanner is a software tool that systematically analyzes computers, servers, services, network devices, and applications for known security vulnerabilities. The scanner compares the discovered software versions, configurations, and behaviors against a continuously updated database of known vulnerabilities (CVEs). In some cases, the vulnerability scanner attempts to actually exploit known vulnerabilities to confirm a vulnerability is exploitable rather than simply theoretically present. For example, the free open-source tool nmap[9] and the OpenVAS scanner mentioned in this tip can be configured to perform exploits to confirm their presence and report on such.

Important: All software, including the below recommendations, entail risk of their own, such as supply chain contamination. Thus, it is advisable to obtain these tools from known trusted sources and to execute them inside sandbox environments (such as virtual machines or containers) with restricted access. Vulnerability scanners that perform attempted exploits should run with the same protections as penetration testing exercises discussed in Tip 9, including written scope and pre-approval, oversight, and use of test environment targets.

IIS use of a vulnerability scanning tool or technique, like Kali Linux, Parrot OS, nmap running exploit scripts, or OpenVAS with exploit tests enabled, essentially spans scanning

[9] For an example of using the free open-source nmap tool to exploit vulnerabilities, please see https://nmap.org/nsedoc/categories/exploit.html.

and pen testing. As with any pen testing tool or technique, testing should be done in close coordination with the IIS IT team and all impacted system owners and should have formal written approval (scope, allowable tests, etc.).

Example of scanner operation: The scanner identifies that a web server is running the React web development framework version X.Y. It checks its vulnerability database and finds that version X.Y has a known remote code execution vulnerability (CVE-XXXX-XXXX). It then attempts to exploit that vulnerability in a controlled, safe manner. If the exploit succeeds, the scanner confirms that the vulnerability is real and active, not just theoretical, and reports it with a severity rating and remediation guidance.

Example: Before a production release, the designated IIS staff member scans the IIS application environment and identifies unsupported web application user interface software libraries that contain known vulnerabilities, such as React. The findings are provided to the IIS vendor for remediation before deployment.

See also Tip 4, which discusses a free external vulnerability scanning service provided by the federal government.

Cost: Staff time is required to investigate and operate open-source vulnerability scanner(s) to identify internal and external vulnerabilities and follow up on issues identified. Commercial vulnerability scanners from vendors cost approximately **$1,500/month** or more for licensing. For small IIS, this is a significant ongoing expense. Several free resources exist, including:

- **Open Vulnerability Assessment System (OpenVAS)** is a multi-licensed open-source vulnerability scanner maintained by Greenbone Security. The open-source licensed product is available without registration at https://github.com/greenbone/openvas-scanner with community build documentation found here: https://greenbone.github.io/docs/latest/background.html#csrk-service. OpenVAS is fully functional and actively maintained and provides a comprehensive scanning capability comparable to many commercial tools. OpenVAS also has commercial offerings that provide more capabilities for annual fees starting at $2,500 per year. Additionally, security-oriented Linux distributions, such as those discussed below, often include OpenVAS as a supported component of their official distribution repositories. Virtual machine, containerized, and native installation options are also available.
- **Nikto Web Application Vulnerability Scanner (Nikto)** is a popular and mature web application vulnerability scanner with over 25 years of development. Nikto is available with many mainstream Linux distributions. Source code is located here: https://github.com/sullo/nikto.
- **Zed Attack Proxy (ZAP)** acts as a proxy (an intermediary) to assess websites for misconfigurations and vulnerabilities. ZAP includes a plug-in marketplace and can

be automated. For more information, please visit the ZAP website at https://www.zaproxy.org/. For the source code, visit https://github.com/zaproxy/.

- **Kali Linux** and similar products such as Parrot OS (https://parrotsec.org/) offer a complete stand-alone Linux operating design specifically for penetration testing (also known as ethical hacking). Kali Linux features a large number of open-source vulnerability scanners.

  These cybersecurity-focused Linux distributions can boot from a USB drive and be installed on an unused IIS laptop or desktop giving IIS internal security teams a formidable set of supported cybersecurity testing software tools that are not normally available on standard Linux, Mac, or Windows computers without significant effort. Kali Linux, for example, includes over 600 cybersecurity testing tools to check web applications, Wi-Fi connectivity, firewall configuration, signal intelligence (information leakage), social engineering, Windows directory security, and many more similar software packages.
  The use of Kali or Parrot by an IIS staff member that is fractionally focused on pen testing, can be a formidable, very cost-effective, and effective means to increase the cybersecurity of an IIS.
  For more information, visit https://www.kali.org/. To download or learn more about the source code, visit https://www.kali.org/get-kali/#kali-platforms or https://gitlab.com/kalilinux.
- **OWASP(r) Foundation** maintains a curated directory of source code analysis tools. These can be employed by jurisdiction and vendors developing IIS software.
  This class of tools reviews software source code to look for known vulnerabilities introduced during product development. The directory can be found at https://owasp.org/www-community/Source_Code_Analysis_Tools.
  As mentioned above in Tip 4, CISA offers a free cybersecurity vulnerability scanning service. The service is administered by CISA and is not open-source.

  Other notable open-source security scanners include:
  - nmap: https://nmap.org/
  - Metasploit: https://www.metasploit.com/
  - sqlmap: https://sqlmap.org/
  - Brutespray: https://github.com/x90skysn3k/brutespray
  - ClamAV: https://www.clamav.net/
  - Logging Made Easy: https://www.cisa.gov/resources-tools/services/logging-made-easy

In summary, the open-source and commercial cybersecurity community has released countless free or dual-licensed (free/commercial) products to assist defenders. The resource-constrained IIS will likely find multiple no-cost cybersecurity products that can be used to help give them an edge over attackers.

Tip 12. Use Firewall Pinholes (IP Address Allowlists) – Reduce Attack Surface Area

Recommendation for IIS: Work with your IIS IT and/or security teams to confirm or configure firewall pinholes to restrict data exchange to authorized partners and ensure policies and technical procedures are in place to request, approve, implement, and review allowlisted IP addresses. Review allowlisted IP addresses whenever a new data exchange partner is onboarded, an existing partner changes network infrastructure or service providers, or a data exchange relationship ends. Blocking an IP address for a no longer authorized data exchange partner helps prevent an attacker from attempting to "spoof" the IP address of the former partner and thus obtaining some access or information. This risk is small. Most health care transactions use a low-level protocol called "TCP" which cannot be spoofed. Attacks are limited to an uncommon protocol called "UDP" or even lower-level protocols which provide less risk. Conduct periodic reviews to remove obsolete entries and verify that only current, authorized partners retain access. A simple, automated process can be used to confirm the partner technical connection can be reached and to create an alert if this "reachability" test fails.

Description: A firewall pinhole (also called IP whitelisting or allowlisting) is a firewall configuration that restricts inbound network connections to only those originating from a specific, pre-approved list of IP addresses. Any connection attempt from an IP address not on the approved list is silently dropped. A rejection response is not sent, making that part of the IIS infrastructure "appear" completely nonexistent to unauthorized parties.

This is a key security property. The IIS infrastructure does not just block unauthorized connections; it becomes effectively *invisible* to them. From an attacker's perspective, there is no server to attack, no port to probe, no surface to exploit. Even attackers using VPNs to masquerade as US-based traffic cannot easily bypass this control, because the firewall does not block based upon geolocation. It only accepts packets from specific authorized IP addresses.

IIS typically exchange immunization data with hospitals via HL7 v2 interfaces. The hospital pushes patient vaccination records to the IIS, or the IIS responds to queries. These interfaces need to be accessible to the hospital but to no one else. A firewall pinhole configured to accept HL7 v2 traffic only from the hospital's known IP address achieves exactly this; the interface is fully functional for the hospital and almost completely "invisible" to everyone else.

Some IIS programs already implement IP address allowlisting. Check with your IT team to see if your IIS has this protection in place.

The concept of using IP address-based filtering generally applies only to exchange between enterprises, agencies, and clinical sites. Access to the IIS web interface is normally unrestricted or is restricted to IP addresses in the United States.

Examples:

- A ransomware operator based in a hostile country routes its traffic through a US-based VPN and attempts to attack an IIS in New York. Without IP allowlisting, the IIS's servers respond to connection probes and are discoverable. The attacker can begin mapping the infrastructure and looking for exploitable vulnerabilities. With IP allowlisting, the IIS's infrastructure simply does not respond. To that attacker, the IIS "does not exist" on the internet.
- The interoperability analyst maintains a list of approved data-exchange partners and works with IIS IT team to allow only HL7 v2 message traffic from authorized provider sites, hospital systems, pharmacies, HIEs, etc. During onboarding, approved trading partner IP addresses are added to a firewall allowlist (whitelist), while connections from unknown or unauthorized IP addresses are automatically blocked.
- When AIRA tests pre-production IIS systems via the Aggregate Analysis Reporting Tool (AART), multiple jurisdictions use firewall pinholes and require AIRA to provide its specific IP address before any testing traffic is accepted.

Cost: Staff time is required to configure firewall pinholes and to manage allowlisted IP addresses.

Tip 13. Perform an IIS Cybersecurity Self-Assessment / Internal Security Audit – Cost-Effective Security Program Implementation

Recommendation for IIS: Work with your IIS IT and/or security teams to conduct regular, structured, internal security assessments of the IIS, using the *2026 Security Guidance for IIS* and the companion *Security Controls for IIS* spreadsheet to identify applicable controls, assess current implementation status and maturity, and prioritize improvement activities. As your IIS increases its cybersecurity maturity, expand the list of implemented controls using the additional "weight 8 and 9" controls from the Secure Controls Framework Controls Catalog spreadsheet.

Description: A comprehensive external cybersecurity audit conducted by a specialized third-party firm is valuable but expensive and may be cost-prohibitive for resource-constrained IIS. For most organizations that are seeking to improve their cybersecurity posture, a cost-effective method is to start with an internal self-assessment rather than undertaking a full external third-party cybersecurity audit. In most cases, a self-assessment provides an excellent foundation for any future external assessments.

Internal assessments are effective because:

- The internal assessor knows the organization's specific systems, vendors, workflows, and quirks – context that external auditors spend significant time (and billing hours) discovering.

- The process of conducting the self-assessment itself builds evidence that the IIS is focused on cybersecurity and is exercising due care and due diligence.
- This process increases IIS staff cybersecurity awareness.
- Findings can potentially be acted upon immediately without waiting for an external report and remediation planning cycle.
- The self-assessment can be repeated regularly (quarterly or annually) at minimal incremental cost.

The Security Controls Framework includes assessment guides. Using AIRA's *2026 Security Guidance for IIS* document together with the SCF is intended to serve as a structured checklist for this self-assessment process, enabling IIS to systematically evaluate their posture across all relevant control domains.

Example: The IIS manager, business analyst, and IIS IT team member complete an SCF-based self-assessment starting with AIRA-recommended IIS controls documented in the *2026 Security Guidance for IIS* and *2026 Security Controls for IIS*. During the assessment, they discover that service accounts used for data exchange have not been reviewed in several years. A new quarterly review process is established, the new process is tested, and evidence is captured to show that the new process is effective.

Cost: Staff time is required to determine applicable controls, assess control implementation and maturity status, and identify areas for improvement.

Tip 14. AI Defensive Tools (Llama Guard) – AI Risk Reduction

Recommendation for IIS: If your IIS uses AI inference within its IIS software system for the customer-facing website, for internal operations, or for any other purpose, then consider the use of AI firewalls, such as Llama Guard, with your IIS IT and/or security teams.

Description: Developed by Meta AI, Llama Guard is a free AI safety model that is, as of this writing, considered the open-weight[10] benchmark for AI content safety. It is specifically designed to act as a firewall between end users and an organization's AI system. It inspects all inputs (user prompts going into the AI) and all outputs (AI responses coming back) against a set of safety policies, blocking any content that violates those policies before it reaches or leaves the core AI model.

Additional free, open-weight, AI firewalls to consider include WildGuard, ShieldGemma, Granite Guardian, and NeMo Guardrails. Note that a solution like Llama Guard also entails risk itself; an attacker could potentially overcome Llama Guard's defenses and compromise it. Micro segmentation, using least privilege, monitoring, and other similar controls

[10] An "open-weight" AI is similar to the term "open-source" but is a specific reference to the availability of the AI inference engine under non-proprietary license. Other aspects of an open-weight AI model, such as training data, are often not available.

discussed elsewhere in this document and the companion *2026 Security Guidance for IIS* can be useful ways to mitigate the risks of using AI and AI firewalls.

As AI adoption grows across industries, IIS are beginning to explore AI-powered tools such as customer-facing chatbots for member inquiries, internal assistants for staff productivity, and automated data analysis. When an organization deploys an AI model (such as a customer service chatbot), the most common attack vector is prompt injection, which is when a malicious user crafts input designed to override the AI's instructions and make it behave outside its intended boundaries. Examples include getting a chatbot to reveal confidential system instructions, generate harmful content, bypass safety filters, or perform actions it was explicitly instructed not to perform. This is the AI equivalent of a SQL injection attack.[11]

The following steps demonstrate using Llama Guard with an IIS chatbot:

1. A consumer (e.g., a parent checking their child's vaccination records) interacts with the IIS's customer service chatbot.
2. Before the input reaches the IIS chatbot AI, it passes through Llama Guard.
3. Llama Guard evaluates whether the input is within acceptable boundaries. If it is, the input passes through normally. If it is a harmful or malicious prompt, Llama Guard blocks it and returns an appropriate refusal.
4. Similarly, all chatbot outputs pass through Llama Guard before being returned to the user, blocking any accidentally harmful AI responses.

Example: The outreach and training coordinator wants to pilot an AI-powered chatbot to answer provider onboarding questions. Before deployment, the IIS manager works with the IIS IT team to implement an AI firewall that prevents users from retrieving patient data or system configuration details through prompt manipulation.

Cost: The primary cost is staff time to deploy, configure, and maintain the model. Most uses do not require licensing fees. Computing resource requirements are generally modest, depending on deployment and traffic volume. For IIS already running AI infrastructure, adding Llama Guard may require little additional infrastructure investment.

Tip 15. Add Behavioral Analysis Rules to Existing Operational Monitoring – Augment Monitoring

Recommendation for IIS: Implement business rules to monitor query/data access activity, to ensure (1) access is by authorized organizations and (2) access is in line with expected behavior, based on the organization's provider profile. For example, query volume seems reasonable given provider type and historical query trends.

[11] NIST defines this as "Attacks that look for web sites that pass insufficiently-processed user input to database back-ends." The impact of a SQL Injection attack includes data corruption, database unavailability, data exfiltration, and even system compromise.

Description: Operational monitoring for IIS may underutilize the existing logging and monitoring infrastructure. IIS programs should work with their IIS IT teams to add additional rules to make sure high-volume operations pass plausibility checks. These new rules normally do not require development effort or new product procurement; they normally entail adding new configuration options for additional rules and alerts.

Example: One contemporary example is data acquisition from IIS via HL7 QBP/RSP messaging. This message exchange pattern introduces a tradeoff. The ability of an IIS to respond to a large number of serial or parallel queries can help public health goals, but it also makes "batch" QBP/RSP messages potentially exploitable by bad actors. To compensate for this, the IIS could apply new rules based on their existing logging and monitoring infrastructure to ensure QBP/RSP activity falls within defined limits. For example, queries from approved organizations may be expected to occur during certain time windows and remain within expected query volumes based on an organization's size, type, and historical messaging patterns. Unusual or anomalous activity should generate alerts for investigation and follow-up as needed.

Cost: Staff time is required to create and implement new configuration rules and monitor corresponding data and alerts.

Tip 16. Evaluate and Select a Prescriptive Risk Management Framework – Not All RMFs Are the Same

Recommendation for IIS: IIS programs should adopt a cybersecurity framework that

1. Is actively maintained and updated on a predictable schedule
2. Is mapped to recognized frameworks and standards, such as the Secure Controls Framework (SCF), NIST Cybersecurity Framework (CSF), and applicable NIST Special Publications (e.g., SP 800-53, SP 800-171, or SP 800-66)
3. Is available under an open or otherwise permissive license that allows use without significant licensing restrictions
4. Provides detailed, actionable security controls and implementation guidance
5. Supports a risk-based approach to cybersecurity governance and decision making
6. Addresses health care or health care IT considerations, where applicable
7. Can be tailored to the unique operational, technical, and regulatory needs of an IIS

Description: During the development of the *2026 Cybersecurity Guidance for IIS*, AIRA and the broader IIS community established a set of evaluation criteria and applied those criteria to a wide selection of candidate cybersecurity frameworks. (See *2026 Cybersecurity Guidance for IIS*, Appendix B: Control Selection Methodology, for additional information.) IIS programs should use similar evaluation criteria when considering a framework for internal self-assessments or external third-party evaluation programs.

Cybersecurity frameworks differ in their purpose, scope, and capabilities. For example, AICPA SOC2 is not a risk management framework. SOC's purpose is to provide reporting

and attestation on controls. In contrast, NIST RMF, SCF, and HITRUST provide true risk management frameworks. SCF and HITRUST go further by defining prescriptive control objectives, governance expectations, and assessment criteria that support evaluating control effectiveness while requiring controls to be managed through activities such as documentation, periodic review, measurement, and ongoing maintenance. Using a prescriptive risk management framework, such as SCF, provides a structured and repeatable approach to assessing and reducing cybersecurity risks within an IIS.

Example: An IIS program implemented SCF, including the detailed SCF control objectives, conformance validation cadence, evidence list, and AIRA control documentation template. The IIS was formally reviewed by its state agency internal security audit department, which found that the IIS documented and managed controls and the captured evidence met or exceeded state requirements.

Cost: Staff time is required to create and implement the SCF program as discussed in the *2026 Security Guidance for IIS*.

For More Information

CISA and the Global Cyber Alliance have compiled and maintain a list of no-cost cybersecurity tools:

- https://www.cisa.gov/resources-tools/services
- https://gcatoolkit.org/smallbusiness/

Critical infrastructure recommended fields for asset management inventory fields:

- https://www.cisa.gov/resources-tools/resources/foundations-ot-cybersecurity-asset-inventory-guidance-owners-and-operators#AppA