WHITE PAPER

Protection Level Agreements for Managed Patching

Why Mid-Market Firms Need Humans and AI Working 24/7, and How BTI Delivers Enterprise Outcomes Without Enterprise PLA Pricing

BTI Communications Group

Managed Patching | Vulnerability Management | vCISO

October 2026

For client and prospect review

Los Angeles 323-510-8003 | Chicago 312-492-2613 | Phoenix 480-272-8728 | www.btigroup.com

Executive summary

Most mid-market organizations still treat patching as a monthly chore. Patch Tuesday arrives, a busy IT team works through what it can, and the rest slides to next month. Nobody writes down how long a critical system may stay exposed, so nobody is accountable when it stays exposed for a quarter.

Gartner has given executives a better way to think about this. Paul Proctor, Distinguished VP Analyst at Gartner, argues that leaders should stop counting attacks and start measuring the number of days it takes to patch critical systems, because the entire value of patching is to limit the number of days a vulnerability is available for hacking.[1] He calls the resulting business commitment a Protection Level Agreement, or PLA: an agreed, funded, measurable level of protection that leadership can defend to customers, regulators, insurers, and the board.[1,2]

Proctor also puts a price on the choice. In his public commentary, 30-day patching of critical systems is illustrated at roughly $1 million per year all-in, 15-day patching at roughly $2 million, and 7-day patching at roughly $5 million.[2,3] Those are illustrative enterprise figures, not quotes, but they make the point: speed costs money, and the board should choose the number on purpose.

This paper makes four arguments:

1. **A real program beats ad hoc Patch Tuesday.** Managed patching, vulnerability management, and virtual CISO (vCISO) leadership should run as one operating model with a written patch interval.
2. **Patch intervals belong in the agreement.** A PLA puts days-to-patch in writing, separates control failure from accepted business risk, and supports a defensible standard of due care.[1]
3. **The work requires humans and AI, 24/7.** Exploitation does not wait for business hours. Skilled technicians supported by governed AI agents, working from a SOC 2 Certified NOC/SOC, can watch, prioritize, patch, and verify around the clock.
4. **Mid-market firms can get PLA-style outcomes without seven-figure budgets.** BTI Communications Group delivers measurable patch intervals, verified closure, and an executive risk picture to multi-site organizations in Los Angeles, Chicago, and Phoenix at a price well below the illustrative $1 million to $5 million enterprise ladder, while helping recover money lost to tool sprawl, overtime, SaaS waste, and downtime.

The question every board should answer

Proctor has a blunt way of opening the conversation with executives. In his talks he suggests that security leaders ask the CEO: "How many days would you like your systems available for hacking?"[3]

The first answer is always "zero." Then the room learns what zero costs. Patches have to be researched, tested, scheduled, deployed, and confirmed, and some vendors have not shipped a fix yet. Zero days is not realistic at any budget, so the real question becomes: how many days will we accept, and what will we fund to get there?

That is the heart of a Protection Level Agreement. Gartner describes PLA decisions as treating cybersecurity "in a manner that reconciles measurable levels of protection with the needs of the business," which "creates a defensibility of their implementation."[1]

PLA versus SLA: not the same promise

Most managed service contracts are built on Service Level Agreements. SLAs are useful. They tell you how fast someone will answer the phone, acknowledge a ticket, or restore a failed service. But an SLA measures the provider's responsiveness. It does not tell you how protected you are.

A PLA measures protection itself. Jason Hart, writing for Forbes Technology Council, describes the shift from SLAs to PLAs this way: like an SLA, a PLA defines a specified level of service delivery, but it moves away from the purely technical focus of traditional SLAs and ties key performance indicators to protection level objectives, including timelines for resolving critical vulnerabilities and named accountable owners on both sides.[4]

	Traditional SLA	Protection Level Agreement (PLA)
What it measures	Provider activity (response time, ticket handling, uptime)	Delivered protection (for example, days to patch critical systems)
Typical metric	"Critical tickets acknowledged within 1 hour"	"Critical exploited vulnerabilities remediated and verified within 7 days"
Who sets the target	Usually the vendor's standard terms	Business leadership, as a funded risk decision
What a breach of the agreement means	Service credit or escalation	Control failure, distinct from accepted business risk
Board value	Limited; describes effort	High; describes how protected you are and what it costs

Proctor draws the accountability line clearly. If an organization's PLA is 30-day patching of critical systems and a system is hacked through an unpatched vulnerability after 35 days, "that is a control failure: security and IT have failed to deliver on the agreement." If the same vulnerability is exploited on day 25, that is the result of a business risk decision, which Proctor calls "a concrete, measurable, enforceable assertion of risk-appetite."[1]

That distinction matters enormously after an incident. One version of the story says, "We had no standard and we missed it." The other says, "We set a defensible standard, funded it, met it, and can prove it."

Why patching intervals belong in the agreement

Outcome-driven metrics

Gartner calls these measures outcome-driven metrics, or ODMs. On the FBI's Ahead of the Threat podcast, Proctor explains that an ODM measures "a delivered protection level," meaning that "when the metric gets better, the organization is measurably better protected."[2]

His favorite example is the one this paper is about. "A very common example I use in vulnerability management," he says, "is number of days that it takes to patch critical systems." He admits it sounds boring, "except that it's what rules everything," because all of your investment in vulnerability management is meant to accomplish one goal: "reduce the amount of time that a vulnerability is available for exploitation."[2]

Contrast that with the metrics many organizations still report: number of attacks blocked, number of open tickets, number of patches deployed this month. Gartner's guidance is direct. Instead of reporting on the number of attacks an organization receives, executives should report on the number of days to patch critical systems.[1]

Due care and defensibility

Proctor also argues that "there is no such thing as perfect protection," and that the real question is whether leaders "can defend the choices they have made on cybersecurity to their key stakeholders."[1] Gartner says the emergence of protection level benchmarks "is a critical step in the development of a cybersecurity standard of due care."[1]

For a mid-market company, those stakeholders are concrete: the insurance carrier asking about patch cadence at renewal, the enterprise customer's security questionnaire, the auditor sampling change records, and the owners who will ask what happened if something goes wrong. A written, measured, verified patch interval answers all of them.

The illustrative enterprise cost ladder

PLAs work because they attach a price to a protection level. Proctor's public examples give a simple ladder. On the FBI podcast he describes spending "$1 million, all-in, to achieve 30-day patching," then asks, "do we want to spend $2 million and do 15-day protection?"[2] In conference talks, as reported by attendees, he extends the ladder to roughly $5 million per year for 7-day patching, an option he notes most executives will consider very expensive for a single problem.[3]

Table 1. Illustrative enterprise PLA cost bands for critical-system patching (Gartner / Paul Proctor public commentary)

Agreed patch interval for critical systems	Illustrative annual cost (enterprise, all-in)	Source
30 days	About $1 million	FBI Ahead of the Threat, S1 E7[2]
15 days	About $2 million	FBI Ahead of the Threat, S1 E7[2]
7 days	About $5 million	Proctor conference commentary, as reported[3]

These figures are illustrative examples used by Gartner's Paul Proctor to explain how protection levels trade against cost in large enterprises. They are not BTI prices, not Gartner benchmarks for any specific organization, and not a quote for your environment.

Two lessons come out of the ladder. First, faster patching is not free; each step takes more people, coordination, and after-hours work. Second, in a PLA executives do not dictate how the protection level is achieved, only that it is achieved.[3] That opens the door to a better delivery model.

Proctor also points out that different parts of the business can carry different intervals. In manufacturing, for example, operational technology is often not patched in 30 days while internal IT systems are patched faster, so investment can be aligned to each part of the business.[2] A good PLA reflects that reality instead of pretending one number fits everything.

The mid-market gap

The enterprise organizations in Proctor's examples have security teams, change advisory boards, and seven-figure budgets for a single control. Most mid-market firms in Southern California, Chicagoland, and Arizona have none of those things. They have a capable but stretched IT lead, perhaps one or two technicians, a stack of vendor portals, and a business that cannot afford downtime.

Lean IT, many sites

Multi-site organizations feel this most. Picture a distributor with warehouses in the Inland Empire and a branch in Phoenix, or a professional services firm with offices in the Loop and the suburbs: endpoints, servers, firewalls, VPN portals, and line-of-business applications spread across locations, patched when someone has time. The internal team is excellent at keeping the business running. It is rarely staffed to watch for newly exploited vulnerabilities at 2 a.m. on a Saturday.

Exploit timelines have outrun the monthly cycle

The threat environment no longer runs on a monthly calendar. The U.S. Cybersecurity and Infrastructure Security Agency (CISA) maintains a Known Exploited Vulnerabilities (KEV) catalog, a public list of vulnerabilities confirmed to be exploited in the wild.[5] Under Binding Operational Directive 22-01, federal civilian agencies are generally required to remediate newer KEV entries within two weeks, with shorter

timelines possible when the risk is grave.[6] That is a useful signal for private-sector leaders: the federal government has decided that, for vulnerabilities attackers are actively using, a monthly cycle is too slow.

The exposures that matter most tend to sit on the edge of the network and in everyday tools: firewalls and VPN portals, email gateways, remote access tools, browsers, and Windows endpoints. When a flaw in one of those is being exploited, the difference between a 7-day and a 90-day response is the difference between a closed door and an open one for nearly three months.

Tool sprawl and fragmented accountability

Many mid-market firms fill the gap by buying tools: one for patching, another for scanning, a separate firewall vendor, a compliance spreadsheet, and a vCISO report once a year. Each tool may be good. Together they create a different problem.

- **No single owner of the outcome.** The scanner finds the issue, the patch tool may or may not fix it, and nobody is accountable for proving it is gone.
- **Duplicate spend.** Overlapping licenses, unused SaaS seats, and agents that do the same job twice.
- **Reopened findings.** Patches that fail silently, devices that drop off the management agent, and vulnerabilities that reappear on the next scan.
- **No executive picture.** Leadership sees dashboards full of counts but cannot answer Proctor's question: how many days are our critical systems available for hacking?

This is the gap a real program closes.

What a real program looks like

A managed patching and vulnerability management program that can support a PLA is a closed loop, not a calendar event. At BTI we run it as one operating model across five stages.

1. Inventory what you actually run

You cannot set a patch interval for systems you do not know about. The program starts with continuous discovery of endpoints, servers, network devices, and cloud workloads across every site, classified by business criticality so the PLA's "critical systems" means something specific.

2. Prioritize what is actually being exploited

Severity scores alone produce long lists and burned-out teams. The program prioritizes using exploitation evidence such as the CISA KEV catalog, vendor advisories, and threat intelligence, plus the exposure of the affected system.[5] The output is a short, defensible list of what must move first.

3. Patch what matters, on an agreed interval

Patches are tested, scheduled, and deployed against the agreed interval, with maintenance windows that respect each site's hours. Where a patch cannot be applied, the program applies a documented

compensating control, such as isolation or restricted access, and records an exception with an owner and expiration date.

4. Rescan and verify closure

This is the stage most ad hoc programs skip. A patch job that reports "success" is not proof. The program rescans the affected systems and confirms the finding is gone before the clock stops. Days-to-patch in a PLA should mean days to verified closure, not days to a deployment attempt.

5. Give leadership a vCISO risk picture

A vCISO translates the operational data into an executive view: days-to-patch by system tier, PLA attainment, open exceptions and owners, trends, and the cost and benefit of a faster interval. This is where Proctor's board conversation actually happens, and where leaders decide, on purpose, whether to stay at 30 days or move critical exploited issues to 7.

Why 24/7 humans plus AI matter

A PLA is only as good as the team that delivers it. Two things have changed that make PLA-style outcomes achievable for mid-market firms.

Around-the-clock watch from a SOC 2 Certified NOC/SOC

BTI operates a SOC 2 Certified network operations center and security operations center with 24/7 monitoring. The team that watches your network, identity, and endpoints sees new vulnerability signal as it arrives, so a newly exploited flaw on a firewall or VPN portal does not wait until Monday to be triaged. BTI is also ISO 27001 compliant, so the processes behind the program are documented and auditable.

One signal, not ten portals

BTI's AI-powered RMM (remote monitoring and management) and vulnerability scanning feed the same NOC/SOC. Patch status, scan results, network surveillance, and security telemetry come together in one view instead of living in separate vendor consoles. When it helps, that signal can be correlated with your phone systems and physical security as well. Shared signal is what makes verified closure possible: the team that patches is the team that rescans and the team that reports.

Governed AI agents working alongside technicians

AI changes the economics of speed. On the enterprise ladder, moving from 30-day to 7-day patching multiplies cost largely because it multiplies human hours. Governed AI agents can take on much of that repetitive load around the clock, for example:

- Watching advisories and exploitation feeds and matching them to your actual inventory.
- Grouping and prioritizing findings so technicians start with what is exploited and exposed.
- Preparing change records, maintenance windows, and rollback notes for human review.

- Checking rescan results and flagging any finding that failed to close.
- Drafting the executive summary and exception list for the vCISO review.

The word that matters is governed. At BTI, AI assists and accelerates; it does not act unchecked. Humans approve changes to production systems, own exceptions, and stay accountable for the outcome, and AI use follows documented policy, access controls, and logging. Skilled people plus governed AI is how a faster interval becomes affordable without burning out a team.

Local teams who know your sites

Automation does not replace being there. BTI has served clients for over 40 years with local teams in Southern California, Chicagoland, and Arizona. When a patch needs hands on a device or a coordinated window at a branch, there are people nearby who know your environment.

Costs and savings: a mid-market scenario

The following scenario is illustrative. It is a composite of patterns we commonly see in multi-site mid-market organizations, not a specific client, and it uses relative cost rather than dollar figures so you can apply it to your own numbers.

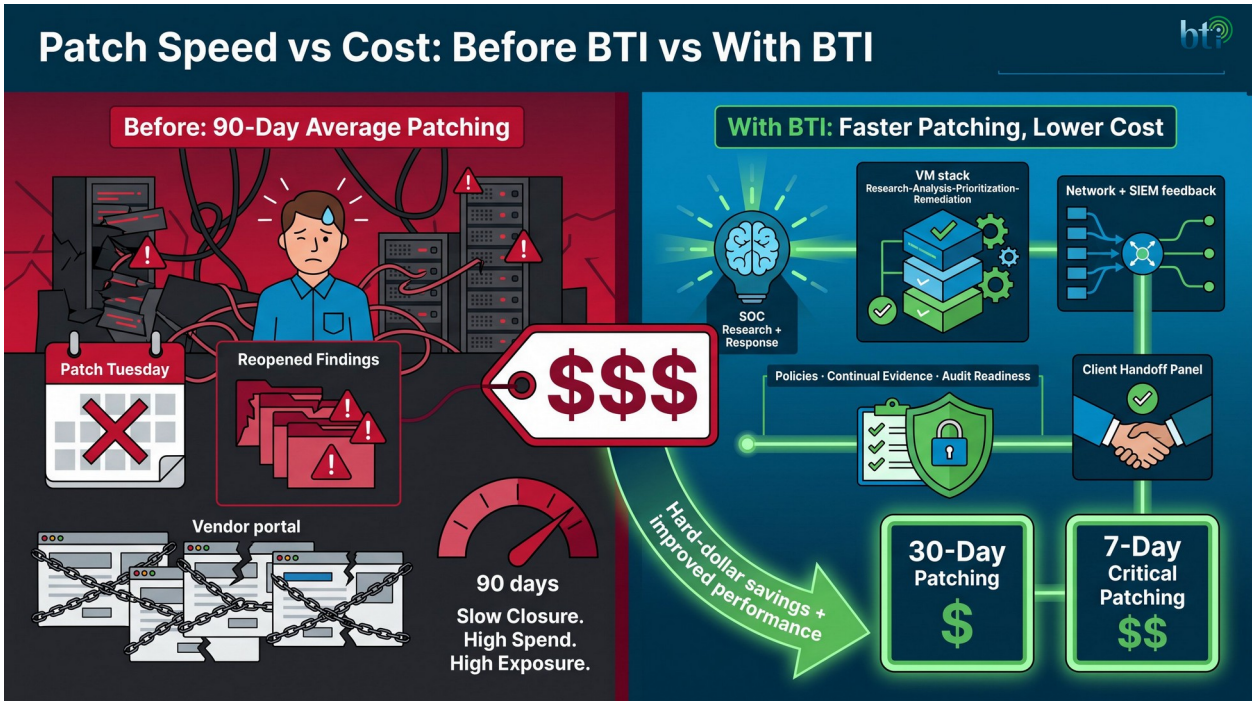


Figure 1. Illustrative before-and-after: 90-day average patching ($$$) versus 30-day patching ($) and 7-day critical patching ($$) with BTI. Relative cost only; not a price quote.

Before: 90-day average, high spend

Picture a company with several locations, a small IT team, and a collection of point tools. Patching happens monthly when it happens at all, devices that need reboot windows are deferred, and findings

reopen because nobody verifies closure. In practice, the average critical vulnerability stays open around 90 days.

That organization is not saving money by going slow. It is paying in less visible ways:

- **Tool sprawl.** Separate licenses for patching, scanning, endpoint, compliance tracking, and reporting, often with overlapping features.
- **Overtime and after-hours labor.** Internal staff working nights and weekends to catch up on deferred patches or respond to emergencies.
- **SaaS waste.** Unused or duplicate subscriptions that nobody has time to audit.
- **Downtime.** Unplanned outages from rushed patching, failed updates, or incidents that exploit known, unpatched flaws.
- **Audit and insurance friction.** Hours assembling evidence for questionnaires and renewals.

On a Proctor-style scorecard, this company has an implicit 90-day protection level, which is well outside the 30-day starting point of the enterprise ladder, at a total cost that is often higher than leaders realize.

With BTI: 30-day standard, 7-day critical

With a BTI program in place, the same organization operates against written intervals: a 30-day standard for routine patching and a 7-day target for critical, actively exploited vulnerabilities on exposed systems, each measured to verified closure.

- **30-day patching ($).** Routine patching becomes predictable, automated where possible, and verified. For most organizations this costs less than the "before" state because it replaces overlapping tools and catch-up labor with one managed program.
- **7-day critical patching ($$).** Moving critical exploited issues to 7 days adds some cost for faster research, after-hours change windows, and tighter verification. It is still typically below the fragmented "before" spend, and it moves the organization from Proctor's least protective band to his most protective band for the vulnerabilities that matter most.

Where the savings come from

Savings lever	How a unified BTI program helps	How to estimate it in your environment
Tool consolidation	One RMM, scanning, and reporting stack feeding one NOC/SOC	Sum current licenses for patching, scanning, and reporting tools
Labor and overtime	24/7 coverage and governed AI reduce after-hours catch-up work	Review last 12 months of IT overtime and contractor spend
SaaS waste	vCISO-led review of overlapping or unused subscriptions	Compare active users to paid seats across security and IT tools
Downtime avoided	Planned windows, tested patches, and faster closure of exploited flaws	Estimate hourly cost of downtime per site multiplied by recent outage hours
Audit and insurance effort	Continual evidence of patch intervals and verified closure	Hours spent per year on questionnaires, audits, and renewals

The goal is not just to spend less. It is to spend less and be measurably better protected, with a number on the scorecard that leadership can defend.

How BTI compares to the Gartner PLA cost bands

Proctor's ladder describes what large enterprises spend to build these capabilities internally. A mid-market firm does not need to recreate an enterprise security department. It needs the same discipline, delivered as a service.

PLA element	Enterprise approach in Proctor's examples	BTI managed approach for mid-market
Measurable patch interval	Internal team staffed to hit 30, 15, or 7 days	Written 30-day standard and 7-day critical intervals, measured by BTI
Verified closure	Internal scanning and change control	Rescan and verify before the clock stops
24/7 coverage	Internal SOC or multiple vendors	SOC 2 Certified NOC/SOC with 24/7 monitoring
Speed at scale	More headcount for faster intervals	Technicians plus governed AI agents
Executive risk picture	CISO reporting to the board	vCISO reporting on PLA attainment, exceptions, and trends
Illustrative annual cost	About $1 million (30-day) to about $5 million (7-day)	Lower than the illustrative enterprise ladder; scoped to your environment

The claim is straightforward. BTI delivers comparable protection outcomes (PLA-style measurable patch intervals, verified closure, and an executive risk picture) at prices well below the seven-figure enterprise ladder, because the cost of the people, the 24/7 operations center, the tooling, and the governed AI is shared across a managed service rather than built from scratch inside each company.

Your actual investment depends on your sites, devices, and the intervals you choose. We scope it with you during an assessment and show it side by side with what you spend today on the tools, labor, and downtime the program replaces.

Recommended PLA wording to take to your MSP

Whether you work with BTI or another provider, ask for patch intervals in writing. The sample clauses below are starting points for discussion with your provider and your legal counsel. They are not legal advice and should be adapted to your environment.

1. Definitions.

> "Critical System" means any asset designated by Client as critical in the asset register, including all internet-facing systems. "Critical Exploited Vulnerability" means a vulnerability listed in the CISA Known Exploited Vulnerabilities catalog, or otherwise confirmed by Provider as actively exploited, that affects a

> Critical System. "Verified Closure" means a successful rescan or equivalent technical validation confirming the vulnerability is no longer present.

2. Protection levels.

> Provider shall achieve Verified Closure of (a) Critical Exploited Vulnerabilities within seven (7) calendar days of the vendor fix becoming available or Provider's identification, whichever is later; (b) other critical and high-severity vulnerabilities on Critical Systems within thirty (30) calendar days; and (c) all other in-scope vulnerabilities within the interval set out in the Service Schedule.

3. Compensating controls and exceptions.

> Where a fix is unavailable or cannot be applied within the protection level, Provider shall implement a documented compensating control within the same interval and record an exception with a named owner, business justification, and expiration date. Exceptions require written Client approval.

4. Measurement and reporting.

> Provider shall report monthly on days-to-Verified-Closure by system tier, percentage of findings closed within the protection level, open exceptions, and trends. Measurement begins at vendor fix availability or Provider identification and ends at Verified Closure.

5. Coverage.

> Provider shall monitor for newly published and newly exploited vulnerabilities affecting in-scope assets 24 hours a day, 7 days a week, 365 days a year.

6. Governance review.

> Provider's vCISO shall meet with Client leadership at least quarterly to review protection level attainment, residual risk, and whether the agreed intervals remain appropriate as the business changes.

7. Use of automation and AI.

> Provider may use automation and AI tools to support monitoring, prioritization, and verification, provided such use is governed by documented policy, logged, and subject to human approval for changes to production systems.

8. Shortfall remedy.

> If Provider fails to meet a protection level, Provider shall notify Client promptly, deliver a root-cause analysis and remediation plan, and apply the remedies set out in the Service Schedule.

Notice what these clauses do. They turn Proctor's question into a contract term. They separate a control failure (we missed the interval) from accepted business risk (it happened inside the interval we chose). And they give your leadership the evidence it needs to show due care.

Call to action

Every organization already has a protection level. The only question is whether anyone chose it. If your answer to "How many days would you like your systems available for hacking?" is "we don't know," that is the place to start.

BTI offers a **free cybersecurity assessment** for mid-market and multi-site organizations. We will inventory what you run, measure your current days-to-patch on critical systems, identify exploited vulnerabilities that are open today, and show you what a 30-day standard and 7-day critical program would look like, side by side with what you spend now on tools, labor, and downtime.

Talk to a BTI IT sales team near you:

- Los Angeles and Southern California: **323-510-8003**
- Chicago and Chicagoland: **312-492-2613**
- Phoenix and Arizona: **480-272-8728**
- Online: **www.btigroup.com**

Hear it from the source. Paul Proctor's thinking on outcome-driven metrics and Protection Level Agreements is worth your leadership team's time:

- Gartner newsroom: "Gartner Says CFOs Must Treat Cybersecurity as a Business Decision"[1]
- FBI Ahead of the Threat Podcast, Season 1, Episode 7, with Paul Proctor[2]

About BTI Communications Group

BTI Communications Group has served businesses for over 40 years, delivering managed IT, cybersecurity, communications, and network services to organizations across Southern California, Chicagoland, and Arizona. BTI operates a SOC 2 Certified NOC/SOC with 24/7 monitoring and is ISO 27001 compliant. Our managed patching, vulnerability management, and vCISO services run as one governed operating model, combining experienced technicians, AI-powered RMM and vulnerability scanning, and governed AI, so mid-market organizations get measurable protection they can defend.

Gartner is a registered trademark of Gartner, Inc. References to Gartner and to Paul Proctor reflect publicly available commentary and do not imply endorsement of BTI Communications Group or its services. Cost figures attributed to Gartner commentary are illustrative examples from public remarks and are not BTI prices.

Notes

1. Gartner, Inc., "Gartner Says CFOs Must Treat Cybersecurity as a Business Decision," press release, National Harbor, Md., May 21, 2024 (URL dated 2024-05-22). Quotes from Paul Proctor, Distinguished VP Analyst. https://www.gartner.com/en/newsroom/press-releases/2024-05-22-gartner-says-cfos-must-treat-cybersecurity-as-a-business-decision

2. Federal Bureau of Investigation, "Ahead of the Threat Podcast: Season 1, Episode 7 - Paul Proctor," video and transcript. Source of the outcome-driven metric definition, days-to-patch example, $1 million for 30-day patching, $2 million for 15-day patching, operational technology example, and PLA funding commitment. https://www.fbi.gov/video-repository/ahead-of-the-threat-podcast-s1-e7-paul-proctor/view

3. Cybersecurity Compass, "Cybersecurity: A Strategic Choice and Balance of Trade-Offs for Business Success," summarizing Paul Proctor's Gartner conference remarks, including the question "How many days would you like our systems available for hacking?" and the illustrative figure of about $5 million per year for 7-day patching. https://cybersecuritycompass.substack.com/p/cybersecurity-a-strategic-choice . Gartner's own press

release uses the closely related phrase "limit the number of days a vulnerability is available for hacking" (see note 1).

4. Jason Hart, "Maximizing Cybersecurity Impact With Protection Level Agreements," Forbes Technology Council, August 2, 2023. https://www.forbes.com/councils/forbestechcouncil/2023/08/02/maximizing-cybersecurity-impact-with-protection-level-agreements/

5. Cybersecurity and Infrastructure Security Agency (CISA), "Known Exploited Vulnerabilities Catalog." https://www.cisa.gov/known-exploited-vulnerabilities-catalog

6. CISA, "BOD 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities" (default remediation within two weeks for KEV entries with CVE IDs assigned in 2021 or later, and within six months for earlier CVEs). https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities