The VPN Under CMMC Level 2:

Remote Access, Connected Facilities, and the Tunnels Nobody Documented

David W. Koran
CyberAB Registered Practitioner Advanced
July 2026

The Most Common Device in the Defense Industrial Base

Nearly every organization in the Defense Industrial Base operates at least one VPN, and most operate several without thinking of them as several. An engineer connects from home to reach the ERP system. A second facility reaches the first over a site-to-site tunnel. A machine tool vendor holds standing remote access to the equipment it services, and a managed service provider maintains its own path into the firewalls it administers. Each of those is a VPN; each carries or protects Controlled Unclassified Information or the systems that handle it, and each is treated differently under the 110 security requirements of NIST SP 800-171 Revision 2, which a CMMC Level 2 assessment scores.

The VPN also carries more misplaced confidence than any other device in the environment. Contractors routinely present the existence of a VPN as the answer to questions about encryption, remote work, facility connectivity, and vendor access, as though the word itself were a control. It is not. A VPN is a road with a locked surface, and what matters at assessment is who may enter the road, what may travel on it, where it terminates, and whether the cryptography on the surface is the kind the government accepts. This paper walks through the three tunnels most contractors actually run: remote access for the workforce, site-to-site between facilities, and third-party connections that rarely appear on any diagram, and states what each must satisfy, with a worked example on common commercial equipment.

One distinction frames everything that follows, so it comes first.

What a VPN Does and Does Not Do

A VPN protects the confidentiality and integrity of traffic in transit. It prevents a party positioned on the path, on the wire, in a carrier network, on a hotel wireless segment, from reading or altering what passes. That is the whole of it, and it is genuinely valuable, which is why the requirements demand it.

What a VPN does not do is restrict what travels through it. To an intruder who has compromised a machine on one end, the tunnel is simply a road, and an encrypted road is still a road. Ransomware propagating from a remote laptop into the plant network does not care that its traffic was beautifully encrypted on the way in. The control that limits movement is not the tunnel but the rule set governing what the tunnel may carry: which systems, which ports, which directions, with everything unnamed denied. Throughout this paper, every tunnel is therefore evaluated twice: once for the quality of its cryptography and once for the discipline of its flows, because an assessment does the same and because an adversary cares only about the second.

The FIPS Question Every Tunnel Must Answer

CUI in transit must be protected against unauthorized disclosure, and requirement 3.13.8 permits either cryptography or alternative physical safeguards to provide that protection. For the tunnels in this paper, which cross carrier networks and the public internet where no physical safeguard is possible, practical control is cryptography; when cryptography is used to protect the confidentiality of CUI, requirement 3.13.11 makes FIPS validation the gating issue. This pair of requirements yields more findings against VPNs than any other because the word "validated" does precise work. Three details decide the outcome.

First, the cryptographic module in the device or client must map to a certificate in the NIST Cryptographic Module Validation Program whose status is acceptable for the system's use case, current for a new procurement, and, at minimum, documented if Historical on an existing deployment, and the assessor will ask for the certificate number, not the vendor's marketing page. The certificate is specific to a module, a model family, and a firmware or software version, which means an upgrade can move a device out of its validated configuration as easily as neglect can leave it behind one.

Second, the device must be operating in its FIPS-approved mode. On most commercial platforms, this is a deliberate configuration state that must be enabled, restricts the available algorithms, and disables features that cannot operate within the approved set. A device that is capable of FIPS mode but not running in it does not satisfy the requirement, and this single checkbox is among the most common gaps found in otherwise well-built environments.

Third, private transport does not waive the obligation. A leased line, an MPLS circuit, or a carrier-managed wide-area service is not under the contractor's physical protection between its endpoints, and CUI crossing it unencrypted is CUI transmitted without the required protection. The safe engineering answer is uniform: encrypt CUI in transit with FIPS-validated cryptography, regardless of what the underlying transport claims to be.

The September Transition

A calendar item now sits atop the FIPS discussion, and it lands squarely on the devices discussed in this paper. FIPS 140-2 certificates remain valid through September 21, 2026; after that date, the Cryptographic Module Validation Program places every remaining 140-2 certificate on its Historical list. On November 10, 2026, Phase 2 of the CMMC program begins, and Level 2 certification requirements start appearing in applicable new solicitations, with the Department retaining discretion to defer the effective requirement within a specific contract, such as to an option period. Seven weeks separate the two dates, and a large share of the firewalls and VPN appliances deployed across the Defense Industrial Base hold FIPS 140-2 certificates today.

Precision about the mechanics matters, because much of the commentary on this transition is written by companies selling FIPS 140-3 migrations and overstates what happens. Historical is not revoked. A revoked certificate means the validation itself has been withdrawn and the module may no longer be used where validated cryptography is required. A Historical certificate remains the record of a completed validation; NIST's transition guidance directs that new procurements use FIPS 140-3 modules while permitting continued use of 140-2 modules in existing systems until 140-3 replacements become available. Nothing is revoked at the transition; the modules do not stop working, and the requirement's operative word, validated, describes an event that occurred and is not unwound by a list change. What changes is what the certificate lookup returns on assessment day, and a 140-2 certificate in an SSP after September invites a question the contractor must be prepared to answer.

The prepared answer is a documented posture, written before anyone asks: the module was validated under 140-2; the system is an existing deployment; NIST's transition guidance permits its continued use; and the refresh path to a 140-3-validated module is identified, with the vendor's timeline attached. That paragraph in the SSP is the difference between a documented position and a discovered one, and the scoring methodology itself shows the program has room for it, since the DoD Assessment Methodology applies a three-point deduction for encryption that is not FIPS-validated, rather than the full five for no encryption at all. One caution belongs alongside that

reading: as of this writing, neither the Department of Defense nor the Cyber AB has published definitive guidance on how assessors must treat Historical status, and individual assessment teams may reach different conclusions until one of them does.

The purchasing rule that follows is the simplest conclusion available. The validation queue offers no rescue, since industry analyses of CMVP processing data place average validation times near a year under 140-2 and closer to a year and a half under 140-3, and a vendor entering the pipeline now will not hold a certificate before the sunset, and compatibility claims, products marketed as FIPS inside on the strength of someone else's embedded module, are only as good as a certificate number and conformant use under that module's security policy. So from this point forward, no new cryptographic equipment enters a CUI environment without an active FIPS 140-3 certificate verified against the CMVP database before the purchase order is issued. The installed base is a documented transition. A new purchase that ignores the transition is a self-inflicted finding.

The Workforce Tunnel: Remote Access

The remote access VPN is the tunnel with the largest population and the longest list of requirements, because an entire family of controls exists for exactly this case. Requirement 3.1.12 obligates the contractor to monitor and control remote access sessions; 3.1.13 requires cryptographic protection of those sessions, which is the FIPS discussion above applied to the client and concentrator; 3.1.14 requires that remote access route through managed access control points, and 3.1.15 requires specific authorization before privileged commands run remotely. Above all of them sits 3.5.3, which requires multifactor authentication for network access, and remote access is the definitive case of network access. A remote-access VPN that accepts a username and password alone fails at the front door, whatever else is behind it.

The managed access control point requirement deserves translation, because it quietly prohibits the most convenient shortcuts. Remote users must enter through a defined, contractor-controlled gateway where monitoring, authentication, and flow control are performed. Direct exposure of internal systems, a remote desktop port opened on the firewall, and a database published to the internet for a traveling salesperson are precisely what the requirement forbids, and assessors find those exposures with a port scan.

Split tunneling is addressed by name. Requirement 3.13.7 obligates the contractor to prevent a remote device from maintaining a connection to the organizational network while simultaneously communicating with external networks over another path. A laptop connected to the plant over the VPN, with its other traffic flowing directly to the

home internet, serves as a bridge between the CUI environment and an unmanaged network, and the configuration must prevent this. Full tunnel policies cost bandwidth and buy compliance, and the choice must be visible in the client configuration the assessor reads.

Three refinements separate a workforce tunnel that merely connects from one built for assessment. The connection should not outlive its use: requirement 3.13.9 requires network connections to terminate at the end of a session or after a defined period of inactivity, and 3.1.11 applies the same discipline to the user session itself, so an idle laptop left connected overnight is a finding waiting for a timestamp. Arrival should not mean the whole network: under the least privilege principle of 3.1.5 and the transaction scoping of 3.1.2, the concentrator should place each remote user in a segment scoped to the role rather than on the flat internal network, which is the remote access version of the defined flows this paper demands of every other tunnel. And the managed device rule needs a technical enforcer: a device certificate or posture check at connection, verifying that the machine is the contractor's, patched, and running its protections, is what turns the policy that only managed devices connect into a control the assessor can watch operate.

Finally, the endpoint at the far end of the workforce tunnel is in scope, where remote access scoping is actually decided. The employee's home router, the hotel network, and the coffee shop are not in the assessment boundary. The laptop is used because CUI is processed on it, which means it must be a contractor-managed device that carries the full weight of the requirements: configuration-managed, encrypted at rest, patched, logged, and inventoried. A personal machine reaching the CUI environment over a compliant VPN is not a compliant arrangement. It is an unmanaged CUI asset with excellent transport security. The one recognized narrowing is virtual desktop infrastructure: under the program's published guidance, an endpoint used strictly as a window into a hosted environment, configured so that CUI is never processed, stored, or transmitted on the device beyond keyboard, video, and mouse, with local transfer, printing, and caching disabled, can fall outside the boundary, and that configuration must be demonstrable rather than assumed.

The technical controls above also need a written counterpart, because the assessment objectives behind 3.1.12 ask whether the types of permitted remote access are identified and whether usage restrictions are defined, and enforcement cannot be evaluated against rules that were never written down. A remote access policy states who may connect, from which devices, and what may be done during a session: CUI is handled only within the approved systems, the session is not left unattended on an unlocked screen, the connection is not shared with anyone else in the household, and work does not migrate onto personal email, personal cloud storage, or the family printer, since a

document printed at home has just left the assessment boundary by the paper tray. The policy must match the configuration it governs; users must acknowledge it; and the acknowledgment records must be added to the evidence file, because a rule nobody signed is a rule the assessor cannot credit. The technical control prevents what it can, the policy governs what technology cannot see, and the assessment expects both.

The Facility Tunnel: Connecting Two Buildings

The second tunnel connects the contractor's own facilities, and here the VPN stops being an access question and becomes a scoping question. Under 32 CFR part 170, the assessment boundary follows where CUI is processed, stored, or transmitted, and which assets can affect those assets. A site-to-site tunnel between two buildings is a statement about which assets can affect which; that statement is made by the rule set, not the tunnel.

If the tunnel carries all traffic without restriction, the two facilities constitute a single environment. Every asset in both buildings falls inside the boundary, the shipping computer in one building and the front office machines in the other alike, because nothing prevents CUI from reaching them, and the System Security Plan must account for the full population. A network diagram that draws a dotted line around a CUI zone is a claim, and the assessment tests claims against enforcement: if the rule set between the zones shows an any-to-any permit, the separation does not exist for assessment purposes. Logical separation is real only where a control actually denies traffic.

Two architectures survive assessment cleanly. The unified enclave defines one CUI environment spanning both facilities, joined by the encrypted tunnel into a single logical enclave with enforced boundaries to everything else at each site; it suits contractors whose CUI work genuinely spans both buildings, and its cost is that both buildings' CUI assets are in scope and both facilities must independently satisfy the physical protection requirements of the 3.10 family. The segmented interconnection maintains a bounded CUI environment at each facility and treats the tunnel as an external interface, permitting only defined flows, specific systems, ports, and directions, and denying the rest; it suits contractors whose cross-site traffic is narrow, and it adds two boundaries to document and maintain.

Whichever architecture is chosen, the tunnel is a boundary under requirement 3.13.1, external in the segmented model and a key internal boundary in the unified one, and requirement 3.1.3 asks the same question from the data's point of view: the contractor must be able to describe how CUI moves between the facilities, through which systems, and what prevents it from moving anywhere else. A flow the contractor cannot describe is a flow the contractor is not controlling. And the requirements that apply per

facility do not merge because the networks did: physical protection, wireless authentication, and audit collection must each stand up in both buildings, with clocks synchronized so that events at the two sites correlate into one reviewable record.

A Worked Example: Two Carriers, One Tunnel

The abstract requirements become concrete in the topology most two-facility contractors actually run. One facility connects to the internet over a business coax circuit from a regional cable carrier such as Spectrum. The other sits on fiber from a metro provider such as Lightpath. A SonicWall firewall at each site terminates a site-to-site IPsec tunnel between the two sites, and ERP, file replication, and log collection traffic rides inside it. The vendor and the carriers are illustrative; the same analysis applies to any equivalent build.

The carriers themselves are out of scope. Commodity internet transport is not an External Service Provider under the CMMC scoping guidance, and the boundary at each facility is the contractor's own firewall, not the carrier modem and not the demarc. Everything between the two firewalls, both carrier networks and whatever routing sits between them, is assumed hostile, and that assumption is what makes the topology clean: because the path is unambiguously untrusted, there is no argument about whether encryption is required. It is.

On the SonicWall, the FIPS obligation resolves into three checkable facts. The appliance model and firmware version must correspond to a CMVP certificate whose status matches the deployment's posture, and the certificate number must be included in the evidence file. FIPS mode must be enabled in SonicOS, a deliberate state that restricts the device to approved algorithms; a unit capable of FIPS mode but running outside it does not satisfy the requirement. And the tunnel proposals must use approved settings: IKEv2 with AES 256, SHA 256, and a Diffie-Hellman group of 14 or higher, which the configuration export will show or fail to show.

The trap specific to this build quietly rebuilds the flat network inside the tunnel. When a site-to-site VPN is defined between the two LAN subnets, SonicOS by default creates access rules that permit all traffic between those networks in both directions, as described in SonicWall's own configuration documentation. The tunnel comes up, everything works, and the two buildings are once again one unrestricted environment, encrypted but flat. The correct build suppresses the automatically created rules and replaces them with explicit permits on the VPN zone, the ERP server on its ports, the replication path, the log collection flow, and a default deny beneath them. The assessor evaluating flow control does not read the tunnel status page. The assessor reads the access rules on the VPN zone, and those rules either name the flows or they do not.

One field note completes the example. Coaxial business circuits frequently carry dynamic public addresses, while fiber circuits carry static ones, which shapes the configuration: the fiber side serves as the stable peer, with the dynamic side initiating, or dynamic DNS covers the coax address. That is an engineering nuisance rather than a finding, but the network diagram must reflect whichever answer was chosen, because an assessor tracing the topology will ask what the tunnel terminates on, and the diagram, the configuration export, and the spoken answer need to match.

The Tunnels Nobody Documented

The third category is the one that damages assessments, because it is the one that surprises the contractor. Real environments accumulate connections nobody mentions: a persistent tunnel to a sister company, a machine tool vendor's standing remote access to the equipment it services, a managed service provider's management plane into the firewalls, a legacy link to a former division that nobody ever tore down.

Every one of those is an external system connection, and each demands the same treatment. It must be identified and documented, its traffic must be controlled at the boundary, and the contractor must be able to state what the connected party can reach. A persistent unrestricted tunnel to an affiliated company that is not itself inside the assessment boundary does not merely create a finding. It potentially extends the boundary, because assets on the far end that can reach CUI assets without restriction become part of the security problem the assessment must evaluate, and the contractor cannot produce evidence for an environment it does not control. The remediation is almost always the same and almost always resisted: the connection is removed or reduced to defined flows through an enforced boundary before the assessment is scheduled. There is no documentation strategy that substitutes for either.

Vendor remote access deserves its own sentence, because the maintenance requirements address it directly: nonlocal maintenance sessions require multifactor authentication and must be terminated when the work is complete, which prohibits the standing, always-on vendor tunnel in its most common form. And if an external provider manages the firewalls, tunnels, or the logging infrastructure described in this paper, that provider is an External Service Provider under the CMMC program; its services are in scope, and the System Security Plan must describe the division of responsibility. In scope describes the contractor's assessment rather than a burden on the provider: the services are evaluated within the contractor's assessment scope, which is distinct from the provider needing to hold a CMMC certification of its own. A tunnel cannot be assessed if nobody at the table can explain who controls it.

The Appliance Is the Perimeter

Every tunnel in this paper terminates on a device, and those devices are the most consequential assets in the environment, because internet-facing VPN and firewall appliances, across every major vendor, have become a recurring, high-value intrusion path into small and mid-sized organizations, a risk emphasized in government guidance on protecting network edge devices. The adversary who owns the appliance owns the tunnel, the rule set, and everything behind both. The discipline matters more than the brand, and it is the same short list everywhere: management access disabled on the WAN side and restricted to the inside, multifactor authentication on administrative accounts, firmware inside the patching program with short timelines for critical advisories, and every appliance enrolled in vulnerability scanning and central logging like the Security Protection Asset it is. An organization that encrypts perfectly through a device it does not patch has built a bank vault door into a tent.

If It Is Not Logged, It Did Not Happen

Every claim made so far in this paper- the tunnel is restricted, the users are authenticated, the vendor session was terminated- is a claim about events, and under CMMC the only acceptable proof of an event is a log record. The 3.3 family requires that system audit logs be created, retained, protected, and reviewed; requirement 3.3.2 requires that logged actions trace to individual users, and requirement 3.1.12 requires that remote access sessions be monitored and controlled, which is impossible for a session that left no record. A VPN without comprehensive logging is a set of controls that cannot be demonstrated, and at assessment, a control that cannot be demonstrated scores the same as a control that does not exist.

What must be captured follows the three tunnels. For remote access, the record is every authentication event, successes and failures alike, with the multifactor result, the source address, the session start and end, and the identity behind each, because failed authentications are where attacks appear first, and per-user traceability is what 3.3.2 demands. For the site-to-site link, the record is tunnel establishment and teardown, the traffic the rule set permitted, and critically the traffic it denied, because a deny rule that does not log cannot prove it ever enforced anything. For vendor and provider tunnels, the record is session start and session termination, since the maintenance requirements make the ending of a nonlocal session an event the contractor must be able to evidence, along with what the connected party reached while inside. And beneath all three sit the appliances themselves: administrative logons, failed administrative attempts, and every configuration change, because the rule set is only as trustworthy as the record of who has altered it.

Where the records live matters as much as whether they exist. VPN appliances hold small local log buffers that roll over in days or sometimes hours, and a buffer that overwrites itself is not retention; it is amnesia with a delay. The requirements resolve only through central collection: every appliance and concentrator forwarding to a log platform that retains records across the full review period, with clocks synchronized to an authoritative source under 3.3.7 so that an authentication in one building correlates with a file access in the other, with the audit information itself protected from unauthorized access, modification, and deletion under 3.3.8 and management of the logging function limited to authorized privileged users under 3.3.9, and with an alert when a device stops sending, since requirement 3.3.4 makes a failure of the logging process itself a reportable event. Silence from a firewall should page someone, because silence is either a dead log source or a covered track, and neither can wait for the weekly review.

The review is the half that organizations skip. Requirement 3.3.3 requires that logged events be reviewed, and 3.14.6 and 3.14.7 require monitoring for attacks and unauthorized use, which together mean that somebody must actually look, and that the looking must leave its own evidence: the review notes, the tickets opened for anomalies, the account disabled after repeated failures from a foreign address. This is also where the timeline argument of this paper closes, because logs are the one control that cannot be remediated retroactively. A contractor can rebuild a rule set in a week, but an assessor sampling the review period will ask for records spanning months, and the 120-day question, can you produce your VPN authentication and deny logs for the last four months and show they were reviewed, is answered by the calendar or not at all. The organization that turns on comprehensive VPN logging today is buying evidence it cannot purchase at any price the week before the assessment.

What the Assessor Will Actually Test

An assessment team evaluates every claim about a VPN against configuration it can inspect. The evidence set is predictable, and a contractor should be able to produce every item in the table below without assembly panic.

Evidence	What It Must Show	Primary Requirements
Network diagram, all facilities	Every tunnel: remote access concentrators, site-to-site links, and all third-party connections, with termination points	3.13.1, 3.13.2
CUI data flow diagram	How CUI enters, moves across tunnels, and leaves, system by system	3.1.3

Evidence	What It Must Show	Primary Requirements
VPN configuration exports and CMVP certificate numbers	FIPS validated modules, FIPS mode enabled, approved algorithms on every tunnel	3.13.8, 3.13.11, 3.1.13
Firewall and VPN zone rule sets	Defined permits and default deny governing what each tunnel may carry; no any-to-any permits between zones claimed as separate	3.13.1, 3.1.3
Multifactor authentication configuration	MFA enforced for every remote access user and every administrative account	3.5.3
Remote access client policy	Full tunnel enforcement preventing split tunneling, managed access control points as the only entry, and inactivity timeouts on connections and sessions	3.13.7, 3.1.14, 3.13.9
Remote endpoint inventory	Every device permitted on the workforce tunnel is contractor managed, encrypted, patched, and logged	3.1.12, 3.4.1
Remote session logs	Monitoring of remote sessions, and termination records for vendor maintenance sessions	3.1.12, 3.7.5
VPN authentication and deny logs	Success and failure events traceable to individual users, denied traffic at the tunnel boundary, central retention across the review period, and evidence of review	3.3.1, 3.3.2, 3.3.3, 3.3.5
External connection inventory	Every third-party tunnel, its purpose, its restrictions, and its owner	3.13.1, 3.12.4

Table 1. The VPN evidence set, mapped to the NIST SP 800-171 Revision 2 requirements, shows which items each item supports.

The pattern across the table is uniform. The word VPN answers no assessment question by itself. Every claim about protection, access, or separation must correspond to a certificate number, a configuration state, or a rule the team can read, and the contractor that passes is the one whose exports say the same thing its diagrams do.

Reading Your Own Tunnels

For an executive who wants to know where the organization stands, the audit takes one meeting and five questions. How many tunnels do we have, including the vendor and provider connections that nobody drew? For each one, can we provide the CMVP certificate number and show that FIPS mode is enabled? Does every remote user pass multifactor authentication through a managed gateway on a company-managed device with split tunneling prevented? For the site-to-site links, can we name every flow the rule set permits and show the deny beneath them? And who controls each termination point, us or a provider, and where is that written down?

An organization that answers all five cleanly has done the work. An organization that answers with the word VPN has found its remediation plan, and the honest version of that plan is measured in months, because rule sets must be rebuilt, endpoints brought under management, and the operational evidence the logs and reviews and the assessment samples accumulated over time rather than assembled the week before. The tunnel can be encrypted in an afternoon. The discipline around it cannot.

About the Author

David W. Koran holds the CyberAB Registered Practitioner Advanced (RPA) credential and is the founder of a consulting practice serving Defense Industrial Base contractors and the legal counsel who support them. His work focuses on CMMC Level 2 readiness, enablement, and implementation. He is an Associate Member of the American Bar Association Section of Public Contract Law and the author of The CMMC Decision. He can be reached at dkoran@davidkoran.com and (802) 335-2662.

References

Cybersecurity Maturity Model Certification (CMMC) Program, 32 CFR part 170, including the asset categories and scoping provisions of section 170.19.
https://www.ecfr.gov/current/title-32/subtitle-A/chapter-I/subchapter-D/part-170

National Institute of Standards and Technology, Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations, NIST SP 800-171 Revision 2.
https://csrc.nist.gov/pubs/sp/800/171/r2/upd1/final

National Institute of Standards and Technology, Assessing Security Requirements for Controlled Unclassified Information, NIST SP 800-171A.
https://csrc.nist.gov/pubs/sp/800/171/a/final

Department of Defense, CMMC Level 2 Scoping Guidance, CMMC Level 2 Assessment Guide, and CMMC Program Frequently Asked Questions.
https://dodcio.defense.gov/CMMC/Documentation/

National Institute of Standards and Technology, Cryptographic Module Validation Program (CMVP), including the FIPS 140-3 Transition Effort and the Validated Modules, Historical, and Modules in Process lists.
https://csrc.nist.gov/projects/cryptographic-module-validation-program

Federal Information Processing Standards Publication 140-3, Security Requirements for Cryptographic Modules. https://csrc.nist.gov/pubs/fips/140-3/final

Department of Defense, NIST SP 800-171 DoD Assessment Methodology.
https://www.acq.osd.mil/asda/dpc/cp/cyber/safeguarding.html

Defense Federal Acquisition Regulation Supplement, DFARS 252.204-7012, Safeguarding Covered Defense Information and Cyber Incident Reporting.

https://www.acquisition.gov/dfars/252.204-7012-safeguarding-covered-defense-information-and-cyber-incident-reporting.

SonicWall, SonicOS IPSec VPN Administration, VPN Auto Added Access Rule Control. https://www.sonicwall.com/support/technical-documentation/docs/sonicos-7-0-0-0-ipsec_vpn/Content/ipsec-vpn-types-auto-added-acess-rule.htm

SafeLogic, What Happens on September 21, 2026? (2025), on CMVP validation processing times. https://www.safelogic.com/blog/what-happens-on-september-21-2026

Cybersecurity and Infrastructure Security Agency, Guidance and Strategies to Protect Network Edge Devices. https://www.cisa.gov/resources-tools/resources/guidance-and-strategies-protect-network-edge-devices