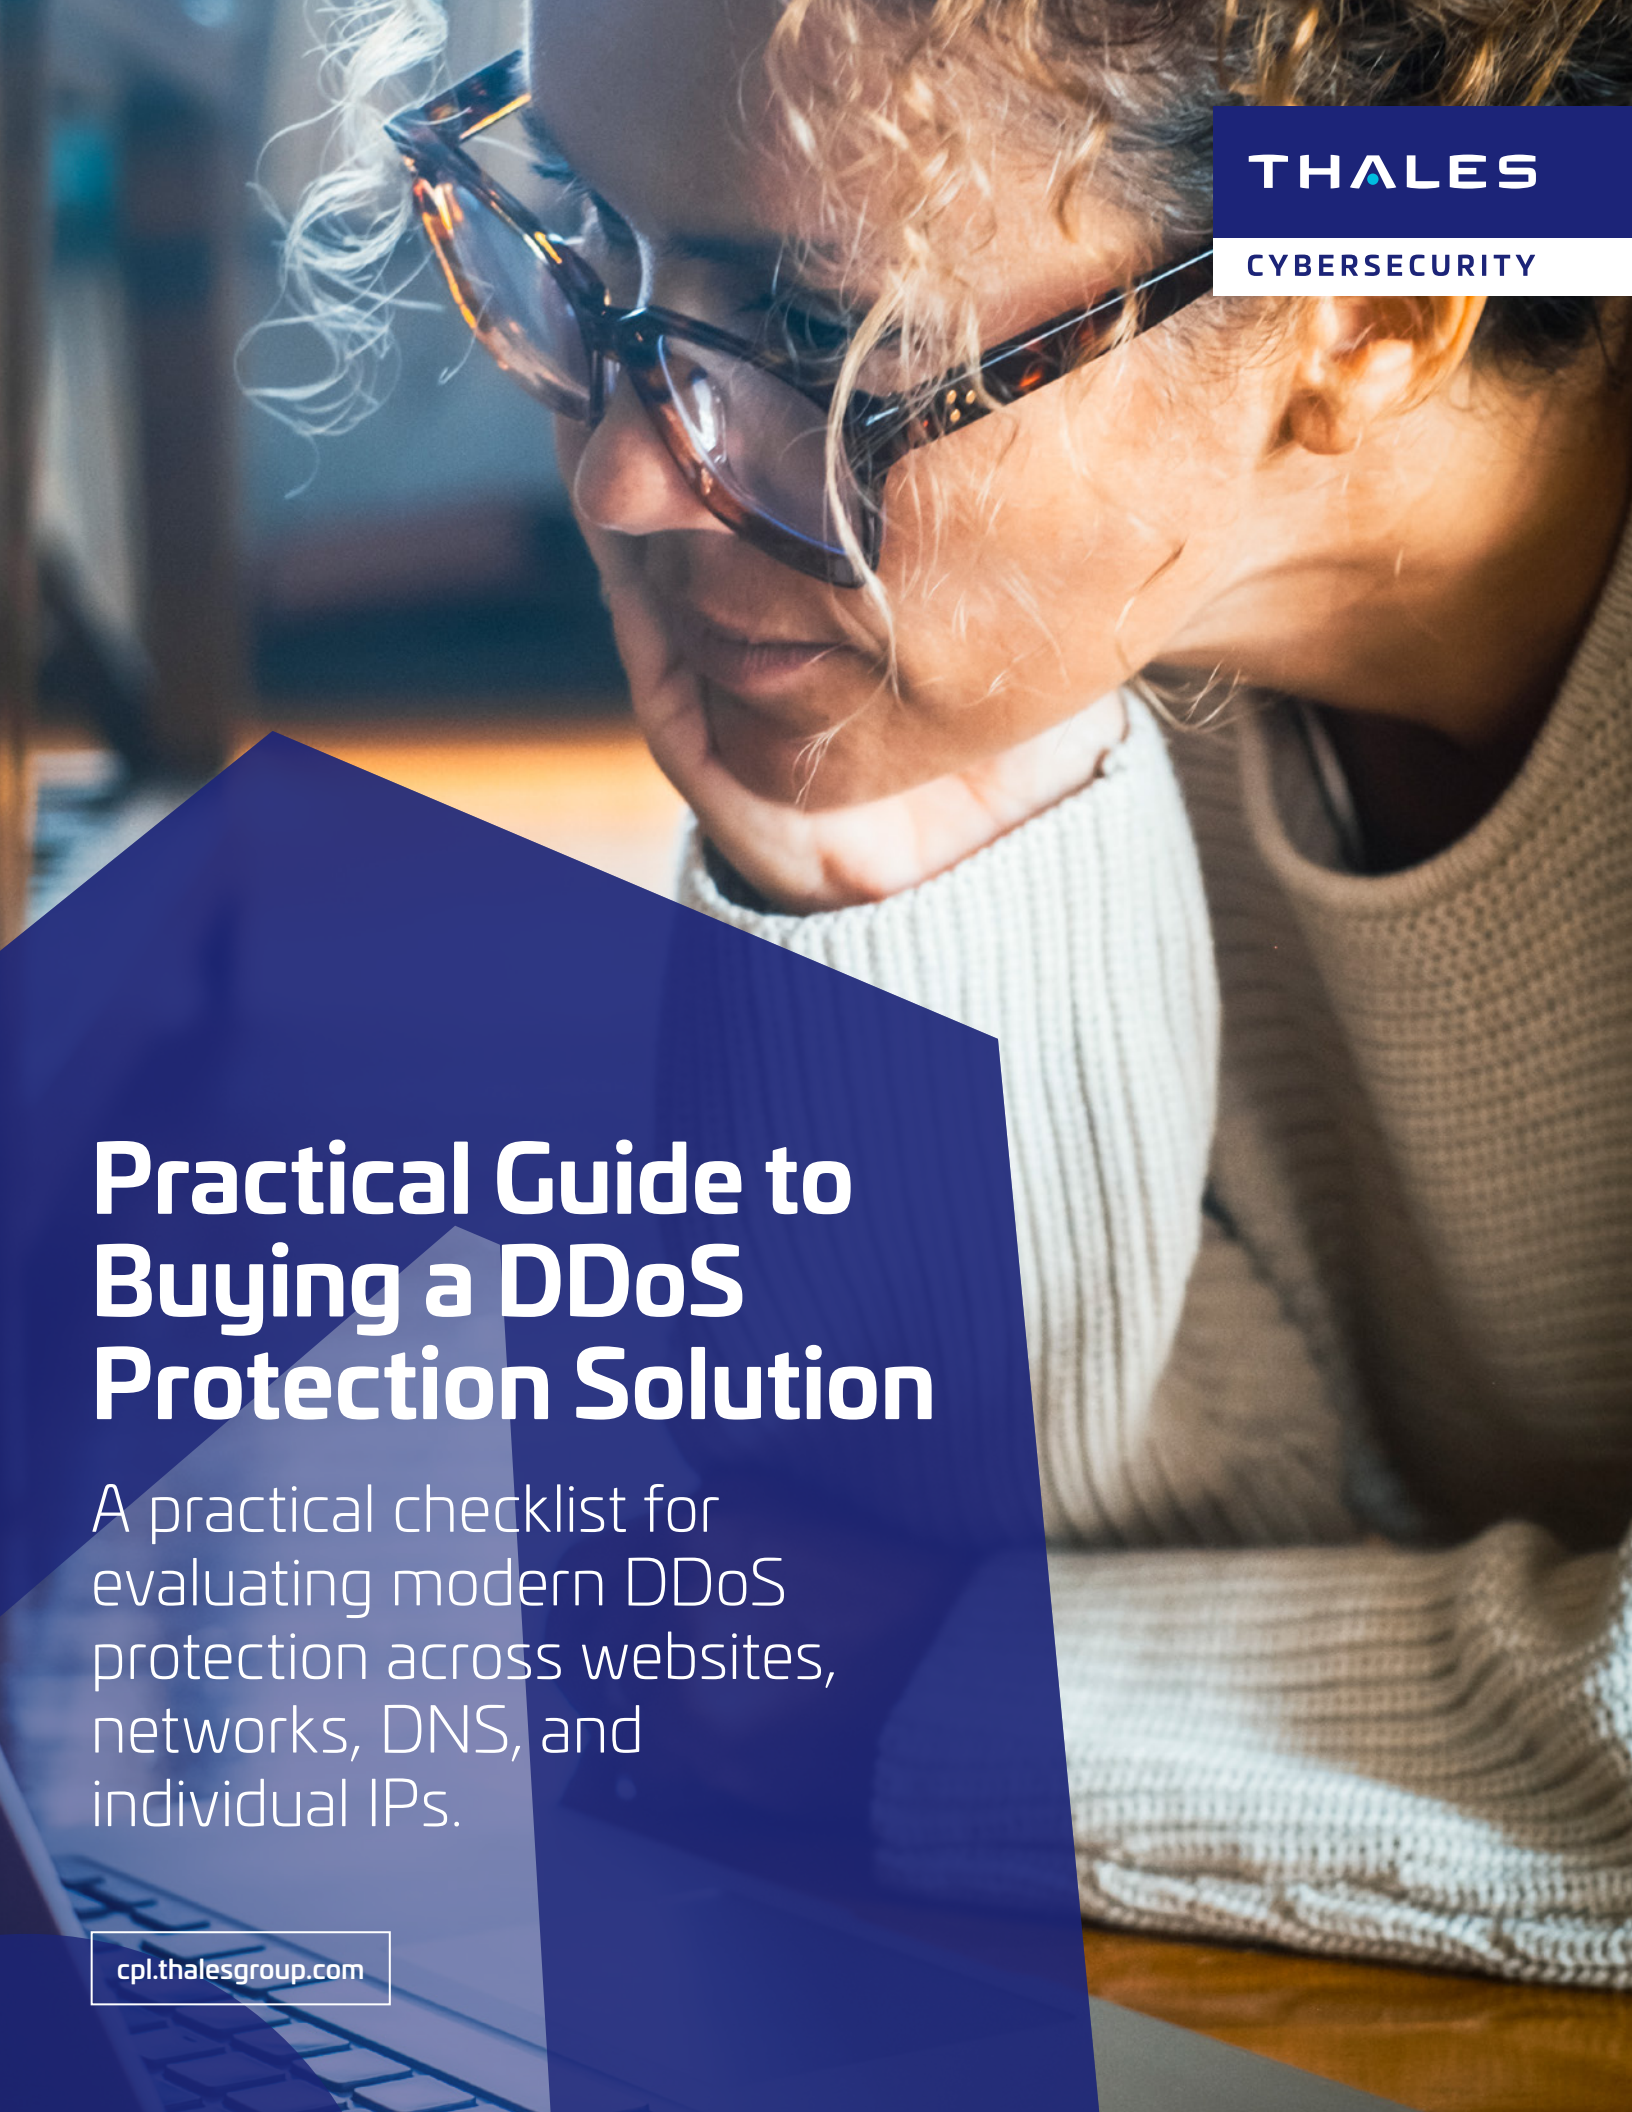

THALES

CYBERSECURITY

Practical Guide to Buying a DDoS Protection Solution

A practical checklist for evaluating modern DDoS protection across websites, networks, DNS, and individual IPs.

cpl.thalesgroup.com

Executive summary

DDoS attacks have evolved from simple bandwidth floods into coordinated, multi-vector campaigns that blend volumetric traffic with application-layer and API abuse, often hitting websites, networks, DNS services, and individual IPs simultaneously.

As attacks become more automated and adaptive, buyers should move beyond point solutions and evaluate DDoS protection as a resilience capability: global scale, fast mitigation, transparent user experience, and clear visibility across every exposed service. This guide provides a step-by-step framework you can use to define requirements, compare vendors, and plan deployment.

- Understand what "good" DDoS protection looks like today: coverage across web, APIs, DNS, networks, and IP-based services.
- Learn the key evaluation criteria to compare vendors consistently (capacity, time-to-mitigate, accuracy at layers 3–7, automation, analytics, and support).
- Leave with a clear set of RFP questions and implementation checks to reduce downtime risk and avoid disruption to legitimate users.

How to use this guide

1. Build a complete inventory of internet-exposed services (web, APIs, DNS, networks, individual IPs) that must remain available during an attack.

2. Use the checklist in Step 3 to turn availability goals into clear requirements and to score providers consistently.

3. Copy/paste the RFP questions into your evaluation document, then validate short-listed vendors with a controlled test of routing, dashboards, escalation, and rollback.

THALES

CYBERSECURITY

Step 1: Define what you need to protect

Before you compare vendors, create a short inventory of your exposed services and how they are reached from the internet. This prevents buying a website-only service when the true outage risk is DNS or public IP infrastructure or selecting a network-only scrubbing service when Layer 7 applications and APIs are the most exposed services.

- Websites & web applications: domains, key paths, peak traffic, performance/latency sensitivity.
- APIs: public endpoints, authentication patterns, rate limits, mobile/partner dependencies.
- DNS: authoritative DNS provider(s), zones, anycast support, DNSSEC usage
- Non-web services: VPN, VoIP, SMTP, FTP, gaming, IoT, proprietary TCP/UDP services.
- On-prem infrastructure: public networks, servers and internet connectivity.
- Individual public IPs: single servers, cloud workloads, or services where you don't control the full network routing.

Step 2:
Evaluate technical capabilities

1) Global scale and service coverage

Start with two fundamentals: (1) can the provider mitigate the attacks you're likely to face, and (2) can it protect every exposed service—not just your website.

- Documented network size and global distribution (number of PoPs and their locations)
- Protection scope across web applications, APIs and non-web applications.
- Ability to mitigate volumetric, protocol, and application-layer/API attacks.

2) Choose a protection model (always-on vs on-demand)

For network DDoS protection, organizations typically choose between always-on protection (traffic continuously routed through the mitigation network) and on-demand protection (traffic is rerouted only during an attack). This choice is generally relevant only for protecting internet-routable network address space, typically a publicly advertised /24 block or larger, and for environments focused primarily on network-layer protection rather than website, API, DNS, or individual IP protection.

Model	Best when...	What to validate
Always-on	You want continuous protection for internet-routable network space, immediate mitigation, and minimal operational dependency during an attack.	Confirm that your address space can be continuously advertised through the provider, mitigation happens automatically, operational overhead is low, and any latency or routing impact is acceptable.
On-demand	You have internet-routable network space and prefer direct routing to your datacentre during business-as-usual conditions, with traffic diverted only when a network-layer attack occurs.	Validate how diversion is triggered, whether activation is automatic or manual, how quickly traffic is rerouted, what testing and rollback options exist, and whether the service can handle large attacks without congestion.

3) Time to mitigate and SLA clarity

In DDoS, minutes matter. Validate how quickly the service detects and mitigates attacks, and ensure any SLA is specific about attack layer, conditions, and measurement method.

- Is there a time-to-mitigate SLA for network-layer volumetric attacks? What is it measured against?
- Does the SLA cover all potential attack vectors, or only those that have been pre-configured by applying specific signatures?
- Is mitigation automatic (always-on) or dependent on manual activation/threshold triggers?
- Is DDoS mitigation performed at all PoPs, or is traffic rerouted to dedicated scrubbing PoPs? This may affect time to mitigate and add latency.
- How does the vendor handle hit-and-run attacks that start/stop rapidly?

4) Application-layer and API attack accuracy (without breaking real users)

Layer 7 and API-focused DDoS often look like legitimate traffic. The key is accurate detection that stops bots and abusive clients while keeping real users moving without forcing widespread challenges that damage conversion and customer experience.

- How does the solution differentiate bots vs humans and good bots vs bad bots?
- What techniques are used for API abuse (credential stuffing-style bursts, token replay, endpoint exhaustion, expensive queries)?
- What is the approach to transparent mitigation versus challenge-based controls (for example, CAPTCHA)?
- How are false positives measured, tuned, and reviewed during an incident?

5) DNS resilience is non-negotiable

If attackers can take down your DNS, they can take down everything, regardless of how strong your web or network defenses are. Treat DNS DDoS as part of your core availability strategy.

- Does the vendor protect authoritative DNS against volumetric and protocol attacks?
- Are protections always-on and globally distributed?
- Can you get attack analytics for DNS events alongside other DDoS activity?

6) Onboarding and traffic routing options

As you evaluate vendors, prioritise solutions that can be deployed quickly and safely, with routing models that match each asset type. For websites and individual IPs, look for onboarding that can be activated through a simple DNS change. For networks and infrastructure protection, validate BGP-based routing that keeps routing control in your hands, as well as diverse connectivity options such as GRE over the internet, direct cross-connects, or network fabrics.

- Web onboarding: reverse proxy model, origin IP masking, TLS certificate handling
- Network onboarding: BGP peering over GRE and direct connects and network fabric; flexible prefix advertisement controls for network engineers and security operators.
- Individual IPs: reverse proxy for TCP apps and/or tunnelling for any application traffic.
- Flexible Onboarding using management portal, APIs or infrastructure as a code platforms.
- Time to deploy: can new assets be onboarded and protected in minutes, even during an attack.

7) Automation and modern operations

Strong DDoS protection reduces dependency on manual intervention. Look for automated operations that continuously profile traffic to improve detection accuracy and automatically reroute traffic during attacks to avoid congestion and interference. This is especially important when you can't guarantee expert responders are available every time an attack begins.

- Is mitigation fully automated, or does it require manual intervention during an attack?
- Are security policies automatically updated as normal traffic patterns shift across both network-layer and application-layer protocols?
- How does the service avoid "noisy neighbor" impact in multi-tenant environments?
- What controls exist for change management (approvals, APIs, role-based access)?

8) Visibility, analytics, and integrations

DDoS events generate high volumes of signals. You want analytics that summarise what's happening (vectors, targets, timelines, mitigations) and help you correlate DDoS with other attack activity so responders can prioritise the right actions.

- Real-time dashboards with clear impact and recommended actions.
- Correlation of DDoS with related security events (for example, WAF, bot, API security)
- Reporting suitable for executives (availability impact, business risk) and for responders (technical details)
- Export/streaming to your monitoring stack (for example, SIEM) and incident response workflows.

9) Support model: 24/7 operations and incident readiness

Even with automation, you should validate the vendor's operational support: coverage, escalation paths, and what "managed" really means during a high-severity incident.

- 24/7 operations centre availability and escalation SLAs.
- Pre-attack readiness: runbooks, tests, onboarding validation, and periodic reviews.
- Post-incident reporting and recommendations to reduce repeat impact.

RFP questions

Threat coverage and scale

- How do you mitigate different type of volumetric/ multi-Tbps DDoS attacks?
- How do you differentiate and keep the DDoS attack away from the protected networks while continuing to route legitimate user traffic?
- Which attack types do you mitigate by default: volumetric, protocol, application-layer, API?
- Which assets can you protect: websites, networks, DNS, individual IPs? Describe limitations.

Mitigation performance and SLAs

- Do you offer a time-to-mitigate SLA? For which layers and under what conditions?
- How do you detect and mitigate hit-and-run attacks with short (and repetitive) spikes?
- How do your security policies stay up to date as legitimate network patterns increase or decrease?
- How does your network handle congestion caused by flood traffic or high packet-per-second rates?
- How does the solution detect advertisement changes or sudden reductions in good traffic that may indicate routing changes outside the organization?
- How does the solution protect against "noisy neighbour" risk, where customers that are not the attack target are still impacted as collateral damage?
- What incident data do you provide, such as attack vectors, timelines, and mitigations applied?
- Also ask how the data is presented and how quickly it becomes available after an attack begins.
- Do allow and block lists exist, and can the customer manage them? Can traffic be blocked or allowed based on ASN, IP ranges, geolocation, network attributes, or other network intelligence?

Application, API, and user experience

- How do you mitigate application-layer and API attacks without introducing user friction?
- Does the solution provide adaptive Layer 7 DDoS protection, including automatic detection of abnormal application traffic patterns and dynamic mitigation without requiring pre-built signatures or manual rule creation?
- Does the offering include advanced DDoS capabilities that announce, detect, and mitigate application-layer attacks in real time, and how are these capabilities surfaced to security and network teams during an incident?
- What is your approach to distinguishing good bots from bad bots?
- How do you measure and tune false positives during and after incidents?

Deployment, routing, and operations

- Which onboarding models do you support: DNS change, reverse proxy, BGP, GRE, and cross-connect?
- Describe your automation for rerouting traffic during attacks (always-on vs on-demand).
- What integrations do you support for logging and event streaming (for example, SIEM)?
- Describe your 24/7 support model and escalation process.

Key Features & Benefits: Imperva DDoS Protection

Below is an overview of Imperva DDoS Protection, included as an example of how a full-stack DDoS offering can map to the evaluation criteria in this guide.

Imperva DDoS Protection

- Mitigation speed: a 3-second mitigation SLA for network-layer volumetric attacks, with mitigation typically occurring in under one second.
- Always-on and on-demand options: always-on mitigation at the edge to avoid delays, plus on-demand routing for latency-sensitive environments.
- Transparent user experience: mitigates attacks without relying on broad CAPTCHA challenges that disrupt users.
- Integrated platform approach: each PoP combines DDoS scrubbing with cloud WAF and advanced bot mitigation on a single, software-defined platform.
- Analytics: attack analytics condense high-volume events into clear insights and correlate DDoS with related activity to focus response.
- Automated operations: self-learning traffic profiling and automated rerouting to reduce manual effort and help prevent congestion and interference.
- Expert backing: 24/7 operations centre and threat intelligence from Imperva Research Labs.

Next steps

- Confirm your scope: list the specific websites, APIs, DNS zones, networks, and IPs that must remain available during an attack.
- Shortlist vendors: eliminate options that can't cover your full scope or can't demonstrate fast, automated mitigation.
- Validate with testing: run a controlled exercise (tabletop and technical) to confirm routing, dashboards, escalation, and rollback.
- Operationalise: integrate logging into your monitoring stack, align on incident roles, and schedule periodic reviews.

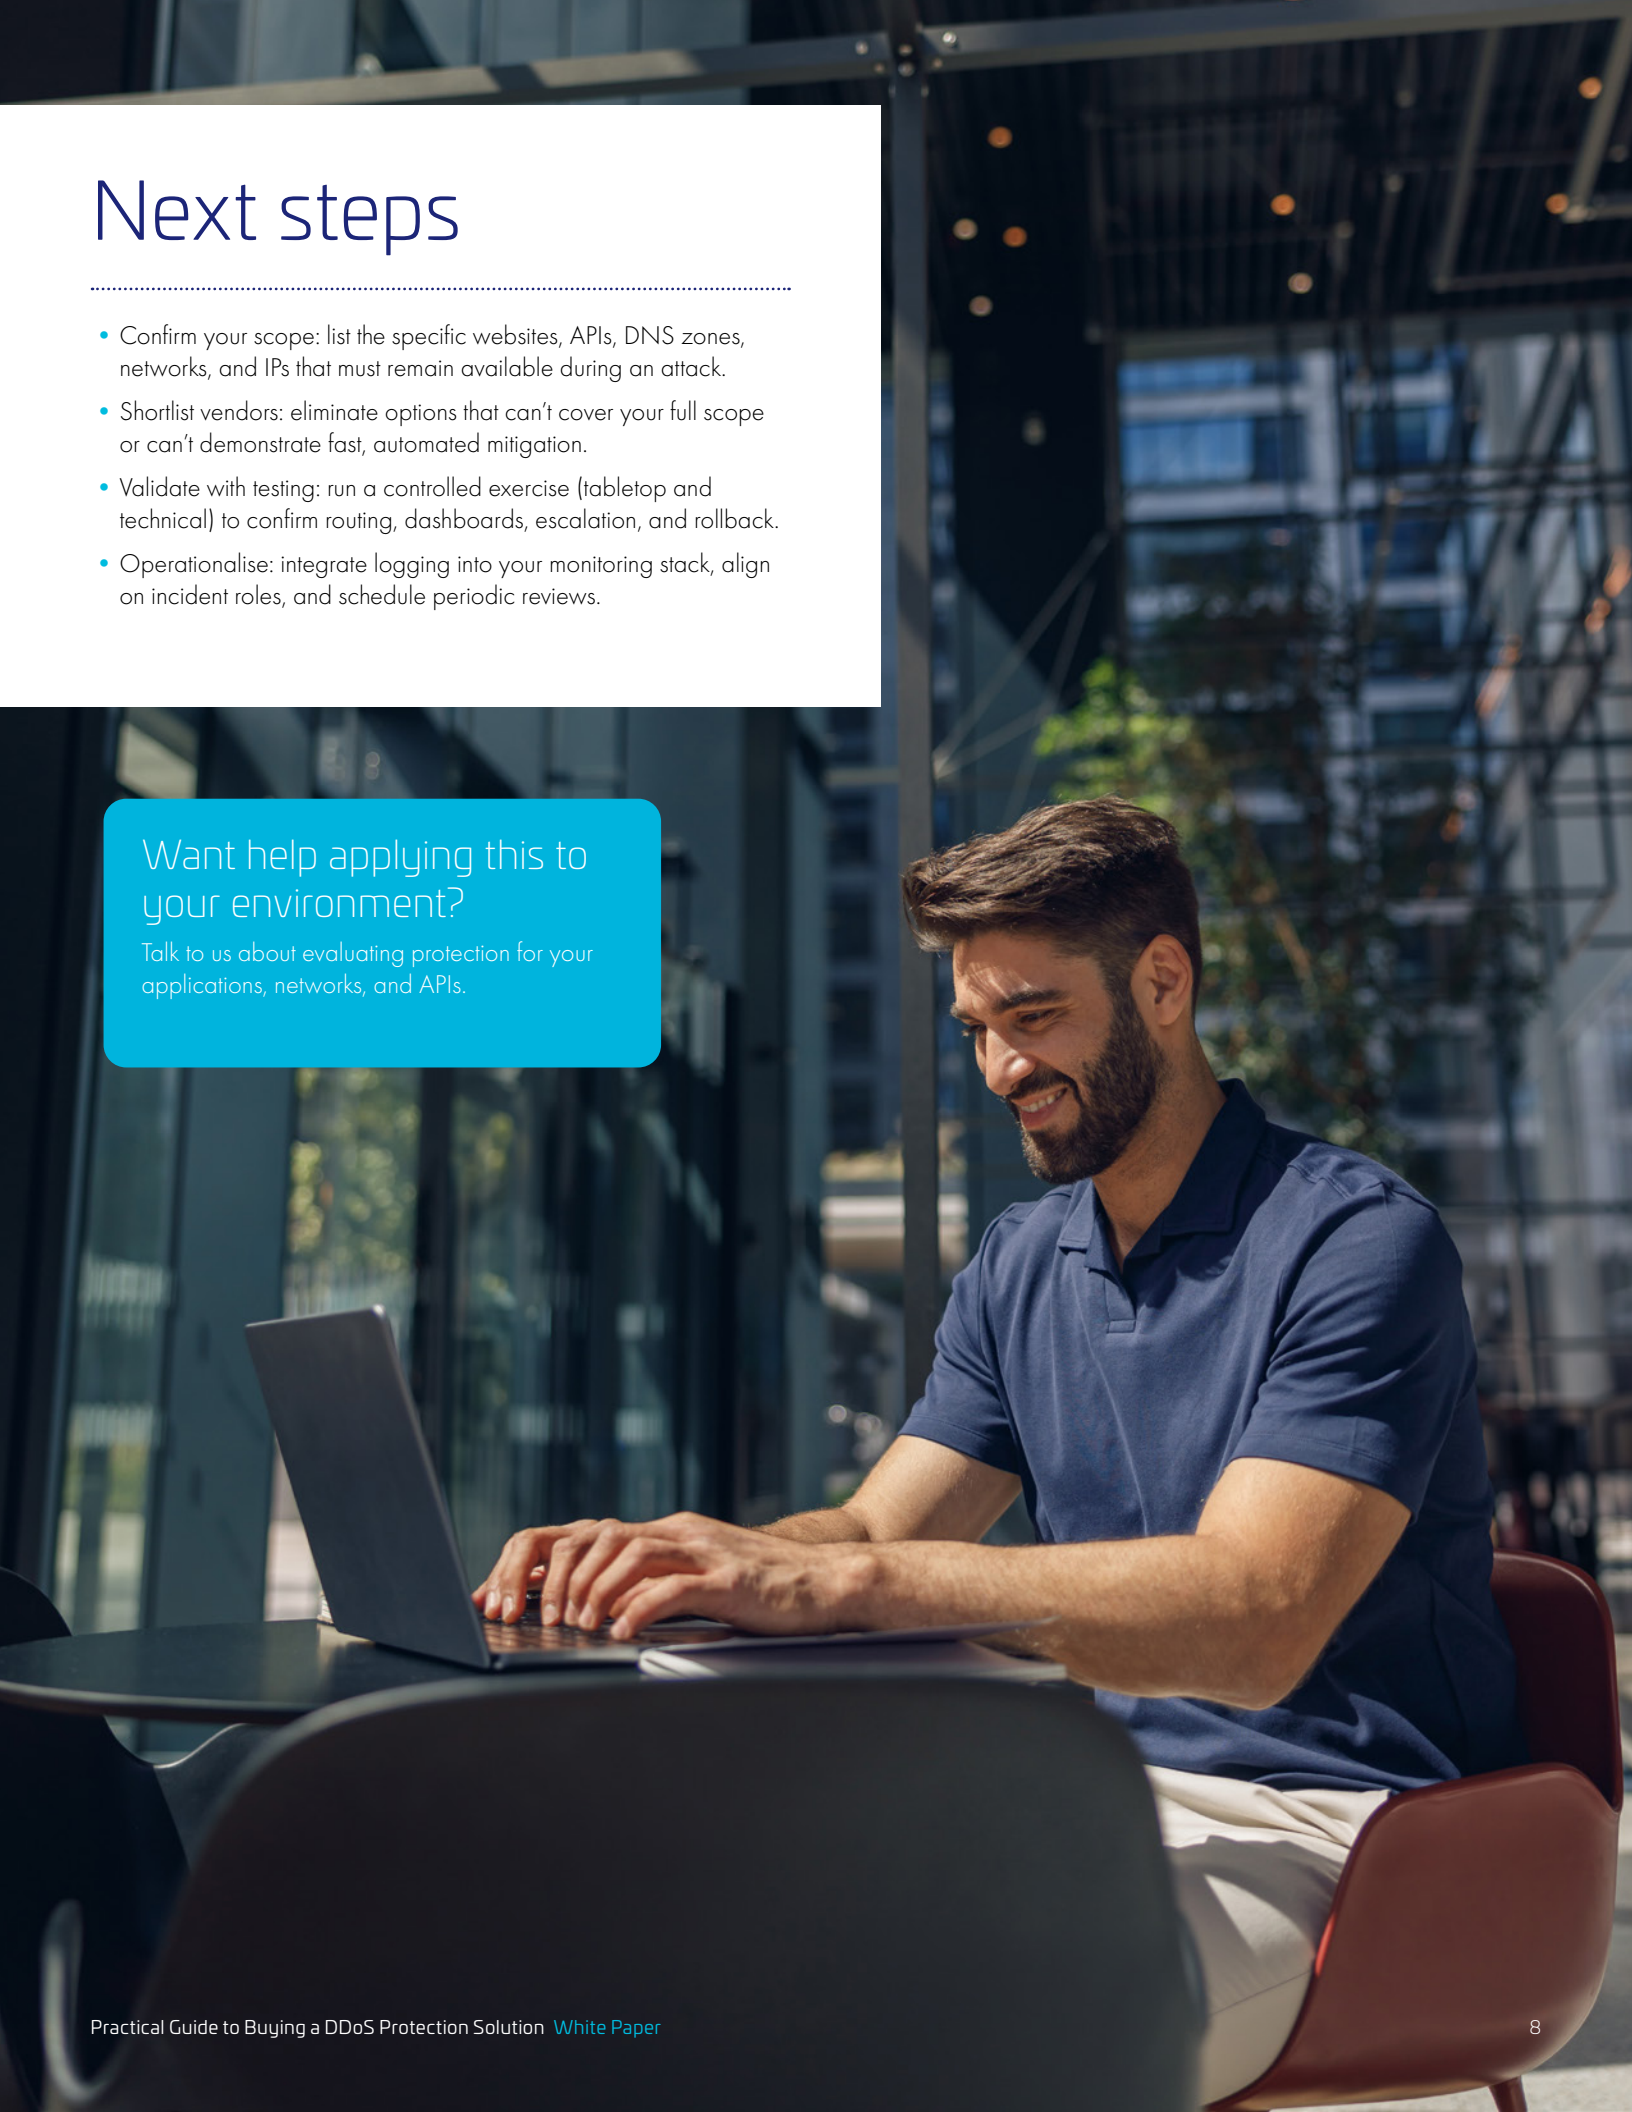

Want help applying this to your environment?

Talk to us about evaluating protection for your applications, networks, and APIs.

THALES

CYBERSECURITY

Contact us

For all office locations and contact information,
please visit cpl.thalesgroup.com/contact-us

cpl.thalesgroup.com