CAPACITY ENHANCEMENT GUIDE: VOLUMETRIC DDOS AGAINST WEB SERVICES TECHNICAL GUIDANCE

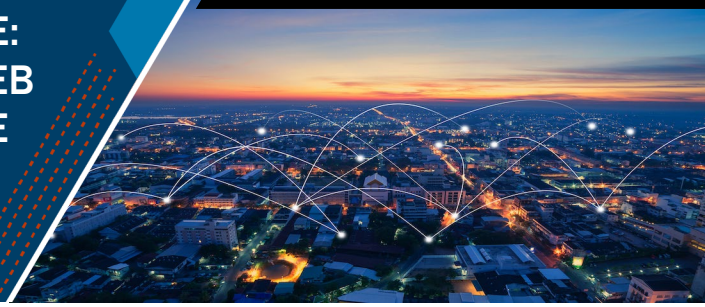

PURPOSE

The Cybersecurity and Infrastructure Security Agency (CISA) is releasing this Capacity Enhancement Guide (CEG) to support Federal Civilian Executive Branch (FCEB) agencies in making risk-informed decisions about the procurement and use of Distributed Denial of Service (DDoS) mitigations to address large-scale volumetric attacks against web services. Section 1 of this guide provides agencies with guidance to prioritize DDoS mitigations based on mission and reputational impact. Section 2 provides detailed descriptions of various DDoS mitigation services to assist agencies as they make risk-informed tradeoff decisions on how to use available resources most effectively. Although this guidance is created and intended for use by FCEB agencies, all organizations are encouraged to review and adopt these recommendations to reduce the risk of volumetric DDoS attacks.

AT-A-GLANCE RECOMMENDATIONS

- Review Section 1: Impact Analysis and consider the impact a successful DDoS attack could have on web services in your agency.
- Consider which mitigation technique will provide acceptable coverage based on risk and impact.
- Create a ranked list of services to inform agencies' use of limited resources to apply DDoS mitigations where they will be

AUDIENCE & SCOPE

CEGs support CISA's role to reduce the risk to the nation's cyber and physical infrastructure by sharing high-priority recommendations, best practices, and operational insights in response to systemic threats, vulnerabilities, and risks. This guide is designed to assist FCEB agencies in evaluating and mitigating the risk of volumetric DDoS attacks against their websites and related web services, including by informing investment decisions by agency leadership. These attacks target specific websites with the goal of exhausting the target system's resources, rendering the target unreachable or inaccessible, and denying users access to the service.

This guide addresses just one type of DDoS attack and should not be considered comprehensive to protect against all types of DDoS attacks[1].

RECOMMENDATIONS
SECTION 1: IMPACT ANALYSIS

This section provides guidance for agencies to assess the impact to their organization of a successful DDoS attack against various web services.

Agencies can use this guide to document risk decisions made in alignment with the National Institute of Standards and Technology (NIST) Risk Management Framework (RMF). For example, agencies can choose to reference this approach when conducting risk assessments[2] on DDoS enterprise risks and when validating whether chosen DDoS-related security controls[3] sufficiently address the risks to organizational operations and assets, individuals, other organizations, and the nation that prompted selection of these controls.[4] This impact analysis is provided as an example of the analysis agencies should be conducting in support of their risk management responsibilities as prescribed by the Federal Information Security Modernization Act. Agencies'

[1] Additional DDoS Guidance for Federal Agencies, CISA, October 2022
[2] NIST, "SP 800-30"
[3] NIST, "SP 800-53"
[4] NIST, "SP 800-37"

decisions should be documented, and risks managed at the appropriate level regardless of the specific risk assessment processes used.

Step 1: Inventory

Create a comprehensive inventory of all agency-owned or -operated web services, including URLs hosted on non-.gov domains.

Step 2: Analyze

Analyze the impact to the agency of a successful DDoS attack against each service identified in Step 1. For purposes of this analysis, agencies should not consider any mitigating factors that may affect the likelihood of a DDoS attack against being successful against a particular service, but only the resulting business and mission impact should an attack occur.

To determine the overall impact, agencies should assign a score in each of the five categories of impact: (1) public transactions, (2) public access to information, (3) government and industry partnerships, (4) internal agency operations, and (5) reputation. These categories are presented in no specific order, and different agencies may have different perspectives on each category's relative importance.

For each web service, agencies should consider:

- The **impact on public transactions**. Does the site enable members of the public to carry out important transactions with the federal government?
    - **4 (Very High):** The site facilitates transactions that members of the public are legally required to carry out (e.g., submitting a tax return or transactions that are critical for delivering agency's core mission). Any amount of downtime will have a tangible impact on the public's ability to complete these transactions.
    - **3 (High)**: The site facilitates transactions that are important to a widely used government service (i.e., shipping packages).
    - **2 (Moderate)**: The site facilitates public transactions, but downtime will not significantly impact members of the public.
    - **1 (Low)**: The site is not used for public transactions.

- **The impact on public access to information**. Does the public use this website to receive updated information or to understand a government service?
    - **4 (Very High):** The site provides time-sensitive information that significant portions of the American public rely upon, or delay may result in physical danger to any population (e.g., emergency alerts, weather forecasts, or statistical data).
    - **3 (High):** The site provides information valuable to the public and/or explains a crucial government service but is not time-sensitive (i.e., health guidelines).
    - **2 (Moderate)**: The site provides information that explains a government service.
    - **1 (Low)**: The site does not provide or facilitate public access to information.

- **The impact on government and industry partnerships** and collaboration activities (e.g., regulatory action, grants coordination, or joint alerting/reporting). Is access to the service needed to carry out interagency activities, or deliver products or services?
    - **4 (Very High)**: The site is needed to support interagency efforts that may be time-sensitive, mission-critical for at least one of the participating agencies, or which impact national security or public safety (i.e., coordinated public alerts).
    - **3 (High)**: The site provides information that is valuable, though not essential, to the activities coordinated between government agencies, or between government and industry.
    - **2 (Moderate)**: The site provides information that supports coordination between government agencies, or between government and industry.
    - **1 (Low)**: The site does not support coordination between government agencies, nor between government and industry.

- **The impact on activities necessary to carry** out the day-to-day operations of the agency?
  - **4 (Very High)**: System has been designated a High Value Asset or is mission-critical, or other critical systems or processes rely on this system to function.
  - **3 (High)**: Downtime will have a significant impact on the agency's ability to perform essential operations.
  - **2 (Moderate)**: Downtime will have modest impact on the agency's ability to perform essential operations.
  - **1 (Low)**: Downtime will have minimal impact to internal operations. The agency can maintain functional operations without access to this system.

- The **reputational impact** of the site's availability becoming degraded. For instance, is the site in question the primary agency website (i.e., *agency.gov*), for which a successful DDoS attack would be highly visible?

  - **4 (Very High)**: The URL is a homepage of the agency's primary website, is a designated High Value Asset with high usage from external customers or is otherwise critical for delivering the agency's mission (i.e., a High Impact Service Provider portal). Any downtime would bring reputational damage to both the agency and the federal government.
  - **3 (High)**: The webpage provides a core feature of a government service (i.e., recalls of federally regulated products) or receives a high number of visits relative to other webpages on the domain, but is not mission critical. Downtime is likely to be noticed and will contribute to a negative public perception of the agency.
  - **2 (Moderate)**: Downtime will impact the agency's reputation.
  - **1 (Low)**: Downtime will have minimal reputational risk to the agency and is unlikely to contribute to a negative public perception due to its lack of availability because the site is obscure and not associated with the government by the public.

Step 3: Calculate

Agencies should determine the relative importance, or "weight," of each impact category, based on that agency's mission and risk tolerance. Agencies that depend on public perception for the successful execution of their mission (i.e., agencies with a law enforcement function) may choose to give more weight to scores in the reputational impact category, whereas agencies that are reliant on partnership with scientific or academic organizations may choose to weight the government and industry partnerships category more heavily.

Agencies can assign a percentage value to each impact category to represent its weight, ensuring that the values add to 100%. For example, an agency that considers all impact categories to be equally important would weight each at 20%:

Impact on Public Transactions (20%)
+ Impact on Public Access to Information (20%)
+ Impact on Government and Industry Partnerships (20%)
+ Internal Impact (20%)
+ Reputational Impact (20%)
= Total Impact (100%)

Alternatively, a different agency might determine that its impact on public transactions and public access to information are more important categories and worthy of additional weight:

Impact on Public Transactions (30%)
+ Impact on Public Access to Information (30%)
+ Impact on Government and Industry Partnerships (20%)
+ Internal Impact (10%)
+ Reputational Impact (10%)
= Total Impact (100%)

Once the overall weights are determined, agencies can calculate a total impact score for each web service between 1 (Low) and 4 (Very High):

Impact on Public Transactions (Service SCORE x Category WEIGHT)
+ Impact on Public Access to Information (Service SCORE x Category WEIGHT)
+ Impact on Government and Industry Partnerships (Service SCORE x Category WEIGHT)
+ Internal Impact (Service SCORE x Category WEIGHT)
+ Reputational Impact (Service SCORE x Category WEIGHT)
= **Total** (SCORE)

Step 4: Prioritize

After calculating a total impact score for each web service, agencies should create a ranked list in order of impact of a successful DDoS attack. This ordered list can be used to prioritize the implementation of DDoS mitigation methods discussed in Section 2. This list should also inform risk management decisions within the Department, including formal documentation of any risk acceptance by the appropriate authority, if the impact analysis and risk mitigations deployed indicate significant residual risk.

Agencies should use this methodology as a guide, customizing it as necessary for their own missions and operating environments. Agencies may find that they need added impact categories to fully capture the impact of a DDoS attack on their services, or that a different approach to calculating impact scores produces results that more closely match their organizational priorities. The goal for this process is to produce a ranked list of services to inform agencies' use of limited resources to apply DDoS mitigations where they will be most valuable to the organization.

SECTION 2: RISK MITIGATIONS

This section provides technical guidance for agencies to consider when mitigating DDoS attacks on web services. When considering protections that can be applied to web servers and web applications, Section 1 enabled agencies to conduct an impact analysis to inform the prioritization of protections for various assets.

DDoS protections can vary in cost and capability, with some protections providing more coverage and guaranteed availability than others. This section compares various approaches to mitigating DDoS attacks, so agencies can select the appropriate mitigation methods.

Content Delivery Network (CDN)

Service Overview

A properly provisioned CDN can effectively mitigate a volumetric DDoS attack against web services. CDNs are uniquely positioned between the agency web server workload and the end user on the internet, allowing them to function as an agency's first line of defense against DDoS attacks. Properly configured CDNs can provide website protection against DDoS attacks at the application level. CDNs can be implemented without procuring any additional hardware, as they reside in geographically distributed cloud service provider data centers. In addition to DDoS mitigation, CDNs cache website content (e.g., files, JavaScript, images, and videos) at data centers around the globe. CDNs reduce bandwidth use charges associated with website hosting and result in lower latency, faster page load times, and improved overall end user experience.

Technical Considerations

Designed to handle large amounts of traffic in a distributed manner, CDNs are typically the most effective mitigation against DDoS attacks. For high-value and high-visibility internet-facing web assets, federal agencies should make use of CDNs. Regardless, agencies should evaluate a CDN's advertised defense capabilities against DDoS attacks, as the primary use case of a CDN service may be geared towards optimizing the delivery of website content. As a result, not all CDNs may necessarily be prepared to mitigate all types of DDoS attacks. Other mitigation options are identified below.

Certain CDN services may provide real-time response to attempted DDoS attacks to avoid any perceptible downtime, while others place the burden on the customer to detect a DDoS attack and notify the CDN. Agencies should default to CDNs that can detect and automatically begin responding to DDoS attacks, to avoid unnecessary downtime.

Additional CDN features:

- Capable of rate limiting by counting web requests and blocking potentially malicious source IPs.
- Scales to meet demand.
- Uses DNS or Anycast based mapping to forward requests to the global CDN node geographically closest to the public user requesting the resource, potentially improving the user experience.
- Automatically redirects requests to the next nearest node if degraded performance on any one node.

Cost Factors

Pricing for CDN services can vary widely based on the nature of the agency's public-facing services, including the geographic location of the agency's users, data stored and transmitted by the service or application, and the agency's tolerance for downtime.

Agency Next Steps

- Discuss DDoS mitigation capabilities available with existing agency CDN services and associated costs.
- Modify contracts as appropriate to incorporate DDoS-related CDN services.
- Exercise DDoS protections at a minimum annually, consistent with the requirements of any applicable CDN Terms of Service agreements.

Internet Service Providers (ISP) & Upstream Providers

Service Overview

ISPs and upstream providers (including packet scrubbing solutions) with sufficient compute and bandwidth may mitigate some volumetric DDoS attacks against agency web services as well as non-web services. This capability should be validated, as not all upstream providers are resilient against this category of attack. ISPs provide "last mile" internet transport to an agency's data centers (i.e., they can redirect internet traffic inbound to an agency's network when malicious activity is detected). If an agency hosts websites in a legacy data center or on-premises environment or is concerned about the high availability of a mission critical website, they should consider the DDoS protection services offered by their ISP. ISP may offer protection against the following types of attacks in particular:

- Open Systems Interconnection (OSI) Layer 3: Typically, a volumetric attack in which the attacker sends thousands of packets to an agency's public IP addresses.
- OSI Layer 4: Often a TCP SYN packet flood attack designed to overwhelm a server.
- OSI Layer 7: HTTP GET or POST requests, or other maliciously crafted application layer request.

Services providing the most uptime in offered Service Level Agreements (SLAs) often have the highest associated costs. Agencies may achieve cost savings through accepting the risk of downtime caused by a DDoS attack by using an on-call mitigation service. This means that once an agency Network Operations Center (NOC)/Security Operations Center (SOC) has detected an attack, the ISP is engaged to begin DDoS mitigation procedures. In this on-call model, the agency web resource under attack will be down or impaired until the agency NOC/SOC engages the provider to start mitigation efforts.

This model requires some upfront configuration and planning on the part of both the agency and the ISP; however, it is significantly less expensive than an always-on model where all traffic is inspected regardless of whether an active attack is occurring. Also, beyond *scrubbing* the traffic, various services may use other technical approaches, including *null routing* (or black hole routing) which can create a high rate of false positives, *sink holing* which uses a blocklist of known malicious source IPs, and *IP masking* which hides the origin IP of your server. Agencies should consider subscribing to reputation-based source IP threat intelligence feeds if ISPs offer this as an option with DDoS protection managed services.

Technical Considerations

Agencies should select a provider that has the capacity to scale and withstand large volumetric DDoS attacks. Agencies should also understand their role and the role of the provider if targeted by a DDoS attack**.** Note the two consumption models previously identified for DDoS mitigation services – always-on and on-call/on-

demand. In an always-on model, all traffic always passes through the mitigation provider's service (which may add latency if the distance between the customer internet circuit and mitigation service are high). Always-on can provide instant protection, but agencies should always validate time-to-mitigation of any proposed solution. The on-demand consumption model only sends traffic to scrubbing centers when directed to do so via human intervention during an attack. Agencies must communicate with their provider to understand which protections are available, the protections that are included in the existing contracts, and those offered à la carte. For services that require manual activation, agencies must understand each organization and individual's roles, as well as develop, maintain, and test the activation procedures for best response.

As with CDNs, agencies should default to mitigations that automatically provide protections (i.e., always-on)

Additional features of provided services:

- DDoS protection is provided by Domain Name System (DNS) re-routing or Border Gateway Protocol (BGP) re-advertisement of a public facing customer URL or subnet. DNS or BGP reconfiguration and advertisement push all agency web traffic to an ISP-managed cloud hosted scrubbing center where DDoS attack traffic is discarded or "scrubbed," and only clean traffic is returned to the customer network, removing resource constraints from the agency servers resulting from the attack.
- Scrubbing facilities are typically geographically distributed to reduce latency between customer edge (traffic destination) and any redirection for scrubbing.
- ISPs may leverage a combination of scrubbing services and inline network sensors at the customer edge to mitigate attacks.

Cost Factors

Pricing varies and is based on bandwidth, number and size of protected IP Prefixes, number of connections between scrubbing center and customer network, and number of individual IPs or DNS names to be protected.

Agency Next Steps

- Work with ISPs or upstream providers (including packet scrubbers) to fully understand the Service Level Agreements (SLA) and the roles of both parties within the agreement. Verify the SLA meets the required services needed to mitigate DDoS attacks, maintain functional operations and any other specific requirement as it pertains to the agency.
- Update contracts with providers to include provisions where 'manual activations' are required, additional fees, etc.
- Update Continuity of Operations Planning (COOP) and Disaster Recovery Plans (DRP) to ensure 'manual activations' are included, and activation procedures and points of contact (POC) are documented.
- Exercise DDoS protection readiness at the minimum annually, consistent with the requirements of any applicable ISP Terms of Service agreements.

Cloud Service Provider (CSP) Hosted Services

Service Overview

Cloud Service Providers with sufficient compute and bandwidth can provide effective mitigations to a volumetric DDoS against web services. DDoS protections are offered as managed services by CSPs (e.g., AWS, Azure, GCP, and Oracle). Agencies should consider CSP hosted DDoS protection services when public facing resources are accessible through CSP provided internet access. A key advantage of CSP offerings is their ability to scale out dynamically to meet demand – which means customers are not paying for dedicated hardware to support DDoS prevention. If an agency can accept some amount of downtime risk from a DDoS attack, the on-call/on-demand protection services detailed in the ISP section above may be acceptable.

Agencies should always consider automated protections, versus protections that require agency personnel to manually initiate DDoS protections.

Technical Considerations

If agencies have multiple points of ingress to their publicly hosted web sites, they should consider CSP offerings to protect web resources accessible via CSP internet circuits. Evaluation of risk and mitigating service offerings available at other ingress points should not be overlooked when building a defense in depth DDoS mitigation strategy. DDoS protection services may not be included in existing CSP contractors. Agencies must communicate with their CSPs to understand which protections are available, the protections that are included in the existing contracts, and those that are offered à la carte. For services that require manual activation, agencies must develop, maintain, and test the activation procedures and are advised to generate a well-developed and documented standard operating procedure.

Additional features of CSP hosted solutions:

- CSP DDoS mitigation solutions typically employ a combination of CDN and WAF (Web Application Firewall) as a managed service to protect web facing resources.
- Can be automatically configured to protect against DDoS attacks based on known good traffic patterns via machine learning.
- Agencies should consider subscribing to reputation-based source IP threat intelligence feeds if CSP's offer this as a la carte with DDoS protection managed services.
- Rate based blocking services are configurable.

Cost Factors

Pricing is often a flat monthly charge which includes a set number of public IP resources with overage charges if additional resources need to be protected. Some CSPs also charge for data transfer related to DDoS mitigation services.

Agency Next Steps

- Work with CSPs to fully understand built in protections in their services and what is available by default vs. what is à la carte vs. what needs 'manual' activation, scope, limitations, SLAs, etc.
- Update contracts with CSPs to include provisions where 'manual activations' are required, additional fees for traffic surges, etc.
- Update COOPs and DRPs to ensure 'manual activation' procedures are included, and POCs are documented.
- Exercise DDoS protection readiness at the minimum annually, consistent with the requirements of any applicable CSP Terms of Service agreements.

On-Premises Solutions

Service Overview

On-premises solutions are highly unlikely to have sufficient compute and bandwidth to provide effective mitigation of a volumetric DDoS against web services. On-premises solutions are unable to scale and handle large volumetric DDoS attacks in the same way as previously discussed mitigations.

Exclusive reliance on on-premises solutions represents acceptance of the risks associated with having no protection against DDoS and should only be considered for the lowest impact URLS as described in Section 1 with documented risk acceptance as part of a website's authorization to operate (ATO).

Technical Considerations

A major disadvantage to hardware or virtualized on-premises solutions is their inability to scale to meet high volume DDoS attacks. The previously discussed CDN, ISP, and CSP solutions offer DDoS protections at a much greater scale to better handle the largest known DDoS attacks, while on-premises solutions are limited to the configured local internet circuit bandwidth and any associated hardware limitations. Note that industry leading firewall vendors do provide limited DDoS protections, yet still recommend subscribing to additional DDoS services as described above in the ISP or CSP sections for the most complete protection.

Agencies should consider subscribing to publicly available threat intelligence feeds. CISA's Shared Cybersecurity Services (SCS) provides federal civilian agencies with no-cost access to commercial Cyber Threat Intelligence. CISA also offers the Automated Indicator Sharing (AIS) feed for agencies. These feeds provide reputation-based intelligence about source IPs including geographic location, organization name, known/past attack types from the source IP, service type (e.g., anonymous proxies, phishing sites), a risk scoring mechanism, etc. Next Generation Firewalls and WAF appliances can translate this feed data into dynamic access control lists (ACLs) that block any processing of traffic from malicious source IPs based on this reputation data. Agencies should confirm any new and existing hardware at the internet edge is resilient enough to support this capability as part of due diligence.

Additional features of an on-premises solution:

- Appliance based solutions are not scalable and face hardware limitations when processing DDoS attack traffic locally.
- Some Next Generation firewall vendors offer DDoS protection either as a separately licensed module or included in a standard license. These protections are constrained by the bandwidth and compute resources available for on-premises solutions and will likely not scale to meet this category of attack.
- Web Application Firewalls (WAFs) can more closely inspect communications between website users and servers to identify malicious intent. Technologies used include device fingerprinting, detection of SQL injection attacks, cross-site scripting, and other customer queries which can compromise an agency web resource. Hardware and circuit bandwidth are limiting factors in a WAF solution.

Cost Factors

Typically, these on-premises solutions come at a lower cost than those with remote scrubbing. Vendors often sell licensing by bandwidth protection scope; prices vary by vendor.

Agency Next Steps

- Work with hardware vendors to fully understand built in protections/limitations in their services and what is available by default vs. what is a la carte vs. what needs 'manual' activation, scope, limitations, SLAs, etc.
- Update contracts with hardware vendors to include provisions where 'manual activations' are required, additional fees for traffic surges, etc.
- Update COOPs and DRPs to ensure 'manual activation' procedures are included, and POCs are documented.
- Exercise DDoS protection readiness at the minimum annually.

SUMMARY

This technical guidance provides four categories of impact metrics to be used when evaluating the impact of volumetric DDoS attacks on agency web services: **(Very High, High, Moderate, and Low)**. The impact analysis within this document is an example of the risk assessment/analysis that agencies should conduct in support of risk management strategies consistent with NIST RMF and the appropriate security controls.

Risk mitigation technologies vary with their effectiveness against volumetric DDoS attacks. This guidance provides agencies with a technical overview of existing mitigation technologies currently used to mitigate volumetric DDoS attacks. CDN mitigations provide the highest degree of protections. Both ISP and CSP are sufficient if service providers can provide the proper compute and bandwidth resources. On-premises solutions are highly unlikely to provide sufficient compute and bandwidth due to its inability to scale; CDN solutions are highly advised.

REPORTING

Agencies should follow all relevant CISA protocols and OMB guidance when reporting events, incidents, breaches, and major incidents. This includes CISA's current Federal Incident Notification Guidelines, CISA's Federal Government Cybersecurity Incident and Vulnerability Response Playbooks, OMB M-22-05 and OMB M-17-12, and Presidential Policy Directive 41 (PPD-41).

CONTACT INFO

For questions about this guidance and other CISA services available to federal agencies, please contact cyberliaison@cisa.dhs.gov.

RESOURCES

- National Institute of Standards and Technology (2012), SP 800-30 Rev. 1: Guide for Conducting Risk Assessments, Department of Commerce, Washington, D.C., September 2012
- National Institute of Standards and Technology (2020) SP 800-53 Rev 5: Security and Privacy Controls for Information Systems and Organizations, Department of Commerce, Washington, D.C., September 2020
- National Institute of Standards and Technology (2018) SP 800-37 Rev 2: Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy, Department of Commerce, Washington, D.C., December 2018
- Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, and Multi-State Information Sharing and Analysis Center, Understanding and Responding to Distributed Denial-of-Service Attacks, October 2022, Washington, D.C.