RECOMMENDED PRACTICES

The F5 DDoS Playbook: Ten Steps for Combating DDoS in Real Time

To the uninitiated, a distributed denial-of-service (DDoS) attack can be a scary, stressful ordeal. But don't panic. Follow these steps to successfully fight an attack.

David Holmes
Senior Technical Marketing Manager, Security

Contents

Introduction

Distributed denial-of-service (DDoS) attacks are a top concern for many organizations today. A DDoS attack saturates a website, renders its services inoperable, and prevents legitimate clients from being able to connect to it. For the uninitiated, this attack can be a scary and stressful ordeal.

DDoS attacks are usually coordinated across a large number of client computers (which may have been set up for that purpose). Even more likely, a client computer may have been infected with a virus that allows an attacker to remotely control the computer, making it participate in the attack.

DDoS attack frequency

Both financially and politically motivated, DDoS attacks are becoming more prevalent. Although a first attack can happen randomly, it often occurs when an attacker with specific knowledge of your high-value service decides to take it offline. This can cause panic and instigate costly decisions, including the payment of ransom, to triage and stop the attack.

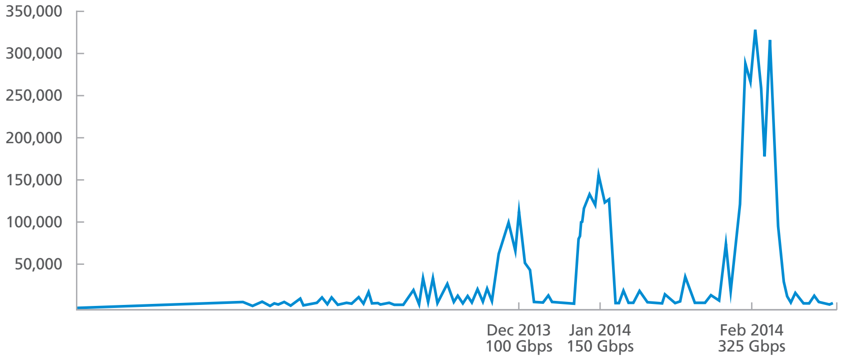

Figure 1: Volumetric attacks increased sharply in 2014.

An objective DDoS combat method

Organizations that have defended against multiple DDoS attacks understand the importance of having a procedural method to assist in combating them.

What is their solution? A DDoS Playbook. This document can be the basis for a procedural document that guides an operations team through DDoS attacks, large or small, frequent or infrequent. The five Quick Reference sheets enclosed, *when completed in advance*, will assist you in repelling a DDoS attack.

- Quick Reference 1: Contact List. Fill this out as you initiate contacts (page 19).
- Quick Reference 2: Whitelists. Map your partners, users, and services (page 20).
- Quick Reference 3: Application Triage. Know your own applications (page 21).
- Quick Reference 4: Device Map. Create a device map (page 22).
- Quick Reference 5: Attack Log. Note the attack details (page 23).

The completed references can be kept in your data center and used for documentation and attack mitigation. If you have not recorded this information prior to your first attack, record it as you collect it to better prepare for a future attack.

Regulatory compliance

Your organization may be subject to regulatory statutes that require a level of reporting around cyber-attacks, breaches, or even DDoS attacks. Quick Reference 5, the Attack Log, can assist you in this situation, as you can track and refer to the log later during the reporting process.

Preparing for a DDoS Attack

If you are fortunate enough to be reading this document prior to being attacked, there are steps that you can take now to make your applications, networks, and processes DDoS-resilient.

Study a DDoS-resilient architecture

After you have filled out the quick reference sheets in this playbook, obtain the *F5 DDoS Recommended Practices* document so you can consider how to align your network architecture defenses.

F5 recommends a multi-tiered approach where layer 3 and layer 4 DDoS attacks are mitigated at the network tier with firewalls and IP reputation databases. See Figure 2.

- The application tier handles high-CPU security functions such as SSL termination and web-application firewall functionality.
- To combat DDoS, modern organizations need a cloud-based DDoS scrubbing tier. These service offerings can scrub hundreds of gigabytes per second and return “clean” traffic to the data center.
- DNS is handled in the DMZ and partially protected by the network tier.

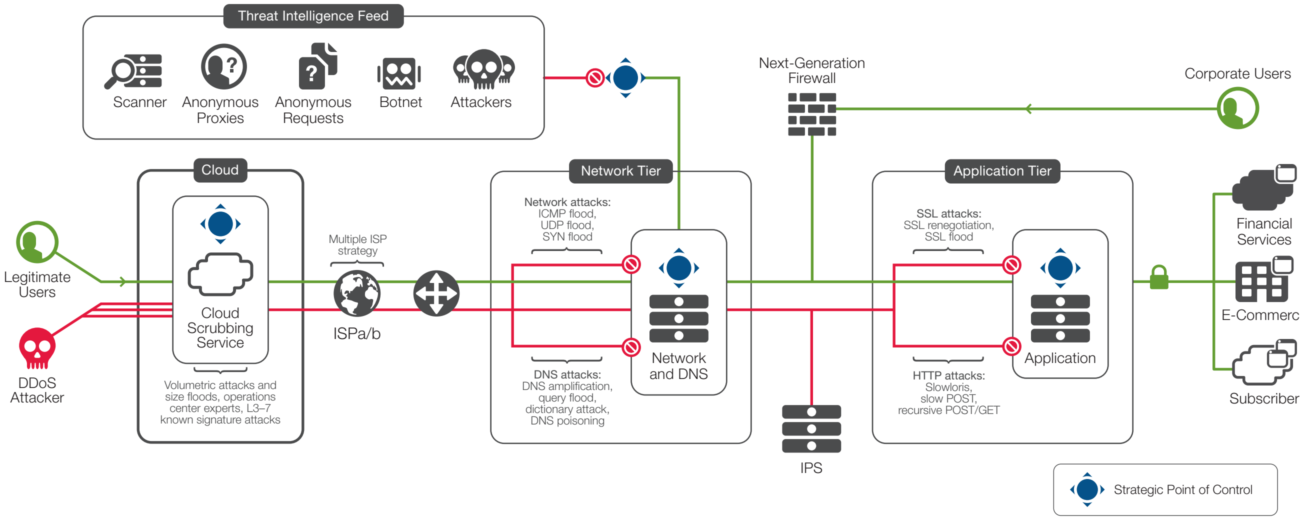
Threat Intelligence Feed
Scanner
Anonymous Proxies
Anonymous Requests
Botnet
Attackers
Next-Generation Firewall
Corporate Users
Cloud
Cloud Scrubbing Service
Volumetric attacks and size floods, operations center experts, L3–7 known signature attacks
Legitimate Users
DDoS Attacker
Multiple ISP strategy
ISPa/b
Network Tier
Network attacks: ICMP flood, UDP flood, SYN flood
DNS attacks: DNS amplification, query flood, dictionary attack, DNS poisoning
Network and DNS
IPS
Application Tier
SSL attacks: SSL renegotiation, SSL flood
HTTP attacks: Slowloris, slow POST, recursive POST/GET
Application
Financial Services
E-Commerce
Subscriber
Strategic Point of Control

Figure 2: F5 recommends a multi-tiered DDoS approach to your architecture.

This multi-tiered approach can:

- Defeat TCP connection floods.
- Overcome SNAT port exhaustion.
- Turn back SSL floods.

These are just a few of the recommended practices and considerations in the comprehensive *F5 DDoS Recommended Practices* document.

DDoS Mitigation Steps

If you appear to be suffering a volumetric attack, it can help to have a historical sense of your own traffic patterns. Keep a baseline of normal traffic patterns to compare against.

If you have determined that you are under a DDoS attack, record the estimated start time. (See Quick Reference 5: Attack Log)

Monitor volumetric attacks. Remember to keep a monitoring web page open to indicate when the attack may be over (or mitigated).

You will need to follow (up to) 10 steps for your DDoS mitigation:

- Step 1: Verify the attack.
- Step 2: Contact team leads.

- Step 3: Triage applications.
- Step 4: Identify the attack.
- Step 5: Protect remote users and partners.
- Step 6: Evaluate source address mitigation options.
- Step 7: Mitigate specific application attacks.
- Step 8: Increase application-level security postures.
- Step 9: Constrain resources.
- Step 10: Manage public relations.

Step 1: Verify the attack

Not all outages are caused by a DDoS attack. DNS misconfiguration, upstream routing issues, and human error are also common causes of network outages. You must first rule out these types of non-DDoS attacks and distinguish the attack from a common outage.

Rule out common outages.

The faster you can verify the outage is a DDoS attack, the faster you can respond. Even if the outage was not caused by a misconfiguration or other human error, there may still be other explanations that resemble a DDoS attack.

For instance, the Slashdot Effect occurs when a particular page on your site is featured on a very popular forum or blog. Your investigation must rule out such possibilities.

Check outbound connectivity.

Is there outbound connectivity? If not, then the attack is so severe that it is congesting all inbound and outbound traffic. Check with your usual diagnostic tools (such as traceroute, ping, and dig) and rule out all such possibilities.

Rule out global issues.

Check the following Internet weather reports to determine if the attack is a global issue:

- Internet Health Report
- Internet Traffic Report

Check external network access.

Attempt to access your application from an external network. Services and products that can perform this kind of monitoring include:

- Keynote testing and monitoring.
- HP SiteScope agentless monitoring.
- SolarWinds NetFlow Traffic Analyzer.
- Downforeveryoneorjustme.com.

Confirm DNS response.

Check to see if DNS is responding for your website. The following UNIX command resolves a name against the OpenDNS project server.

```
% dig @208.67.222.222 yourdomain.com
```

Step 2: Confirm the DDOS attack

Contact team leads.

Once the attack is verified, contact the leads of the relevant teams. If you have not previously filled out Quick Reference 1: Contacts List, fill it out now.

When an outage occurs, your organization may hold a formal conference call including various operations and applications teams. If your organization has such a process, use the meeting to officially confirm the DDoS attack with team leads.

Contact your bandwidth service provider.

One of the most important calls you can make is to the bandwidth service provider. The number for your service provider should be listed in Quick Reference 1: Contacts List. The service provider can likely confirm your attack, provide information about other customers who might be under attack, and sometimes offer remediation.

Contact your fraud team.

It is especially important to *invoke the fraud team as soon as the attack is verified*. DDoS attacks can be used as cover to hide an infiltration. Logs that would normally show a penetration may get lost during a DDoS attack. This is why high-speed, off-box logging is so important.

Step 3: Triage applications

Once the attack is confirmed, triage your applications.

When faced with an intense DDoS attack and limited resources, organizations have to make triage decisions. High-value assets typically generate high-value online revenue. These are the applications you will want to keep alive.

Low-value applications, regardless of the level of legitimate traffic, should be purposefully disabled so their CPU and network resources can be put to the aid of higher-value applications. You may need the input of team leads to do this.

Ultimately, these are financial decisions. Make them appropriately.

Quick Reference 3: Application Triage takes only a few minutes to fill out, and it will greatly assist you in making tough application decisions while combating an actual DDoS event. If you have not done this yet, now is the time.

Decide which applications are low priority and can be disabled during the attack. This may include internal applications.

Record your choices in Quick Reference 3.

Step 4: Protect partners with whitelists

Whitelist partner addresses.

Very likely you have trusted partners who must have access to your applications or network. If you have not already done so, collect the IP addresses that must always be allowed access and maintain that list. Quick Reference 2: Whitelists includes a template for your whitelist collection.

You may have to populate the whitelist in several places throughout the network, including at the firewall, the Application Delivery Controller (ADC), and perhaps even with the service provider, to guarantee that traffic to and from those addresses is unhindered.

Protect VPN users.

Modern organizations will whitelist or provide quality-of-service for remote SSL VPN users. Typically this is done at an integrated firewall/VPN server, which can be important if you have a significant number of remote employees.

Step 5: Identify the attack

Determine the nature of the attack.

Now is the time to gather technical intelligence about the attack. The first question you need to answer is "What are the attack vectors?"

Four DDoS attack types

You are trying to determine the nature of the attack. Is it:

- Volumetric—flood-based attacks that can be at layers 3, 4, or 7?
- Asymmetric—designed to invoke timeouts or session-state changes?
- Computational—designed to consume CPU and memory?
- Vulnerability-based—designed to exploit software vulnerabilities?

By now you should have called your bandwidth service provider with the information on Quick Reference 1: Contacts List. If the attack is solely volumetric in nature, the service provider will have informed you and may have already taken steps at DDoS remediation.

Even though well-equipped organizations use existing monitoring solutions (such as NetScout) for deep-packet captures, you may encounter cases where you have to use packet captures from other devices, such as the ADC, to assist in diagnosing the problem. These cases include:

- SSL attack vectors. If the attack is launched over SSL, there may be no other way to diagnose it other than at the ADC. Capture the packet streams either at the ADC or elsewhere, and then use the ssldump utility to decrypt the stream file.
- FIPS-140. If your ADC is using a FIPS-140 hardware security module (HSM), then you can often still use ssldump to decode the file capture.
- Use of a mirror-port or clone pool. One way to capture packets is to mirror them from the ADC. This high-performance method allows data to flow through the ADC and also to an external device without interruption.

Step 6: Evaluate source address mitigation options

If Step 5 has identified that the campaign uses advanced attack vectors that your service provider cannot mitigate (such as slow-and-low attacks, application attacks, or SSL attacks), then the next step is to consider the following question: “How many sources are there?”

If the list of attacking IP addresses is small, you can block them at your firewall. Another option would be to ask your bandwidth provider to block these addresses for you.

Geoblocking

The list of attacking IP address may be too large to block at the firewall. Each address you add to the block list will slow processing and increase CPU. But you may still be able to block the attackers if they are all in the same geographic region or a few regions you can temporarily block.

For example, if the majority of your attacks appear to be coming from Southeast Asia, evaluate the revenue you will lose if you block all traffic from that region. Be deliberate about geoblocking. The decision to block entire regions via geolocation must be made as a business decision.

Finally, if there are many attackers in many regions, but you don’t care about any region except your own, you may also use geolocation as a defense by blocking all traffic except that originating from your region.

Mitigating multiple attack vectors

If there are too many attackers to make blocking by IP address or region feasible, you may have to develop a plan to unwind the attack by mitigating “backwards”—that is, defending the site from the database tier to the application tier, and then to the web servers, load balancers, and finally the firewalls.

You may be under pressure to remediate the opposite way—for example, mitigating at layer 4 to bring the firewall back up. However, be aware that as you do this, attacks will start to reach further into the data center.

As you identify the different mix of attack vectors, check this table for remediation specific to individual attacks.

Attack Vector	Firewall	On-Premises DDoS	Application Delivery Controller	Cloud Scrubber
SYN Flood	×	×	×	×
ICMP Flood	×	×	×	×
UDP Flood	×	×	×	×
TCP Flood			×	×
DNS Flood		×	×	×
Apache Killer		×	×	
Slowloris		×	×	
Keep Dead		×	×	
HTTP Recursive GET		×	×	

Step 7: Mitigate specific application-layer attacks

If you have reached this step, the DDoS attack is sufficiently sophisticated to render mitigation by the source address ineffective. Attacks that fall into this category may be generated by tools such as the Low Orbit Ion Cannon, the Apache Killer, or the Brobot. These attacks look like normal traffic at layer 4, but have anomalies to disrupt services in the server, application, or database tier. (To learn about common attack tools and mitigation strategies, see “The Taxonomy of Application Attacks” in the *F5 DDoS Recommended Practices* document.)

To combat these attacks, you must enable or construct defenses at the application delivery tier.

Mitigating specific attack tools

Once you have analyzed the traffic in Step 4, if the attack appears to be an application-layer attack, the important questions are:

- Can you identify the malicious traffic?
- Does it appear to be generated by a known attack tool?

Specific application-layer attacks can be mitigated on a case-by-case basis with specific F5 countermeasures. Attackers today often use multiple types of DDoS attack vector, but most of those vectors are around layers 3 and 4, with only one or two application-layer attacks thrown in. We hope this is the case for you, which will mean you are nearly done with your DDoS attack.

Step 8: Increase application-level security posture

If you have reached this step in a DDoS attack, you've already mitigated at layers 3 and 4 and evaluated mitigations for specific application attacks, and you are still experiencing issues. That means the attack is relatively sophisticated, and your ability to mitigate will depend in part on your specific applications.

Asymmetric application attack

Very likely you are being confronted with one of the most difficult of modern attacks: the asymmetric application attack. This kind of attack can be:

- A flood of recursive GETs of the entire application.
- A repeated request of some large, public object (such as an MP4 or PDF file).
- A repeated invocation of an expensive database query.

If you implemented some of the architectural recommendations discussed in the introduction, you may be able to make use of those defenses now.

Leveraging your security perimeter

The best defense against these asymmetric attacks depends on your application. For example, financial organizations know their customers and are able to use login walls to turn away anonymous requests. Entertainment industry applications such as hotel websites, on the other hand, often do not know the user until the user agrees to make the reservation. For them, a CAPTCHA might be a better deterrent.

Choose the application-level defense that makes the most sense for your application:

- A login wall
- Human detection
- Real browser enforcement

Login walls

A login wall is a logical defense that requires a client to be logged in as a known user before that client can access any high-value asset or run a database query. Login walls can be implemented at the level of the service provider, a web application firewall, or an ADC.

The drawback to this otherwise perfect solution is that not every application has tight integration with known users. For example, hoteliers must serve room availability applications that do not require the user to log in.

Human detection

Human detection is the second-best approach. Validating that the client connection is being controlled by a human (instead of a malicious bot) can go a long way to turning back a layer 7 DDoS attack. Usually this is done with a CAPTCHA.

CAPTCHA is an acronym for Completely Automated Public Turning test to tell Computers and Humans Apart. It is a challenge used in computing to tell whether or not the requesting entity is human by requiring a specific response. The drawback to CAPTCHAs (and the reason that they do not protect every resource all the time) is that they will turn away some percentage of legitimate users. Flexible applications will allow CAPTCHAs to be turned on during an attack and then off again afterward.

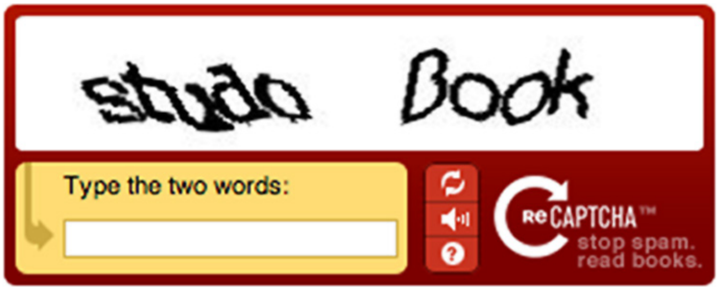

Figure 3: A typical CAPTCHA can help turn back a layer 7 attack.

Real browser enforcement

Some web application firewalls provide this functionality by inserting a JavaScript redirect to new connections and then blacklisting them if they do not follow the redirect. This is a worthwhile approach because it foils the majority of bots without interfering with real users using real browsers.

Step 9: Constrain resources

If all the previous steps fail to stop the DDoS attack, you may be forced to simply constrain resources to survive the attack.

This technique turns away both good and bad traffic. In fact, rate limiting often turns away 90 to 99 percent of desirable traffic while still enabling the attacker to drive up costs at your data center. For many organizations, it is better to just disable or "blackhole" an application rather than rate-limit it.

Rate shaping

If you find that you must rate-limit, you can provide constraints at different points in a multi-tier DDoS architecture. At the network tier, where layer 3 and layer 4 security services reside, use rate shaping to prevent TCP floods from overwhelming your firewalls and other layer 4 devices.

Connection limits

Connection limits can be an effective mitigation technique, but they do not work well with connection-multiplexing features. Application tier connection limits should provide the best protection to prevent too much throughput from overwhelming your web servers and application middleware.

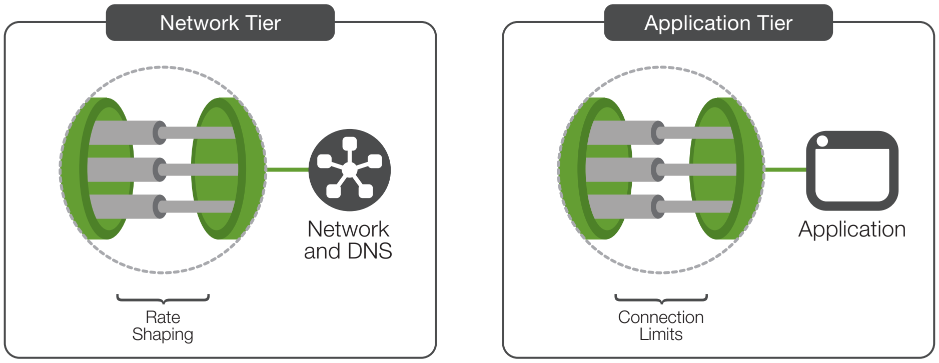

Figure 4: Resource constraints can help you survive the attack.

Step 10: Manage public relations

Hacktivist organizations today use the media to draw attention to their causes. Many hacktivists inform the media that an attack is underway and may contact the target company during the attack.

Financial organizations, in particular, may have policies related to liability that prevent them from admitting an attack is underway. This can become a sticky situation for the public relations manager. The manager may say something like, “We are currently experiencing some technical challenges, but we are optimistic that our customers will soon have full access to our online services.”

Handling reporters

Reporters, however, may not accept this type of hedging, especially if the site really does appear to be fully offline. In one recent case, a reporter called a bank’s local branch manager and asked how the attack was proceeding. The branch manager, who had not received media coaching, responded, “It’s awful, we’re getting killed!”

If the DDoS attack appears to be a high-profile hacktivist attack, prepare two statements:

1. For the press. If your industry policies allow you to admit when you are being externally attacked, do so and be forthright about it. If policy dictates that you must deflect the inquiry, cite technical challenges but be sure to prepare the next statement.

2. For internal staff, *including anyone who might be contacted by the press*. Your internal statement should provide cues about what to say and what not to say to media, or even better, simply instruct your staff to direct all inquiries related to the event back to the PR manager. Include a phone number.

Conclusion

If this information has been helpful, create a custom playbook for your organization.

- Include the reference sheets in the next section. Print them, fill them out, and laminate them.
- Use them to start a physical playbook or post them on the wall in the data center.

As you defend yourself against DDoS attacks, you can refine your playbook and improve the resilience of your applications.

Quick Reference 1: Contacts List

Many different teams may need to come together to fight a large, hectic DDoS attack. Use this form to collect and maintain contact information for the different teams and agencies that might be required during a DDoS attack. Add rows as necessary.

Team	Name	Phone	Email
Network Security			
Threat Intelligence			
Applications Director			
DNS Manager			
F5 Professional Services		1-888-88-BIG-IP	
Reseller Services			
Bandwidth Service Provider			
Public Relations Director			
Fraud Team Liaison			
Financial Comptroller			

Quick Reference 2: Whitelists

Maintain the list of IP addresses that must always be allowed access. Addresses that should be recorded here include:

- External monitoring tools.
- Google and the other search engines you do not want to block.
- Your own global traffic managers (GTMs). These will be monitoring your applications throughout the attack.
- Your DDoS cloud-scrubbers such as F5 Silverline DDoS Protection.
- Your other cloud service providers (this could be large list).
- Business partners.

IP Address Range	Maps to?	External Contact	Internal Contact

Quick Reference 3: Triage Applications

For all applications at the data center:

1. Decide on and record a priority value indicating whether or not it should be disabled.
2. Record a triage decision. (You can use the priority value to assert a decision like "disabling all applications that are priority 3 or lower").
3. Add a column for application owner contact information if necessary.

A defined set of priorities may enable you to automate tasks. For example, you can write a script to disable (and later re-enable) all applications with a priority of 3 or lower.

Application Name		Priority	Triage	Associated Virtual Server	Location
1	Example Application	2	Disable	dc1-rxspc.example.com	BIG-IP2, Rack 5, 192.168.11.5
2					
3					
4					
5					
6					
7					
8					
9					
10					
11					

Quick Reference 4: F5 Device Map

If you engage F5 Professional Services during the DDoS attack to assist with defense, it will be helpful to have a map of available F5® BIG-IP® devices in the data center, along with their serial numbers and the other information here. This information will guide those advising you on defensive strategies. The command below provides the serial number and the platform type:

```
%tmsh show sys hardware
```

Both of the F5 configuration management solutions, F5 Enterprise Manager™ and the F5 BIG-IQ® system, gather the device information (other than the location) for you and may assist you in filling out this table. Keep this information with Quick Reference 1: Contacts List. Better yet, keep all of your quick references in one place.

F5 Device		Model	Modules	Serial Numberr	Location
1	bigip2-dmz.dcxnet.com	1600	GTM	f5-wtax-exgw	Data Center 1, DMZ, Rack 5, 192.168.11.5
2					
3					
4					
5					
6					
7					
8					
9					
10					

Quick Reference 5: Attack Log

Information recorded here can be useful for after-action reporting, lessons learned, and regulatory reporting requirements. Print out several copies of this page and use it as a cover sheet for notes taken during the attack.

DDoS Attack Log	
Attack Started	Date & Time
Attack Stopped	Date & Time
Fraud Team Alerted	Date & Time
Intrusion Detected	Date & Time
Assets Exposed, If Any	
DDoS Attack Vectors (circle one)	ICMP UDP TCP DNS HTTP HTTPS
Attribution or Attackers Identified	

Source addresses may be turned over to the authorities. If the addresses are isolated to a specific country, the attack may be mitigated via geolocation (see Step 6 in the F5 DDoS Playbook).

Source Address Analysis
Geolocation:
Source Address:

Once the attack is over, provide a summary that includes a description of the attack, the mitigations that worked, and those that did not work. Include services that were disabled and their weaknesses. Use that information to evolve your services before the next attack.

Attack Summary
Geolocation:
Source Address:

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan K.K.
f5j-info@f5.com

Solutions for
an application world.