radware

Choosing the Right DDoS Solution

A guide to DDoS mitigation solutions and how to select the optimal one for your organization

Executive Summary

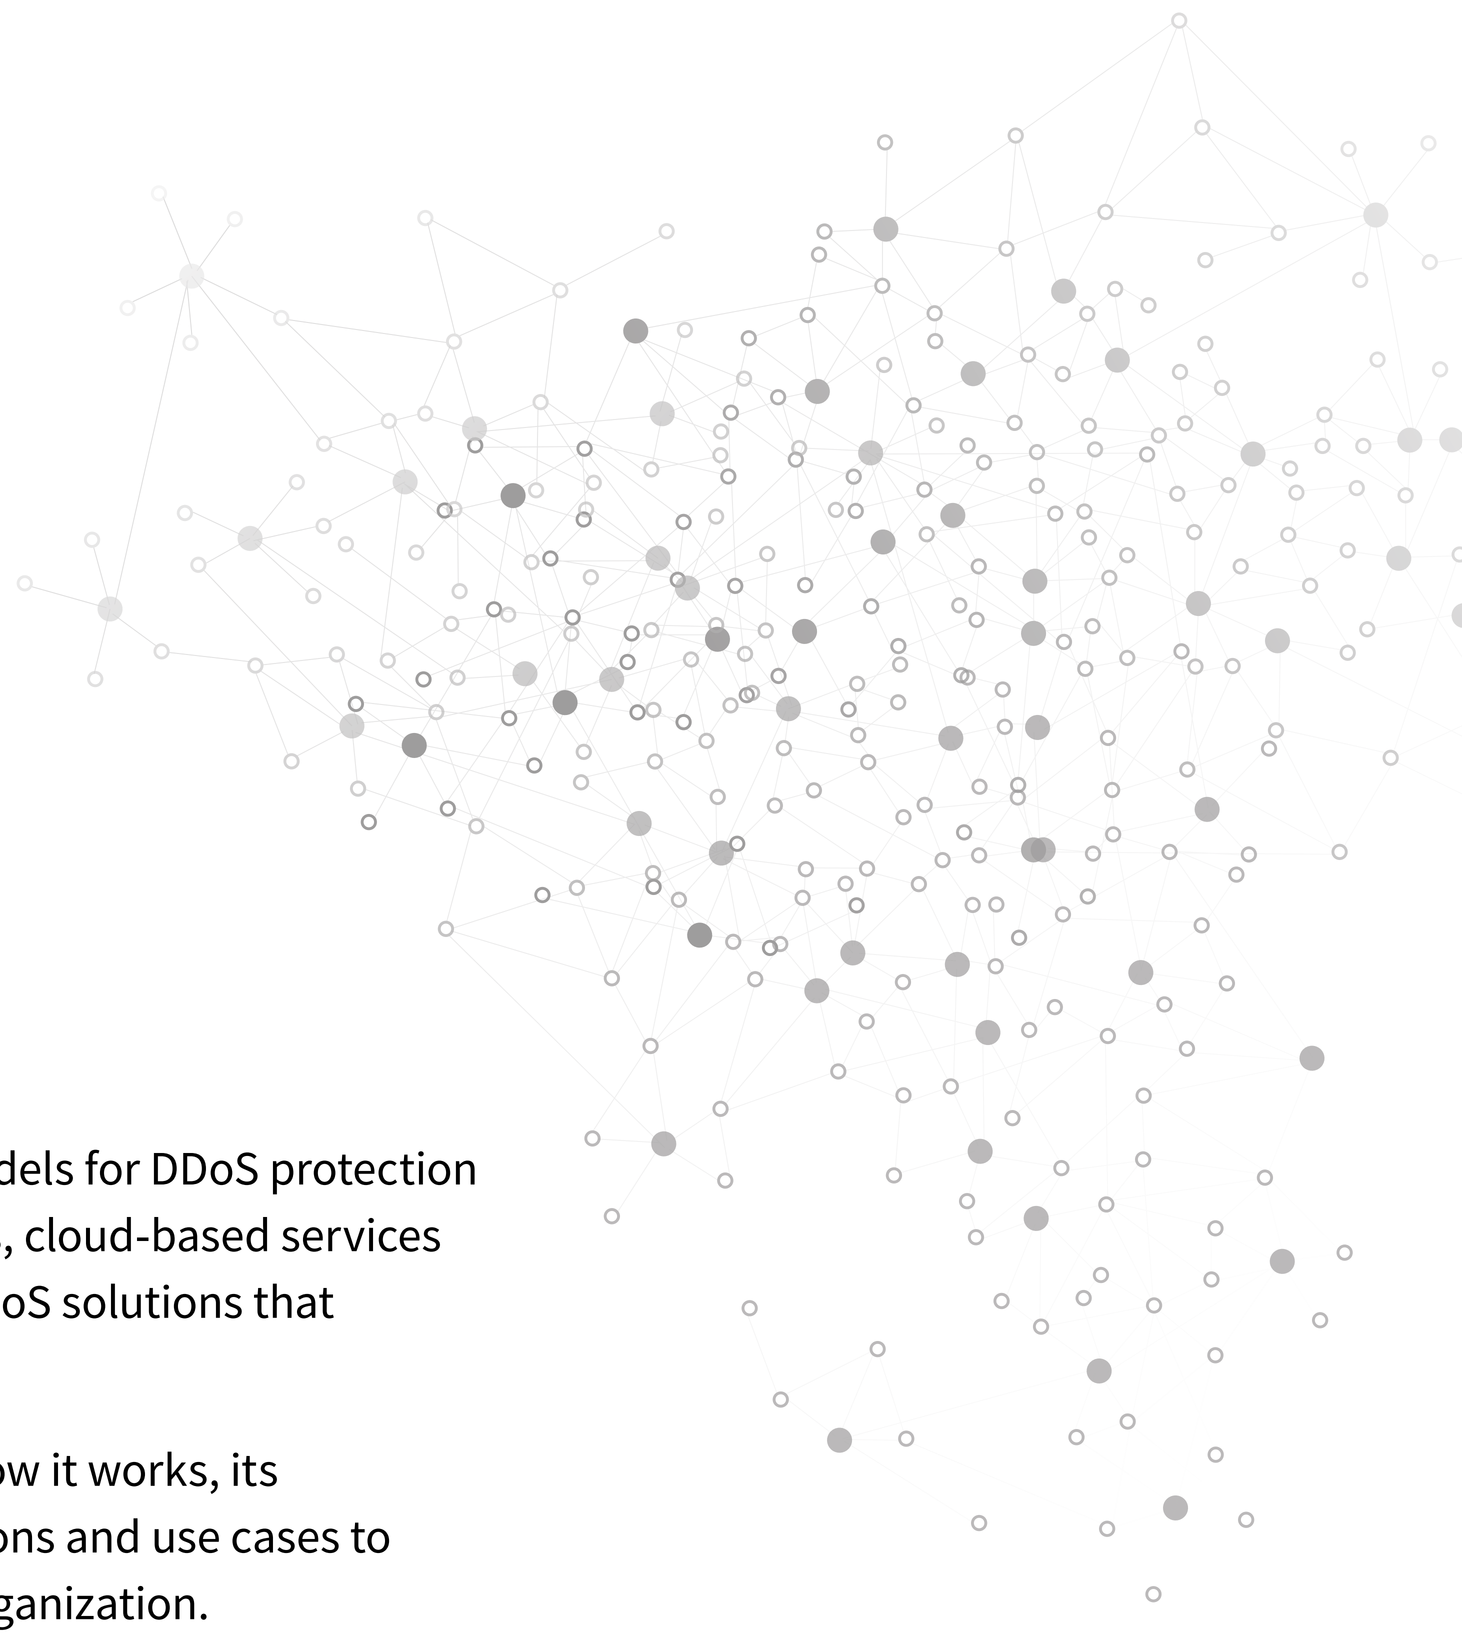

DDoS protection is not a one-size-fits-all fixed menu; rather it is an a la carte buffet of many choices. Each option has its unique advantages and drawbacks, requiring each organization to select the optimal solution that best fits its needs, threats and budget.

Overall, the DDoS mitigation market continues to be influenced by the transition to the cloud, with two major trends underscoring this. First, applications continue their migration our of private data centers and into the cloud. Secondly, nearly 75% of surveyed global security decision-makers want to consume DDoS protection as a service instead of racking hardware themselves, according to Forrester Research.

Hybrid DDoS protection combines both premise- based and cloud-based components. It provides both low latency and uninterrupted protection in addition to the high capacity required to mitigate large-scale volumetric DDoS attacks. This is best for customers seeking data center protection as well as customers running mission- critical and latency-sensitive applications.

Always-On cloud service provides constant, uninterrupted cloud-based DDoS protection. However, since all traffic is routed via the provider's scrubbing network, it may add latency to requests. This is best for applications hosted on public clouds or customers who come under attack frequently.

On-Demand cloud service is activated only when organizations come under DDoS attack. However, detection and diversion usually take longer than with other models, meaning that customers may be exposed for longer periods. This is best for customers who are infrequently attacked (or not at all) or otherwise have limited budgets.

This guide examines the various deployment models for DDoS protection and reviews premise-based hardware appliances, cloud-based services (both on-demand and always-on), and hybrid DDoS solutions that combine cloud and premise-based components.

For each deployment model, this guide covers how it works, its advantages and drawbacks, and key considerations and use cases to determine if the model will work best for your organization.

< Prev   Next >

DDoS Criteria and Considerations

Before evaluating DDoS protection solutions, it is important to assess the needs, objectives and constraints of the organization, network and applications. These factors will define the criteria for selecting the optimal solution.

1 WHAT ARE YOUR DATA CENTER PLANS?

Many organizations are migrating their data center workloads to cloud-based deployments. The decision of whether to invest in new equipment or to use a cloud service depends heavily on this consideration. Organizations that are planning to downscale (or completely eliminate) their data centers might consider a cloud service. However, if you know for sure that you are planning to maintain your physical data center for the foreseeable future, then investing in a DDoS mitigation appliance could be worthwhile.

2 WHAT IS YOUR THREAT PROFILE?

Which protection model is best for you also depends heavily on your company's threat profile. If a company is constantly attacked with a stream of nonvolumetric DDoS attacks, then a premise-based solution might be effective. However, if it faces large-scale volumetric attacks, then a cloud-based or hybrid solution would be better.

3 ARE YOUR APPLICATIONS MISSION CRITICAL?

Some DDoS protection models offer faster response (and protection) time than others. Most applications can absorb short periods of interruption without causing major harm. However, if your service cannot afford even a moment of downtime, that should factor heavily into the decision-making process.

Continued on next page

< Prev Next >

Continued from previous page

4 HOW SENSITIVE ARE YOUR APPLICATIONS TO LATENCY?

Another key consideration is the sensitivity of the organization and its applications to latency. Cloud- based services tend to add latency to application traffic, so if latency is a concern, then an on-premise solution — either deployed in-line or out of path

5 IS YOUR ORGANIZATION HEAVILY REGULATED?

Some organizations are within regulated industries that handle sensitive user data. As a result, they're prevented from — or preferring to avoid — migrating services/data to the cloud. In such cases, there may not be an alternative to using an on-premise appliance.

6 HOW IMPORTANT IS CONTROL?

Some organizations place a big emphasis on control, while others prefer to delegate the burden. Physical devices will provide organizations with more control but will also require additional overhead. Others might prefer the lower overhead usually offered by cloud services.

7 OPEX VS. CAPEX?

Solutions which include hardware devices (such as a premise-based DDoS appliance) are usually accounted for as a capital expenditure (capex), whereas ongoing subscription services (such as cloud DDoS protection services) are considered operating expenses (opex). Depending on accounting and procurement processes, organizations will have a preference for one type over the other.

8 WHAT IS YOUR BUDGET?

Finally, when selecting a DDoS protection solution, many times the decision comes down to costs and available funds. That's why it is important to be cognizant of the total cost of ownership (TCO), including added over- head, infrastructure, support, staff and training.

< Prev Next >

DDoS Deployment Options

Fundamentally, there are four models of DDoS protection solutions from which customers can select to protect their assets.

PREMISE-BASED APPLIANCE	ON-DEMAND CLOUD SERVICE	ALWAYS-ON CLOUD SERVICE	HYBRID PROTECTION
A hardware-based appliance located directly in the customer's data center	A cloud-based service that is only activated when a DDoS attack is detected	A cloud-based service that diverts traffic through the DDoS protection provider at all times	Combines cloud-based and hardware components, enjoying the benefits of both worlds

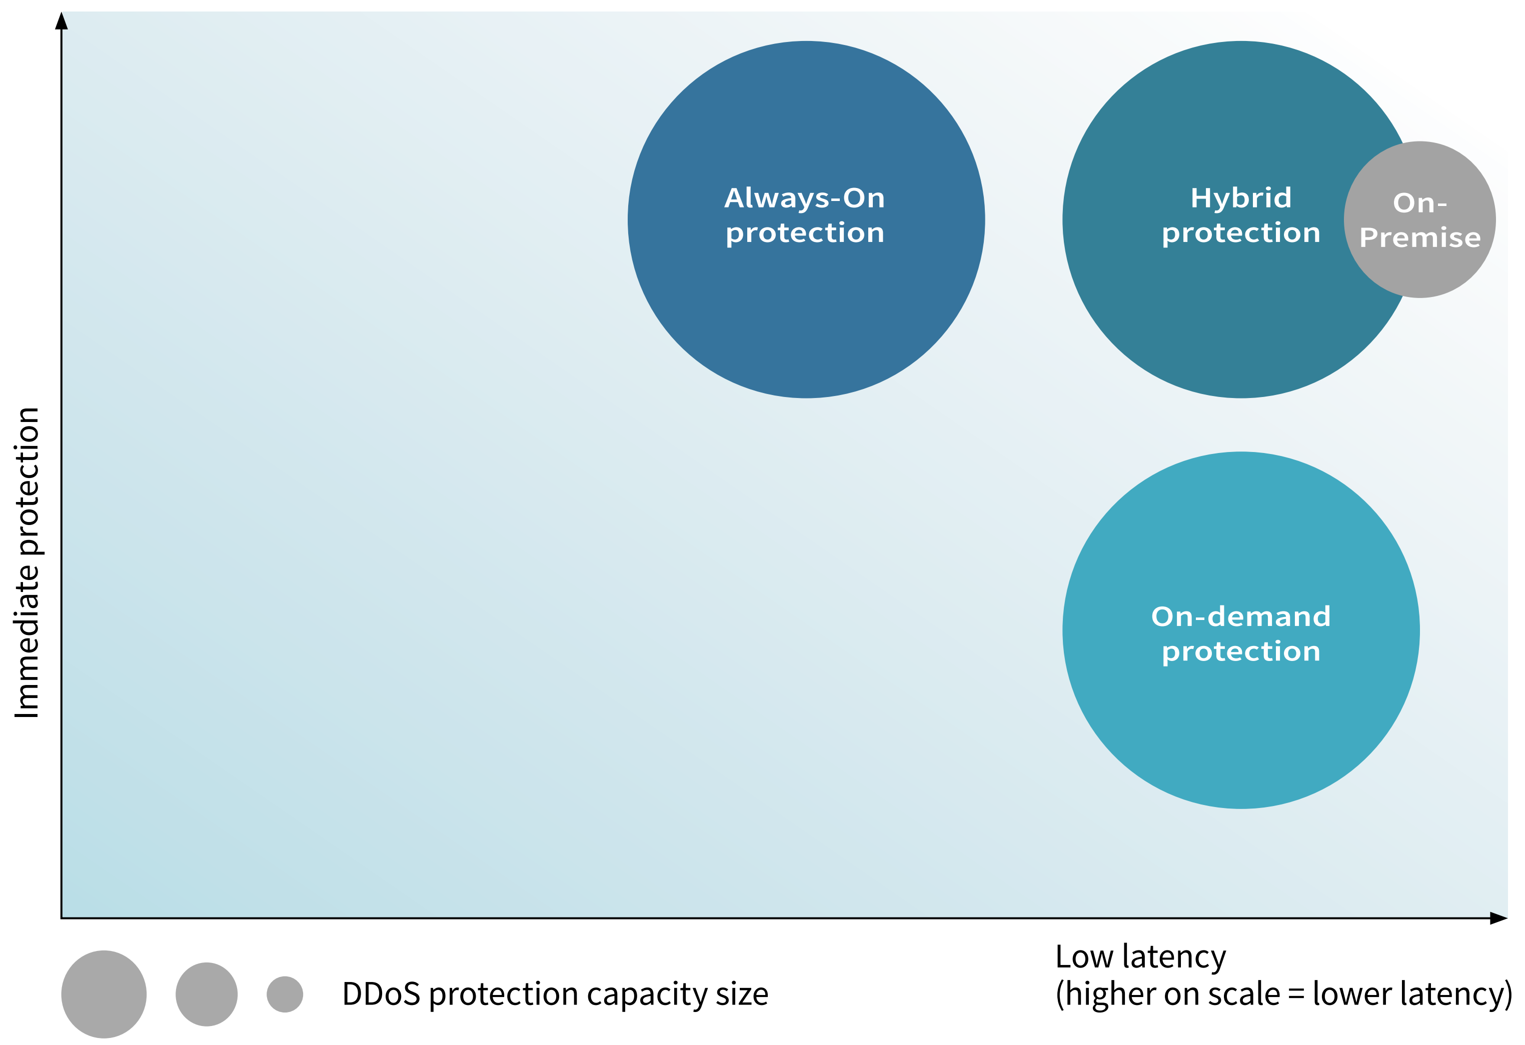

< Prev   Next >

Premise-Based Appliances: High Control But Limited Capacity

Premise-based appliances were the first form of DDoS protection, starting in the early 2000s in response to the first generation of DDoS attacks. These devices are deployed on-site at the customer's data center alongside other networking equipment such as firewalls, switches and routers.

The usage of stand-alone premise-based appliances has diminished in recent years, as the scale of volumetric DDoS attacks has outpaced the capacity of most appliances. Many customers have migrated DDoS protection to cloud-based services or hybrid solutions because they offer the bandwidth capacity necessary to deal with large-scale DDoS attacks.

Nonetheless, stand-alone DDoS appliances can still be found in many organizations, particularly those that have specific requirements that mandate the usage of such devices or are otherwise constrained from migrating to the cloud.

Understanding the merits and constraints of such solutions is helpful in understanding the relative advantages (and drawbacks) of cloud-based solutions.

ADVANTAGES AND DRAWBACKS

Low latency: One of the key advantages of the premise- based appliance is the low latency that it permits. The device is located directly in the data center, close to the application servers, with minimal or no latency. Moreover, some on-premise appliances can also be deployed in an out-of-path deployment, meaning there is no added latency at all during peacetime.

Control: Another key reason for selecting a premise- based DDoS protection device is control. Many organizations (and network managers) put a high premium on control and having their own device directly in the data center allows for maximum control.

Regulation: Finally, some organizations are in regulated industries, such as healthcare or finance, and are constrained by industry regulations to migrate their IT workloads to the cloud.

Continued on next page

< Prev   Next >

Continued from previous page

However, there are also certain drawbacks to deploying a premise-based appliance.

Capacity: While DDoS attacks continue to increase in size, premise-based DDoS appliances are constrained by their size and available bandwidth that they can handle.

Limited pipe size: Beyond the bandwidth capacity of the devices itself, stand-alone appliances are limited by the size of the network pipe going into the network. Large-scale volumetric attacks can quickly fill the network pipe, regardless of the capacity of any mitigation devices at the end.

Cost: A key consideration for many organizations is their available budget for a DDoS mitigation solution. The cost of a DDoS mitigation appliance can range from several tens of thousands of dollars for an entry-level device to hundreds of thousands of dollars for carrier-grade devices. Moreover, there are frequently associated costs for support and maintenance as well as dedicated staff needed to manage the equipment, which may impact the overall TCO.

Management overhead: With great responsibility comes additional overhead. Premise-based equipment frequently requires dedicated staff to manage the devices, in addition to utilities overhead, such as power, networking and cooling.

WHO IS IT BEST FOR?

Looking at the relative merits and drawbacks of stand-alone on-premise DDoS appliances, there are several categories of customers for whom it makes sense to explore such solutions.

Service providers that have a large install base and provide services to end customers using their data centers

Organizations that own existing data centers and are planning on maintaining them in the foreseeable future

Organizations in regulated industries that are unable to migrate workloads to the cloud

Latency-sensitive critical applications that require low latency and a high degree of control

Stand-alone on-premise solutions are less suited for certain organizations.

Applications hosted in the cloud that cannot be protected by premise-based equipment

Organizations migrating to the cloud that are planning to scale down their data center footprint

Price-sensitive customers who don’t have large budgets

Organizations frequently breached by large volumetric attacks, which can saturate the connection pipe or overwhelm the device

Premise-based solutions provide high-quality protection against a wide array of DDoS threats, as well as a high degree of control for organizations that put a high premium on it.

However, capacity is limited by the size of the appliance and the bandwidth of the network pipe. Customers who require low latency and the high degree of control offered by premise- based appliances should consider hybrid solutions, which add high-scale capacity against large-scale DDoS attacks.

< Prev Next >

Cloud-Based DDoS Protection

More organizations have begun to move away from hardware appliances toward subscription- based cloud services. There are numerous advantages to moving to a cloud-based protection service.

Protecting cloud-based applications: Applications that are hosted in the cloud cannot be protected by premise-based equipment and therefore require cloud-based protection.

Larger capacity: As volumetric DDoS attacks become bigger, many attacks can surpass the capacity of typical enterprise-grade DDoS mitigation appliances. In such cases, a cloud service will be able to provide backup capacity that can absorb these attacks.

Less management: Using a cloud service frequently requires less management overhead and staff than a premise-based device.

Lower cost: Whereas DDoS mitigation appliances require large upfront capital costs (capex), cloud-based DDoS mitigation services tend to be lower cost and can be purchased as an ongoing subscription model. This allows the customer to expand (or contract) their service based on operational needs. Moreover, such expenditures are usually classified as operating expenses (opex), which for many companies are easier to allocate.

It should be noted that the convenience of the cloud is tapered by some drawbacks.

Lower level of control: Since the services are not managed by the customers and not installed on-premise, it will allow for a lower degree of control. For customers (or network managers) for whom control is important, this might be a challenge.

Conflict with regulatory requirements: Certain regulatory requirements may limit the organization's ability to migrate/move data to the cloud.

< Prev Next >

On-Demand Cloud Service: DDoS Protection When You Need It

The first model of cloud-based DDoS protection is an on-demand model, in which traffic flows directly to the host in peacetime (when not under attack). However, once a DDoS attack is identified, traffic is rerouted to the cloud DDoS mitigation service, which scrubs the attack traffic and passes only clean traffic to the origin server. As its name implies, this type of protection is activated — on demand — only in times of need.

ADVANTAGES AND DRAWBACKS

No latency in peacetime: One of the big advantages of an on-demand service is that there is no latency during “peacetime” when an organization is not under attack. Traffic is diverted only during times of attack, for the duration of the attack.

Lower cost: On-demand services tend to be cheaper than purchasing dedicated DDoS mitigation appliances or always-on cloud services. This allows for effective protection for customers who don’t have a large budget.

Simplicity: On-demand cloud-based services are simple to maintain and require no management during normal times.

There are certain drawbacks to the on-demand model.

Detection time: Perhaps the biggest drawback is that it does not provide protection 100% of the time. Most on-demand services detect DDoS attacks based on volumetric traffic thresholds. Protection will be activated only when a certain traffic threshold is reached, and it may take a few minutes to accumulate and analyze the data. During this time, the server may be exposed.

Continued on next page

Diversion time: After the diversion is initiated, it may take some time — usually a few minutes — until the diversion is complete. Diversion time consists of two factors: the time it takes to initiate the diversion and the time it takes for the diversion to propagate through BGP or DNS tables. While diversion time can be minimized using automatic or programmatic (API-based) diversion techniques, propagation time is usually outside of the provider’s direct control.

Latency during diversion: Once traffic has been diverted, all requests to the origin server flow through the network of the cloud DDoS mitigation provider, which may add latency to transactions. The amount of latency can depend on the location of the scrubbing center, the distance from the origin server and the quality of connectivity. However, this latency continues only while the diversion is taking place and returns to normal once the diversion is over.

< Prev Next >

Continued from previous page

WHO IS IT BEST FOR?

Considering the relative merits and drawbacks of the on-demand cloud DDoS protection model, there are several types of customers (or applications) for whom this model will work best.

Infrequently attacked: Companies that are not frequently attacked and do not need constant coverage

Latency sensitive: Applications that are very sensitive to latency and therefore for which an always-on solution will not be suitable

Price sensitive: Organizations that do not have a large budget to spend on DDoS protection and wish to have cost-effective protection.

There are certain organizations and application types for which this solution is less suited.

Constantly attacked: Organizations or applications that constantly come under attack, resulting in traffic being constantly diverted. In these cases, an always-on or a hybrid solution will probably be more suitable.

Mission-critical applications: Mission-critical applications must always be available and cannot afford any downtime. Since on-demand DDoS protection usually takes a few minutes to detect and divert, this may result in short interruptions to availability. If this is a major issue, an always-on or hybrid solution will be better.

The on-demand cloud DDoS protection model is a cost- effective solution for organizations that do not require constant protection. For organizations that do, cloud always-on and hybrid solutions can offer optimal protection.

< Prev Next >

Always-On Cloud Service: Uninterrupted Cloud-Based Protection

DDoS protection solutions using an "always-on" model work by constantly routing all customer traffic through the network of the DDoS mitigation provider. Customers change their routing advertisements (usually BPG or DNS) to the network of their DDoS mitigation provider, which then routes all traffic through its scrubbing centers. Communications are then scrubbed for malicious traffic, and only clean traffic is forwarded to the customer.

The difference between the on-demand model and the always-on model is that, in the on-demand model, traffic is diverted through the provider's network only for limited durations when an attack has been detected, whereas in the always-on model, traffic is diverted through the provider's network at all times.

ADVANTAGES AND DRAWBACKS

Using an always-on DDoS protection service provides key benefits.

Uninterrupted protection: One of the biggest benefits of the always-on model is you are protected at all times against DDoS attacks.

No protection gaps: There are no protection gaps during the detection and diversion stages. Most on-demand models detect attacks based on volumetric traffic thresholds. Only once the threshold has been reached will the diversion be initiated. The detection and diversion steps may take up to several minutes, during which time the application is still exposed. In the always-on model, traffic is constantly routed through the DDoS mitigation provider, and therefore no gap exists.

Low management overhead: An always-on deployment usually requires low management overhead. Once the initial configuration of the service is complete, there is no need for additional overhead since traffic is constantly routed.

There are downsides to this model.

Latency: Since all traffic is routed through the network of the DDoS mitigation provider, this will inevitably lead to additional latency to traffic. The amount of latency will depend on the location of the provider's scrubbing center, the distance from the customer's host and the connectivity.

Cost: Since traffic is always routed through the scrubbing center, always-on deployments use more bandwidth than on-demand services. As a result, always-on service tends to be noticeably more expensive than on-demand service.

Continued on next page

< Prev Next >

Continued from previous page

WHO IS IT BEST FOR?

There are numerous use cases for which this model is particularly suitable.

Critical applications: Mission-critical applications that cannot afford any downtime at all. The always-on aspect of the service will ensure the application is constantly protected.

Frequently attacked: Companies that frequently come under attack. In this case, an on-demand service doesn't make sense since it will constantly be diverting on or off.

Low-latency sensitivity: Applications that are not sensitive to the minor added latency usually incurred by such services.

There are also use cases for which such a solution is less suited.

Latency-sensitive applications: Real-time applications with high sensitivity to latency. In this case, an on-premise or hybrid solution will probably be more suitable.

Price-conscious customers: Always-on services tend to be more expensive due to the added traffic surcharges and additional overhead incurred by the service providers. Therefore, customers who have limited budgets might consider on-demand services.

The always-on model provides effective protection for applications that require constant protection against DDoS attacks and cannot afford any downtime. However, this added security comes at the cost of additional latency.

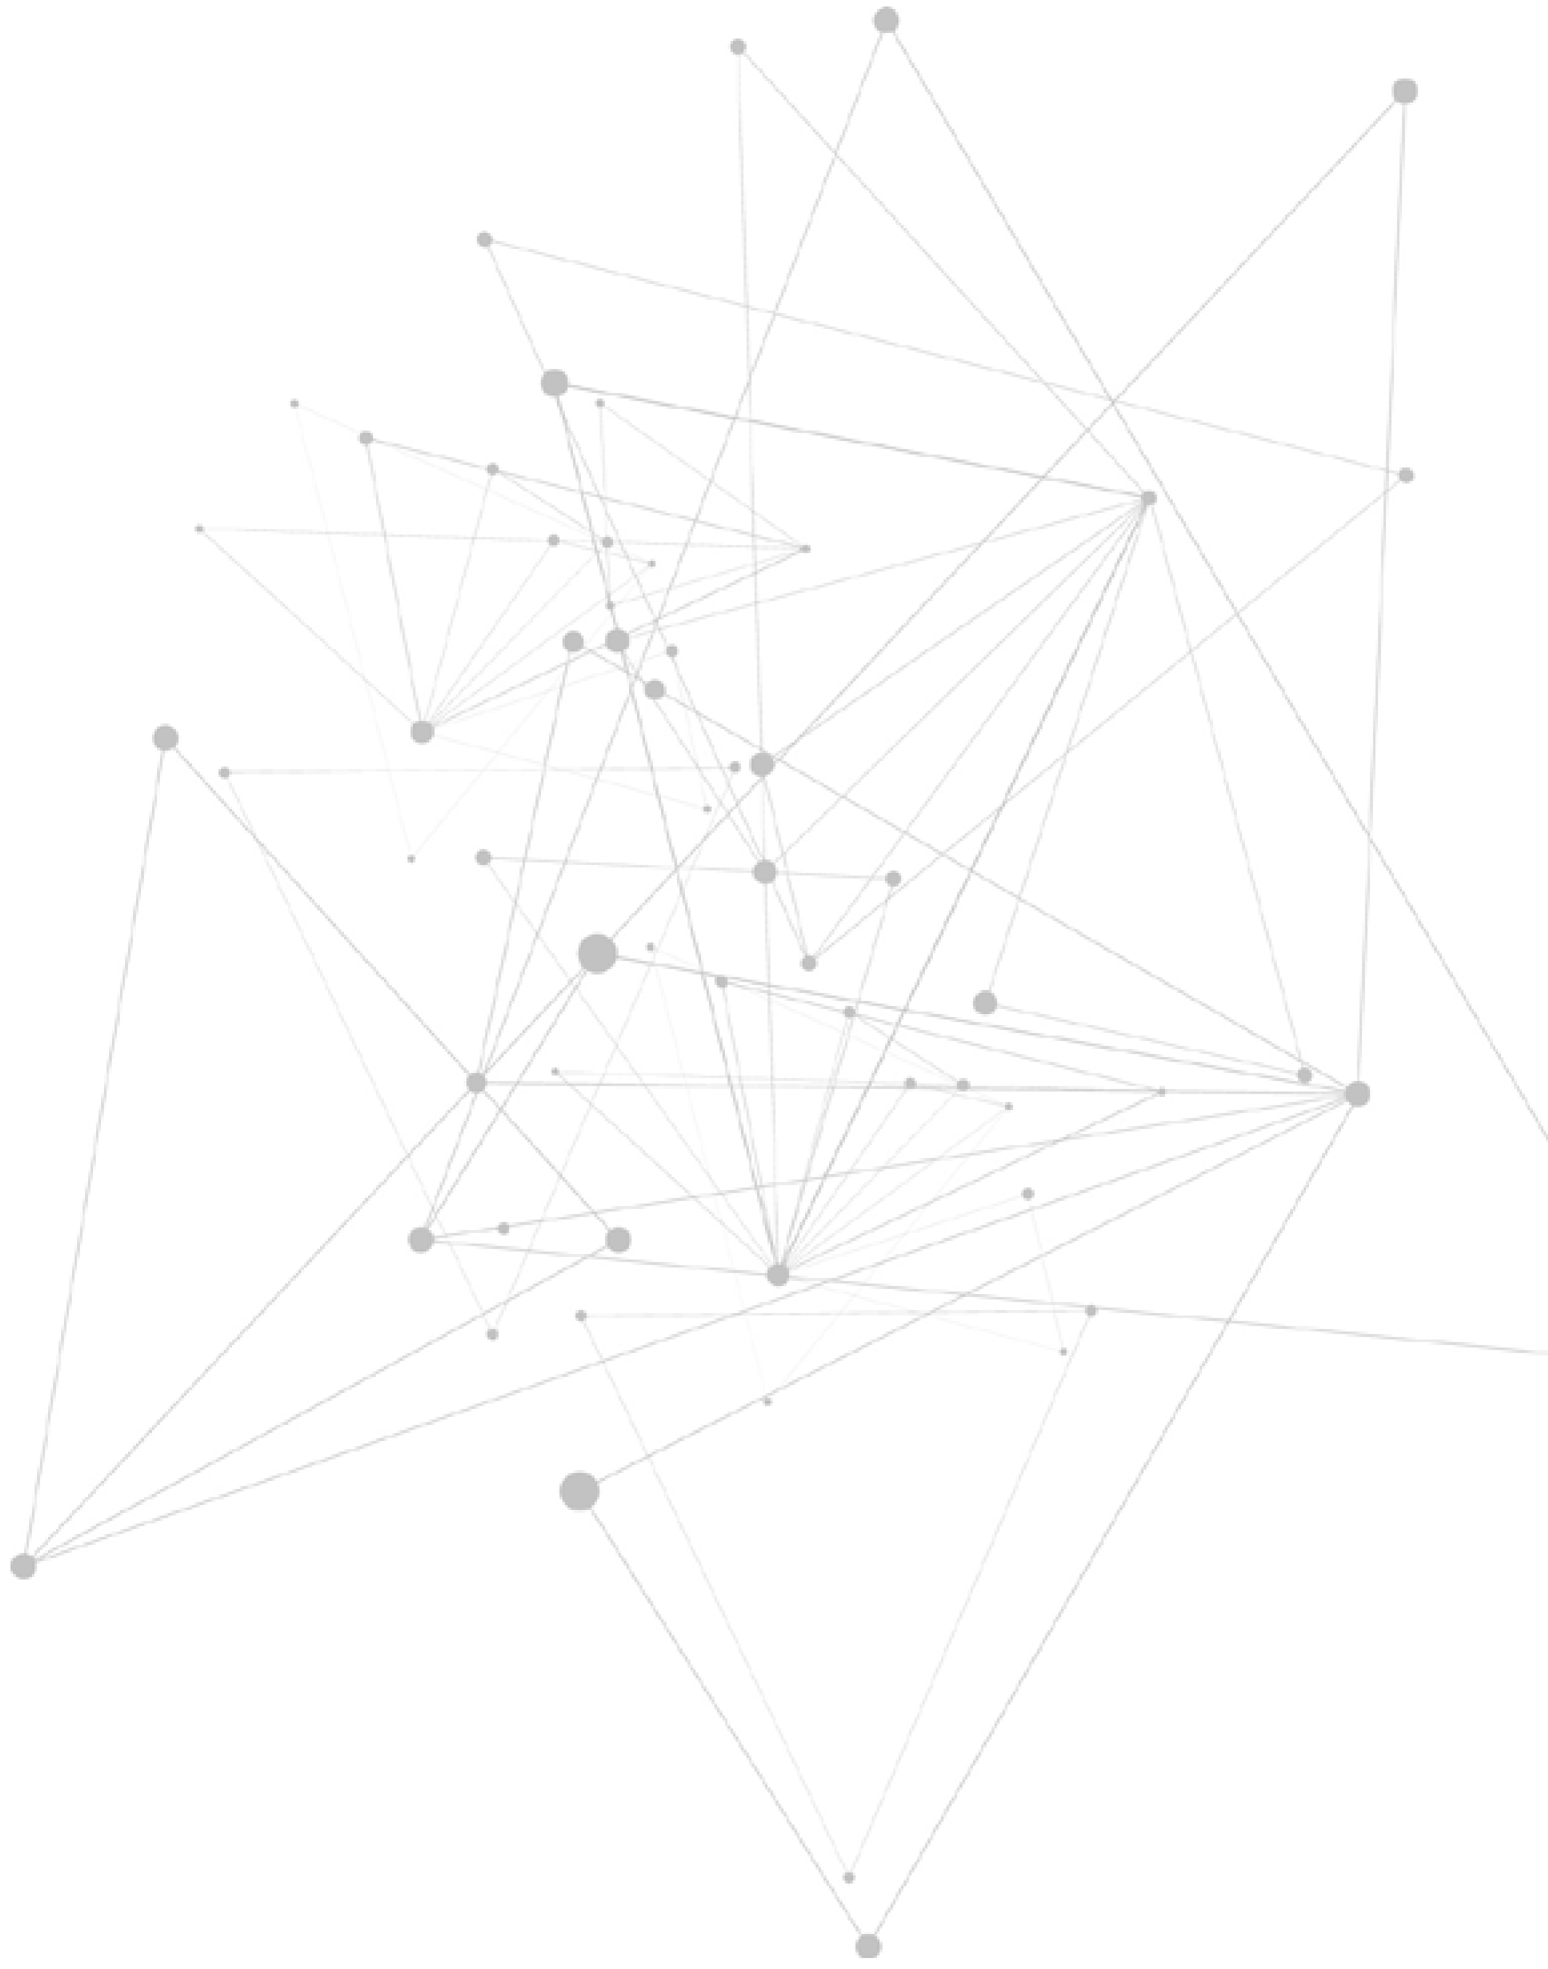

< Prev   Next >

Hybrid Protection: The Best of Both Worlds

Whereas a premise-based solution relies strictly on a local hardware appliance and on-demand and always-on solutions are purely cloud based, a hybrid model combines a local hardware appliance with expandable capacity in case of a large volumetric attack.

During the normal course of business, traffic flows directly to the data center. The premise-based appliance will inspect for attack traffic and block most attacks. If a large-scale attack, which may overwhelm the device (or even completely saturate the pipe), is detected, traffic is diverted to a cloud scrubbing center. The scrubbing center will block attack traffic and send only clean traffic to the customer. Once the attack is over, traffic is diverted back to the device.

Hybrid DDoS protection allows organizations to enjoy the best of both worlds: the low latency and high control of premise-based solutions together with the scalable capacity of cloud solutions.

Continued on next page

ADVANTAGES AND DRAWBACKS

There are certain advantages and drawbacks to choosing a hybrid DDoS protection solution.

Best quality of protection: Hybrid protection is the recommended best practice by most security analysts as it combines both low latency and high capacity for protection of mission-critical services.

Immediate detection: Since traffic flows through the local appliance at all times, attacks can be detected immediately by the appliance. This is an advantage over cloud on-demand services, which usually have a detection and protection gap until the diversion is initiated.

Flexible capacity: The availability of flexible mitigation capacity in case of large-scale volumetric attacks. Such attacks can overwhelm stand-alone hardware appliances and even saturate the internet pipe leading to the data center. Having backup cloud capacity allows customers to handle any attack, regardless of size.

Low latency: Hybrid solutions allow for low latency as day-to-day protection is handled by on-premise appliances directly in the data center. Only in times of attack will traffic be diverted to the cloud. This is an advantage compared to always-on cloud solutions, which usually add some latency to communications, even during peacetime.

Regulation: Companies in regulated industries, such as finance or healthcare, are frequently constrained in their ability to migrate services to the cloud. Therefore, a hybrid solution could be useful in providing on-premise protection most of the time while still allowing for backup capacity in case of large-scale attacks.

Control: Having an on-premise device allows for greater control and configurability, especially for organizations with unique network topologies or specific needs.

< Prev   Next >

Continued from previous page

The hybrid DDoS protection model also entails various drawbacks.

Management overhead: Having a premise-based solution usually incurs higher management overhead and staff requirements, while keeping premise-based and cloud-based defenses synchronized and aligned at all times.

Cost: Since a hybrid solution combines both a hardware appliance and cloud service, the combined cost usually tends to be higher than a cloud service alone.

WHO IS IT BEST FOR?

There are several types of customers (and use cases) who would benefit from this model.

Data center protection: Customers who have extensive data center infrastructure and services that require protection

Mission-critical applications: Mission-critical applications that require both constant protection and 100% availability

Latency sensitive: Services that require fast (or real- time) responsiveness and have low latency tolerance

Regulated industries: Companies in regulated industries that cannot migrate workloads and data to the cloud

There are also certain use cases in which a hybrid solution is not optimal.

Cloud-hosted applications: Applications that are hosted solely on public clouds (such as AWS or Azure). For such applications, a cloud-based solution is required.

Price sensitive: Organizations that don't have the budget to allocate for such comprehensive solutions. For such organizations, an on-demand cloud solution is usually best.

The hybrid protection model has been recommended by market analyst firms as the best practice for organizations looking to protect mission-critical applications against DDoS attacks.[1]

For customers (and applications) who need both constant protection and low latency, a hybrid solution combining both premise-based equipment and scalable cloud service is the best option.

[1] The Forrester Wave: DDoS Mitigation Solutions, Q1 2021. hyperlink: https://www.radware.com/ddos-wave-2021q1/

< Prev Next >

Summary: A Buffet, Not a Fixed Menu

DDoS protection is a buffet, not a fixed menu. There are many DDoS protection providers, who provide varying levels of protection and cost. Every model has its relative merits and disadvantages; there are many options, and it is up to each customer to choose the optimal solution for their particular use case.

On-premise DDoS solutions provide a high degree of control but have limited capacity when facing volumetric DDoS attacks. On-premise solutions are best for service providers that are building their own capacity or customers who are constrained by regulation from using a cloud service.

On-demand cloud services provide cost-effective, low-latency protection against volumetric DDoS attacks but take more time to initiate compared to other protection models. This is best for organizations that have a limited budget and/or are not attacked frequently.

Always-on cloud services provide constant protection but may add latency since traffic is routed through the DDoS protection provider at all times. This is best for applications hosted in public cloud environments as well as organizations that are constantly attacked.

Hybrid DDoS protection provides the best of both worlds with the low latency and control of on-premise solutions and the scalability of cloud solutions. This is best for data centers, mission-critical services and applications sensitive to latency.

Continued on next page

< Prev   Next >

Continued from previous page

See the table below for a comparison of the different deployment models

	ON-PREMISE	ON-DEMAND	ALWAYS-ON	HYBRID
Detection	Volumetric + nonvolumetric	Volumetric only	Volumetric + nonvolumetric	Volumetric + nonvolumetric
Latency	None in peacetime	None in peacetime	Minor added latency	None in peacetime
Protection	Immediate	Few min. until diversion	Immediate	Immediate
Capacity	Limited	High	High	High
Best For	Service providers Regulated industries	Service providers with many protected assets Cost-sensitive enterprises not frequently attacked	Applications on public cloud Organizations constantly attacked	Data center protection Apps sensitive to latency Mission-critical applications

< Prev Next >

About Radware

Radware® (NASDAQ: RDWR) is a global leader of cybersecurity and application delivery solutions for physical, cloud and software- defined data centers. Its award-winning solutions portfolio secures the digital experience by providing infrastructure, application and corporate IT protection and availability services to enterprises globally. Radware's solutions empower more than 12,500 enterprise and carrier customers worldwide to adapt quickly to market challenges, maintain business continuity and achieve maximum productivity while keeping costs down. For more information, please visit www.radware.com.

Radware encourages you to join our community and follow us on: Radware Blog, LinkedIn, Facebook, Twitter, SlideShare, YouTube, Radware Connect app for iPhone® and our security center DDoSWarriors.com that provides a comprehensive analysis on DDoS attack tools, trends and threats.

< Prev