

Why You Should Perform Security Functions at the Load-Balancing Level

WHITEPAPER





Executive Summary

Protecting web applications from cyberattacks is a critical operational imperative for organizations in today's increasingly complex cloud, SaaS and hybrid deployment environments. As the threat landscape evolves, implementing a multi-layered strategy is essential to improve your IT defenses. Load balancers play a pivotal role in this strategy, serving as strategic control points and application gatekeepers.

This whitepaper explores the crucial role of load balancing in cybersecurity and demonstrates how integrating security functions at the load-balancing level can significantly enhance an organization's security posture. Given their position in the user and application server chain, load balancers are ideally suited to inspect access requests before they reach applications and data sources. This enables load balancers to identify and stop malicious traffic from compromising backend servers.

This paper examines how load balancers function and the techniques they use to intelligently distribute access requests across available application servers. It then delves into the specific security technologies and functionalities that you can integrate with load balancers, including Web Application Firewalls (WAFs), pre-authentication methods, intrusion prevention, security certificate management, TLS/SSL security offloading and extensive logging.

We outline how organizations can reap numerous benefits by using the robust security features offered by load balancers. These include improved performance through offloading security processing tasks, centralized security policy enforcement, enhanced visibility and control over network traffic and the ability to proactively mitigate threats before they reach application servers.

This paper also highlights the challenges associated with using load balancers without security and provides best practices for effectively integrating load balancing with security functions to maximize their effectiveness. Real-world case studies, such as ASOS's successful implementation of Progress® Kemp® LoadMaster® to handle 167 million website visits during peak Black Friday sales activity, demonstrate the tangible benefits of deploying load-balancing and the security provided.

As the threat landscape continues to evolve, the need for adaptable, multi-layered defenses remains paramount. LoadMaster is an ideal solution for delivering cybersecurity functionality that complements other security layers across on-premises and cloud applications or hybrid infrastructure deployments. With its scalable security features and the ability to integrate with existing security infrastructure, LoadMaster empowers organizations to fortify their defenses and stay ahead of emerging threats.

Introduction

Protecting web applications from cyberattacks is crucial as cloud, SaaS and hybrid deployment models become more prevalent. The increased use of cloud deployment, plus hybrid deployments across cloud and on-premises infrastructure, has made cybersecurity more complex.


The risk of financially-motivated cyberattacks shows no signs of decreasing. As a result, improving web application protections (and other IT systems) against cyberthreats is a core operational imperative for organizations.

Every organization requires a multi-layered cybersecurity defense strategy to counter ever-present threats. Layered protection adds multiple levels of defense against potential threats. Similar to a castle with multiple walls, moats and gates to keep attackers out, having a layered approach to cybersecurity helps prevent attacks from succeeding. Each layer can act as a barrier, reducing a cyberattack's speed and providing security personnel more time to detect and respond to threats.

By incorporating multiple layers of protection, organizations significantly improve their overall security posture and reduce the likelihood of successful cyberattacks. Deploying components of this layered security functionality on load balancers can play a significant role in broader cyber defense strategies.

Due to their central role in controlling access to web applications, load balancers are ideal for inspecting network traffic before it reaches applications and data sources. After detecting suspicious activity, load balancers can mitigate potentially malicious network traffic from reaching backend servers.

In this paper, we outline what load balancing is, how it works and how deploying additional security protections on load balancers enhances security. As we know it best, we will use LoadMaster in examples and illustrations of the concepts.



What Is Load Balancing?

Load balancing is the process of distributing incoming client access requests to a pool of application servers so that the load is spread evenly over the available servers. Load balancers use multiple intelligent algorithms and network traffic monitoring to accomplish this task. More advanced load balancers continuously monitor the health and status of each server in the pool to prevent requests from getting routed to servers that are too busy or offline. By managing the allocation of requests, the load-balancing process helps maximize application performance and minimize downtime.

A grocery store checkout analogy helps explain what load balancers do at a network traffic level. When the number of shoppers in a store is low, having a few checkouts open is sufficient. However, as the evening rush hour arrives and more people start shopping, the number of people waiting at the checkouts increases. If only a few lanes are operating, long queues build up. Opening additional checkouts adds capacity and allows for an efficient flow of people through the tills.

In the same way, a load balancer spreads out client requests over the available servers. If required, you can spin up additional load balancer instances by adding additional virtual machines or new cloud instances to handle increased demand. And then take them offline again when demand reduces.

We can extend the checkout analogy to highlight how load balancing deals with a server issue. If someone drops a bottle of tomato juice at a checkout, then that single checkout lane will need to be closed. People queuing there will be moved to other available checkouts until cleaning is complete, and the closed checkout reopens. Similarly, if a server or service is unavailable, the load balancers will redirect client traffic to other available servers. It will then start to reuse the temporarily offline server when it comes back online.

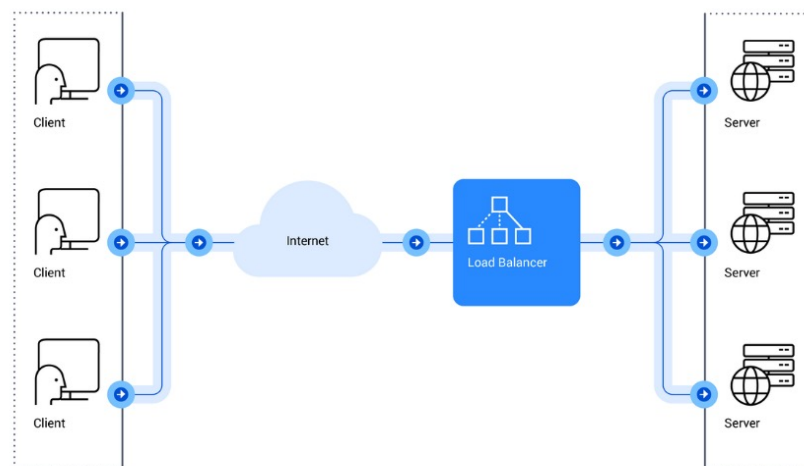



Figure 1: How a Load Balancer works



How Do Load Balancers Work?

There are various techniques and algorithms that load balancers use when distributing access requests among available servers. The method used depends on which load balancer you deploy and the type of service request or application. LoadMaster, for example, can use these methods to load balance access requests to a server pool:

DNS Round Robin

Uses load-balancing servers or provides fault tolerance. This method involves configuring multiple servers with the same services and unique IP addresses that use the same internet domain name. The load balancer allocates requests on a rotating round-robin basis in a continuous loop.

Weighted Round Robin

Improves upon the DNS Round Robin method described above. The network administrator assigns each server in the pool a static numerical weight. The administrator assigns the most efficient and powerful servers with higher weights. The load-balancing algorithm prefers higher-weight servers when deciding where to send requests.

Least Connection

Differs from DNS Round Robin or Weighted Round Robin in that it considers the current server load when distributing requests. Instead of simply rotating requests among servers, it sends the latest request to the server currently servicing the fewest active connections.

Weighted Least Connection

Assigns a numerical value to each server, similar to the Weighted Round Robin method. The load balancer uses these weights to distribute sessions to servers. If two servers have the same number of active connections, the server with the higher weighting gets sent the new request.

Agent-Based Adaptive Load Balancing

Installs an agent on each server in the pool that reports the current load on the server to the load balancer software. This real-time current load information gets used when deciding which server can handle the new access request.

Chained Failover (Fixed Weighted)

Configures a predetermined order of servers in a chain. All requests go to the first server in the chain. When it can't accept anymore, the next server in the chain gets sent all new requests until it can't handle more. Then the third server receives all requests and so on for the length of the server chain. When the last server in the chain can't accept any more requests, the load balancer retries from the start of the chain.

Weighted Response Time

Uses the response time from a server health check to determine the server responding fastest at any particular time. The next client access request then goes to the server that's responding fastest. This means that any servers under heavy load, which will respond more slowly, are not sent new requests, allowing the load to even out on the available servers in the pool over time.

Source IP Hash

Uses an algorithm that takes the client and server IP addresses and then generates a unique hash key. This key gets used to allocate the client to a particular server. As the client and the server pool can regenerate the key if the session breaks, this load-balancing technique can reconnect clients to the same server they were using previously. Load balancers can use this to deliver session affinity for applications with shopping carts so that items placed in carts are still there when a dropped connection gets re-established.

Software Defined Networking (SDN) Adaptive

Combines knowledge of upper networking layers with information about the network's state at lower layers. SDN Adaptive uses information about the status of the servers, the status of the applications running on the servers, the health of the network infrastructure connecting the servers and the level of congestion on the network in the load balancer decision-making.

In reality, the methods outlined above are not used individually but rather in combinations that determine where to send requests.

The load balancer pool and services shared on them can be located in a single data center or geographically distributed over multiple data centers using Global Server Load Balancing (GSLB - see ref 1).

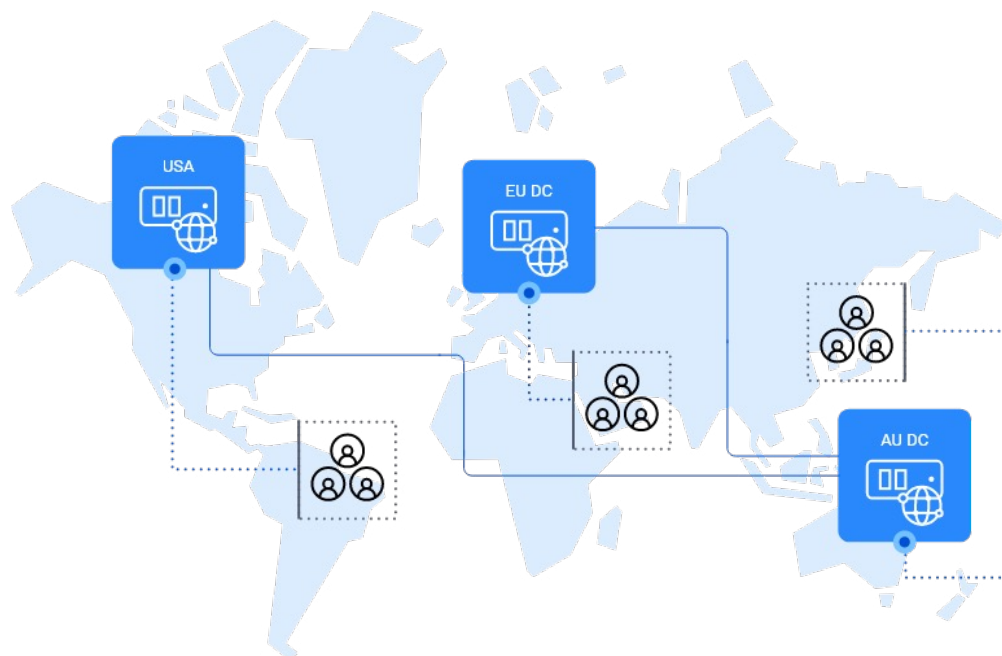


Figure 2: Load balancing across different geographic locations using Global Server Load Balancing

The Role of Load Balancing in Cybersecurity

Modern load balancers like LoadMaster operate as strategic points of control, acting as application gatekeepers. Organizations can enhance their defenses, mitigate risks and deliver improved business continuity by combining specific security functions with load balancing.

As we've noted, load balancing has traditionally been seen primarily as a way to distribute network traffic among servers to enhance reliability and high availability. This is still a core function of load balancers, but in the increasingly complex threat landscape, deploying security measures on load balancers can significantly enhance security protections in combination with other cybersecurity defenses such as firewalls, zero-trust networking, Network Detection and Response (NDR) solutions, endpoint protection and more.




Figure 3: Implementing Zero Trust principles for security

Enhancing Security through Load Balancing

Integrating security functions with an application load balancer can significantly enhance an organization's cybersecurity posture. Solutions such as LoadMaster can have multiple built-in and optional additional installable cybersecurity components. Cybersecurity professionals and System Admins can choose what security components to deploy on each load balancer instance in a way that best meets their unique needs.

The cyber security technologies and functionality typically available via deployment of load balancers are as follows.



Web Application Firewall (WAF)

Deploying an enterprise-class WAF to protect applications from vulnerabilities is a common practice when rolling out load balancers. WAFs also support the creation of per-application security profiles to enforce source location-level filtering, have pre-integrated rulesets for common attack vectors (including rulesets such as the OWASP ModSecurity Core Rule Set) and also include the creation of custom security rules. The out-of-the-box deployment of a WAF delivers protection against common threats without changing your applications or infrastructure. Every organization can tailor the default security with custom rules to meet its unique needs.

LoadMaster has an industry-leading WAF. The diagram below shows where the WAF on LoadMaster logically sits within a typical application deployment infrastructure. Other implementations will have a similar topology. The WAF sits between the applications and clients to mediate access requests, inspect the traffic and provide security via the predefined rulesets and any custom rules deployed.

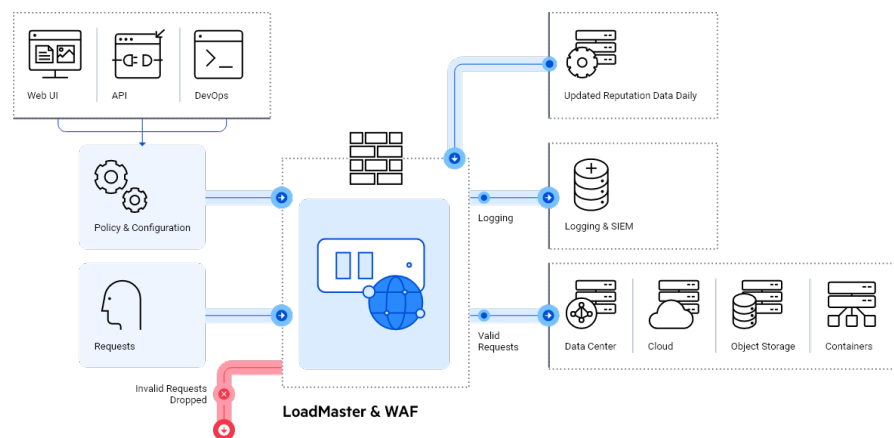



Figure 4: How Kemp LoadMaster WAF integrates security into web application access workflow



Pre-authentication Methods

Load balancers play a crucial role in delivering and enhancing authentication by acting as a central control point for user access to applications and services. They operate as a reverse proxy, sitting in front of application servers and handling incoming client requests. This allows the load balancer to perform authentication before forwarding the request to the appropriate server. By centralizing authentication at the load balancer level, organizations can enforce consistent authentication policies across multiple applications.

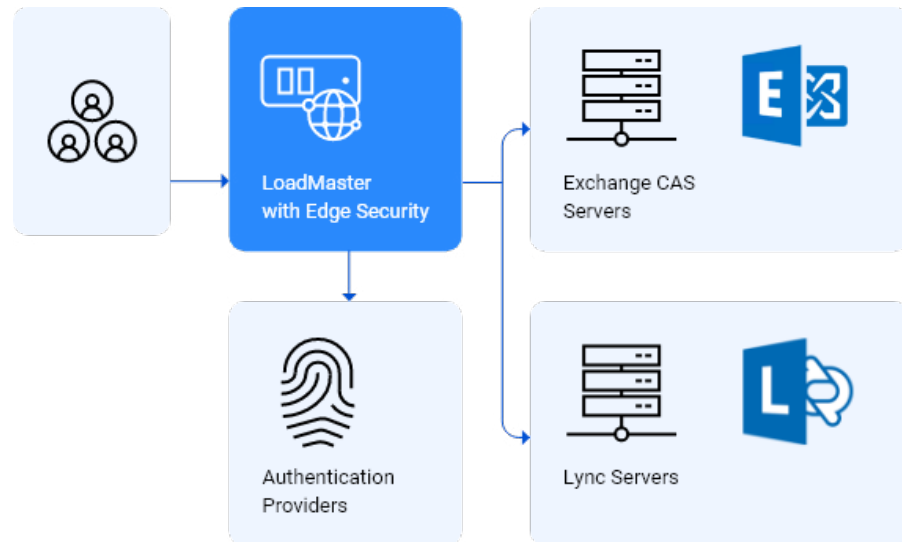



Figure 5: How Kemp LoadMaster provides security through enhanced authentication

User authentication is a common approach where the load balancer validates user credentials against a designated authentication server, such as Active Directory (AD) or RADIUS. The load balancer securely communicates with the authentication server to verify user credentials before granting access to the requested resource. Integration with Active Directory enables the load balancer to access and use an organization's existing user credentials and information for authentication. This simplifies user management and maintains access policies defined in AD are consistently applied across applications. RADIUS authentication allows the load balancer to authenticate users against a RADIUS server, which can be useful for organizations with diverse authentication requirements or legacy systems.

Another key benefit of using load balancers for authentication is single sign-on (SSO) across virtual services. By implementing SSO, users can authenticate once and seamlessly access multiple applications without the need to re-enter credentials. As long as it's implemented in a way that doesn't undermine zero-trust security by enabling SSO across systems where it's not required or appropriate. The load balancer securely shares the authentication context across virtual services where appropriate, enhancing the user experience and reducing the authentication burden.


Load balancers can also integrate with advanced authentication mechanisms like RSA SecurID, which provides two-factor authentication using hardware or software tokens. This adds an extra layer of security, providing valid tokens to only authorized users so they can access protected resources. By using load balancers for authentication, organizations can centralize access control, enforce consistent policies and enhance security across their application infrastructure.

See ref 2 for details of how LoadMaster implements and supports enhanced authentication methods.



Intrusion Prevention

Many load balancers include an Intrusion Prevention System (IPS) to further boost security. LoadMaster includes one that is compatible with rules created in SNORT syntax. Administrators can use an IPS engine to apply Permit/Deny IP by address rules for HTTP or HTTPS traffic with TLS/SSL offloading enabled. Note that the IPS systems on most load balancers (including the LoadMaster IPS system) do not replace a dedicated IPS and that using a WAF is a superior security choice in most cases. See ref 3 for details of the LoadMaster IPS.



Security Certificate Management

One of the most common issues with the cryptographic security needed to encrypt and protect data at rest and in transit over networks is the expiry of a security certificate. When this happens for a certificate being used to provide TLS/SSL security for a web application, unplanned downtime is the result. Expired security certificates have been the root cause of many of the significant service outages in the past few years that have impacted some of the most recognizable technology companies. Renewing certificates before they expire prevents this from happening, but those responsible for this task often fail to do it before it's too late.


Many load balancers take on the management role for certificates and can often automate the renewal process and so minimize unplanned downtime. LoadMaster natively supports the Automated Certificate Management Environments (ACME) protocol. ACME enables load balancers that use security certificates to communicate via an API with Certificate Authorities responsible for issuing trusted security certificates. Using ACME in LoadMaster enables the automatic renewal of security certificates issued by Let's Encrypt, Entrust, DigiCert, SSL dot Com and ZeroSSL.



TLS/SSL Security Compliance


Many industry and government regulatory frameworks require that applications use the latest version of TLS/SSL possible (ideally TLS 1.3, the latest at the time of writing). Legacy applications can have problems with this as they may only support older TLS/SSL protocols or none. If any of these applications are used in organizations that must implement PCI DSS, NIST FIPS 140-2 or other regulations, this needs to be addressed.

Many load balancers can bridge any TLS/SSL gaps that exist due to legacy systems. They can mediate access requests from systems using TLS 1.3 but also support legacy applications that don't support TLS 1.3 and communicate with them using the latest TLS/SSL protocol they support. With each system involved, they communicate with the highest level of TLS/SSL possible and translate between the protocols to deliver a more secure communication channel. It should be stated that the best practice is to upgrade all systems to the latest version of TLS/SSL. However, it's not always feasible and boosting security in the interim via load balancers is useful.



Comprehensive Logging

Load balancers also include in-depth logging of inbound and outbound activities and events. This occurs on a granular per-application basis and sends logged information to third-party SIEM systems. This provides useful data for real-time monitoring and post-event analysis to find root causes for any cyber or other incidents that occur.




The Challenges of Load Balancing Without Security

It may be tempting to deploy load balancers without taking advantage of the security functionality they provide. System administrators may think that they have several layers of security already and that adding the security features available via load balancers into the mix will be just another layer to manage and update—something else added to the IT and threat landscapes.

This represents short-term thinking because it introduces challenges that are larger than any additional management overhead from using the security functionality.

- 1. Increased security risks** - Without harnessing the security features provided by load balancers, organizations expose themselves to various security risks. Not centralizing and combining authentication and access control mechanisms at the load balancer level increases the risk of cybercriminals exploiting authentication weaknesses on individual servers. The lack of SSL/TLS offloading also increases the load on application servers.
- 2. Inconsistent security policies** - Maintaining consistent security policies across multiple applications is challenging. Each application server may have its own authentication and access control mechanisms, leading to a fragmented and inconsistent security posture. This lack of centralized control makes setting and managing security policies difficult. Having load balancers in central roles where you can set and deploy consistent rulesets is very useful.
- 3. Limited visibility and monitoring** - Load balancers provide valuable insights into traffic patterns, user behavior and potential security threats. Without using the security features, organizations miss out on this visibility. Detecting and responding to security incidents becomes more challenging as the detailed logs or alerts that load balancers can provide related to authentication failures, access violations or suspicious activities are not available.
- 4. Inconsistent authentication** - Implementing authentication becomes more complex without using the load balancer's authentication and session management capabilities. Each application may require its own authentication process, forcing users to log in multiple times and leading to a fragmented user experience.
- 5. Compliance challenges** - Many industry and government regulations and security standards require specific security controls, such as encryption, access control and logging. Without the load balancer's security functionality, organizations may find it challenging to meet these compliance requirements consistently across all servers and applications. Centralizing these functions on load balancers using their security functionality simplifies delivering regulatory compliance.

By deciding to use the security functionality that load balancers have available, you can mitigate the challenges outlined here and deliver the benefits outlined below.



Load Balancing Security Benefits

The benefits of implementing security functions at the load-balancing level include:

- Improved performance through offloading security processing tasks
- Centralized security policy enforcement
- Enhanced visibility and control over traffic
- Detected and mitigated threats before they reach your application servers

Security is enhanced when an organization has IT deployments across hybrid and multi-cloud infrastructure. Using load balancers with WAFs in hybrid and multi-cloud environments significantly enhances security across an organization's entire infrastructure. Doing so allows for unified security policies, providing consistent rule enforcement and centralized management across the multiple deployment platforms.

It also provides improved threat visibility through a consolidated view of security events via a unified console, enabling faster detection of distributed attacks and easier responses to incidents. Additionally, it enables advanced analytics, leveraging aggregated data in dedicated third-party analysis systems to provide better threat intelligence.

Load balancers placed strategically across a hybrid infrastructure help to provide a more robust, adaptable and efficient security posture.

Integrating security functions with load balancing offers these additional benefits:

- **DDoS Protection** - Load balancers can mitigate Distributed Denial of Service (DDoS) attacks, absorbing volumetric floods and detecting and filtering malicious traffic patterns.
- **Web Application Firewall** - When load balancers function as WAFs, they safeguard applications against common attack types like SQL injection and cross-site scripting (XSS).
- **Authentication and Authorization** - Load balancers simplify and strengthen authentication security via pre-authentication and single sign-on (SSO) across systems and platforms to control access to sensitive resources.
- **TLS/SSL Offloading** - Load balancers optimize TLS/SSL offloading and encryption/decryption, reducing the load on application servers while enabling strong data protection.
- **Intrusion Prevention System (IPS)** - LoadMaster offers IPS capabilities for detecting and blocking various network-based attacks (although deploying a WAF is a preferred defense mechanism).

Best Practices for Load Balancing and Security Integration

To maximize the effectiveness of integrating load balancing with security functions, organizations should follow best practices. These include:

- **Plan Strategically** - Define security objectives and align load-balancing strategies with the overall cybersecurity posture.
- **Apply Layered Security** - Load balancing security should complement existing firewalls, Network Detection and Response (NDR) tools (such as Progress Flowmon - ref 5) and endpoint security solutions.
- **Regularly Update Policies** - Maintain up-to-date security policies and threat intelligence to protect against evolving threats.
- **Monitoring and Analytics** - Use load balancer logging and analytics tools for proactive threat detection and incident response.
- **Security Certificate Management** - Load balancers such as LoadMaster can use automated security certificate lifecycle management to renew certificate keys before they expire. The unmanaged expiration of security keys is frequently a root cause of unplanned application downtime.

Implementing Security Functions at the Load Balancing Level with LoadMaster

Every organization should have a multi-layered cybersecurity defense strategy to protect against modern threats. Load balancers can add an extra layer of security functionality and become an important part of this strategy.

As they play a central role in controlling access to web applications, load balancers are in a prime position to inspect network activity and help block any potentially malicious traffic from reaching vulnerable data sources and applications. Load balancers can prevent harmful network traffic from reaching backend servers by detecting anything suspicious in access requests.

Sections of the Progress Support documentation thoroughly explain how to deploy and configure LoadMaster cybersecurity features such as WAF, certificate management, TLS/SSL, authentication, logging and integration with SIEM systems (ref 4).

Each organization will have a unique set of applications deployed and a specific set of LoadMaster instances deployed to manage access requests. The deployment of security components across these load balancer instances will be specific to each deployment. Follow the advice in the support articles and contact the Progress support team for additional guidance. Assistance from the Progress professional services team is also available if required.

Real-World Case Studies - Successful Load Balancing Security Implementation

Many organizations are enhancing their cybersecurity posture by deploying LoadMaster WAF, ESP and other security features on LoadMaster instances alongside their other security infrastructure.

Here are some selected examples. You can read more on all of these examples via our Case Studies page (ref 6).

VANQUISH TECH

Vanquish Tech - Vanquish Tech is a UK-based provider of specialized Global IT services. Using LoadMaster, it has strengthened client IT security and created always-on VPNs for companies with hybrid work models.



Wowrack - A managed cloud service provider with data centers in multiple countries aimed to enhance application performance and security. They deployed LoadMaster as virtual machines and used the LoadMaster RESTful API to automate deployment processes.



Texas A&M University - The University's IT office needed a new load-balancing solution for critical student services. After struggling with their previous system, they deployed LoadMaster, which delivered improved high availability, better security, simplified SSL management, and reduced downtime.



Redplaid - A successful managed hosting provider for SMBs, chose LoadMaster to deliver high availability, enhanced security, and scalability for its customers' websites. LoadMaster enabled Redplaid to improve website reliability and performance and manage costs.

Conclusion: The Future of Load Balancing and Cybersecurity

Data from multiple industry sources predict that cybercriminals' threat will not diminish in the near future. See the January 2024 World Economic Forum and Accenture report titled Global Cybersecurity Outlook 2024 for more (ref 8). As the threat landscape evolves, the need for multi-layered and adaptable defenses will remain paramount.

LoadMaster consistently remains an ideal solution for delivering cybersecurity functionality that enhances other cybersecurity layers for on-premises and cloud application and infrastructure deployments.



Find Out More

The LoadMaster Load Balancer - Read more on how virtual, cloud and hardware LoadMaster load balancers can enhance your organization's cybersecurity posture by visiting our product pages (ref 7).

Contact Us for Expert Guidance - If you're looking for tailored advice on integrating load balancing with other security functions, our team of experts is here to help. Contact us to arrange a chat with our expert team.

Learn More About Load Balancing and Security - For more information on load balancing and its role in cybersecurity, visit the LoadMaster Solution Briefs website. Here, you'll find additional overviews of the LoadMaster security components and features discussed in this white paper.

References

1. **Progress Kemp: Global Server Load Balancing (GSLB)**
<https://kemptechnologies.com/global-server-load-balancing-gslb>
2. **Progress Kemp: Protection for Your Applications and APIs**
<https://kemptechnologies.com/solutions/security>
3. **Progress Kemp: Support — How to configure Intrusion Protection on KEMP Loadmaster (IPS+SNORT)**
<https://community.progress.com/s/article/How-to-configure-Intrusion-Protection-on-KEMP-Loadmaster-IPS-SNORT>
4. **Progress: Documentation — LoadMaster**
<https://docs.progress.com/category/loadmaster-documentation>
5. **Progress Flowmon**
<https://www.progress.com/flowmon>
6. **Progress Kemp: Resource Library — Case Studies**
<https://kemptechnologies.com/resources?content-types=case-study>
7. **Progress LoadMaster**
<https://kemptechnologies.com>
8. **World Economic Forum: Global Cybersecurity Outlook 2024**
<https://www.weforum.org/publications/global-cybersecurity-outlook-2024/>



Free Trial

Try a Progress Kemp Load Balancer Free for 30 days






About Progress

Progress (Nasdaq: PRGS) provides software that enables organizations to develop and deploy their mission-critical applications and experiences, as well as effectively manage their data platforms, cloud and IT infrastructure. As an experienced, trusted provider, we make the lives of technology professionals easier. Over 4 million developers and technologists at hundreds of thousands of enterprises depend on Progress. Learn more at www.progress.com

© 2024 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.
Rev 2024/11 RITM0259641

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA01803, USA
Tel: +1-800-477-6473

 facebook.com/progresssw
 twitter.com/progresssw
 youtube.com/progresssw
 linkedin.com/company/progress-software
 [progress_sw_](https://instagram.com/progress_sw_)