GUIDE

Avoiding CMS Disaster: A four-part guide

- Improve website performance
- Prevent website downtime
- Scale WordPress for high traffic
- Raise WordPress security to the next level

VIP | Enterprise WordPress

Is your website and content management system capable of scaling to meet demand on your biggest day? Making that great first impression is everything in the digital age.

But everything from unknown traffic and poor caching to a CMS ill-prepared for the challenge can slow things to a crawl or take your site offline.

Don't let that happen.

For expert advice from the security front lines, whether you're a content developer or content creator, explore our *Avoiding CMS Disaster* guide—and how WordPress VIP, an enterprise-grade CMS, answers each challenge at scale.

Find out how to:

1. **Improve website performance.** When your business is on the line, you can't afford to send traffic elsewhere or tarnish your brand by having your CMS deliver a poor digital experience. Learn how to mitigate five common slowdown culprits.

2. **Scale WordPress for high traffic.** High-traffic days ought to be cause for celebration, not a nightmare for engineers trying to keep a site up and humming—and your reputation intact. Learn four approaches for enabling a WordPress website to handle traffic tidal waves.

3. **Prevent website downtime.** Lack of caching, untested code, PHP errors, slow MySQL database queries, excessive database writes, we explore the common root causes of website downtime and the role continuous monitoring and other factors play in avoiding that.

4. **Raise WordPress security to the next level.** From a data breach to a site collapse, no organization wants to make headlines. Learn how to battle your sites and applications against malicious attacks, protect sensitive customer data, and stay open for business.

AVOIDING CMS DISASTER PART 1

How to Improve Website Performance

No one likes waiting. And waiting. And waiting some more.

Having a site that performs like a digital superhero is a business imperative, especially when you're making a first impression with a potential customer researching your organization.

The following website performance statistics should strike fear in the heart of every site developer and engineer, not to mention your brand marketing team:

- One in two users will abandon a site if it takes longer than four seconds to load. Source: BBC.
- A one-second delay results in a 16% drop in user satisfaction. Source: LoadStorm.
- Fifty-three percent of mobile site visits are abandoned if pages take longer than three seconds to load. Source: Google.

Perhaps you've noticed poor site performance when you check Google's Core Web Vitals performance metrics, but aren't sure how to remedy that?

When your business is on the line, you can't afford to send new business elsewhere and tarnish your brand by having your content management system (CMS) deliver a poor digital experience.

In the first of our *Avoiding CMS Disaster* series, we diagnose five common slowdown culprits and how to improve website performance using an enterprise-grade CMS.

First, what actually happens when you load a web application?

It's a big ask for any browser, which immediately has to kick into high gear, performing a whole host of actions in milliseconds. Here's the skinny courtesy Mozilla:

1. The browser goes to the DNS server, and finds the real address of the server that the website lives on.

2. The browser sends an HTTP request message to the server, asking it to send a copy of the website to the client. This message, and all other data sent between the client and the server, is sent across your internet connection using TCP/IP.

3. If the server approves the client's request, the server sends the client a "200 OK" message, which means, "Of course you can look at that website! Here it is." It then starts sending the website's files to the browser as a series of small chunks called data packets.

4. The browser assembles the small chunks into a complete web page and displays it to you.

Note: Of course, it's not as easy as all this. There are a multitude of smaller steps between each of these. For now, let's dive into five key reasons "slowness" happens and how our WordPress VIP CMS mitigates them.

So, what causes a slow browsing experience?

Culprit 1 **Lack of PoPs and CDN**

More than likely, your business is global. That increases the logistical challenge of ensuring fast, consistent, stable connections for users on a variety of connection types on a variety of devices.

So what does this mean for your application?

When a network request is initiated from Step 2 above, the request goes through multiple network hops. A hop is a computer networking term that refers to the number of routers a packet passes through, from its source to its destination.

As a consequence of these hops from a user's geographical location to your server, there can be added time experienced by your users during loading. This can be mitigated by serving the content closer to your visitor, via a content delivery network (CDN) and points of presence (PoPs).

> ***The WordPress VIP answer:***
>
> WordPress VIP's CDN is a global network of edge and origin PoPs that serve your site to global customers, as quickly, efficiently, and reliably as possible. This localized availability happens automatically as soon as our edge servers are notified of new content from your application. This means there is no additional configuration to be handled by your team, relieving them to focus on other tasks.

Culprit 2 Unoptimized media

Those high-quality images everyone is shooting today with the latest mobile devices are quite beautiful! That is, until they're used as thumbnails on your homepage.

Adding a 4K image and limiting it to a 100×100 pixel container means your users have to load the whole 4K image. A single image isn't much to fret about, but for businesses publishing content at scale, page load can easily explode exponentially the more articles populate an index page.

> ***The WordPress VIP answer:***
>
> Applications on WordPress VIP automatically receive all the benefits of edge cached images, and automatic image resizing on uploaded images. Additionally, images can be resized on the fly.

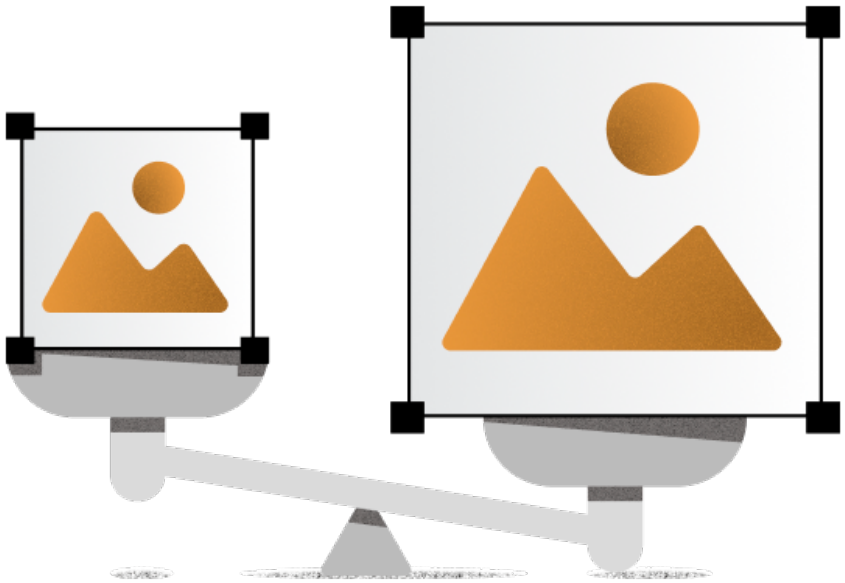

Culprit 3 Unoptimized Javascript and CSS

When loading large files, every bit and every request counts. If you have already optimized the number of network hops taken to download a file, the largest amount of time a file takes to arrive is for the opening and closing of requests.

To address this, modern Javascript and CSS development often uses minification and concatenation.

Minification refers to the removal of all unnecessary characters in a file, normally all unnecessary whitespace. The aim here is to reduce the total number of bits being transferred. Concatenation is the act of gluing multiple files together to produce one larger file. The aim here is to reduce the number of network requests.

> ***The WordPress VIP answer:***
>
> WordPress VIP automatically concatenates JavaScript and CSS files to reduce the number of requests that occur on a single page load. CSS files are minified as well as concatenated. Credit goes to our open-source plugin of choice.
>
> Minified and concatenated files are then cached for 15 days, or until a change is detected.

Culprit 4 Poor cache utilization

Caches are special temporary storage to speed things up. What they speed up depends on the type of cache.

If you are not utilizing a form of caching, you might be seriously impacting your page response speed when your application is under load.

The WordPress VIP answer:

WordPressWordPress VIP employs several different types of caching:

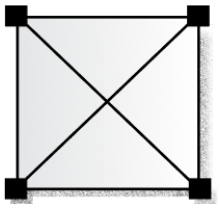

- **Object cache**—used for storing application level data. This helps retrieve the data quickly and efficiently without hitting your database. In short, object cache prevents extensive computations and/or queries from taking up valuable connection time.

- **Query cache**—lightweight memory caching of database queries. In WordPress, any post query that utilizes the WP_Query API greatly benefits from this form of caching.

- **Page cache**—storage for a fully drawn page, stored in each edge PoP location. This is the first level of cache that most users visiting your application encounter. Page cache enables our CDN to serve localized content without your application ever running a line of code.

All of these work in harmony to decrease request response times.

Culprit 5 Database connection saturation

When your application's content and user base grow, those older database queries that were originally performing fine can now degrade performance.

When queries take time to execute, this occupies a seat at the database until the query finishes. Sometimes these seats fill very quickly, subsequently preventing more queries from being executed. This commonly manifests as degraded application performance until the database serves 500 responses.

Connection saturation can occur rapidly and randomly, making the root cause hard to diagnose.

> ***The WordPress VIP answer:***
>
> As part of the WordPress VIP service package, we provide access to New Relic, which helps proactively diagnose and monitor potential problems like those above. This saves your team time, costs, and the trouble of securing additional tools and resources when seconds count.

Final thoughts: how to improve website performance

From network hops and lack of points presence to database content saturation, addressing the root causes of poor website performance should be a P0 priority for any forward-thinking organization and its content management system. After all, making that great first impression with an audience is everything in the digital age. Because speed really does kill.

AVOIDING CMS DISASTER PART 2

Scaling WordPress for High Traffic

"Yikes! We're getting 100 times more requests than normal!"

Why is it happening? Where did it come from? Is our website up for the challenge and capable of scaling to meet the demand?

Everyone knows having a performant website is a business imperative, especially during peak traffic periods in the wake of major marketing campaigns or breaking news.

But in a modern enterprise ecosystem, many other factors, some random and unpredictable, can also cause traffic spikes. Think everything from celebrity endorsements, inflammatory posts, and major events like U.S. Election Night, to deep technical issues like bad deployments of decoupled front ends or mundane bot indexing.

Some events are welcome, driving influxes of new customers and eyeballs. Others can be disastrous for your organization, tarnishing your brand and weakening customer trust. In every case, your engineering team needs to do everything in their power to respond quickly, efficiently, and agilely to keep your site and applications up and humming—and your reputation intact.

In the second of our *Avoiding CMS Disaster* series, we explore four approaches to scaling a WordPress website to handle those high traffic tidal waves.

1. Horizontal and vertical scaling

When you're thinking of scaling your operations as a response to the traffic, you've got two choices: going horizontal or vertical. Each approach has plus points and downsides.

Horizontal scaling

Horizontal scaling is creating more instances of your site or application to cope with traffic spikes. All that is required is adding additional hardware. The goal is to serve more traffic by distributing the load before it reaches your properties.

Key benefits:

- Hardware decisions are significantly easier vs. vertical scaling.
- Eliminates the need to analyze system specifics.
- More servers offers better resilience to traffic.
- Distributing your data across multiple nodes means there is no longer a single point of failure.

What to watch out for:

- Adds significant complexity to your infrastructure setup.
- Increased maintenance and operation costs—one server is much easier to maintain than multiple.

Vertical scaling

Vertical scaling is increasing the server resources allocated to the instance(s) you already have to meet demand. This aims to increase your application's ability to handle more requests by providing more processing power.

Key benefits:

- Easier to decide hardware options, less expensive to implement—upgrading an existing setup is generally cheaper than preparing an entirely new one.
- Less complex configuration vs. horizontal—no load balancer needed to distribute traffic or requirement to synchronize data.

What to watch out for:

- Harder to decide which parts of server infrastructure need to be upgraded or if software needs to be moved to an entirely different server.
- Higher chance of downtime—your application has a single point of failure.
- Limitations to the amount of computing power you can add to a single server.

> ***The WordPress VIP answer:***
>
> On WordPress VIP, your application becomes distributed by default. We employ a worldwide network of engineers and points of presence to safeguard your application, using dynamic auto scaling and our content delivery network (CDN) to regulate the traffic that hits your application or site. All this ensures you stay up and serving content even when demand is spiking.

2. Calling in a third-party CDN

Services like Cloudfront, Cloudflare, and Akamai aim to put a point of presence local to the requester to serve pages. This allows fast response with less network communication, eliminating the need for your application to process the request at all.

The goal is to reduce load, spreading requests evenly to other nodes of the CDN. Some of these services also promise to prevent malicious requests; however, these services are often pricey.

Third party CDNs do provide some protection for your application. But do you have enough time or budget to engage them as traffic ramps up and your application is hitting its limits?

> ***The WordPress VIP answer:***
>
> What if there was a CDN that simply worked for you right out of the box? That's what WordPress VIP's CDN accomplishes. We serve your application as close to the request as possible through our global points of presence, reducing load times and protecting your application from the negative impact of increased traffic.
>
> Still want to use your CDN of choice? No problem. We have experience with all major CDN providers and can assist with connecting to them.

3. Going on the counterattack against attack patterns

How and when do you recognise that your properties are under attack?

Unfortunately, modern attacks—by bots, malware, or malicious, state-sanctioned activity—are more sophisticated, able to tunnel through VPNs or co-opt penetration testing tools to do damage. Worse, attack patterns can't always be mitigated immediately. Even when diagnosed, they can be hard to block.

In the end, they can leave your engineers frustrated and tired, like they're playing an endless game of virtual whack-a-mole.

> ***The WordPress VIP answer:***
>
> WordPress VIP's expertise, infrastructure, and distributed team can proactively control and diagnose attack patterns. With built-in monitoring tools and 24/7 technical support on your side, WordPress VIP helps safeguard your site and application during attacks.

4. Preparing for known traffic by partnering with experts in CMS at scale

What if spikes are expected? Do you ask your engineers to work within budget (and existing resources) to ensure your site and applications remain online? Or do you plan for best- and worst-case scenarios by partnering with experts at scaling sites?

Even the best laid plans and traffic estimations don't always cover the real numbers being served. Consider WordPress VIP customer FiveThirtyEight, "devoted to rigorous analysis of politics, polling, public affairs, sports, science and culture," who shattered their traffic records on the 2020 US Election Night.

Before their site launch, dedicated engineering teams from both WordPress VIP and FiveThirtyEight partnered with 10up (a WordPress VIP development agency) to diligently optimize site performance, including cache efficiency, in anticipation of extraordinary traffic.

All the preplanning paid off. During Election Week, WordPress VIP helped FiveThirtyEight serve an astonishing 1.3 billion page views, hitting a peak of 132,000 requests per second with server response time staying flat at 144 milliseconds under the load.

> ***The WordPress VIP answer:***
>
> Mitigating the downside of known traffic spikes is vital for modern businesses going all in on their digital transformation. During these spikes, WordPress VIP automatically scales your application to meet demand, ensuring a smooth experience for your users. Another safeguard is optional code review by our expert engineers, available as part of our Application Support and higher tiers.

Final thoughts: scaling WordPress for high traffic

Whether it's known or unknown traffic hitting your site and application, preparation and mitigation measures are mission-critical for any organization hoping to harden its properties, user experience, and reputation. From horizontal and vertical scaling to calling in CDNs, today's businesses have a range of options to scale WordPress for high traffic while freeing up their engineering teams from on-call whack-a-mole duties.

AVOIDING CMS DISASTER PART 3

How to Prevent Website Downtime

What does it actually mean for a site to be considered down?

Often that depends on whom you ask.

For a website to be considered down, it may mean a number of different things:

- The website is completely unavailable.
- The website is online but unusably slow.
- The website is giving error messages for certain users or locations.
- The website is working for most visitors, but some simply can't log in to their CMS, for example, to create, edit, or publish content.

No matter the cause or degree, the impact of website downtime can be serious, from lost ecommerce orders and frustrated users to weakened customer trust.

In the third of our *Avoiding CMS Disaster* series, we explore classic root causes of website downtime and the role continuous monitoring and other factors play in avoiding that.

First, the role continuous monitoring plays

We monitor different aspects of a website, so we can tell when something is not working correctly at any of the different layers that make up our fully managed WordPress VIP platform.

Those layers include:

- Network connectivity
- Load Balancers
- Web servers
- Object caching (Memcached)
- Databases
- Elasticsearch
- Files Service (CDN)

We try to spot issues early so that we can anticipate future issues that might affect website stability. Cross-referencing logs from different system components allows us to review periods when a website was reported unstable. Because a combination of factors rather than a single issue might be responsible for downtime, we employ a number of tools to compare data across both systems and applications.

In most cases, website instability is a result of application code, i.e., custom or third-party WordPress theme and plugin code. Here are a few things we look for when investigating an unstable site, and how to mitigate each.

Not enough caching

The most important thing you can do to ensure a site is performant and stable is to make sure any full page that can be cached, is cached. Uncached pages need to be built on the server each time they are requested, which is a slower process and more prone to errors.

> ***The WordPress VIP answer:***
>
> WordPress VIP Platform provides powerful page caching via a global network of edge cache servers, each used to store and serve content closest to an end user. The response time from an edge cache server is almost always a magnitude faster than anything that bypasses page caching and hits the origin servers.

Caching challenges

Because they demand a personalized, fully interactive experience, some sites, particularly ecommerce ones, simply can't be cached at the page-cache level.

Often a compromise can be found whereby a static page is served by edge cache, with dynamic features (e.g., logged-in status, shopping carts) added via JavaScript. Asynchronous requests from JavaScript can then be used to communicate with a WordPress REST API endpoint designed with a much lower overhead than a full page load.

Alternatively, this is where object caching comes into play. The page can remain dynamic but parts of the page and any data used in it can be stored and retrieved in object cache to avoid needing to query the database.

> ***The WordPress VIP answer:***
>
> Each WordPress VIP application environment has its own dedicated Memcached cluster, which stores object cache data in memory for lightning quick and efficient retrieval.

Untested code deployments

This is another common culprit of website downtime and pretty easy to diagnose, based on pure cause and effect.

If your website has just deployed untested code, leading to immediate site issues, there's your likely cause. If you can, revert the suspect code to the previous version ASAP.

The best thing to do to avoid this situation? Thoroughly test every piece of code on a separate development or staging environment before releasing to production.

> ***The WordPress VIP answer:***
>
> Because all our site deployments are via GitHub, WordPress VIP customers can easily revert code themselves, without losing any new code changes, which remain stored safely in the GitHub revision history. Optionally, in emergency situations, we can rollback a customer's website to a previous deployment on their behalf, independently from GitHub.
>
> Regarding environments, all applications hosted on our fully managed service can have a separate development or staging environment. Syncing data there from production is easy, letting you test code against the same amount and same type of data as on your production website.

PHP errors

WordPress uses PHP code on the server. A PHP error might be "fatal," meaning that once the error occurs, the web page, script, or command will stop running. These will almost always surface as visible errors somewhere, and will be recorded in the PHP logs.

Note: Some PHP warnings in PHP 7 become fatal errors in PHP 8, so it's important to take these errors seriously.

> ***The WordPress VIP answer (plus helpful advice):***
>
> Our platform automatically logs all PHP errors, making them available to WordPress VIP customers in their dashboard and to our engineers.
>
> **Pro tip:** Address and fix all PHP errors—even if a site appears to be working fine. Routinely, we see logs full of PHP errors, even fatal ones, on a site that appears stable. However, that doesn't necessarily mean the site is working correctly. Keeping PHP logs clear by addressing minor errors and warnings makes it easier to find more serious errors during debugging.

Slow MySQL database queries

Every WordPress website uses a database to store website content and configuration data. Database queries fetch that content data for web pages, but sometimes those queries are written inefficiently. They may work fine for sites with only a few hundred pages, but stall when handling large amounts of data (some websites on our platform have millions of stored records).

A slow query ties up database resources, potentially impacting site stability—not just for the page, script, or command running the SQL, but across the whole application. Sites often struggle because single or multiple database queries are slow, e.g., any query that takes longer than 0.75 seconds to execute.

> ***The WordPress VIP answer:***
>
> WordPress VIP helps mitigate database bottlenecks by providing each application with a dedicated database cluster featuring a primary database, where all database write queries occur, and one or more read-only replica databases.
>
> This increases the number of simultaneous database queries that can take place, spreading the resource load when a site is under pressure.
>
> That said, slow database queries can't always be resolved simply by adding additional database resources.
>
> That's why we advise customers to monitor slow database queries by using Query Monitor and New Relic (provided by our platform). These highlight where queries originate in the database, so your development team can refactor them to optimize performance.
>
> Finally, our Application Support and Premier Engineers can also help your team find and analyze these queries, and suggest ways to improve them for speed and efficiency.

Excessive database writes

Sometimes a feature, such as custom logging or tracking code, updates the database on every request. This can lead to instability for two reasons:

- **Foregoing database replicas:** All write queries are directed to the primary database; subsequent database queries for the same table (or tables) in the same page request will also be directed there. By not taking advantage of database replicas, this limits the scalability of the site.
- **Bypassing page caching:** For a database write to happen on every page request, page caching must be bypassed. But doing so means the first (and best) line of defense has been compromised.

> ***The WordPress VIP answer:***
>
> In these circumstances, we advise refactoring the feature. For example, content analytics is usually best delegated to an external service that uses a snippet of JavaScript in the page rather than server-side code, which doesn't work well with caching and may result in excessive database writes.

Other known causes of downtime and how to avoid them

Plugins

There are thousands of popular, helpful third-party plugins in the WordPress ecosystem that provide fantastic features and functionality. Some, though, have challenges scaling, potentially leading to downtime issues when added to a website with tons of content and traffic.

> ***The WordPress VIP answer:***
>
> As good ecosystem stewards, we regularly reach out to vendors with suggestions to make their plugins perform better in high-traffic environments. We can also suggest alternative plugins that have been tried and tested at scale on our platform.

Custom logging

Custom logging is a powerful debugging tool, often the only viable method to track down a bug or issue that seems to happen only on a production site. On numerous occasions, however, we've seen custom logging built in PHP on a high-traffic site slow down things or put a site in danger of downtime through excessive database writes.

> ***The WordPress VIP answer:***
>
> For customers, we provide access to standard PHP logs in the Health panel of the WordPress VIP Application Dashboard. There they can log custom errors (and also to New Relic), which will not negatively impact the database.

Remote API calls

Some websites take advantage of server-side REST API calls to other applications or services. These are pretty fast under normal circumstances, but sometimes the underlying application code leads to a slow response, times out, or throws an error.

> ***The WordPress VIP answer:***
>
> To minimize these issues, we advise "defensive coding." It depends on the purpose of the remote call, but often when a remote request fails, it's possible to fall back on a cached response from a previous request—or at least "handle the error gracefully," so that the rest of the page can still load. We provide a number of helper functions to handle these scenarios. Keeping a low timeout also means PHP resources are freed up quicker if the remote API is not responding.

Final thoughts: how to prevent website downtime

Website downtime costs more than lost business—it tarnishes brand and reputation in the eyes of customers. Fortunately, whether it's lack of caching, PHP errors, or slow MySQL database queries at fault, today's enterprise organizations have a fulsome toolkit of technical options—continuous monitoring and edge cache servers to defensive coding and tools like GitHub and Query Monitor—to mitigate those root causes of a site going offline.

AVOIDING CMS DISASTER PART 4

Raising Your WordPress Security to the Next Level

No one wants to be in the headlines for a data breach or catastrophic site collapse.

Threats online today exist in various forms—from bots scanning for known vulnerabilities to script kiddies, cyber gangs, and even nation state actors.

For beleaguered IT teams, ensuring an online presence isn't compromised can feel like more than a full-time job, "running just to stand still." Unfortunately, that often leaves precious little bandwidth for the real task of preparing a site, its infrastructure, and business for its future state.

In the fourth of our *Avoiding CMS Disaster* series, we explore five ways organizations and their IT guardians can battle-harden their sites and applications against malicious attacks, protect sensitive customer data, and keep their business always open for business.

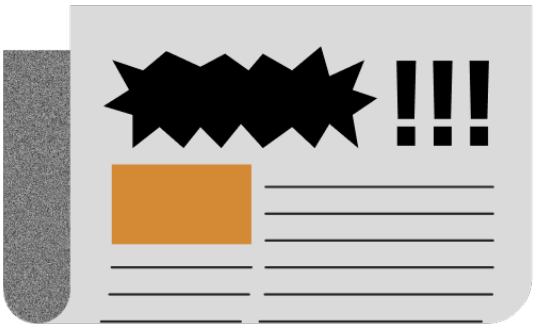

1. Vulnerability management

Identifying and mitigating software vulnerabilities can be an overwhelming task, even for the largest IT teams.

Vulnerability management requires on-going identification, classification, prioritization, managing, and remediation of vulnerabilities within your organization's infrastructure.

According to the UK's National Cyber Security Centre, "Exploitation of known vulnerabilities in software remains the greatest cause of security incidents."

Vulnerability scanning

Periodic scanning of your network to monitor for any new vulnerabilities or unintended open access to machines is a must for any organization that wants to minimize their attack surface.

Unfortunately, maintaining scanning software, reviewing the scan results, and actioning them often fall down the list of priorities in a busy IT department.

> ***The WordPress VIP answer:***
>
> Our platform provides security throughout your network—from edge security to protection of data in transit between components. For example, DDoS protection continuously monitors web traffic and takes active mitigation steps when suspicious activity is detected. Network and host-based firewalls with real-time notification processes are there to prevent unauthorized access attempts.

Vulnerability remediation

Keeping the application layer updated is just one step in the process. All the supporting layers and infrastructure need to be kept up to date, too. Sometimes vulnerabilities are exposed that don't have an immediate patch or require mitigation in your codebase.

To keep your applications infrastructure up to date requires multiple steps, including:

- Checking for updates for each piece of software
- Building the new release
- Testing on your non-production architecture
- Ensuring no new issues are introduced and subsequently fixing any that are
- Putting your production application into maintenance mode and rolling out your updates

Note: All these steps should be done with every patch available at every layer of your application.

Specifically for WordPress, it's not just keeping core and underlying infrastructure up to date. Third-party themes and plugins must be updated and patched regularly.

It's also important to recognize the quality of third-party plugins added to your WordPress site. Some might be poorly coded or introduce security vulnerabilities—neglectfully or maliciously. Tools such as WPScan, SonarQube, or PHP_CodeSniffer can help automate your code reviews to catch unwanted exploits being introduced.

The WordPress VIP answer:

Our platform is managed by active members of the WordPress community. When an issue arises, we immediately patch it, often ahead of the fix getting pushed to WordPress core code.

Meanwhile, we proactively alert all customers of: 1) upcoming, automated WordPress updates, checking to make sure they're on the latest version of our platform, and 2) plugin exploits in the wild and attempts to patch these issues at a platform level.

Going deeper, we use automated code scans for pull requests created in an application's GitHub repository. This can identify potential security concerns before code goes into production (and is useful for evaluating plugins from the wider WordPress ecosystem.) Leveraging our Kubernetes orchestration, WordPress VIP provides zero downtime deployments for customers' applications.

Finally, based on years of experience running WordPress at scale, we can mitigate common attack vectors—thanks to continuously testing our infrastructure for vulnerabilities and engaging independent third parties to run penetration tests against it.

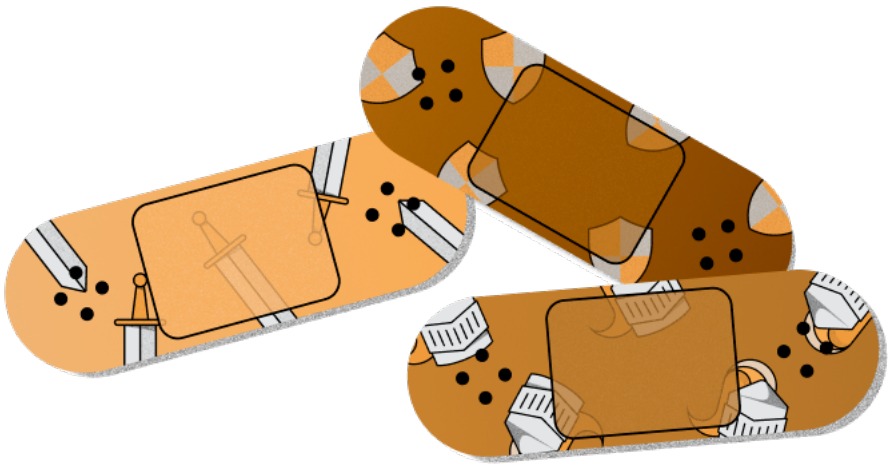

2. Network security

Network security is a vital part of an organization's online presence.

Best-in-class security means managing both perimeter-based security and internal network security. Here multiple factors must be considered and managed to effectively protect users and their data.

Intrusion detection systems

Monitoring and logging all network traffic is essential in identifying malicious or suspicious activity. To prevent unauthorized access, security teams need either automated rules or alert system administrators to review suspicious traffic and take appropriate action.

Firewalls

Understanding what traffic is allowed to traverse your network and how applications communicate within and outside the network is a must to minimize security risks. This means setting and reviewing ingress/egress rules using either software or hardware firewalls, which allow only essential network traffic for your applications to run.

Blocking or allowing the wrong traffic could prevent a vital system from performing or, worse, expose your database to the world.

Physical network security

A network is only as secure as its physical security. The best firewalls, intrusion detection systems, and threat management software can all be circumvented by a malicious actor gaining physical access to your servers.

Data centers require multiple levels of physical security, such as:

- Physical access controls
- Environment monitoring
- Alarms and sensors
- Surveillance
- Backup power

All these also need periodic auditing to ensure they are meeting security best practices.

> ***The WordPress VIP answer:***
>
> To permit only authorized traffic, we monitor and control communications at the external boundary of our infrastructure and at key internal boundaries. Automated alerts and continuous logging at every level of our stack help our experts intervene when needed to keep your site secure. Our data centers also monitor networks of hundreds (even thousands) of sensors used for real-time telemetry, ensuring peak performance from our hardware.

"We wanted to use WordPress, and the fact WordPress VIP was the only WordPress option with FedRAMP authorization to operate (ATO) made it a strong option for us as a government agency."

Kristen Loflin
Public Relations Manager
Marine Corps Marathon Organization

3. Data protection

No company wants their data leak showing up on haveibeenpwned.com for all the world to see. Users lose trust fast with companies that don't protect their data to the level they expect. Ensuring only authorized, role-specified users can access sensitive customer data requires multiple layers of protection.

Data encryption

Communication between your application and your users needs to be encrypted in transit to prevent data being intercepted or tampered with by a third party. Transport Layer Security (TLS) is generally used to encrypt the data. This requires creating TLS certificates and ensuring they are renewed.

Because data can also be encrypted at rest, this requires protecting data kept on storage media such as backups. If a malicious actor does gain access, they will still need the data encryption key to actually use the data.

Key management is an important part of data encryption. Keys have several stages in their lifecycle: generation, distribution, use, backup, rotation, and destruction. At each stage, there are best practices to follow to keep your data safe.

Audit trails

Keeping a chronological record of all activities that occur within your application at every level of the stack is essential for any enterprise. Audit logs are required for forensic investigations, detecting security breaches and their impact, and understanding systems errors.

Audit trails need to collate logs for multiple layers and applications, secure enough so they can't be altered by users. And they must ensure chronological accuracy. This requires knowing what actions need to be logged, connecting multiple systems into an ELK tool such as Kibana, synchronizing systems using network time protocol (NTP) so timestamps are meaningful, and managing access to logs.

Automated alerting

Knowing what defines a security incident, what requires manual assessment, and how they should be managed is an art in itself.

Automated log analysis can flag suspicious behavior at an early stage (if you know what to look for). Tools with predefined or custom rules can be created specific to your application. This requires setting parameters within your log analysis to know when:

- Automated actions should be executed to protect against attacks and malicious traffic.
- The Systems team should manually intervene to examine the patterns and try to determine whether it's falsely flagged benign behavior or action is required.

All this relies on well-configured tools and an experienced security team that understands your applications' usage patterns. No enterprise wants a piece of content to go viral, only for their system to flag it as a DDoS attack and block the traffic.

> ***The WordPress VIP answer:***
>
> We maintain separate containerized database infrastructure for every client and application, each with their own unique authentication. This mitigates the risk of unauthorized access between applications and protects each customer's data and reduces the risk of attack.
>
> We provide database, file system, application, and data center security, as well as hourly encrypted backups. And our origin data centers meet the International Organization of Standardization (ISO), International Electrotechnical Commission (IEC) 27001 certification, Standards for Attestation Engagements (SSAE) No. 18 (SOC1) and SOC2 Type 2.

Case study:

For a media outlet with the global reach and influence of Al Jazeera, it was essential to migrate its properties to a CMS platform hardened against malicious actors. Read why they chose WordPress VIP.

"Security and privacy are things that need to be top of mind—we have a big bullseye on us, which forces us to be hypervigilant."

David "Hos" Hostetter
Digital CTO
Al Jazeera Media Network

4. Access and authentication

According to Telesign, more than half of consumers use five or fewer passwords across their entire online life and almost half of consumers rely on a password that hasn't been changed for five years.

Gaining access to a user's account may be one of the easiest ways to access a secured system. That's why granular access control, multifactor authentication, and/or single sign-on are so important for security-conscious organizations.

Access controls

Granular access controls and implementing a policy of "least privilege" is vital to keeping your data safe and reducing attack surfaces for your application. The policy of least privilege states that a user should be given only those permissions needed to complete their task. Ensuring every user does not have administrator privileges, for example, means that if a malicious actor does gain a user's credentials, the likelihood of them being able to do significant damage is limited.

Single sign-on (SSO)

With SSO, users login to multiple services via one set of login credentials. If a user does not exist in a specific service, they can sometimes be provisioned on the fly by utilizing user mapping from the service's Identity Provider. Services like Azure AD, Google Apps, AuthO, or OneLogin provide SSO functionality.

SSO helps IT departments set centralized rules for users, reduce time recovering lost passwords, and remove the need for manually provisioning and deprovisioning users during onboarding/offboarding.

Multifactor authentication

Using MFA provides a further layer of protection against your organization's users being compromised.

MFA requires a combination of at least two methods of authentication to login. Generally it will be configured with a username and password as the first layer of authentication followed by a time-based authentication token generated via a hardware device or software like Google Authenticator. The benefit of this process is that even if the username and password are compromised, a user can't login without the authentication token and vice versa.

> ***The WordPress VIP answer:***
>
> WordPress VIP is built on a foundation of granular access controls and permissions, including multifactor authentication, brute force protection, data access audit trails, and physical security. These provide an extra layer of protection against compromised passwords, prevent unauthorized employees or contractors from accessing customer data, and dynamically apply restrictions at the network level when unnatural behavior is detected.

5. Breach recovery

Automated backups and hardware redundancy is vital to the smooth running of your day-to-day online business operations.

Backups

Backups are vital for data loss prevention, preventing ransomware attacks, and quick recovery from outages. There are a number of backup best practices every organization should follow to ensure they have full control and redundancy of their data.

1. **Regular backups.** The more frequent the better in terms of reducing your Recovery Point Objective (RPO) and minimizing data loss.

2. **Backup redundancy.** Storing backups in multiple locations (e.g., offsite) ensures you can still access them, even if you lose access to your main server.

3. **Encrypted backups.** Even if your backup storage is compromised, the data will be useless without the encryption key.

4. **Regular testing.** Regularly extract your backups and test them in a non-production environment to ensure your team can actually restore your site with them.

Hardware redundancy

Having backups available is little to no use without backup hardware to restore to.

This requires redundant hardware within and outside your primary data center. No matter if the issue is with a single server or the entire data center, your team will be able to access this hardware to quickly get back online.

> ***The WordPress VIP answer:***
>
> In the unlikely event of a breach, we help customers quickly recover and get back to business, thanks to multiple levels of backup (origin datacenter and offsite locations), plus disaster recovery and security breach procedures.
>
> We also provide the ability to automatically ship your backups to your own S3 storage to ensure you can set your own data retention policies on them or even run automated recovery testing. Utilizing multiple levels of redundant storage, we can reconstruct data in its original or last-replicated state before the moment it was lost. WordPress VIP also has multiple origin data centers that sites can be migrated to in the improbable event of a single data center failure.

Hear more from our experts: Watch the webinar *Four Strategies to Run WordPress at Scale*

Dive deeper into exclusive insights and best practices inspired by our *Avoiding CMS Disaster* series for creating better customer experiences on the enterprise web.

Watch on-demand

Final thoughts: raising your WordPress security to the next level

From vulnerability management to breach recovery, working WordPress VIP gives organizations the opportunity to leverage years of experience keeping high-profile, high-scale WordPress-based sites online and secure in the face of threats.

Built with multiple levels of security controls and protection—including edge protection, secure networking, robust access controls, continuous security monitoring, and code scanning—WordPress VIP meets the most exacting security requirements.

That's why it's trusted by customers in high-risk industries such as banking, pharmaceuticals, public utilities, and government. We're also the only WordPress platform to achieve FedRAMP Authority to Operate (ATO).

The best companies run the web with WordPress VIP.

WordPress VIP combines the ease and flexibility of WordPress—the CMS that runs 43% of the web—with unmatched scalability and security for the enterprise. Our solutions are trusted by iconic media titans, major brands, and government agencies like CNN, Salesforce, News Corp, The White House, NBC Universal, Capgemini, and Bloomberg. With WordPress VIP, brands can scale their web presence, enable their teams to produce more web content, and use data to continuously improve content performance, eliminating wasted effort while maximizing ROI.

Learn more at wpvip.com.