

Basics of Hosting

Cal Evans
cal@getpantheon.com
<http://getpantheon.com>





Each year as his team gathered together for the first time, Vince Lombardi held up a football and said those famous words...

“Gentlemen, this is a football.”



Fundamentals are everything. Professional football players don't need to be introduced to a football, they have been playing the game all their lives. In showing the football, and later the field, the goal posts, the boundaries, Mr. Lombardi was reinforcing the importance of the basics of the game. Once you understand the basics, you can move on to the more advanced topics. If you don't get the basics, you will never get the advanced topics correct.

Ladies and gentlemen, this is hosting.

The Basics - Part 1

- RAM
- Filesystem
- CPU
- Bandwidth
- Operating System



No, your website does not live in “the cloud”. It lives on a server somewhere. That server has a finite amount of resources available to it. Your website and all of it’s moving parts consume those resources. You can’t decide if you need more resources though, till you know what they are and how they affect your site.

RAM - Random Access Memory

This may be the most important resource available to your website. If you have too little, you are starving your processes and causing things to have to swap to the hard disk. If you have too much...well, let’s be honest, there’s no such thing as too much RAM. When considering a hosting partner for your website, make sure you ask how much RAM will be dedicated to your site.

File System

This is also called “Persistent Storage”, or simply hard disk. it is where you store things long term. Important things like pictures, media files, etc. There are two important things to ask about when talking Filesystem, how fast and how much. You want to make sure that you have enough room to store everything. you also want to make sure that the web server can get files off of the File System quickly. These days speed usually means SSD or “Solid State Disks”. When considering a host for your website, look for SSD. If you can’t find it, ask before laying out your money.

CPU

The heart of any server is the CPU or the Central Processing Unit. CPU used to be the absolute most important factor in hosting, these days however - especially with the ease at which you can spread the load around between servers, it takes a back seat to RAM. That having been said, make sure you know how fast the CPU of your server is and how much of it is guaranteed to you. Most shared hosting providers will not be able to give you a guarantee of a specific portion of the CPU’s cycles - the unit of measure for how much of the CPU is dedicated to you. Virtual Private Servers will be able to guarantee you a specific percentage of the CPU’s cycles. Of course, if you purchase your own physical server, you get 100% of the cycles.

Bandwidth

Bandwidth is the lifeblood of any website. This is the “pipe” your website travels through to get to your customers. You will hear vendors talk about “fat pipe” meaning that they think they have a fast connection to the Internet and enough bandwidth to move your site to your customer quickly. When talking about bandwidth, there are two important factors, “transit” and “transfer”. Transit is how fast file can be moved out of the datacenter - where the server lives - and to your viewer’s browser. Transfer is how much can be transferred at one time. Most hosts will discuss transfer, but not transit. It is important for your site to have both fast transit and high transfer. Sacrifice either of these and the performance of your site will suffer.

Operating System

What operating system your site will reside on is important not only to the overall performance of your site but also to the development of your site. If your developer is working on Windows, make sure that you host your site on Windows. Yes, you can easily develop a website on Windows and host it on a Linux based computer. There are however, differences between the two that your developer will need to know about. In general it is much easier to make sure that your development, testing, and production sites are all running on identical environments than it is to remember to keep everything straight.

The Basics - Part 2

- Your Domain Name
- Your Registrar
- DNS



There are several other basic facts that you need to consider before selecting a host. These aren't quite as important as Part 1 but they are very important, none the less.

Your Domain Name

One of the most critical parts of any website is the domain name. It is how people find you. I see so many people however getting this one wrong. Many site owners allow someone else to register - and own - their domain name. I had a customer once, a very famous singer in his genre, who did not own his domain name. The person who built his first website did. It was never a problem to get a change made, but he was always at the mercy of this person. Should this person get mad at him, or worse yet, should this person pass away, changes would be a lot more difficult.

If you allow your host - or anyone else for that matter - to register your domain, then you do not own your domain, they do. It is very important, so very vitally important, that you register your own domain name. You must register it, you must own it. Anything else is wholly unacceptable and will cause you problems down the road. If you own and control your domain, most other problems can be solved. If someone else owns it, you are at their mercy.

Registrar

Your host is not your Registrar. Your Registrar is where you register your domain. This is also where you put in important information like who owns your site, who is the technical contact, and what DNS servers are responsible for your site. There are a lot of good registrars out there, there are a lot of bad ones too. Ask around if you are not sure. If all else fails, email me and ask me. I'll share with you who I use and who I do not...and why.

DNS

Domain Name Service is how the internet finds your domain name and points a visitor to your little piece of the cloud. When you register your domain, you will have to enter the DNS servers that are "authoritative" for your domain. Whomever you choose to host your site will give you this information, you will enter it into your Registrar's system for your domain name. As long as you own your domain name, you control which servers are authoritative for your domain. This means that no matter what, no one can take your site away from you. Your host may be able to shutdown your account because of a disagreement, but you can always point it to a new host and get SOMETHING back up and running.

Hosting Options

- Shared Hosting
- Virtual Private Servers
- Containers
- Physical Servers



Shared Hosting

This is bad; so very, very bad. Shared hosting is how the web started. Back in the day, when websites were static HTML files, you could host a lot of them on a single server. It was energy efficient, and easy to manage. Then along came scripting and things changed. Regardless of what language you are using, the fact is that your web server is running the script and if it has permission to see your site and another site, the potential for mischief and damage is there. In the PHP community we have partially solved this with security patches like suhosin. Suhosin helped protect sites, but not all hosts implement it. Also, it isn't foolproof. The fact of the matter is that if you are on shared hosting, your site is a ticking time bomb waiting to explode. As soon as another site on your server gets compromised, all other sites on that server are vulnerable.

Security isn't the only problem though with shared hosting. Resource allocation and "noisy neighbors" are a serious concern, as well as your site accidentally being associated with sites that you may not want to be associated with.

Noisy Neighbors

In shared hosting arrangements, the host can strictly control the amount of the file system that any given site can use. It is much harder however for them to control the amount of bandwidth or CPU that a given site uses since everyone is sharing the same instance of the web server. Many hosts will shut a site off after they have exceeded their bandwidth transfer cap, but this is at best a stopgap. Until a site is shutdown it can consume the lion's share of the transfer available to the server.

Noise neighbors can also take the shape of a runaway script. Whether by planning - or by poor planning - some scripts can start consuming resources - RAM and CPU - until it either brings the server to a halt, or someone kills it, assuming someone is watching at all.

Guilty by Association

If you are on shared hosting and you are not paying for your own private IP address, your site is hosted on the same IP address as other websites. Sometimes that may be one or two, other times that may be hundreds. The problem is you don't have control over what other sites are on your same IP address or even on the server with you. If one of the websites sharing your website's IP address is a bad actor, you could very quickly find your IP address blacklisted and all of your emails bouncing. Getting into an email blacklist is easy, getting out of one is difficult. It is a situation you do not want to be in. If you are hosting on shared hosting, you may not have any options though, other than to move hosts. Sometimes a host will assign you to a new IP address without you having to move, but if it is still shared, you run the risk of having it happen all over again.

Shared hosting was a good idea once, a long time ago. These days however, it is a supremely bad idea for every website. If you are only motivated by the price of the hosting then you will get exactly what you pay for. There are many web hosts that will charge you \$5/month for unlimited everything. There is at least one that will charge you \$5/year for unlimited everything. Since resources cost, the only way they can stay in business is to "oversubscribe" servers and hope that all the websites they are hosting are incredibly low traffic sites. They also hope that they don't rent to users who have no serious interest in hosting a website but are looking to compromise other websites on the same server. Either way, it's a crap shoot and all you can do is watch the dice roll.

Better (?) Options

- Managed Hosting
- Platform as a Service



I am old-school, I manage my own VPS. I used to own my own server but I got tired of disposing of them in an environmentally responsible way. Now it's someone else's problem. I do not recommend this though unless you just truly love compiling software and editing config files. It is a thankless job that requires you to get up at 2:13 AM some days and fix the *&^% database server. You can however, still have a rock solid platform for your applications and get a good night's sleep. There are a couple of solutions for you.

Managed Hosting

Managed Hosting providers will usually give you a VPS or a physical server for your application. They will manage everything about it, including things like updates to your core application. In our space this means that if you host with a Managed Hosting provider who does WordPress, they are responsible for your application from the core, down to the wire. You are responsible for everything else. Themes, plugins, custom code, etc. Since your Managed Hosting partner is responsible for your infrastructure and part of your code, there may be limitations on what you can and cannot do. This is a reasonable requirement given that they are responsible if anything goes wrong. Before you select Managed Hosting, make sure you understand any and all limitations that will be put on you. More importantly, look at your existing application - and your growth plan - and make sure you can live with the limitations.

Platform as a Service

In the last couple of years, another option has arisen, Platform as a Service. Companies that provide you with the infrastructure you need to get the job done without you having to manage - or even worry - about that infrastructure are considered PaaS providers. The advantage of a PaaS provider is you are divorced from the infrastructure. Not just the server, but the entire infrastructure. Most PaaS providers working in the WordPress space take care of everything.

- The server
- The network infrastructure
- Load balancing
- Caching
- File systems
- etc.

When running a web infrastructure, there are a lot of moving parts. Things that most people don't want to bother with. You really only have a few choices.

1. Learn to manage all of it and become an expert
2. Hire someone to do it for you
3. Go with a managed solution
4. Go with a PaaS

PaaS is different from managed solution. PaaS providers - in most cases - stick with the underlying infrastructure. The line of demarcation is usually your application.

DTAP

- **Development**
- **Testing**
- **Acceptance**
- **Production**



Half Way Point

There is more to hosting your site than just the final server and resources

Some do it all on the same server

Some do it on different servers on the same provider

Some do it on different providers

Consistency across all platforms is what is important.

SFTP

- How do you move files to the server?
- FTP vs. SFTP



Source Code Control

- GIT
- SVN
- Others



Database

- What RDBMS does the host run?
- Do you have access to the database?
 - Direct Access?
 - ssh tunnel?
 - PHPMysqlAdmin?



There are hosts that do not give you access to the database at all. You have to decide if you need it. If not, don't worry about it. If you just need it for reporting purposes, use a backup tool to pull it down and install it in a local server. Sometimes however, you will need access. Make sure you ask the question.

SSL

- Do you really need SLL?
 - Encryption
 - Identification
 - SEO Ranking
- What Type of SSL do you need?
 - Encryption level?
 - Wild Card?
 - Self-Signed?
- Vendors



SSH

- **YAGNI** - You Ain't Gonna Need It
- You may think you need it, you do not.
- Tunnels are the exception



Secure Shell

Old school developers like myself can't imagine using a server that I can't ssh into; it's just not natural. That being said, the biggest problems I have ever had with a site can be traced directly back to me doing something stupid while being sshed into the server.

These days ssh access is not necessary to run a complex website, regardless of what you are using to power it. WordPress, properly installed, has everything you need to install plugins and themes. If you need ssh for something, you probably need to reconsider what you are doing.

More than that though, if you are using Managed Hosting or a PaaS, you have turned over your infrastructure to someone else. If they are responsible for your infrastructure then you do not need need access. More importantly, you do not deserve access. It is their domain now, not yours. Just because your site is running on top of it does not mean that you somehow deserve access to the underlying infrastructure.

The one exception to this is "tunneling". SSH tunneling allows you to connect a local program to a remote service, securely. For instance, tunneling a connection to MySQL is a good thing. It allows me to use my existing MySQL tools to manage my database

Recommendations

- Do your homework
- Know your limitations
- Research your options
- Ask someone



Do your homework

Before you even start to try and figure out where to host your site, make sure you understand what you need. You can't make an intelligent decision without understanding your needs. Selecting a hosting partner without first determining your needs leads to selecting them based on Logo, swag, or other irrelevant factors. Know what your app needs now, know what you will need in the future. Then start the process.

Know your limitations

Don't select physical server if you've never managed a server. Also, don't select managing your own infrastructure if it is just you. No matter how good you are, without a backup, you are on call 24x7x365. Kiss that vacation goodbye, along with your significant other, and any other aspect of your life. Being the sole "DevOps" person in a company is worse than being married, you are attached to the server.

If you are going to manage it yourself, make sure you understand what that means, and make sure you have at least 2 other people working with you on managing the server/infrastructure.

Research your options

In this talk, we've covered several options. These are the big categories, there are an infinite number of subcategories of hosting providers. There are a lot of options for you to consider. To pick the right one you can either Do your homework, research your options, and pick the one you think best matches you. Or you can throw darts at a wall full of logos. If you choose the latter, you get what you deserve.

Ask Someone

It is always a good idea to get a second opinion. Most of us know someone involved in web development. Don't assume you know everything you need to know. Take a developer out for a coffee and get their opinion. If you don't even know where to start, take a developer to lunch and ask them. Developers usually respond positively to bribes of coffee or food. sometimes even adult beverages, if appropriate. We are usually willing to offer advice, help go over options, and generally guide the way. be respectful of your developer's time. A \$4 coffee will usually get you a half-hour consultation. A \$20 lunch will get you a full hour. If you want more, be prepared to pay. If you want someone to make the choice for you, and be responsible for the decision, yes, you are going to pay full rate.

Checklist

- How many visitors do you expect each month?
- Does the Host specialize in WordPress?
- Are there restrictions on plugins or themes?
- Can you install an SSL certificate?
- Do you get SFTP access?
- Do you get SSH access? (Do you really **really** need it?)
- Do you get direct access to the database?
- Other than WordPress, do you need to run something else?



Thank you



Cal Evans

b. blog.calevans.com

t. [@calevans](https://twitter.com/calevans)

w. getPantheon.com

e. cal@calevans.com



Questions?