

NGINX



© 2012 NGINX, Inc.

What is NGINX?

NGINX provides a unique combination of web server, caching proxy and load balancing solution to any web site that just wants to be consistently efficient. Because of its design and architecture NGINX has already enabled more performance, scalability, reliability and security to many organizations across the world. Today NGINX is the 2nd most popular open source web server on the Internet.

History of NGINX

NGINX was written by Igor Sysoev, a Russian system and software engineer. From the very beginning NGINX was created with a vision on how the pressure for content and the increasing number of web and mobile users will transform the essential infrastructure elements used by the Internet focused companies.

Igor started offering NGINX to the open source community in 2004 and since that time NGINX has already become a vital part of the Internet DNA.

Who uses NGINX

NGINX has been really successful with customers in a variety of market areas which include Service Providers, Media & Entertainment, Online Retail, Social Media and more.

Today over 10% of the top 1,000,000 Internet web sites and 25% of the top 1,000 web sites run on NGINX, including Groupon, LivingSocial, Playdom, Zappos, Hulu, TechCrunch, Dropbox, Yandex, WordPress and many others.

Hosting providers like DreamHost, Media Temple, CloudFlare, Engine Yard and Linode have been deploying NGINX as part of their infrastructures to achieve more density, scalability and predictable performance and to ensure best quality of service for their users.

Overall NGINX already powers over 60,000,000 domains and this number is growing.

Challenges of today's Internet

The Internet has evolved from a simple HTML information browsing medium into a sophisticated, always-on information superhighway used by more than 2 billion users worldwide for social media, content sharing, video on demand, online transactions, telephony and more.

With the ongoing consumerization and with the proliferation of permanently connected mobile devices, tablets, laptops and even set-top boxes, Internet landscape is rapidly changing and the global economies become digitally wired.

Today people tend to continuously expand their daily Internet habits towards exponentially growing demand for instant, proactive and responsive online experiences.

Younger generations are live on Internet almost 24-hours a day, machine to machine communications steadily improve, and more and more major businesses and institutions start to deliver online services.

All of the factors above create ever increasing pressure on web-based businesses.

Currently applications used on the Internet strive to maintain the hard balance between delivering content and services quickly and using server infrastructure cost efficiently. Permanent readiness to serve the users today requires efficiency and scalability of the server infrastructure more than ever.


NGINX helps Internet companies to protect, develop and propel their web infrastructure into the future. NGINX encompasses years of experience in facilitating dynamic infrastructure growth. We believe that the agility of your web servers today is what enables the existence of your online business tomorrow.

In fact, NGINX has been a software-to-the-rescue for thousands of installations worldwide. But it is also a logical step towards using the next generation of Internet infrastructure software regardless of whether you're trying to make your legacy setup more secure and scalable or you're selecting the best of breed solution for your next project.

How old is your web server?

Technology moves forward and software does too. Surprisingly, with the Internet infrastructure technology and software, there are still things that haven't been changing very dynamically.

Around 65% of web installations across the world are still running on legacy software engineered in the 90's. Back then servers were different, the Internet was different and most importantly—the operating systems were different. Server software was often architected and written with the idea to overcome many historical limitations of the underlying OS. However, many of such limitations have since been addressed and eliminated.



Is NGINX different?

NGINX differs from the legacy web server software primarily by its architecture. NGINX is clever in its behavior and it doesn't try to outthink a well-engineered modern operating system when handling the network and disk load with ultimate efficiency in memory and CPU.

NGINX achieves that by utilizing the maximum of performance enabling mechanisms available with all of the major OS platforms.

With such combination of a modern operating system and a capable web server software like NGINX system architects can make the most out of the available server and network infrastructure.

NGINX architecture

NGINX optimizes the usage of the operating system and the hardware resources with its modular, event-driven, asynchronous, non-blocking architecture.

Using event notifications and asynchronous handling of a variety of consequent actions associated with accepting, processing and managing network connections and content retrieval, NGINX provides hints to the operating system and gets timely feedback in regards to when expect an inbound or outbound traffic, when check disk operation, when refresh content and so on.

Traditional process-based or thread-based model of handling concurrent connections involves processing each connection with a separate process or thread, blocking on network and input/output operations until completion. Depending on the application it can be extremely inefficient in terms of memory and CPU consumption.

Spawning a separate process or thread requires preparation of new runtime environment, including allocation of new stack and memory areas. Additional CPU time is also spent on creation. If your web site is aimed at handling many thousands of connections, but still uses the old web server architecture, you may quickly find that the resource starvation is blocking your growth.

NGINX runs two types of processes in memory. There are worker processes that accept and handle connections, and there is a single master process that starts and stops workers and controls configuration. It is now possible to handle many thousands of concurrent connections and requests per second with each worker versus handling just one connection per process or thread.

When handling network events NGINX workers are able to elaborately use every possible modern notification mechanism invented for an operating system, like for instance *kqueue*, *epoll* or *event ports*, taking only half a kilobyte of memory per keepalive connection.

The model of having separate worker processes also allows NGINX to fully utilize multi-core hardware architectures and eliminates most of the issues with blocking on disk I/O. The latter problem is specifically addressed by tight integration with a number of disk performance enablements like *sendfile*,

AIO and *direct I/O* which are already available in a number of operating systems today.

Scenarios of use

What is remarkable about using NGINX in different setups is that you can utilize its features as standalone enabling mechanisms or you can use them in flexible combinations.

Simple use which would instantly enable you to offload serving concurrent connections, is to put NGINX as a reverse web or mail proxy in front of your busy backend servers. You can enable more efficiency here with NGINX caching functionality. You can even offload SSL to NGINX and be able to process thousands of handshakes per second on a generic server hardware.

If you happen to serve a lot of static content like HTML, CSS, JavaScript, FLV or MP4 videos, images and alike you can turn on the web server functionality of NGINX and stream your static content from the same software to really skyrocket your web sites performance.

Through the combination of NGINX features you can achieve extreme performance for static content and videos, conserve server and network resources by caching repetitive content, handle concurrent connections with ultimate efficiency, offload your backend infrastructure from unnecessary tasks, securely integrate your applications servers using a number of interfaces and protocols, and more.

Benefits of using NGINX

Essentially, NGINX has the lowest memory footprint possible and optimizes CPU usage while delivering maximum performance even on a very inexpensive server hardware.

More importantly, NGINX is able to continuously take more connections while retaining compact memory footprint.

NGINX supports all the major standards and protocols, providing the set of features necessary today to build a scalable web setup.

NGINX configuration is logical and flexible and allows more efficiency and scalability in day-to-day engineering operations.

All of the above is what really makes NGINX so much more resource efficient in comparison to the older web server architectures.

Being very flexible and ridiculously efficient in hardware resource utilization NGINX means not just performance and scalability, but also more density and a greener environment too.

Conclusion

For a service provider or any Internet organization trying to optimize ROI and total cost of ownership deploying NGINX edge web server can serve as a great foundation to develop more competitive advantages upon.