



Why Xen?

Author: Stephen Spector and Xen.org Community

Previous documents – [What is Xen?](#) and [How Does Xen Work?](#) highlight important aspects of the Xen hypervisor but fail to address a fundamental question, Why? This document presents the case of selecting the open source Xen hypervisor as your virtualization platform. I will address the benefits of the Xen hypervisor and why they matter in selecting a hypervisor.

Xen Compared to Other Hypervisors


The following points compare Xen to VMware ESXi, Hyper-V, and KVM.

- Xen has “thin hypervisor” model
 - No device drivers and keeps domains/guests isolated
 - 2 MB executable
 - Relies on service domains for functionality
- VMware ESXi – similar to Xen
 - Contains device drivers and base of management stack
 - Hardware support depends on VMware created drivers
- Hyper-V – modeled on Xen hypervisor
- KVM – unlike Xen model
 - Linux kernel as hypervisor

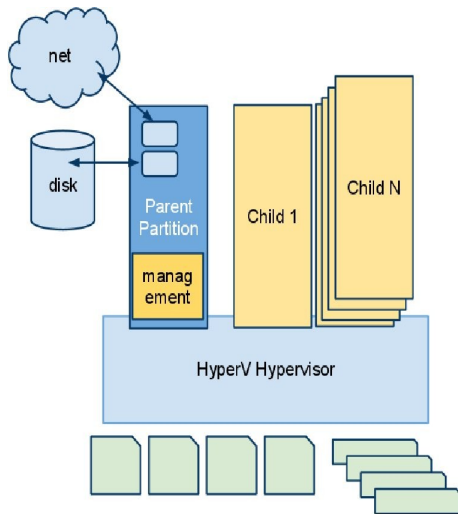
Hypervisor Architecture Diagrams

Xen Architecture


Linux Service domain “Domain0” can be Solaris, BSD, etc.




Hyper-V Architecture



VMware ESXI Architecture



KVM Architecture



Xen Comparison Highlights

- Xen separates the hypervisor execution from management OS, management stack, device drivers, and guests (components)
- Components are interchangeable – choose the best OS to support your needs
- Strong isolation between all components – assisted with modern hardware and domains can restart without taking out full system
- Scalability

Operating System Neutrality

A critical benefit of the Xen Hypervisor is its neutrality to the various operating systems. Due to its independence, Xen is capable of allowing any operating system (Linux, Solaris, BSD, etc) to be the Domain0 thereby ensuring the widest possible use case for customers. For example, many hardware manufacturers leverage NetBSD as their OS of choice for Domain0 and are able to deploy Xen in the manner of their choosing.

This separation of hypervisor from the Domain0 operating system also ensures that Xen is not burdened with any operating system overhead that is unrelated to processing a series of guests on a given machine. In fact, more are beginning to break up the Domain0 from a single guest into a series of mini-OS guests each with a specific purpose and responsibility which drives better performance and security in a virtualization environment.

As for DomainU, guests, the Xen community has led the way in paravirtualization technology which is now part of mainstream Linux and available to everyone creating a Linux distribution. The Xen community also provides industry leading paravirtualization drivers for Windows guests as well to ensure that the broader enterprise computing industry is able to deploy their OS of choice as a guest on Xen.

Security and Reliability Built-In

A critical aspect of building a hypervisor is ensuring that the solution is secure, especially when the solution is deployed in enterprise and cloud computing environments. Xen ensures a high level of security via a variety of methods/features:

- Guest Isolation – every DomainU guest is isolated from other DomainU guests with no way to access each other's memory or networking connections
- Privileged Access – only the Domain0 or single purpose control guests are given the ability to communicate with the hardware via the hypervisor
- Small Code Base – the Xen hypervisor contains a “tiny” code footprint which limits the areas for attack
- Operating System Separation – by separating the hypervisor from an operating system, the Xen hypervisor cannot be used to attack an operating system; e.g. Xen cannot attack the host operating system as there is no host operating system to attack

The Xen.org community also works closely with the [The Invisible Things Lab](#) who focus exclusively on cutting edge security needs for computing infrastructures. Working together, the Xen hypervisor is attacked by leading security experts who work directly with Xen.org to close the loop on any security hole found. In fact, The Invisible Things Lab has developed a complete, secure open source operating system, [Qubes OS](#), leveraging the Xen hypervisor as their virtualization foundation.

Performance

It is difficult to credibly present hypervisor performance statistics as there are multiple benchmarks available which are not universal in the methodologies and metrics. Instead, I will highlight some of the ways that the Xen hypervisor is fundamentally lowering processing time, memory usage, and networking efficiency.

Paravirtualization (enlightenment), created by the founders of the Xen hypervisor allows the guest operating system to co-operate with the hypervisor to improve overall performance for I/O, CPU, and memory virtualization. By being aware that the operating system is running in a virtualized platform, the modified operating system is able to assist the hypervisor in a variety of tasks. All Linux distributions currently support paravirtualization out of the box and Windows paravirtualization drivers are available in the Xen.org community.

Pass-through technology allows a guest domain to speak directly with a specific piece of hardware directly without having to send communication to and from the Domain0. Allowing a guest domain direct access to hardware significantly improves time to response for a guest, lowers processing time by eliminating the Domain0 middleman, and reduces load on the Domain0 queue. Of course, security is maintained as the guest is restricted in what hardware it can access thereby preventing guest interaction.

As stated previously, having the hypervisor separate from the operating system also ensures maximum performance. Any operating system will have a series of tasks that must be scheduled and processed during normal operation. The majority of these tasks are not related to processing the virtualized guests and thus can potential impact overall performance. The Xen hypervisor is able to process the virtualized guests without any operating system overhead and can even be tuned specifically to

maximum guest processing based on user demands and requirements for a given guest. The scheduler within Xen is also customized for a virtualized environment thereby ensuring that a Xen infrastructure is capable of meeting the highest user expectations.

Cutting Edge Features

As an open source community with many leading technology vendors participating (e.g. AMD & Intel), Xen.org is able to ensure that the Xen hypervisor supports both paravirtualized (or enlightened) guests and fully virtualized guests, so Xen users can take advantage of both the latest in software and hardware virtualization technology. By working closely with hardware manufacturers including networking equipment vendors, the Xen hypervisor becomes the testing ground for new hardware ensuring that Xen is always a cutting edge solution.

The Xen hypervisor also benefits from a variety of university research which is developed and tested before integrating into the final Xen hypervisor solution. A recent example of this close relationship is Project Remus from the University of British Columbia. Remus provides transparent, comprehensive high availability to ordinary virtual machines running on the Xen virtual machine monitor. It does this by maintaining a completely up-to-date copy of a running VM on a backup server, which automatically activates if the primary server fails.

The Cloud

The Xen hypervisor is the most used virtualization platform in the cloud computing space. With leading vendors such as Amazon, Cloud.com, GoGrid, and Rackspace all using Xen, the community is able to ensure that their scalability and performance needs are met. In fact, the largest virtualization deployments in the world are primarily running the Xen hypervisor in a cloud computing environment. To better support these cloud providers, the Xen.org community created a new project in 2009, Xen Cloud Platform.

Xen Cloud Platform offers ISVs and service providers a complete cloud infrastructure platform with a powerful management stack based on open, standards-based APIs, support for multi-tenancy, SLA guarantees and detailed metrics for consumption based charging. More details in the [Xen in the Cloud](#) document.

To follow XCP, join the Xen.org [xen-api](#) mailing list.

Proven Technology

The Xen hypervisor has been available for enterprise deployment since 2004 and is the first open source hypervisor to successfully be deployed by industry leading Linux vendors, software corporations, and global customers taking advantage of the virtualization revolution. The Xen hypervisor is currently available in solutions from Avaya, Cisco, Citrix, Fujitsu, Lenovo, Novell, Oracle, Samsung, VALinux, and others. Cloud providers including Amazon, Cloud.com, GoGrid, and Rackspace are amongst the many cloud solutions using Xen as their virtualization foundation.

For sample case studies on the Xen hypervisor go [here](#), [here](#), and [here](#). These case studies from Citrix, Oracle, and the Xen.org community demonstrate the global, secure, and scalable capabilities of the Xen hypervisor.

Community

The open source Xen hypervisor is built by a fully-engaged, global developer community consisting of corporate developers, university researchers, independent developers, and virtualization technologists. With over 250 patches (on average) submitted to the community monthly, the Xen hypervisor continues to lead the industry in hardware support and new features. Supported by the Xen Advisory Board (Citrix, Fujitsu, HP, IBM, Intel, Novell, Oracle, and VA Linux Systems Japan) and a large customer base, the Xen hypervisor has a solid future in being a leader in hypervisor technology.

Become a part of the Xen.org community...

- Developers? Join the [xen-devel](#) mailing list
- Users? Join the [xen-users](#) mailing list
- General Interest?
 - Follow the [Xen.org community blog](#)
 - Follow the [Xen.org community on twitter](#)