

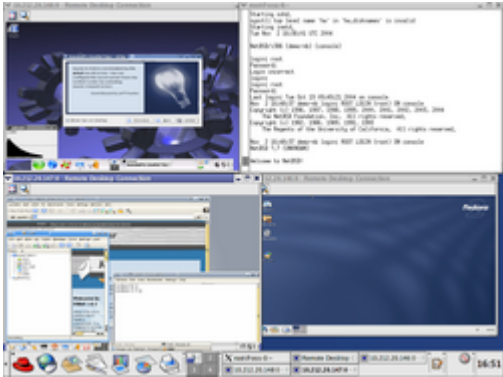
Xen


By Contributors to Wikimedia projects

Published: 2004-04-11 · Archived: 2026-04-05 20:14:53 UTC

From Wikipedia, the free encyclopedia

This article is about the virtualization software. For other uses, see [Xen \(disambiguation\)](#).

Xen Project	
	
	
Xen running NetBSD and three Linux distributions	
Original authors	Keir Fraser, Steven Hand, Ian Pratt , University of Cambridge Computer Laboratory .
Developers	Linux Foundation Intel
Initial release	October 2, 2003; 22 years ago ^{[1][2]}
Stable release	4.21 ^[3] ✎ / 19 November 2025; 4 months ago
Written in	C

<u>Platform</u>	<ul style="list-style-type: none"> • x86 • ARM • RISC-V • PowerPC
<u>Type</u>	<u>Hypervisor</u>
<u>License</u>	<u>GPLv2</u>
<u>Website</u>	<u>xenproject.org</u>
<u>Repository</u>	<ul style="list-style-type: none"> • <u>xenbits.xen.org/gitweb/?p=xen.git</u> 

Xen (pronounced) is a [free and open-source](#) type-1 [hypervisor](#), providing services that allow multiple computer [operating systems](#) to execute on the same [computer hardware](#) concurrently. It was originally developed by the [University of Cambridge Computer Laboratory](#) and is now being developed by the [Linux Foundation](#) with support from [Intel](#), [Citrix](#), [Arm Ltd](#), [Huawei](#), [AWS](#), [Alibaba Cloud](#), [AMD](#), [Bitdefender](#) and [EPAM Systems](#).

The Xen Project community develops and maintains Xen Project as [free and open-source software](#), subject to the requirements of the [GNU General Public License](#) (GPL), version 2. Xen Project is currently available for the [IA-32](#), [x86-64](#) and [ARM instruction sets](#).^[4]

Software architecture

[\[edit\]](#)

Xen Project runs in a more privileged CPU state than any other software on the machine, except for [firmware](#).

Responsibilities of the hypervisor include memory management and CPU scheduling of all virtual machines ("domains"), and for launching the most privileged domain ("dom0") - the only virtual machine which by default has direct access to hardware. From the dom0 the hypervisor can be managed and unprivileged domains ("domU") can be launched.^[5]

The dom0 domain is typically a version of [Linux](#) or [BSD](#). User domains may either be traditional operating systems, such as [Microsoft Windows](#) under which privileged instructions are provided by hardware virtualization instructions (if the host processor supports [x86 virtualization](#), e.g., [Intel VT-x](#) and [AMD-V](#)),^[6] or [paravirtualized](#) operating systems whereby the operating system is aware that it is running inside a virtual machine, and so makes hypercalls directly, rather than issuing privileged instructions.

Xen Project boots from a [bootloader](#) such as [GNU GRUB](#), and then usually loads a [paravirtualized](#) host operating system into the host domain (dom0).

Xen originated as a research project at the [University of Cambridge](#) led by [Ian Pratt](#), a [senior lecturer](#) in the [Computer Laboratory](#), and his PhD student Keir Fraser. According to [Anil Madhavapeddy](#), an early contributor, Xen started as a bet on whether Fraser could make multiple Linux Kernels boot on the same hardware in a weekend.^[7] The first public release of Xen was made in 2003, with v1.0 following in 2004. Soon after, Pratt and Fraser along with other Cambridge alumni including Simon Crosby and founding CEO Nick Gault created XenSource Inc. to turn Xen into a competitive enterprise product.

To support embedded systems such as smartphone/ IoT with relatively scarce hardware computing resources, the Secure Xen ARM architecture on an ARM CPU was exhibited at Xen Summit on April 17, 2007, held in IBM TJ Watson.^{[8][9]} The first public release of Secure Xen ARM source code was made at Xen Summit on June 24, 2008^{[10][11]} by [Sang-bum Suh](#),^[12] a Cambridge alumnus, in Samsung Electronics.


On October 22, 2007, [Citrix Systems](#) completed its acquisition of XenSource,^[13] and the Xen Project moved to the xen.org domain. This move had started some time previously, and made public the existence of the Xen Project Advisory Board (Xen AB), which had members from [Citrix](#), [IBM](#), [Intel](#), [Hewlett-Packard](#), [Novell](#), [Red Hat](#), [Sun Microsystems](#) and [Oracle](#). The Xen Advisory Board advises the Xen Project leader and is responsible for the Xen trademark,^[14] which Citrix has freely licensed to all vendors and projects that implement the Xen [hypervisor](#).^[15] Citrix also used the Xen brand itself for some proprietary products unrelated to Xen, including [XenApp](#) and [XenDesktop](#).

On April 15, 2013, it was announced that the Xen Project was moved under the auspices of the [Linux Foundation](#) as a Collaborative Project.^[16] The Linux Foundation launched a new trademark for "Xen Project" to differentiate the project from any commercial use of the older "Xen" trademark. A new community website was launched at xenproject.org^[17] as part of the transfer. Project members at the time of the announcement included: Amazon, AMD, Bromium, CA Technologies, Calxeda, Cisco, Citrix, Google, Intel, Oracle, Samsung, and Verizon.^[18] The Xen project itself is self-governing.^[19]

Since version 3.0 of the [Linux kernel](#), Xen support for dom0 and domU exists in the mainline kernel.^[20]

Version	Release date	Notes
1.0	October 2, 2003 ^{[1][2]}	
2.0	November 5, 2004 ^[21]	Live migration of PV guests.
3.0	December 5, 2005 ^{[22][23]}	<ul style="list-style-type: none"> Supports the Intel VT technology for HVM guests. Support for the Intel IA-64 architecture. <p>The releases up to 3.0.4 also added:</p> <ul style="list-style-type: none"> Support for the AMD SVM virtualization extensions.^[24]

Version	Release date	Notes
		<ul style="list-style-type: none"> • Support for the PowerPC architecture.^[25] • Graphical framebuffer support for paravirtualized guests.^[26]
3.1	May 18, 2007 ^[27]	Live migration for HVM guests, XenAPI
3.2	January 17, 2008 ^[28]	PCI passthrough and ACPI S3 standby mode for the host system.
3.3	August 24, 2008 ^[29]	Improvements for the PCI passthrough and the power management. Xen ARM hypervisor source code released for ARM CPU support
3.4	May 18, 2009 ^[30]	Contains a first version of the "Xen Client Initiative", shortly XCI.
4.0	April 7, 2010 ^[31]	Makes it possible to use a dom0 Linux kernel, which has been implemented by using PVOps. A Linux kernel of version 2.6.31 has been modified for this purpose, because the official Linux kernel actually does not support the usage as dom0 kernel (date July 2010). ^[32]
4.1	March 25, 2011 ^[33]	Some of the improvements: Support for more than 255 processors, better stability. Linux kernel v2.6.37 and onward support usage as dom0 kernel. ^[34]
4.2	September 8, 2012 ^[35]	XL became the default toolstack. Support for up to 4095 host processors and up to 512 guest processors.
4.3	July 9, 2013 ^[36]	Experimental ARM support. NUMA-aware scheduling. Support for Open vSwitch .
4.4	March 10, 2014 ^[37]	Solid libvirt support for libxl, new scalable event channel interface, hypervisor ABI for ARM declared stable, Nested Virtualization on Intel hardware. ^{[38][39]}
4.5	January 17, 2015 ^[40]	With 43 major new features, 4.5 includes the most updates in the project's history. ^[40]
4.6	October 13, 2015 ^[35]	Focused on improving code quality, security hardening, enablement of security appliances, and release cycle predictability. ^[35]
4.7	June 24, 2016 ^[41]	Improved: security, live migrations, performances and workload. Hardware support (ARM and Intel Xeon). ^[42]
4.8.1	April 12, 2017 ^[43]	

Version	Release date	Notes
4.9	June 28, 2017 ^[44]	Xen Project 4.9 Release Notes
4.10	December 12, 2017 ^[45]	Xen Project 4.10 Release Notes
4.11	July 10, 2018 ^[46]	Xen Project 4.11 Release Notes
4.12	April 2, 2019 ^[47]	Xen Project 4.12 Release Notes
4.13	December 18, 2019 ^[48]	Xen Project 4.13 Release Notes
4.14	July 24, 2020	Xen Project 4.14 Release Notes
4.15	April 8, 2021	Xen Project 4.15 Release Notes
4.16	December 2, 2021	Xen Project 4.16 Release Notes
4.17	December 14, 2022	Xen Project 4.17 Release Notes
4.18	November 23, 2023	Xen Project 4.18 Release Notes
4.19	July 29, 2024	Xen Project 4.19 Release Notes
4.20	March 5, 2025	Xen Project 4.20 Release Notes
	This article needs to be updated . Please help update this article to reflect recent events or newly available information. <i>(May 2016)</i>	

[Internet hosting service](#) companies use hypervisors to provide [virtual private servers](#). [Amazon EC2](#) (from August 2006 to November 2017),^[49] IBM [SoftLayer](#),^[50] Liquid Web, [Fujitsu Global Cloud Platform](#),^[51] [Linode](#), [OrionVM](#)^[52] and [Rackspace Cloud](#) use Xen as the primary VM hypervisor for their product offerings.^[53]

Virtual machine monitors (also known as hypervisors) also often operate on [mainframes](#) and large servers running IBM, HP, and other systems.^[*citation needed*] Server virtualization can provide benefits such as:

- Consolidation leading to increased utilization

- Rapid provisioning
- Dynamic [fault tolerance](#) against software failures (through rapid bootstrapping or rebooting)
- Hardware fault tolerance (through migration of a virtual machine to different hardware)
- Secure separations of virtual operating systems
- Support for legacy software as well as new OS instances on the same computer

Xen's support for virtual machine live migration from one host to another allows [load balancing](#) and the avoidance of downtime.

Virtualization also has benefits when working on development (including the development of operating systems): running the new system as a guest avoids the need to reboot the physical computer whenever a bug occurs.

[Sandboxed](#) guest systems can also help in computer-security research, allowing study of the effects of some [virus](#) or [worm](#) without the possibility of compromising the host system.

Finally, hardware appliance vendors may decide to ship their appliance running several guest systems, so as to be able to execute various pieces of software that require different operating systems. [\[citation needed\]](#)

Types of virtualization

[\[edit\]](#)

Xen offers five approaches to running the guest operating system: [\[54\]](#)[\[55\]](#)[\[56\]](#)

- **PV** (paravirtualization): Virtualization-aware Guest and devices.
- **HVM** (hardware virtual machine): Fully hardware-assisted virtualization with emulated devices.
- **HVM with PV drivers**: Fully hardware-assisted virtualization with PV drivers for IO devices.
- **PVHVM** (paravirtualization with hardware virtualization): PV supported hardware-assisted virtualization with PV drivers for IO devices.
- **PVH** (PV in an HVM container): Fully paravirtualized Guest accelerated by hardware-assisted virtualization where available.

Xen provides a form of virtualization known as paravirtualization, in which guests run a modified operating system. The guests are modified to use a special hypercall [ABI](#), instead of certain architectural features. Through paravirtualization, Xen can achieve high performance even on its host architecture (x86) which has a reputation for non-cooperation with traditional virtualization techniques. [\[57\]](#)[\[58\]](#) Xen can run paravirtualized guests ("PV guests" in Xen terminology) even on CPUs without any explicit support for virtualization. Paravirtualization avoids the need to emulate a full set of hardware and firmware services, which makes a PV system simpler to manage and reduces the attack surface exposed to potentially malicious guests. On 32-bit x86, the Xen host kernel code runs in [Ring 0](#), while the hosted domains run in *Ring 1* (kernel) and *Ring 3* (applications).

CPUs that support virtualization make it possible to run unmodified guests, including proprietary operating systems (such as Microsoft Windows). This is known as [hardware-assisted virtualization](#), however, in Xen this is known as hardware virtual machine (HVM). HVM extensions provide additional execution modes, with an explicit distinction between the most-privileged modes used by the hypervisor with access to the real hardware (called "root mode" in x86) and the less-privileged modes used by guest kernels and applications with "hardware"

accesses under complete control of the hypervisor (in x86, known as "non-root mode"; both root and non-root mode have Rings 0–3). Both Intel and [AMD](#) have contributed modifications to Xen to exploit their respective Intel VT-x and AMD-V architecture extensions.^[59] Use of [ARM](#) v7A and v8A virtualization extensions came with Xen 4.3.^[60] HVM extensions also often offer new instructions to allow direct calls by a paravirtualized guest/driver into the hypervisor, typically used for I/O or other operations needing high performance. These allow HVM guests with suitable minor modifications to gain many of the performance benefits of paravirtualized I/O. In current versions of Xen (up to 4.2) only fully virtualized HVM guests can make use of hardware facilities for multiple independent levels of memory protection and paging. As a result, for some workloads, HVM guests with PV drivers (also known as PV-on-HVM, or PVH) provide better performance than pure PV guests. Xen HVM has device emulation based on the [QEMU](#) project to provide I/O virtualization to the virtual machines. The system emulates hardware via a patched QEMU "device manager" (qemu-dm) daemon running as a backend in dom0. This means that the virtualized machines see an emulated version of a fairly basic PC. In a performance-critical environment, PV-on-HVM disk and network drivers are used during the normal guest operation, so that the emulated PC hardware is mostly used for booting.

Administrators can "live migrate" Xen virtual machines between physical hosts across a LAN without loss of availability. During this procedure, the LAN iteratively copies the memory of the virtual machine to the destination without stopping its execution. The process requires a stoppage of around 60–300 ms to perform final synchronization before the virtual machine begins executing at its final destination, providing an illusion of seamless migration. Similar technology can serve to suspend running virtual machines to disk, "freezing" their running state for resumption at a later date.

Xen can scale to 4095 physical CPUs, 256 VCPUs^[clarification needed] per HVM guest, 512 VCPUs per PV guest, 16 TB of RAM per host, and up to 1 TB of RAM per HVM guest or 512 GB of RAM per PV guest.^[61]

Xen implements [shared memory](#) using *grant tables*, which allow domains to specify [memory pages](#) to share with other selected domains.^[62]

The Xen hypervisor has been ported to a number of processor families:

- Intel: IA-32, IA-64 (before version 4.2^[63]), x86-64
- PowerPC: previously supported under the XenPPC project, no longer active after Xen 3.2^[64]
- ARM: previously supported under the XenARM project for older versions of ARM without virtualization extensions, such as the Cortex-A9. Currently^[when?] supported since Xen 4.3 for newer versions of the ARM with virtualization extensions, such as the [Cortex-A15](#).
- [MIPS](#): XLP832 experimental port^[65]

Xen can be shipped in a dedicated virtualization platform, such as [XCP-ng](#) or XenServer (formerly Citrix Hypervisor, and before that Citrix XenServer, and before that XenSource's XenEnterprise).

Alternatively, Xen is distributed as an optional configuration of many standard operating systems. Xen is available for and distributed with:

- [Alpine Linux](#) offers a minimal dom0 system ([Busybox](#), [musl libc](#)) that can be run from removable media, like USB sticks.
- [Arch Linux](#) provides the necessary packages with detailed setup instructions on their Wiki.^{[66][67]}
- [Debian Linux](#) (since version 4.0 "etch") and many of its derivatives;
- [FreeBSD](#) 11 includes experimental host support.^[68]
- [Gentoo](#) has the necessary packages available to support Xen, along with instructions on their Wiki.^[69]
- [Mageia](#) (since version 4);
- [NetBSD](#) can function as domU and dom0.^[70]
- [OpenSolaris](#)-based distributions can function as dom0 and domU from Nevada build 75 onwards.
- [openSUSE](#) 10.x to 12.x:^[71] only 64-bit hosts are supported since 12.1;
- [Qubes OS](#) uses Xen to isolate applications for a more secure desktop.^[72]
- [SUSE Linux Enterprise Server](#) (since version 10);
- [Solaris](#) (since 2013 with [Oracle VM Server for x86](#), before with [Sun xVM](#));
- [Ubuntu](#) (since [8.04 "Hardy Heron"](#), but no dom0-capable kernel in [8.10 "Intrepid Ibex"](#) until [12.04 "Precise Pangolin"](#).^{[73][74]})

Guest systems can run fully virtualized (which requires hardware support), paravirtualized (which requires a modified guest operating system), or fully virtualized with paravirtualized drivers (PVHVM^[75]).^[76] Most operating systems which can run on PCs can run as a Xen HVM guest. The following systems can operate as paravirtualized Xen guests:

- [Linux](#)
- [FreeBSD](#) in 32-bit, or 64-bit through PVHVM;^{[77][78]}
- [OpenBSD](#), through PVHVM;^[79]
- [NetBSD](#)
- [MINIX](#)
- [GNU Hurd](#) (gnumach-1-branch-Xen-branch)
- [Plan 9 from Bell Labs](#)

Xen version 3.0 introduced the capability to run Microsoft Windows as a guest operating system unmodified if the host machine's processor supports [hardware virtualization](#) provided by Intel VT-x (formerly codenamed Vanderpool) or AMD-V (formerly codenamed Pacifica). During the development of Xen 1.x, [Microsoft Research](#), along with the University of Cambridge Operating System group, developed a port of [Windows XP](#) to Xen — made possible by [Microsoft's](#) Academic Licensing Program. The terms of this license do not allow the publication of this port, although documentation of the experience appears in the original Xen [SOSP](#) paper.^[80] James Harper and the Xen open-source community have started developing free software paravirtualization drivers for Windows. These provide front-end drivers for the Xen block and network devices and allow much higher disk and network performance for Windows systems running in HVM mode. Without these drivers all disk and network traffic has to be processed through QEMU-DM.^[81] Subsequently, Citrix has released under a BSD license (and continues to maintain) PV drivers for Windows.^[82]

Third-party developers have built a number of tools (known as Xen Management Consoles) to facilitate the common tasks of administering a Xen host, such as configuring, starting, monitoring and stopping of Xen guests. Examples include:

- The [OpenNebula](#) cloud management toolkit
- On openSUSE [YaST](#) and virt-man offer graphical VM management
- [OpenStack](#) natively supports Xen as a Hypervisor/Compute target
- [Apache CloudStack](#) also supports Xen as a Hypervisor
- Novell's [PlateSpin](#) Orchestrate also manages Xen virtual machines for Xen shipping in [SUSE Linux Enterprise Server](#).
- [Xen Orchestra](#) for both [XCP-ng](#) and Citrix Hypervisor platforms

Commercial versions

[[edit](#)]

- [XCP-ng](#) (Open Source, within the Linux Foundation and Xen Project, originally a fork of XenServer)
- XenServer^[83] (Formerly Citrix Hypervisor^[84] until 2023 and formerly Citrix XenServer until 2019)
- [Huawei](#) FusionSphere^[85]
- [Oracle VM Server for x86](#)
- Thinsy Corporation
- [Virtual Iron](#) (discontinued by Oracle)
- [Crucible \(hypervisor\)](#) by Star Lab Corp.^[86]

The Xen hypervisor is covered by the GNU General Public Licence, so all of these versions contain a core of free software with source code. However, many of them contain proprietary additions.

- [CloudStack](#)
- [Kernel-based Virtual Machine](#) (KVM)
- [OpenStack](#)
- [Virtual disk image](#)
- [tboot](#), a TXT-based integrity system for the Linux kernel and Xen hypervisor
- [VMware ESXi](#)
- [Qubes OS](#)

1. [^] [Jump up to: ^a ^b "Xen"](#). *SourceForge.net*. October 2, 2003. Retrieved October 18, 2012.
2. [^] [Jump up to: ^a ^b Jonathan Corbet \(October 2, 2003\). "The first stable Xen release"](#). *Lwn.net*. Retrieved October 18, 2012.
3. [^] ["Xen Project Delivers Xen 4.21, a Modernized Hypervisor with Broader Architecture Support and Improved Performance"](#). November 19, 2025. Retrieved November 22, 2025.
4. [^] [jgross \(April 2, 2019\). "What's New In XEN 4.12"](#). *xenproject.org*. Retrieved May 6, 2019.
5. [^] ["Xen Overview"](#). Retrieved April 5, 2015.
6. [^] ["OSCompatibility - Xen Project Wiki"](#). *Wiki.xenproject.org*. February 8, 2007. Retrieved June 8, 2013.

