

What Is High-Performance NVMe Storage?

什么是高性能 NVMe 存储？

Carol Platz

作者: *Carol Platz*

This blog post is an introduction to software-defined storage, NVMe, and private/edge/public clouds. If you are already familiar with all of the above, feel free to skip it and check out <https://www.lightbitlabs.com/blog/introducing-lightbits-on-aws/> to get your hands dirty with Lightbits awesomeness on the public cloud. If, however, you want to learn more about software-defined storage, NVMe, private/edge/public clouds, and how Lightbits fits in, read on.

本篇博客文章介绍了软件定义存储、NVMe 和私有/边缘/公有云。如果您已经熟悉上述所有内容，请跳过此文，可直接访问 <https://www.lightbitlabs.com/blog/introducing-lightbits-on-aws/>，在公有云上亲身体会 Lightbits 解决方案的奇妙之处。但是，如果您想了解更多关于软件定义存储、NVMe、私有/边缘/公有云以及 Lightbits 为此所提供的解决方案，请继续阅读。

Overview

概览

What is High-Performance NVMe Storage?

什么是高性能 NVMe 存储？

What is High-Performance Storage?

什么是高性能存储？

Why Is High-Performance Storage Important?

为什么高性能存储很重要？

High-Performance Software-Defined Storage

高性能的软件定义存储

What is Software-Defined Storage?

什么是软件定义存储？

Why Is Software-Defined Storage Important?

为什么软件定义存储很重要？

High-Performance NVMe Storage

高性能 NVMe 存储

High-Performance NVMe for Private and Edge Clouds

用于私有云和边缘云的高性能 NVMe

High-Performance Storage for VMware

用于 VMware 的高性能存储

High-Performance Storage for Kubernetes

用于 Kubernetes 的高性能存储

Lightbits High-Performance, NVMe-native Storage

Lightbits 高性能的 NVMe 原生存储

What is High-Performance NVMe Storage?

什么是高性能 NVMe 存储？

Your data may now be stored in dozens of locations. Making sure your applications have immediate access to data, regardless of the type of cloud the data is stored on, is now essential for every organization and can be achieved with cloud-native, NVMe-optimized storage.

您的数据现在可能存储在几十个地方。确保您的应用程序能够立即访问数据，无论数据存储在哪种类型的云上，这对于每个组织来说都是至关重要的，并且可以通过云原生、NVMe 优化的存储来实现。

The old saying “less is more” may be true sometimes, but not when it comes to data storage. As the amount of data we produce every day continues growing, there will always be a need for more capacity, more flexibility, more secure data storage, and faster access to data. Cloud-native applications need their data as quickly as possible. That’s where high-performance storage comes into play.

“少即是多”这句老话有时可能是正确的，但在数据存储方面却并非如此。随着我们每天生成的数据量不断增长，我们总是需要更大的容量、更高的灵活性、更安全的数据存储以及更快速的数据访问。云原生应用程序需要尽可能快地访问数据，这就是高性能存储发挥作用的地方。

What is High-Performance Storage?

什么是高性能存储？

Storage performance is the measure of how quickly an application can store or retrieve its data from a storage device or system. Technically, a storage device or system’s performance refers to the speed at which an application can read

data from it and write data to it, measured in Inputs/Output Operations per Second (IOPS).

存储性能是衡量应用程序在存储设备或系统中存储或检索数据的速度的一個标准。从技术上讲，存储设备或系统的性能是指应用程序可以从存储设备或系统中读取数据和向其中写入数据的速度，以每秒输入/输出操作次数（Inputs/Output Operations per Second, IOPS）衡量。

You can also measure storage performance in terms of data transfer rates, which is concerned less with the number of distinct I/O operations but rather with the total amount of data transferred over a period of time. A network-attached high-performance storage system can transfer petabytes of data at hundreds of gigabytes per second. That means complex data or unstructured data, like video files, can move across the network at lightning speed.

您还可以根据数据传输速率来衡量存储性能。数据传输速率与不同 I/O 操作的数量无关，而是与在一段时间内传输的数据总量有关。网络连接的高性能存储系统能够以每秒数百 GB 的速度传输 PB 量级的海量数据。这意味着复杂的数据或非结构化的数据（如视频文件），可以在网络中以闪电般的速度进行传输。

Why Is High-Performance Storage Important?

为什么高性能存储很重要？

High-performance storage is important because the performance of the larger systems of which storage is a part is often dictated by the performance of the underlying storage system. You can have a blazingly fast CPU and network, but if the storage is not performing, the entire system will be sluggish. Mission critical, I/O-intensive applications always need high-performance storage.

高性能存储非常重要，因为存储是大型系统的一部分，其性能通常由底层存储系统的性能所决定。您可以拥有运行速度非常快的 CPU 和网络，但如果存储性能

不佳，整个系统运行将变得迟缓。任务关键型、I/O 密集型应用程序始终需要高性能存储。

High-performance storage is ideal for primary storage workloads such as big data analytics and highly transactional workloads. High-performance storage may cost more than conventional storage, so it's not ideal for backup use cases. In addition, high-performance storage would not be the most economical solution for storing cold data—data sets that need to be stored long-term for compliance or other regulatory reasons or archived on a permanent basis—so it's not ideal for archive use cases.

高性能存储非常适合大数据分析等主存储工作负载和高事务性工作负载。高性能存储可能比传统存储成本更高，因此对于备份用例来说并不是理想选择。此外，高性能存储不是冷数据（由于合规性或其他监管原因需要长期存储的数据集或永久存档的数据集）存储的最经济实用解决方案，因此，它不是存档用例的理解选择。

High-Performance Software-Defined Storage

高性能的软件定义存储

Software-Defined Storage (SDS), which abstracts storage software from the underlying hardware and delivers all expected storage capabilities while running on any hardware, is available in high-performance variants. Some workloads are better when implemented using SDS with high-performance capabilities.

软件定义存储（SDS）从底层硬件中抽象出存储软件，并在任何硬件上运行时，都可提供所有预期的存储功能，以提供高性能存储。一些工作负载使用高性能的软件定义存储执行时，其效果会更好。

What is Software-Defined Storage?

什么是软件定义存储？

Traditional data storage solutions were delivered using proprietary hardware and software, with the software highly customized and only able to run on the exact hardware platform it was delivered on. Software-defined storage provides the same functionality and capabilities but runs on any hardware platform. This removes the dependency on proprietary hardware and its various limitations. SDS can use standard servers and commodity off-the-shelf SSDs to provide full-featured storage solutions.

传统的数据存储解决方案都是使用专有的硬件和软件来提供的，其中软件是高度定制的，并且只能在提供的特定硬件平台上运行。而软件定义存储提供相同的功能和性能，还可以在任何硬件平台上运行，这消除了对专有硬件的依赖及其各种限制。SDS 可以使用标准的服务器和现成可用的固态硬盘（SSD）来提供功能齐全的存储解决方案。

Why Is Software-Defined Storage Important?

为什么软件定义存储很重要？

SDS is important because data today is no longer confined to on-premises data centers. Workloads can be scattered across the enterprise, in a variety of media, in traditional infrastructure, virtual machines, or in the cloud. Handling and managing today's highly dispersed and dynamic data requires a new level of agility and flexibility. SDS makes this happen.

软件定义存储之所以重要，是因为当今的数据不再局限于企业内部数据中心。工作负载可以分布在整个企业、各种媒介、传统基础设施、虚拟机或云中。现在，处理和管理高度分散的和动态的数据需要更高级别的敏捷性和灵活性。软件定义存储可使之成为现实。

Furthermore, when SDS untethers storage from hardware, your organization

will be able to use more of your enterprise servers and existing hardware for storage and thus enjoy an increase in storage capacity. Utilization of hardware improves, so you won't need to add more hardware when you need more storage.

此外，当 SDS 将存储从硬件中分离出来时，您的组织将能够使用更多的企业服务器和现有硬件进行存储，从而提高存储容量。硬件的利用率提高了，因此当您需更多存储容量时，不需要添加更多硬件。

The benefits of SDS include:

使用 SDS 的好处包括：

No vendor lock-in

无需锁定到某一家供应商

Rapid provisioning of storage resources

快速配置存储资源

Disaggregation of storage from compute

将存储与计算拆分开

Automation and centralized management

自动化和集中化管理

Cloud-like scalability, with improved resource utilization

具有类似于云的可扩展性，提高了资源利用率

Business continuity

业务连续性

Greater availability, resiliency

更高的可用性和弹性

Reduced costs

降低成本

High-Performance NVMe Storage

高性能 **NVMe** 存储

What is NVMe?

什么是 **NVMe**?

Nonvolatile memory express (NVMe®) is storage access and transport protocol that can deliver the fastest response times and highest throughput for next-generation workloads. NVMe is superior to older protocols like SAS, which were originally designed for hard disk drives (HDDs). Yes, SAS protocols do work with solid-state drives (SSDs), but not efficiently or economically. In contrast, NVMe is designed to make the most of flash and next-generation SSDs as well as today's multi-core processors.

非易失性存储器标准 **NVMe®** 是一种存储访问和传输协议，可为下一代工作负载提供最快的响应时间和最高的吞吐量。**NVMe** 优于 **SAS** 等旧协议，后者最初是为硬盘驱动器（**HDD**）设计的。是的，**SAS** 协议确实适用于固态硬盘（**SSD**），但效率和经济性都较差。相比之下，**NVMe** 设计旨在充分利用闪存和下一代 **SSD** 以及当今的多核处理器。

Systems using NVMe drives enjoy high-bandwidth and low-latency access to

storage devices by accessing flash storage through a PCI Express (PCIe) bus. NVMe provides thousands of parallel command queues, unlike older protocols which were serial in nature, with a single command queue. NVMe delivers higher IOPS and lower latency by spreading I/O requests across multiple cores, providing quick access to critical data. NVMe devices also consume less power, thus reducing the total cost of ownership. NVMe devices also support standard security protocols and facilitate scalability for next-generation demands.

使用 NVMe 驱动器的系统通过 PCI Express (PCIe) 总线访问闪存，从而实现以高带宽和低延迟的方式访问存储设备。NVMe 提供了数千个并行命令队列，与传统的串行协议不同，它只有一个命令队列。NVMe 通过将 I/O 请求分布到多个内核，提供更高的 IOPS 和更低的延迟，从而快速访问关键数据。NVMe 设备的功耗也较低，因此降低了总体拥有成本。NVMe 设备还支持标准安全协议，并推动下一代需求的可扩展性。

High-Performance NVMe for Private and Edge Clouds

用于私有云和边缘云的高性能 NVMe

Public Cloud versus Private Cloud

公有云和私有云的对比

Any discussion of the private cloud must begin with the understanding that cloud computing is a loosely defined term. It can refer both to software architecture and to a business model. A cloud computing architecture abstracts software from physical infrastructure using sophisticated management tools and virtualization. With a cloud architecture, an admin can deploy a virtual machine (VM) or virtual storage without any need to know where specifically it is hosted, or on what kind of hardware. It is running “in a cloud,” so to speak.

任何关于私有云的讨论都必须首先理解云计算是一个定义宽松的术语。它既可以

指软件架构，也可以指业务模式。云计算架构通过使用复杂的管理工具和虚拟化将软件从物理基础设施中抽象出来。通过云架构，管理人员可以部署一个虚拟机（VM）或虚拟存储，而无需知道它具体托管在哪里，或在何种类型的硬件上。可以说，它是在“云中”运行。

With this in mind, a hyperscaler, such as Amazon Web Services (AWS), enables users to deploy software and storage onto their platform. AWS takes care of setting up the server and storage hardware, along with the network and so forth. You don't need to know anything about their infrastructure except that it's there and you can spin your VMs up or down as needed. It's a multi-tenant environment where different tenants belong to different organizations. E.g., both Coca-Cola and Pepsi may run their workloads on the same AWS cloud.

考虑到这一点，诸如 Amazon Web Services（AWS）等超大规模云服务提供商（hyperscaler）支持用户将软件和存储部署到他们的平台上。AWS 负责设置服务器和存储硬件以及网络等方面的工作。您不需要了解他们的基础设施的任何信息，只需要知道它就在那里，您可以根据需要增加或减少虚拟机数量。这是一个多租户环境，不同的租户属于不同的组织。例如，可口可乐公司和百事可乐公司可能会在同一个 AWS 云上运行他们的工作负载。

A private cloud, in contrast, is a single organizational environment. While it may have multiple tenants, e.g., both engineering and accounting can both be thought of as different tenants, the tenants all belong to the same organization. A private cloud is typically run on private, on-premises infrastructure or in colocation facilities. A private cloud offers the flexibility of the cloud architecture, but with the security and control, that's only possible when you own the entire system and don't share it with anyone.

相比之下，私有云是一个单一的组织环境。虽然它可能有多个租户，例如，工程和会计都可以被认为是不同的租户，但这些租户都属于同一个组织。私有云通常运行在私有的、企业内部的基础设施上，或在主机托管设施中。私有云提供了云

架构的灵活性，但在安全性和控制性方面，只有当您拥有整个系统且不与任何人共享时，这才可能实现。

Companies may choose a private cloud because their workloads deal with sensitive data including medical records, intellectual property, financial data, or other confidential documents. A private cloud (sometimes called a corporate cloud) combines important cloud features (like scalability and easy delivery of services) with features of on-premises infrastructure (including access control, resource customization, and security).

公司可能会选择私有云，因为他们的工作负载要处理敏感数据，包括医疗记录、知识产权、财务数据或其他机密文件。私有云（有时被称为企业云）将重要的云功能（如可扩展性和易于提供服务）和企业内部基础设施的功能（包括访问控制、资源定制和安全性）相结合。

What is an Edge Cloud?

什么是边缘云？

Today's corporate data centers and hyperscale cloud facilities are sometimes referred to as the "core" of the Internet and corporate networks. That's where the bulk of the data and computing power resides. There is nothing wrong with having most of the data hosted at the core. The difficulty is that the end users or devices requiring this data are often far away, with the resulting latency causing poor user experiences.

今天的企业数据中心和超大规模云设施有时被称为互联网和企业网络的“核心”。这是大部分数据和计算能力所在的地方。将大部分数据托管在核心上并无不妥。困难之处在于，需要这些数据的终端用户或设备通常都在很远的地方，由此产生的延迟导致用户体验不佳。

Edge computing seeks to solve this problem by placing computing and storage

capacity closer to end users, often in small-scale “micro” edge data centers that may be located next to 5G towers.

边缘计算力图通过将计算和存储能力放置在更靠近终端用户的地方来解决这一问题，通常是在可能位于 5G 塔旁边的小规模“微型”边缘数据中心。

An Edge Cloud applies the cloud software architecture to edge computing. System admins can now deploy cloud-based VMs to edge data centers, rather than to public cloud platforms like AWS or private clouds running in core corporate data centers.

边缘云将云软件架构应用于边缘计算。系统管理人员现在可以将基于云的虚拟机部署到边缘数据中心，而不是部署到 AWS 等公有云平台或在核心企业数据中心运行的私有云。

Why Clouds Require High-Performance Storage

为什么云需要高性能存储？

Whether it’s an edge or private cloud, both models have a similar goal and that is to enable people and organizations to work smarter and faster. Much faster. In today’s competitive, budget-tightening environment, milliseconds count because seconds add up to hours and hours add up to big costs. Bringing applications and data closer is important; but in addition, network and storage systems must perform at the highest bandwidth, with the lowest latency, and the tightest security. High-performance NVMe does all this and more.

无论是边缘云还是私有云，这两种模式都有一个相似的目标，那就是支持人们和组织能够更智能、更快速地工作。速度快得多。在当今竞争激烈、预算紧缩的环境中，毫秒都是非常重要的，是因为多秒加起来就是小时，小时加起来就是巨大的成本。让应用程序和数据更靠近是很重要的；但除此之外，网络和存储系统必须以最高的带宽、最低的延迟和最严格的安全性运行。高性能 NVMe 可以做到

这一切，甚至更多。

High-Performance Storage for VMware

用于 **VMware** 的高性能存储

What are VMs?

什么是虚拟机？

VM stands for virtual machine, a software-defined computer that resides within a physical machine. Special software known as a hypervisor splits up the computing capacity of the physical machine and enables it to host one or more VMs on its hardware. A VM can run a complete software stack, with an operating system, application software, database, and more. A VM borrows specific amounts of CPU, memory, and storage from its physical computer or server. But the virtual machine is walled off from the host system. As a result, the VM's software can't impede the physical computer's operations. Each VM does not "know" that other VMs are running on the same underlying hardware. You can have a Windows Server VM and a Linux Server VM running on the same physical server.

VM 代表虚拟机，一种驻留在物理机器中的软件定义的计算机。被称为管理程序的特殊软件可以将物理机器的计算能力分割开来，使其能够在其硬件上托管一个或多个虚拟机。一个虚拟机可以运行一套完整的软件协议栈，包括操作系统、应用程序软件、数据库等。虚拟机从其物理计算机或服务器上借用特定数量的 CPU、内存和存储，但是虚拟机与主机系统是隔离开的。因此，虚拟机的软件不会妨碍物理计算机的运行。每个虚拟机并不“知道”其他虚拟机正在相同的底层硬件上运行。您可以让一个 Windows Server 虚拟机和一个 Linux Server 虚拟机在同一台物理服务器上运行。

Why is High-Performance Storage Important for VMware Environments?

为什么高性能存储对 **VMware** 环境很重要？

VMware environments do well with high-performance storage because VMs may add performance overhead to a system. Storage should not add to this problem. The virtual machine should not be seen as a lesser device, but as a system that performs highly essential tasks which require high-performance storage in order to execute with the fastest possible response times.

VMware 环境在高性能存储方面做得很好，因为虚拟机可能会增加系统的性能开销。存储不应该加剧这个问题。虚拟机不应被视为一个小型的设备，而应将其视为执行任务关键型的系统，这些任务需要高性能存储才能以尽可能快的响应时间执行。

High-Performance Storage for Kubernetes

用于 **Kubernetes** 的高性能存储

What is Kubernetes?

什么是 **Kubernetes**？

Kubernetes (K8s) is an open-source container orchestration system. Its primary function is to automate the deployment, scaling, and management of container-based applications. Containers free application services from physical hardware by making them portable. And when services are separated into containers, they can be independently scaled.

Kubernetes (K8s) 是一个开源的容器编排系统。它的主要功能是对基于容器的应用程序进行自动部署、扩展和管理。容器通过使应用程序服务具有可移植性，将其从物理硬件中释放出来。当服务被分离到容器中时，它们可以独立扩展。

Kubernetes and containers have grown in popularity because they offer a simple way to efficiently scale and manage cloud-native applications. They do

this by dividing applications into a set of loosely coupled micro-services.

Kubernetes 和容器越来越受欢迎，因为它们提供了一种简单的方法来高效地扩展和管理云原生应用程序。它们通过将应用程序划分为一组松散耦合的微服务来实现这一点。

Why is High-Performance Storage Important for Kubernetes?

为什么高性能存储对 **Kubernetes** 很重要？

For the maximum value, Kubernetes requires a high-performance persistent storage solution. It must be as portable as containers but must function with the performance of local flash. The storage solution must also be standards-based and run on standard servers. To optimize storage for K8 and containers, hyperscale efficiency, and flexibility are essential. This can be found with a system of high-performance scale-out and redundant storage that performs like local NVMe SSDs.

为了实现最大价值，**Kubernetes** 需要一种高性能的持久存储解决方案。它必须像容器一样便携，但同时需要具有本地闪存的性能。存储解决方案还必须基于标准，并运行在标准服务器上。为了优化 **K8** 和容器的存储，超大规模的效率和灵活性是必不可少的。这一点可以在高性能扩展和冗余存储系统中实现，该系统的性能与本地 **NVMe SSD** 类似。

Lightbits High-Performance, NVMe-native Storage

Lightbits 高性能的 **NVMe** 原生存储解决方案

Lightbits, the inventors of NVMe/TCP storage, offers an efficient, agile NVMe SDS for private clouds, edge clouds, and the public cloud. It offers high performance, high IOPS and throughput (4.7M IOPS and 22GB/s from a single storage server), and consistent low latency (160µs) for database and analytics

workloads in OpenStack, VMware, and Kubernetes environments. It's simple to consume because it works on standard Ethernet TCP/IP networks and NICs requiring no RDMA confirmation on the network switches. If you are building a cloud and want to offer fast, secure, and resilient services, consider the Lightbits high-performance, NVMe SDS.

作为 NVMe/TCP 存储的开拓者，Lightbits 为私有云、边缘云和公有云提供了高效、灵活的 NVMe SDS。它为 OpenStack、VMware 和 Kubernetes 环境中的数据库和分析工作负载提供了高性能、高 IOPS 和高吞吐量（单台存储服务器可提供高达 4.7M IOPS 性能和 22GB/s 的传输速率），以及一致的低时延（160µs）。NVMe SDS 易于使用，因为它适用于标准以太网 TCP/IP 网络和网卡，无需在网络交换机上进行 RDMA 确认。如果您正在构建云，并希望提供快速、安全、弹性的服务，请考虑使用 Lightbits 高性能的 NVMe SDS 解决方案。

Other related blogs:

其它相关博客:

Accelerate Cloud-Native Applications With NVMe

[使用 NVMe 加速云原生应用程序](#)

NVMe/TCP Storage PoC That Resulted In Great Things

[NVMe/TCP 存储 PoC 带来的巨大效益](#)

Data Storage 101: The NVMe Storage Landscape

[数据存储 101: NVMe 存储环境](#)