Performance on the Edge - Bringing Computing to the Home

The rapid growth of embedded devices – from sensors in your smartphone and self-driving automobile to your home security system – has led to an abundance of data that must be gathered and analyzed in real-time. These devices are intelligent sensors that facilitate everything we do; there is an interconnection

between these devices even as they integrate into our daily lives. The data produced by this intertwined communication eventually travel through and reside on servers in the cloud. However, the interaction of the devices is at the edge of the network – the “last mile” to the customer and the manifold related applications. It is here that the user connects to the larger networking “grid.”

Edge Computing – What Is It?

The user and related devices at the “edge” of the networking structure come in many varieties with the most prominent being nodes or hubs within the so-called Internet of Things (IoT). These are essentially an array of sensors utilized for a wide variety of tasks. This includes augmented and virtual reality, drones and automated vehicles, remote monitoring for security and healthcare, as well as commercial and industrial applications related to safety, agriculture, and much more. The user is also a consumer of multimedia content, such as music and videos, which arrive by means of a content delivery network (CDN) of servers. This most commonly requires constant bandwidth for streaming which, for efficacy’s sake, must have servers located within geographic proximity.

Market Growth

The expectation is for the edge computing market to grow at a Compound Annual Growth Rate (CAGR) of 34.1% globally within the next five years, from 3.6bn to 15.7bn USD between 2020 and 2025, according to

both Expert Interviews and MarketsandMarketsAnalysis. Development will be across all major regions but especially in Asia Pacific (APAC) with Australia leading the way, according to these sources. This includes all major players in the market, but many have seen unexpected growth due to the COVID-19 pandemic with higher demand for online services from entertainment streaming to telemedicine. Users and workers are also staying at home, ensuring that edge computing is more efficient than ever. Regardless of the unexpected the proliferation of IoT, continuous investment in edge infrastructure, and burgeoning opportunities for lightweight autonomous computing devices will widen the market despite looming security issues.

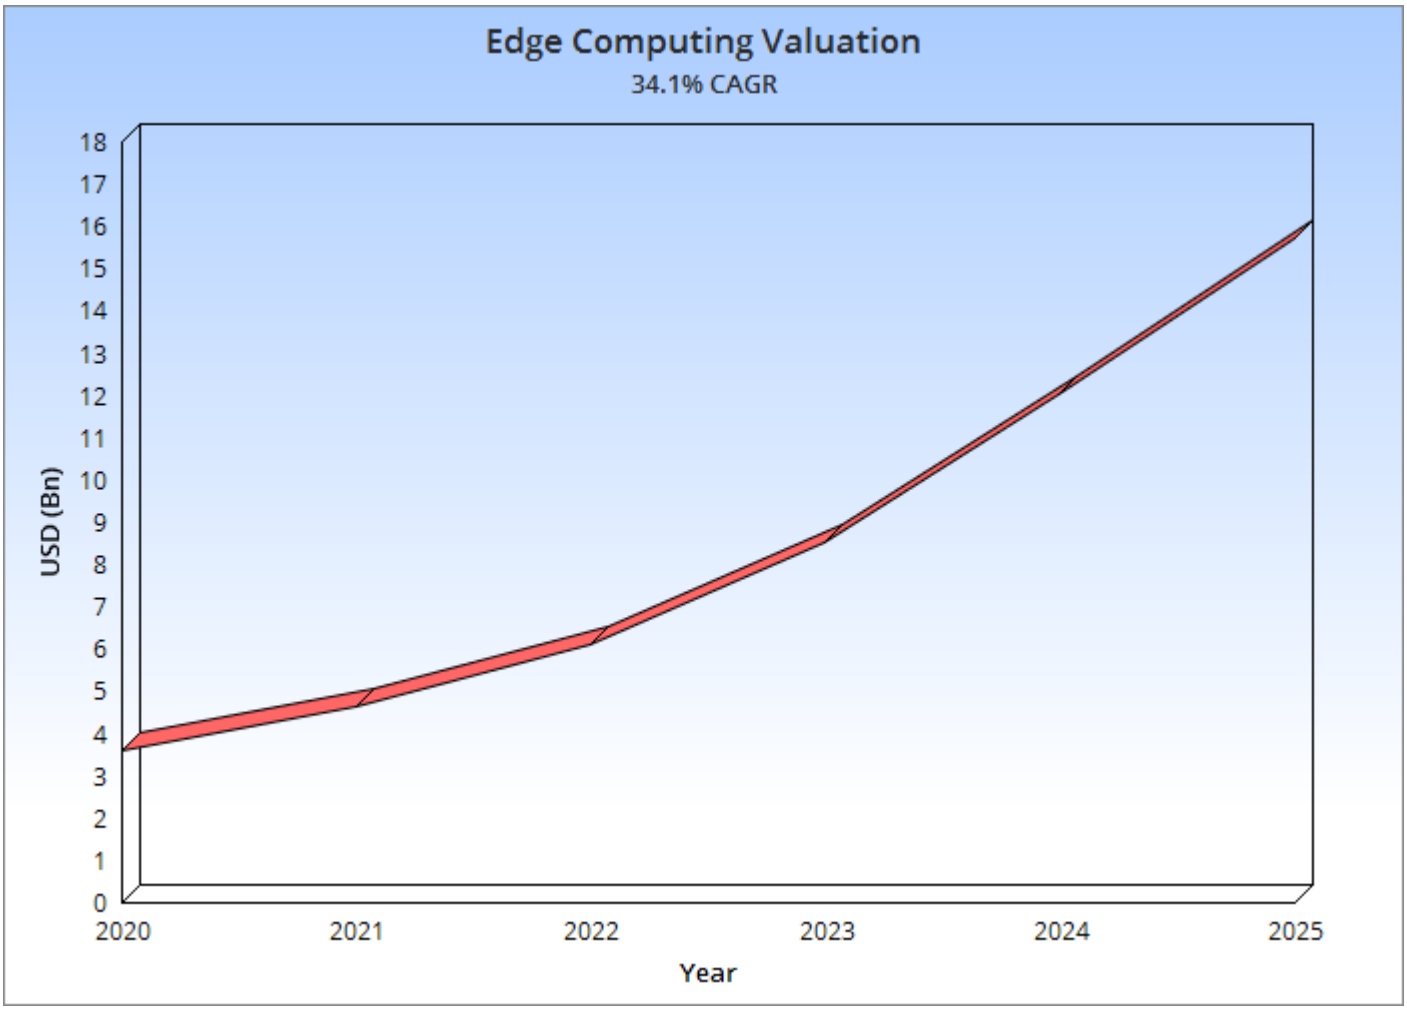

Growth is driven by the Internet of Things (IoT) with demand for low-latency, real-time connectivity.

The cloud and the need for local computation, from cars to security systems, pushes the need for higher efficiency.

Storage Requirements

The Internet of Things by its very nature involves, in aggregate, a large number of devices and sensors that as a whole require high input/output (I/O) throughput. This is not only in terms of raw bandwidth but also for I/O operations per second (IOPS). As these sensors are constantly recording data it is necessary for their related storage to have sufficient write performance. Further, because these devices are at the edge of the network it is necessary to reduce the roundtrip latency as much as possible. The sensors themselves may be able to do some computation with individualized storage but ultimately the data must be located on centralized servers.

These servers in general require a high amount of read throughput as accessing data is constant, whether as an aggregate IoT server or for CDN purposes. Latency remains critical and overall throughput in IOPS is demanding due to the sheer number of devices and simultaneous accesses. Although scalability remains a concern, the primary goal for edge server storage is consistency rather than peak performance. This means predictable performance for both reads and writes with minimal downtime.

Designing and Selecting the Proper Solid State Drive

Selecting the right solid-state drive (SSD) for edge storage involves two areas: software and hardware. For software, we start with the firmware – optimization focuses on steady-state or equilibrium performance rather than peak performance. This includes utilizing artificial intelligence (AI) and machine learning (ML) to better predict and shape I/O for specific workloads. There is also a need for proper file management, classifying data by application, and sensor type based on date as an example. The SSD must also fully support the non-volatile memory express (NVMe™) standard including for the transport control protocol/internet protocol (TCP/IP), remote direct memory access (RDMA), and NVMe™ over Fabrics (NVMe-oF™). The SSD's flash translation layer (FTL) must also intelligently handle external host requests in an optimal way.

Emphasis on the hardware side is also around consistency and reliability rather than peak throughput. Smoothing out performance and latency involves having a large cache, e.g. DRAM, and ample overprovisioning of the flash with no reliance on pseudo-SLC (pSLC). Low latency flash is especially paramount. This means utilizing an enterprise- rather than consumer-class drive. The large amounts of data require a drive of high capacity in the interests of saving physical space. Further, cooling and powering the drives add maintenance costs so power efficiency is important. Speaking of power, the SSD should have a hardware form of power loss protection (PLP) to avoid data loss in the event of outages. Lastly, security is of primary importance for user data – this means end-to-end hardware encryption.

Summary

There is an ever-increasing need for edge capacity as ubiquitous integration of tiny computers and myriad sensors continues unabated. Solid-state drives provide the required low latency necessary to keep up with this constant barrage of information. Of key importance is performance consistency and reliability – the Internet is always on, and mission-critical devices cannot afford to wait. Streaming is at an all-time high due to worldwide conditions, automation and tracking is more important than ever, and so the stage is set for “big data” to take on an even more prominent role in the development of cloud services within the IoT.

Those wishing to extend and improve their services to their customer base – whether that client is consumer, commercial, or industrial – must provide an appropriate storage solution through optimized software and hardware implementations. This means low-latency, high-throughput, read-, write-, or mixed-heavy I/O performance, depending on the application. Predictable performance with the categorization and security of data – through AI, encryption, and PLP, among others – is the primary goal. It is important to select an appropriate enterprise-class drive with sufficient capacity to reduce overall costs while keeping these criteria in mind.

Our SSD Solution

Our ED1 Series is a powerful, high performance SSD made for edge storage applications. It comes in M.2 and U.2 form factors. Contact us today to find out why Tier 1 cloud providers and Fortune 500 Enterprise companies adopted our SSDs.

Specifications and features are subject to change without prior notice. Images are samples only, not actual products. Please check with an authorized SSSTC representative for details.

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com