WHITE PAPER

State of the Enterprise Edge

velocloud™ by Broadcom

TABLE OF CONTENTS

Introduction

According to a Broadcom® survey[1], the driving challenge for scaling edge solutions and AI workloads at the distributed edge is network connectivity issues across locations (57%). When organizations implement these edge solutions, the top benefit they plan to achieve is faster response times for latency-sensitive applications (68%) and improved bandwidth/reduced network congestion (65%). By providing faster bandwidth and more reliable connections at the edge, enterprises can more efficiently process data leading to faster, smarter decision making and further encourage edge compute and AI workload deployments.

This *State of the Enterprise Edge* report is based on VeloCloud SD-WAN Edge telemetry data collected from VeloCloud customer deployments, including more than 5 trillion data points analyzed daily, across more than 300,000 edges and 3700 gateways, representing more than 25,000 customers in 78 countries. It is also based on a survey of 192 respondents across North America, Europe, South America, Asia, Australia/New Zealand and Africa. Broadcom conducted the survey in October 2024.

Market and Deployment Trends

Organizations across a range of industries are exploring edge technologies to enhance efficiency and enable improved business outcomes. Businesses are deploying these solutions at highly distributed sites that may be geographically remote and have limited or no IT staff. Common environments range from offices and hybrid work environments to retail stores, healthcare facilities, factory floors, and emergency medical vehicles.

Rise of the Intelligent Edge

As organizations learn how to deal with edge data, they move through stages, starting with limited connectivity and evolving toward intelligence.

As small language models (SLMs) become more prevalent in applications including retail, manufacturing, and logistics, advanced AI tasks are performed locally while querying large language models (LLMs) deployed in the cloud or data center for additional information. Another class of applications gaining momentum are generative AI apps that make use of

[1]Broadcom, *Survey Data for the State of the Enterprise Edge Report: November 2024*

58% OF CUSTOMERS SURVEYED HAVE DEPLOYED APPLICATIONS AT THE EDGE TO DRIVE BUSINESS OUTCOMES

LLMs. Users of these applications reside at the edge and interact with the LLM models located in the cloud or data center. Common to all these AI applications is the dependency on reliable and secure networks, whether it's metadata, full video streams, or API calls to LLMs.

The organizations surveyed that are investing in AI applications cited these use cases as top drivers:

- 21% for IoT applications such as predictive maintenance
- 17% for supply chain
- 14% for computer vision/video
- 11% for smart factories/robotics

For many organizations, edge solutions offer practical outcomes to increase productivity, improve efficiency, and offer better user experiences. To achieve these goals, organizations need to address connectivity challenges and strengthen security associated with these emerging applications.

New Edge Capabilities Bring New Challenges

Each stage of edge data demonstrates significant potential but also places a burden on network infrastructure, introducing new challenges.

Bandwidth Limitations

By their nature, edge applications generate a tremendous volume of data. Organizations require a solution to bring back metadata that's generated at the edge to the data center, to the cloud for machine learning, for analytics, and other functions. These organizations require fast and secure WAN connectivity for real-time results. This often means broadband becomes the natural choice for edge compute deployments.

38% OF PARTICIPANTS INDICATED THAT FASTER RESPONSE TIME IS THE TOP BENEFIT THEY PLAN TO REALIZE WITH EDGE APPLICATIONS AND WORKLOADS

According to our deployment data on WAN, many organizations are either augmenting or replacing their existing WAN transport options like MPLS with cheaper, faster and widely available broadband connections. Although broadband is faster, it is not as reliable as MPLS. To compensate, organizations provision their edge connectivity for redundancy and higher speed where possible. Analysis of 513,000 active circuits shows the following:

- 92% of links are broadband to address high bandwidth demands with a cost-optimized solution; 75% of customers deploy more than one link for redundancy, resulting in double the WAN bandwidth.
- MPLS adoption remains low at 8%, based on trailing 22-month data.
- 25% of edge deployments have a single link, which is typical for home deployments.

Figure 1: Transport Options

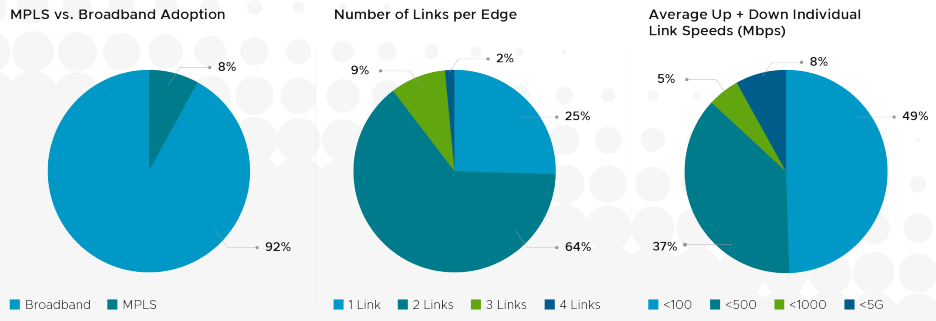

Broadband Quality Limitations

While broadband is widely accessible and affordable, it often suffers from outages, performance degradation, and inconsistent speeds.

Figure 2: Outages and Degradation for Sampled ISPs

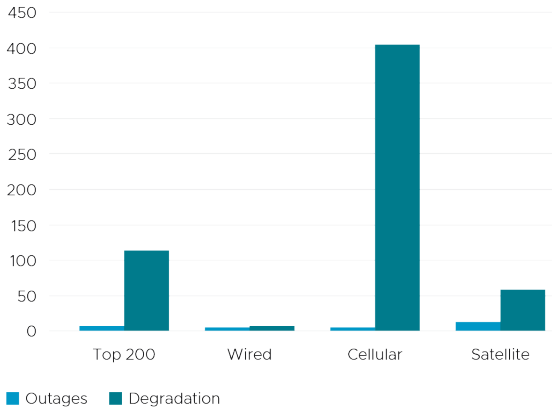

The VeloCloud solution continuously monitors all circuits for packet loss, jitter, and delay. When any of these metrics exceed predefined thresholds, the system flags the circuit as degraded and promptly initiates corrective actions, such as rerouting affected application flows or performing real-time remediation. Shown below is deployment data for combined outages and degradation in hours per connection, per month for the top 200 SPs and representative providers for wired, cellular and low earth orbit satellite:

- Wired connections: 13 hours
- 4G/5G cellular connections: 408 hours
- Satellite connections: 71 hours

Fixed Wireless Access and Satellite Adoption

As shown in Figure 2, the broadband quality issue is more pronounced in wireless access, such as 5G and satellite. However, adoption of fixed wireless access such as 5G/4G LTE and satellite connectivity is growing rapidly, driven in part by the geographic diversity of remote edge locations:

- Cellular adoption from a top ISP jumped 42% over a 24-month period
- Satellite adoption from a top ISP jumped 420% over a 24-month period

Figure 3: Adoption of Cellular and Satellite is Growing Rapidly

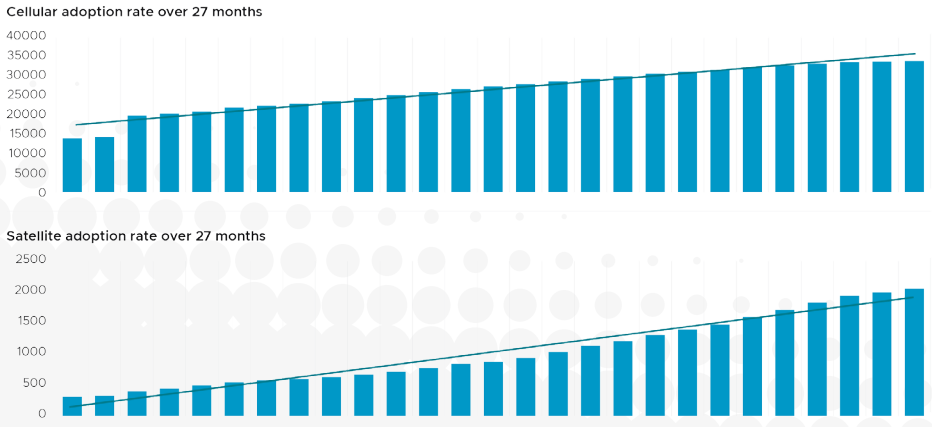

Satellite adoption is gaining strong momentum, yet cellular deployments continue to lead in volume by a wide margin.

Improve the Quality of Connections

Modern WAN solutions must effectively address circuit-related issues like packet loss, latency, and jitter to ensure reliability and performance for business-critical applications. This is especially important given the increased adoption of wireless highlighted above. Technologies that enhance wired connections such as VeloCloud Dynamic Multipath Optimization (DMPO)—which provides significant increases in the ability to handle voice and video and business-critical applications—are extended to satellite and 5G/4G-LTE connectivity using VeloCloud SD-WAN, delivering similar benefits across a lossy medium such as 4G/5G and satellite.

Figure 4: Connectivity Improvements with VeloCloud DMPO

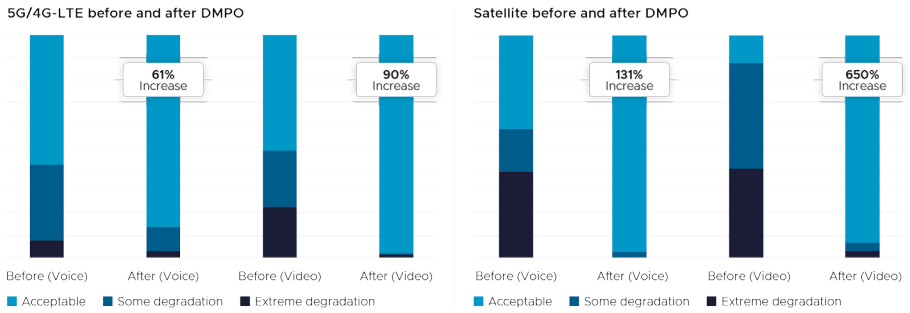

Application Visibility and Traffic Prioritization Challenge

To ensure optimal performance of edge workloads and applications, a reliable infrastructure must be established at the remote location with visibility into the application experience. This includes robust Wi-Fi coverage, as well as stable access to network services like DNS and DHCP. Based on deployment data from approximately 2000 customers over the past year, which analyzed users and devices connected to our solution, we continue to observe significant wireless and mobile deployments at branch offices, with Android and iOS devices accounting for 55% of the total.

In retail environments, 92% of the end devices are wireless, and Wi-Fi coverage is a major issue with 44.5% of companies.

Figure 5: Operating Systems for End Devices in a Retail Environment

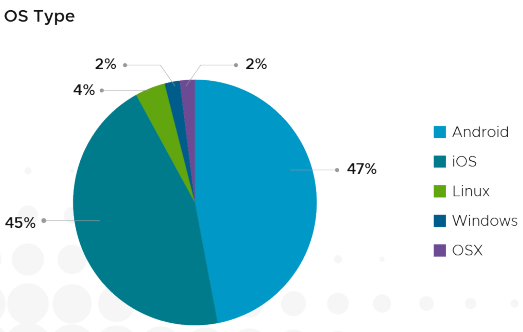

Manufacturing environments utilize more wired connectivity, but 9.2% of companies experienced DHCP issues.

To prioritize business-critical application traffic effectively, first identify the applications that are consuming bandwidth. Once identified, the right policies can be implemented to prioritize essential applications over non-critical ones, ensuring optimal network performance for key business functions.

Figure 6: Prioritize Essential Applications over Non-critical Applications

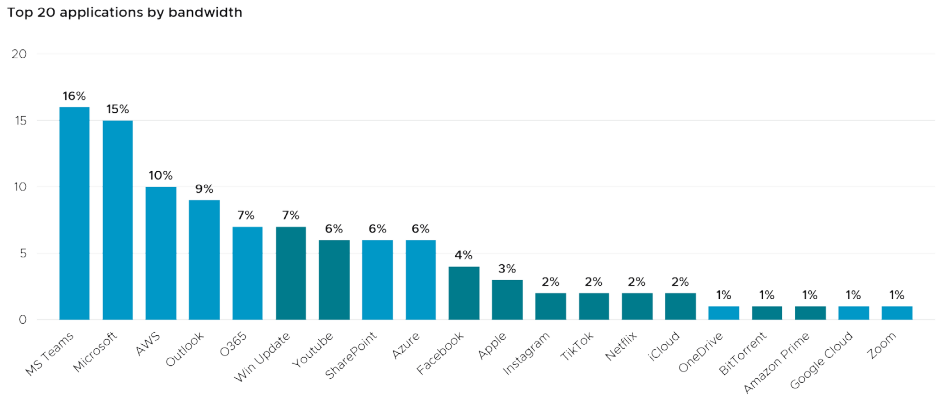

Policy management is a cumbersome process, but with the right solution, complexities can be reduced with application recognition and default policy built into the solution. Data analysis shows that, on average, customers have one business policy to prioritize business-critical applications, simplifying the scale of management. Deployment data shows approximately 30% of bandwidth is consumed by non-business-critical apps in the absence of resource contention.

Evolving Traffic Patterns

With the distributed nature of AI applications and the underlying models, traffic patterns are also undergoing transformation. In the client-server model used in typical applications, a small payload query sent upstream from a client results in a large payload response sent downstream from the application server. With AI applications, the nature of asymmetric traffic profile is reversing: upstream is greater than downstream, as only synthesized data is sent down. For example, when a full video stream is sent upstream from the edge to an LLM, or an SLM at the edge queries an LLM in the cloud, the upstream bandwidth increases significantly.

These changing patterns require a network capable of handling symmetric bandwidth, in addition to being reliable, secure and efficient. Networking and security has to be delivered on an optimal path between the edge, cloud, and the data center. Hair-pinning traffic via a data center will create bottlenecks to onboard AI applications.

CHANGING PATTERNS REQUIRE A NETWORK CAPABLE OF HANDLING SYMMETRIC BANDWIDTH, IN ADDITION TO BEING RELIABLE, SECURE AND EFFICIENT

Figure 7: The Movement Toward Distributed Architecture

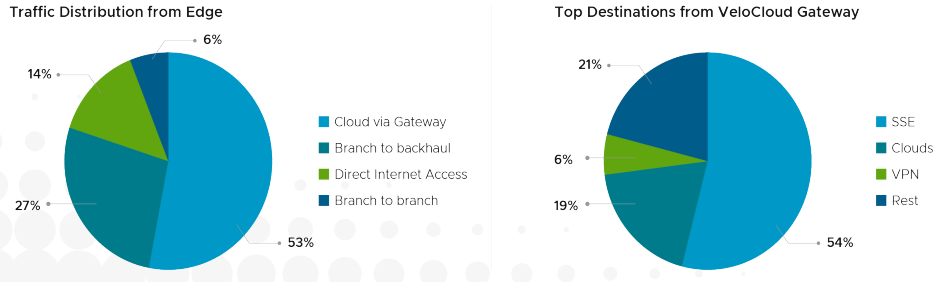

VeloCloud solution deployment data reveals interesting results that show customers moving away from a data center-centric approach and adopting a distributed architecture:

- 20% of the traffic is either direct Internet access or branch to branch.
- 53% of the traffic is sent to the VeloCloud Gateways for cloud apps and security policy enforcement, with half that percentage going to enforcement.
- Traffic sent to the data center is at 33%, with 6% sent to data centers directly, 27% through a hub.

Customers are leveraging the global deployment of VeloCloud Gateways to eliminate data center hair-pinning architectures, while also benefiting from circuit mitigation techniques for both last-mile and gateway-to-data center connections.

Security and Compliance Remain Critical

Even as enterprise architectures transform and extend farther to the edge, cybersecurity and compliance remain top of mind for organizations across many industries.

53% OF THOSE SURVEYED SAID THAT DATA SECURITY AND COMPLIANCE WERE IMPORTANT FOR EDGE COMPUTING APPLICATIONS

Edge sites are often highly diverse and are usually characterized by a mix of industrial machines, such as manufacturing robots, as well as consumer IOT devices such as security cameras or environmental sensors and controls. Because they are exposed to internal users as well as outside personnel, these devices can significantly expand the security threat vector.

Edge modernization requires elaborate protection for edge applications—including server, OS, and applications. IOT and end-user devices also tend to become points of entry for a breach or data loss.

Figure 8 shows vulnerability exploits blocked for 110 million unique devices targeted, based on data for the trailing 18 months from the Symantec® Threat Research and Global Intelligence Network.

Figure 8: Vulnerability Exploits Blocked for 110 Million Unique Devices

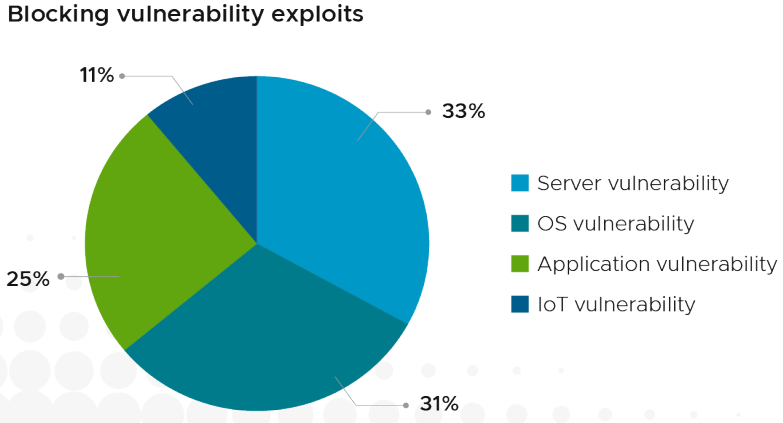

Edge Applications are Vulnerable to a Variety of Attacks

Research reveals that a variety of edge applications are vulnerable to exploits. It is critically important to monitor vulnerabilities in browser-based and non-browser-based applications, including the following:

- AI applications that utilize APIs to communicate
- Edge apps interacting with web resources without a full web browser
- IoT devices communicating with their applications in the cloud

Browser-based attacks reflect an attack that goes through a vulnerability in the browser to reach the intended target. Non-browser attacks are those where the apps utilize APIs and web services to fetch data, make requests, and interact with web-based resources without requiring a full web browser.

Figure 9: Distribution of Attacks Based on Data for the Trailing 18 Months

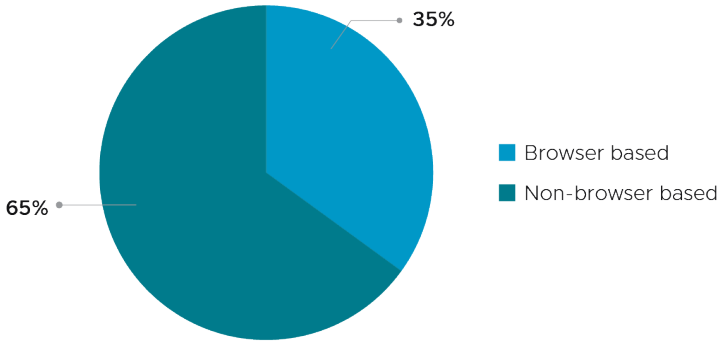

10% OF WEB REQUESTS ARE BLOCKED BY POLICY, COMPLIANCE, SITE RISK LEVEL, INTENTIONAL OR UNINTENTIONAL MALICIOUS ACTIVITY USING SECURE WEB GATEWAY

A Comprehensive, Layered Approach to Security

While organizations strive to keep up with security updates, threat actors tend to exploit exposure to known critical vulnerabilities.

A comprehensive security strategy must offer threat and data protection for systems that are exposed to these known vulnerabilities. Edge locations also have users accessing web applications, and the Internet exposes organizations as a potential entry point for threats.

In Figure 10, analyzed data shows a layered defense using Symantec SSE to prevent threats from reaching the edge.

Figure 10: Layered Defense Using Symantec SSE Prevent Threats from Reaching the Edge

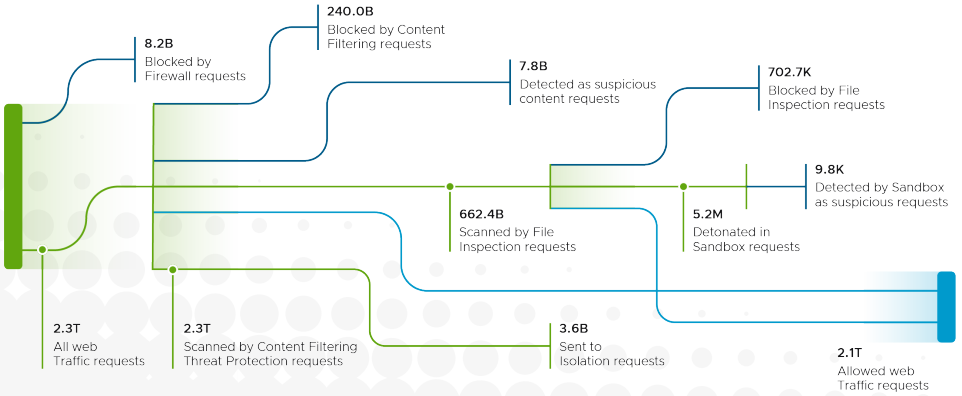

Recommendation

Each stage of edge technology evolution is capable of transforming a variety of industries. The latest stage—the intelligent edge—is on the brink of rapid adoption. Deployment data shows that customers have taken steps to implement an AI-ready infrastructure by adopting cloud-first architecture, optimizing connectivity, and addressing security and compliance requirements. However, as AI workloads at the edge continue to proliferate, enterprises must take this moment to reassess their network infrastructure.

A RESILIENT AND ADAPTIVE NETWORK IS ESSENTIAL FOR CAPITALIZING ON THE FULL CAPABILITIES OF THE INTELLIGENT EDGE

A resilient and adaptive network is essential for capitalizing on the full capabilities of the intelligent edge, particularly in enabling real-time data processing, ensuring seamless connectivity, and maintaining secure operations. Organizations that proactively evaluate and upgrade their network infrastructure to meet these requirements will be better positioned to support scalable, AI-driven outcomes that drive competitive advantage and operational efficiencies.

AI in the overlay can also play a pivotal role in detecting and blocking unauthorized or malicious AI applications on the network through SD-WAN appliances. By continually learning from new network patterns and adapting to emerging threats, AI can improve its accuracy in detecting and blocking unauthorized or malicious AI applications over time. This dynamic threat detection capability will become essential as malicious AI applications evolve.

With millions of circuits under management and extensive experience across hundreds of thousands of customer environments, VeloCloud solutions have helped customers evolve their infrastructure to be performant, available, and resilient to address the needs of each stage of edge transformation.

BROADCOM®
connecting everything®

For more information, visit our website at: www.broadcom.com

VC-SD-WAN-SOTEE-WP100 January 14, 2025