CLOUDBRINK

SOLUTION BRIEF



The New Office is Anywhere

Personal SASE with high perfprmance ZTNA: Enabling enterprises to deliver best-in-class quality of experience and security for their Workers in today's work-from-anywhere world.

Work patterns have changed drastically over recent years, requiring organizations to adapt to users expecting consistent access to applications and data from any device and any location. Their employees are increasingly mobile, often with poor wireless conditions and varying quality of internet connections, all while utilizing immersive digital experience applications like video conferencing. In addition, ZTNA and VPNs were created to extend corporate networks to remote workers, but were never designed with quality of experience, application acceleration, performance, or scalability in mind. More recently SD-WANs have greatly helped multi-branch office topologies, but they still do not scale for or give the quality of experieince needed by the hybrid workerforce. Meanwhile, data breaches have skyrocketed as organizations struggle to implement effective security for their ever-changing application landscape.

The Challenge

Enterprises must increasingly modify and diversify the way they provide application and data access to their employees. Data, applications, and employees previously were co-located in the office, and interconnected via secure hardwired lines. Over the last several decades, applications have begun moving out of local data centers and employee work behavior has shifted to a mobile-first approach, and organizational responses so far have been more stopgap measures than true architectural redesigns. While organizations desire to provide improved quality of experience (QoE) to their employees, to date this hasn't been possible due to an absence of products on the market that directly target QoE.

One of the first types of mobility required by end users came from laptop computers and WiFi. Employees needed the ability to move about their offices and remain connected to data and applications, forcing enterprises to implement WiFi to keep their users productive. Providing reliable, secure wireless access brings its own set of challenges – especially for home users – yet is feasible for enterprises with an adequate budget. Employees also needed the ability to access their on-premises data and applications while at home or on the road, giving rise to VPNs. Remote work increased the need for QoX demanding applications like collaboration tools, all while reducing the quality of the network, resulting in significantly reduced productivity for end users and their organizations.

More recently SD-WANs have revolutionized WAN networks, providing improved security, performance, and resiliency for multi-office topologies. However they do not provide any QoX improvements for end users over the most unreliable portion of the office network – wireless. Nor do they account for home offices unless a proprietary home router is deployed, which requires IT to procure, manage, and support yet another piece of hardware – which lacks QoX features – for every remote worker. Not only are there large cost and scalability issues associated with this, it still doesn't address mobile workers. To provide connectivity to every employee without hardware reliance, organizations frequently revert to legacy VPN technology.

VPNs have worked well for the occasional remote worker, however they have scalability challenges once a significant number of workers have shifted to working remotely. These scalability, performance, and QoX issues are further magnified when customers use firewalls as VPN concentrators. Additionally, once a user has connected to the internal network through the VPN, their device is assumed to be trusted. This means a single compromised device can provide prolonged access to an enterprise's entire network, increasing the probability of a data breach or ransomware attack.
Over the last decade, the data and applications that users access have slowly shifted from on-premises data centers to public clouds. This causes organizations to choose between forcing a large amount of user traffic to be "hair-pinned" through the on-prem data center to the SaaS application – resulting in a poor user experience – or increased administrative and operational challenges around application authentication.

The rise in the work-from-anywhere workforce and SaaS based applications demands a new network and security architecture built specifically with QoE in mind. Only Cloudbrink can deliver this required architecture in a 100% software-only solution, enabling increased productivity and security for your organization.

The Cloudbrink Solution

Cloudbrink's **Personal-SASE Service** with **High-Performance ZTNA** is a software-only solution which provides excellent QoE for all users, even when connected to poor networks, in a highly secure manner. This provides increased productivity for end users, and fewer threat vectors for organizations. Cloudbrink accomplishes Hybrid Access as a Service through the following key features:



FAST Edge Technology – Cloudbrink utilizes machine learning to proactively provision and deprovision software-only edges based on predicted use, ensuring optimal capacity nearest to end users. Since there's no reliance on proprietary hardware like traditional "points of presence," Cloudbrink has the utmost flexibility on where and when these software edges can be deployed, allowing for expansion into new regions within minutes and guaranteeing edges closest to end users.



Brink Protocol – is an overlay protocol that is used to carry all types of traffic from the end user to its destination. It enforces end-to-end security and programmable optimization functions, where the protocol can be programmed to behave in certain optimization modes on a per-stream and per-packet basis. In conjunction with Cloudbrink's FAST edge technology, Brink Protocol is at the core of providing industry leading QoE for end users.



Just-in-Time Security – Cloudbrink shifts application and data security from an IPSEC VPN or SD-WAN based architecture, to a zero-trust architecture through application- and user-centric contexts. Through the use of application and user profiles, administrators can easily and programmatically define security policies based on the principle of least privilege, thereby bringing security directly to the end user rather than a potentially distant firewall. Cloudbrink's Internet security allows for secure web access using built-in policies for ease of deployment.



Access Anywhere – With Cloudbrink users are able to access their applications in a direct manner without hairpins, regardless of their location or application. Administrators can easily define policies to ensure different categories of traffic traverses based on business need, whether it's from the client to its destination over the public internet, or through the fabric of Cloudbrink edges to its destination in the most ideal path possible. On-premises applications can also be accessed in a secure manner by placing a Cloudbrink connector in the on-prem data center, in a more scalable manner.

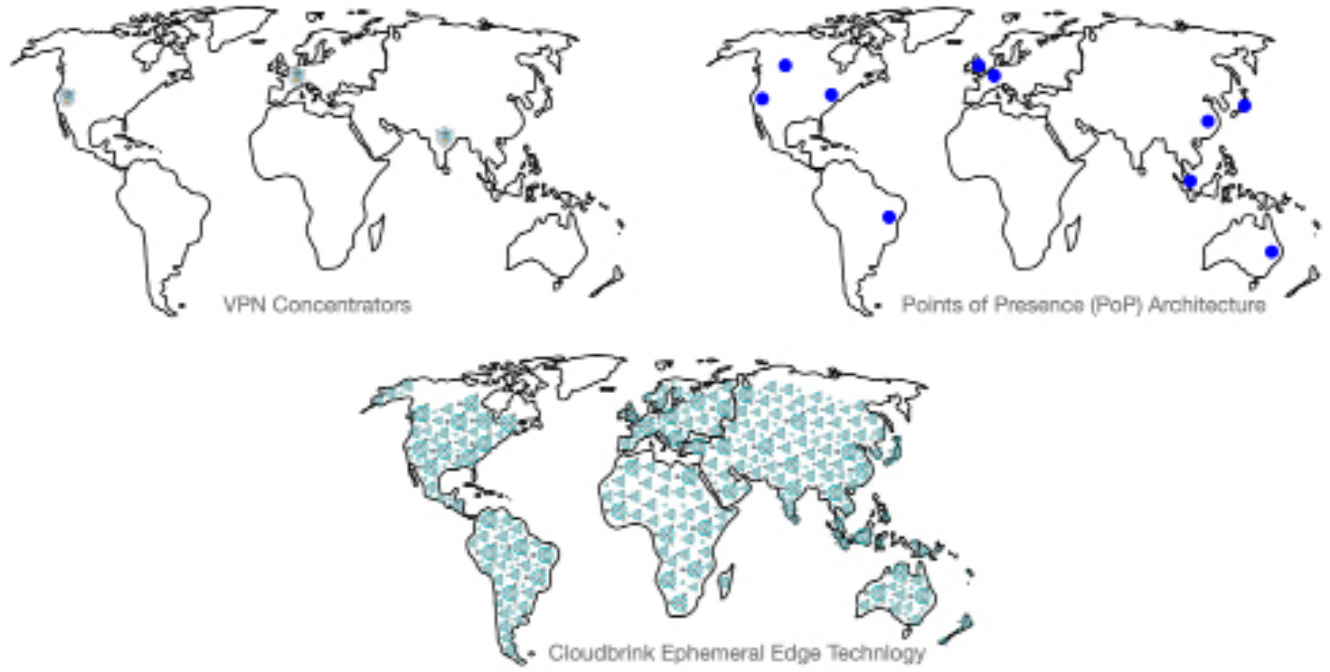
FAST Edge Technology - better than ZTNA or SDP

Traditional VPNs will have a handful of concentrators around the globe, potentially forcing a large number of end users to send traffic through a geographically distant location. As an organization scales, it's possible to add VPN concentrators as additional data centers are opened, however this is a time consuming and expensive process. In today's mobile first world, it's unlikely an organization will ever be able to keep up with the distribution of end users.

More recently, some organizations have switched to a software defined perimeter (SDP) or zero trust network access (ZTNA) architecture, which overcomes many issues present with legacy VPNs. These solutions place points of presence (PoP) across the globe, and are also able to provide efficient and secure access to both on-premises and SaaS applications. Since the SDP or ZTNA provider can focus on deploying and managing the PoPs, they can scale far more effectively than a single organization's VPN concentrators. However, these solutions are often just repackaged VPN, SD-WAN, and SWG (secure web gateway) technologies managed by an external vendor, rather than a true re-architecture. With many of these providers, adding new points of presence still requires deploying proprietary hardware in a new data center, effectively limiting how quickly new PoPs can be deployed.

Cloudbrink's FAST Edge Technology is a software-only solution, allowing new edges to be dynamically created in rapid fashion without human input and without any reliance on proprietary hardware or creation of a new data center. Machine learning is utilized to quickly react to and predict changes in end user demand, programmatically creating or terminating edges as needed. In addition to scaling across regions, edges can be horizontally added within the same region, providing scale that's not possible with traditional VPN concentrators nor ZTNA providers. Cloudbrink can also expand to new regions on the order of minutes rather than months like SDP or ZTNA providers that rely on physical data centers.

This degree of scalability ensures that Cloudbrink edges are nearest to the end user, typically only 5 to 20 ms round trip time away from the end user. This proximity, in conjunction with the Brink Protocol, is what enables Cloudbrink to provide the best in class QoE that end users demand.



Brink Protocol

Whether a user is working from the office, a home office, or on the road, the majority of access today is through wireless technologies. While WiFi and cellular networks have made great advancements in the past few years, they're still subject to a variety of challenges:



Collisions – unlike Ethernet, WiFi utilizes a collision avoidance scheme (CSMA/CA) due to its half-duplex nature. Devices utilizing WiFi will wait for an idle period prior to transmitting data, with exponential back-off times. Collision avoidance works well in ideal environments, however it quickly deteriorates in busy settings, which are typical in offices.



Interference – co-channel WiFi interference happens on "congested" networks when devices operating on the same WiFi channel interfere with each other. Adjacent channel interference is the result of nearby WiFi networks operating in an overlapping channel, which generally happens in dense office or apartment buildings. Finally, since WiFi operates on an unlicensed spectrum, other wireless devices (cordless phones, microwaves, cameras, etc.) can also cause WiFi interference. All these types of interference can cause frame corruption and a variable delay in getting access to the channel. Although cellular operates in licensed bands and the channel planning is centrally coordinated, it is not completely immune to interference.



Fading – due to reflection off of walls and other objects, signals take multiple paths to go from the transmitter to receiver. From the receivers view, these multiple path signals appear to fade, which can cause packet loss. There are built-in compensations to fading, suching as antenna diversity and MIMO (multiple-input and multiple-output). Both WiFi and cellular technologies are subject to fading.



Weak Signal – received signal strength indicator (RSSI) is a measure of the power of the signal received by the end user device. While important, the signal-to-noise ratio (SNR) is a better measure of the quality of the link, as it factors in both the signal strength and the noise due to the limitations previously mentioned. Low SNR at best results in a slow sustained data rate, and at worst results in significant frame error rates.



Roaming – cellular networks use a "make-before-break" approach for voice communication where the new connection is established prior to breaking connection to the old tower. However, this method is not used for data connections, meaning VOIP or Unified Communications as a Service (UCaaS) apps can experience packet loss when roaming to a new tower. WiFi roaming is generally quicker than cellular roaming, however users can still experience packet loss.

There are many built in techniques for both WiFi and cellular technologies that help mitigate these issues. Layer 1 (physical) techniques, such as antenna diversity, stream diversity (MIMO), and modulation and coding schemes (MCS) including block level forward error correction (FEC), help reduce these issues, but they do not eliminate them entirely.

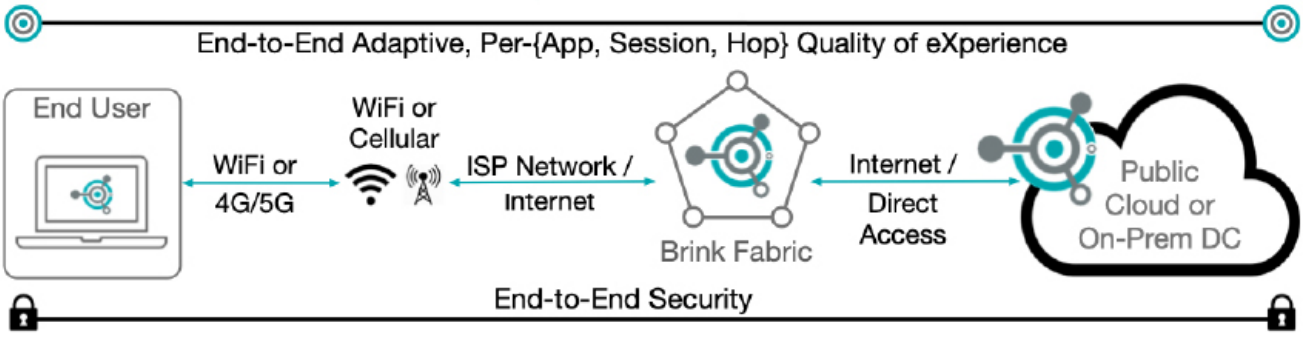
While layer 1 techniques have no or minimal effects on overall delay and delay variation, the same cannot be said for layer 2 (data link) techniques. These techniques include channel access (and the aforementioned exponential backoff), rate adaption (reactive data rate adjustments), and retransmission (including head of line blocking which prevents subsequent frames in a given flow from being delivered until the frame at the head is dropped or delivered), all which introduce significant delay variation.

More recently 5G cellular networks have seen increased deployment and utilization, which have much lower latency and higher throughput compared to 4G networks, and can even have favorable bandwidth compared to traditional home internet. While this will certainly help with demanding applications, 5G networks will still be subject to the same cellular impairments mentioned above, and

has greatly reduced range compared to 4G, which slows down rollout time and causes mixed network complications. Once 5G is ubiquitous, new applications will surely rely on the increased bandwidth of 5G and QoX will continue to be paramount.

Most UCaaS applications have implementations aimed at mitigating the impairments described above by detecting end-to-end delay, loss, and jitter. However these implementations are not segment aware, so there's no way to determine whether the issues are due to real network congestion or wireless impairments. Therefore these implementations are limited to high level FEC, retransmissions, and reducing fidelity or resolution. In many cases, by the time these corrective actions have been taken, the underlying network conditions may have changed.

Brink Protocol



Cloudbrink's Brink Protocol is an overlay protocol developed to specifically address the impairments described above. Brink Protocol can optimize traffic in the most challenging portion of the network (the last mile, or client to nearest edge), presenting a much improved network to the UCaaS application running on the end user device. Brink Protocol is not limited to last mile optimizations; through Cloudbrink's FAST Edge Technology a mesh between several Cloudbrink edges can be created, allowing for per-segment optimization. Traffic between these edges may be carried over the public internet, direct access cloud backhauls, or a combination, depending upon use case. Finally, a Cloudbrink connector can be added in the data center, enabling Brink Programmable Protocol to carry, optimize, and secure traffic end-to-end.

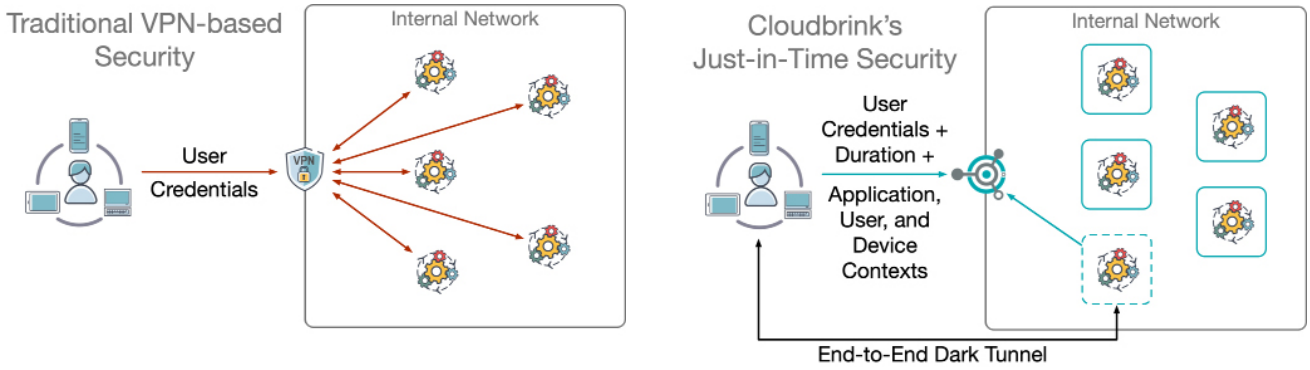
Brink Protocol continuously performs dynamic network profiling, which is responsible for characterizing various network conditions such as available bandwidth, congestion, loss (burst vs. random) and jitter in real-time. With this network profile Cloudbrink optimizes each application session over Brink Protocol on a per-packet basis, where the techniques applied to each packet may vary depending upon underlying network conditions. The techniques used per-packet include, but are not limited to, adaptive retransmissions, multipath transmissions, duplication, packet-level interleaving, FEC, signal/message bundling, data caching, and data compression. Brink Protocol also enforces jitter management and trades off a variety of application requirements in real-time (e.g., reliability vs latency vs jitter) on a per-session basis.

In summary, Brink Protocol provides end-to-end optimization and security, on a per-segment, per-location, and per-user basis, which no other solution provides. In conjunction with Cloudbrink's FAST Edge Technology, the Brink Protocol delivers an industry leading QoE for end users.

Just-in-Time Security

Traditional security architectures with VPNs rely on a simple authentication mechanism based on user credentials. Once authenticated, they're treated as a normal internal user with full access to the network. If a user's device or credentials are compromised, the attacker can easily discover internal applications, data, and servers. Even if the compromised user doesn't have authorization to these internal resources, the attacker can discover them through network scans, and use brute force or known exploits to gain access to these systems. This can result in a data breach or a ransomware attack, costing the organization millions of dollars.

With Cloudbrink's **Just-in-Time** security, in addition to authentication through a user's credentials, policy and rules are implemented at the application, user, and device level. These individual contexts are easily and programmatically configurable by administrators through Cloudbrink's SaaS management portal – which ties into industry standard identity and authentication providers – and then pushed to the device, bringing security directly to the end user. By left shifting functions such as Internet security, network performance and employee productivity is improved.



Assuming the user and their device has the authorization to access an internal application, a per-device end-to-end dark tunnel is established through mTLS authentication, and access is limited to a predefined duration, rather than permanent access. These internal applications and servers are not discoverable within the internal network as all access is handled through an inside-out connection from the application to the Cloudbrink connector. This means if a user's device is compromised, the attack vector is limited to a single device – as all user traffic is segmented in isolated dark tunnels – and to specific applications within a finite period of time.

Access Anywhere

Prior to the popularity of the public cloud and SaaS applications, having a user send traffic to a distant VPN concentrator was generally unavoidable, as the data and applications they were accessing were behind the VPN within the organization's private data center. In today's world where users routinely connect to SaaS applications, forcing all traffic through a full tunnel VPN can result in a network hairpin, and a poor QoE for end users. Split tunnel VPNs can improve performance, but require additional administrative and operational overhead, and potentially increase security challenges.

For example, imagine a user working from New York, either at their home office or a small sales office without a dedicated data center. Their organization is headquartered in California, so the user's full tunnel VPN sends all traffic to California. Our New York user is trying to access a 3rd party SaaS application hosted in a public cloud in Ohio. Their network traffic now has to go from the east coast through headquarters on the west coast, to then traverse back across the country to terminate in the midwest. Even with ideal home internet access, the end user will face high latency when using the SaaS application. When the user doesn't have ideal internet access, or has interference on their WiFi connection, these issues will compound the hairpin problem resulting in an extremely low QoE..







If an organization implements a split tunnel VPN, administrators must decide on where to host authentication and access policies. If stored on-premises, the same network hairpin occurs for our New York user upon authentication. If these policies are also stored in the public cloud provider, the organization trades app access performance gains for reduced operational efficiency due to management overhead, especially if using multiple cloud providers.

Instead let's consider the same New York based end user that's running Cloudbrink software. Rather than connecting to a VPN concentrator in California, they'll instead connect to a Cloudbrink edge located in Northern Virginia (or closer), with a round trip time of under 20 ms. Their congested and high interference WiFi issues have been resolved through Brink Programmable Protocol. Assuming they have privilege to access the SaaS application through Cloudbrink's Just-in-Time security, their traffic will traverse a short distance from Northern Virginia to Ohio, which could also be optimized through the Brink Protocol if a Cloudbrink connector has been deployed in the Ohio data center.

Not only does the network traffic have to traverse a much shorter distance than the VPN examples, any poor network conditions during the traversal can be mitigated by Cloudbrink. Even in instances where the end user needs to access an application still residing in the headquarter data center, the optimizations provided by the Brink Protocol result in a much improved experience for the end user.

Summary

In today's mobile-first era, hybrid workers need the ability to work effectively from anywhere in the world, requiring organizations to securely deliver a quality work-from-anywhere experience if they expect to remain competitive and retain talent. Organizations must also adapt to the shift from entirely on-premises applications to public cloud architectures. This changing landscape requires an entirely new network architecture designed with security and quality of experience first. Only Cloudbrink's Hybrid Access as a Service can deliver on these requirements in a software-only solution and enable organizations to deliver best-in-class security and quality of experience for all of their end users.

For more information and next steps, visit us at www.cloudbrink.com

CLOUDBRINK

Corporate Headquarters
Cloudbrink, Inc.
530 Lakeside Drive, Suite 190
Sunnyvale, CA 94085

© 2024 Cloudbrink, Inc. All rights reserved. Cloudbrink, the Cloudbrink logo, and all product and service names mentioned herein are registered trademarks or trademarks of Cloudbrink, Inc. in the United States