Not all Next-Gen Firewalls are created equal.

How to build your own secure application delivery network.

White Paper

Not all next-gen firewalls are created equal.

Reconciling the discrepancy between real world buyer needs and mainstream firewall industry offering.

By Dr. Klaus Gheri
VP Network Security
Barracuda Networks, Inc.

You might have noticed that typical next-gen firewall conversations have moved on lately. With application control and user awareness having become increasingly commoditized features the firewall industry now tries to differentiate around advanced threat protection capabilities. Advanced threats are in fact a pressing problem but there are a whole range of other novel operational and economic challenges that IT organizations face in a world where applications are increasingly consumed as services.

So what is a firewall? It is both a choke point for security policy enforcement and a router for network communication. To further clarify, there is also a significant **application delivery aspect** to its functions. In this regard, it's similar to the human brain where there are two equally important hemispheres, each with very distinct capabilities and tasks. Only their intimate collaboration allows us to function efficiently.

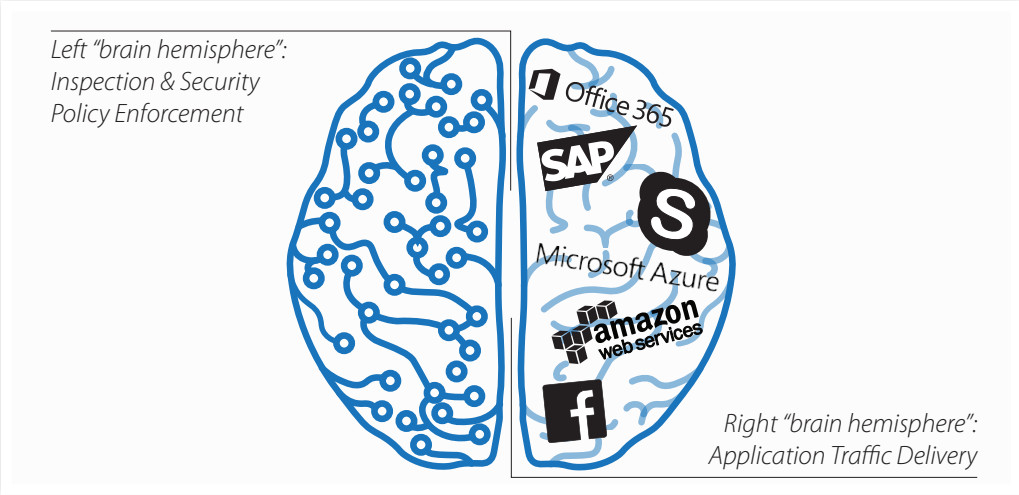

Unfortunately, the typical conversation involving firewall refreshes centers around the deep packet inspection and security enforcement. However, what about favorable and benign traffic that your lines of business rely on? This application traffic needs to be delivered securely with predictable quality of service - Site-to-Site and increasingly Cloud-to-Site. Focusing on security alone and not addressing the delivery aspect properly in your firewall is a bit like a brain that is using only one hemisphere fully while the other remains dormant. While we consider such brains as dysfunctional, we seem to be good with comparable deficiencies in our firewalls.

Let's revisit the past to better understand why this lopsidedness is becoming an increasingly limiting issue. Not long ago, it was a world where the corporate network was predominantly used to access data on internal shares, receive/send email, and leisurely browse the Internet. The two major constituents of corporate infrastructures were the LANs that users resided in, and the WAN, i.e., the network links connecting satellite offices with the main corporate locations.

Centralizing office backend infrastructures and extending applications not exactly designed for high latency environments across the WAN quickly led to performance issues and reduced productivity. This paved the way for the advent of a new breed of WAN optimization products that have since been used to address these issues.

They provide both generic optimization via compression and data deduplication techniques as well as protocol specific enhancements for certain popular applications. Typically, these pure play products would not provide adequate VPN or deep inspection functionality. Too often, customers shied away from the complexity of using firewall and WAN optimization gear side-by-side with the associated more complicated traffic flow patterns through both devices. Thus, they would often resort to MPLS-based WAN structures where all the traffic including surfing the Internet is backhauled to the data center.

At the time, there was little dependence on the accessibility of SaaS offerings or cloud services that resided outside the confines of the corporate LANs. Now, with the wealth of SaaS offerings that are increasingly affecting mainstream needs such as office application offerings (like Microsoft Office365) the backhauling approach is flawed as the MPLS backbone cannot differentiate between the various apps that are now using the same physical line. So, it's possible for an internal backup app or an update agent to totally flood the line, causing use of business-critical online apps to become difficult.

The ultimate solution to this—separate different types of application traffic from each other, either physically or logically, by partitioning the wire into different quality segments. Coming back to our hemispheres analogy, this means that the inspection hemisphere is needed to find out about the nature of a traffic flow and then the other hemisphere is needed to do something smart and deliver this traffic flow differently, from other less important traffic flows.

With a lot of SaaS offerings being used, the best option is to go for local Internet breakouts in the various locations. Of course, any Internet breakout has to be protected by a proper next-generation firewall. This sounds daunting, as now, policies of multiple firewalls need to be kept in sync. This is not the case if the firewall product comes with a powerful central management architecture where a single, or at best, a few policies can be maintained centrally, and then be pushed to the choke points for decentralized enforcement, i.e., implements the classic "manage once run everywhere" (MORE) paradigm.

The added benefit from a local breakout is that now you can operate a hybrid WAN. The hybrid WAN consists of traditional WAN lines complemented by Internet-based VPN links. A more radical approach would be to rely exclusively on Internet-based VPN but using two different ISPs and then create multiple tunnels. Again, an application-aware firewall can do the orchestration and distribute different types of internal traffic across the available links. If individual links fail, the firewall should notice and redistribute the flows accordingly, with no service disruption.

If you run such firewalls and create network redundancies, you will have built what your business is really craving—a secure communication backbone that provides all the deep packet inspection and threat mitigation you need in addition to resilient and performant delivery of business apps to the corporate users. **We call this an "application delivery network" and we help our customers build those every day.**

This solution, the Barracuda NextGen Firewall F-Series is our application-connectivity and delivery -aware, next-generation firewall for dispersed enterprise networks – and recall not all "Next-Gen Firewalls" are created equal. It's best to check how much brain capacity your current firewall is bringing to the table.

What does an application delivery network look like?

It requires advanced next-generation firewalls at each location of the app aware network.

Each location from which access to SaaS offerings is needed has a direct Internet breakout (DIB) that is controlled by this firewall. Each location uses more than one uplink to connect to the other locations. The firewall uses these uplinks in an application-selective fashion but has the capability to quickly react to line outages and reroute traffic with an appropriately adjusted policy – intelligent dynamic path selection.

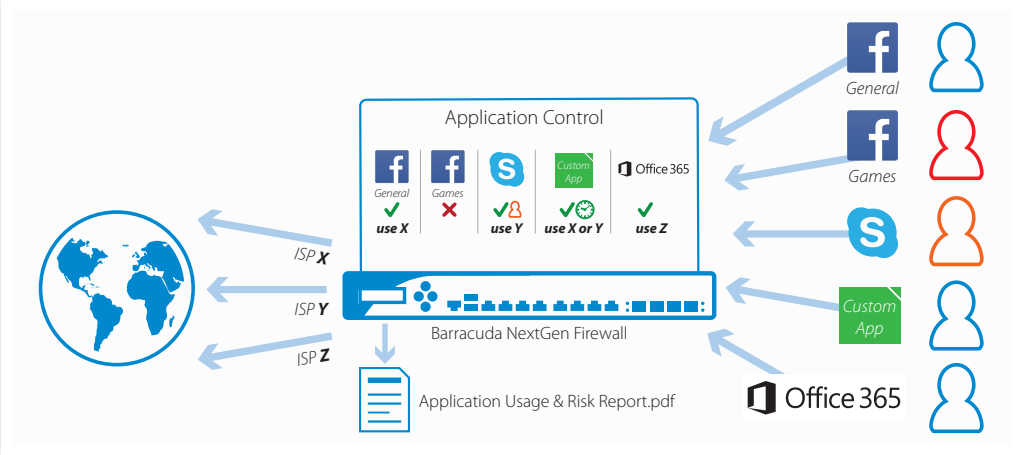

Traffic between locations can be appropriately optimized for latency and throughput via compression, data deduplication, and protocol optimization techniques. The firewall provides these services in order to have a close connection between nature of traffic and enforced delivery action in terms of quality of service, bandwidth, privacy, and delivery path.

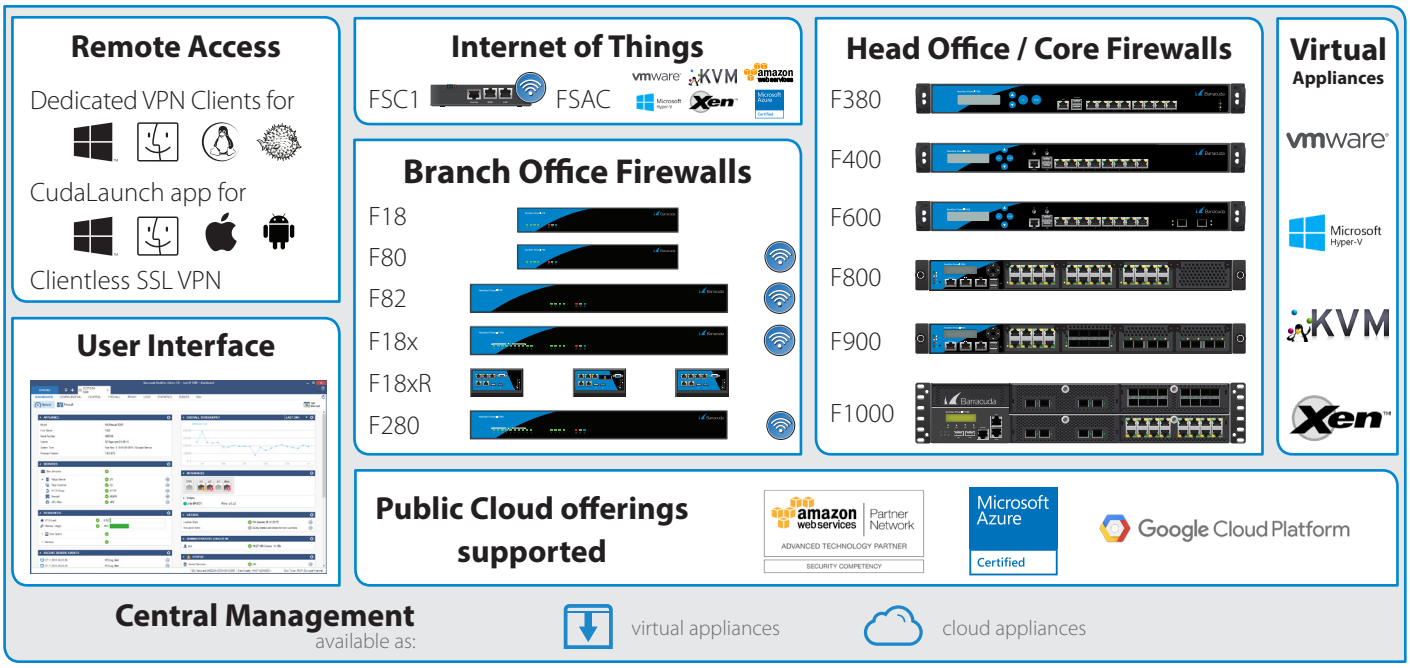

About Barracuda NextGen Firewall F

The Barracuda NextG Firewall F-Series is the ideal enterprise solution for IT administrators seeking to protect vital data in networks made chaotic and vulnerable by the explosive use of mobile and BYOD devices, evasive Web 2.0 applications, and remote network users. Barracuda NextGen Control Center adds a powerful and intuitive centralized management portal that makes it extremely simple to deploy, configure, update, and manage multiple units from a single location, while also providing comprehensive, real-time network visibility and reporting. As a result, it is an ideal solution for enterprises looking to manage large numbers of users, or several sites with few IT personnel, while meeting PCI Compliance requirements.

For questions about the Barracuda NextGen Firewall F, or for a free 30-day evaluation, visit http://www.barracuda.com/products/nextgenfirewall-f or call Barracuda Networks at +1 408-342-5400.

About Barracuda Networks, Inc.

Barracuda provides cloud-connected security and storage solutions that simplify IT. These powerful, easy-to-use, and affordable solutions are trusted by more than 150,000 organizations worldwide and are delivered in appliance, virtual appliance, cloud, and hybrid deployments. Barracuda's customer-centric business model focuses on delivering high-value, subscription-based IT solutions that provide end-to-end network and data security. For additional information, please visit barracuda.com.

US 2.0 • Copyright © Barracuda Networks, Inc. • 3175 S. Winchester Blvd., Campbell, CA 95008
408-342-5400/888-268-4772 (US & Canada) • barracuda.com
Barracuda Networks and the Barracuda Networks logo are registered trademarks of Barracuda Networks, Inc. in the United States.
All other names are the property of their respective owners.

Barracuda Networks Inc.
3175 S. Winchester Boulevard
Campbell, CA 95008
United States

t: 1-408-342-5400
1-888-268-4772 (US & Canada)
e: info@barracuda.com
w: barracuda.com