

OPTIMIZE CLOUD NETWORKING COSTS BY CREATING A FLATTER NETWORK

By Larry Cantwell,
Vice President and CTO

As cloud computing applications evolve, the resulting data center networks grow more and more complex.

With this increasing complexity, there is an understandable temptation to attempt to scale existing cloud networks and data centers by continuously adding new layers of network switching and routing elements. This results in an obvious increase in both CAPEX and OPEX, but can also manifest into network management nightmares! Most techies understand what a flat network represents, both architecturally and structurally. Due to the gradual evolution that tends to occur, the tangled and interconnected web of equipment and resources found in most data centers is generally anything but flat. Many companies have lost sight of the benefits of a flat network, choosing instead to add more elements to their systems in a race against efficiency. These same firms are literally throwing money away in a stand to make their network infrastructure scale.

Typically a flat network is described as *a network in which all devices or stations can reach one other without going through any intermediary hardware devices such as bridges or routers. In essence, a flat network is one single network segment.*

This definition omits some of the more notable benefits that a flat network can provide, including:

- Less network switching 'boxes'
- Better performance in terms of bandwidth and latency
- Lower CAPEX and OPEX costs due to simplification
- Reduced power and cooling expenditures
- Improved flexibility
- Elimination of common network problems such as spanning tree loops

These benefits are direct, desirable, and ultimately provide a lower network TCO. While data center network architectures may be growing out of control, building your cloud network from the ground up, including physical layer switching as part of the architecture, can provide a solid platform upon which to construct a simplified, flatter network. Such a network, if implemented properly, will mitigate bottle necks, security issues, scalability limitations, spanning loops, and performance liabilities of more rigid and hierarchical networks.

To further develop the thought, questions come to mind about the traditional “3-tiered network” approach. For example, what is the aggregation layer really doing? In most cases, data traffic at the aggregation layer is being routed through the core, as opposed to being routed or switched at the edge. This being the case, the aggregation problem could be better solved by a fully non-blocked physical layer aggregation device that offers minimal latency, lower power, higher scalability, and deterministic performance without bottlenecks.

BETTER PERFORMANCE

As in mathematics, systems become more tractable when you remove variables from the equation. In addition to saving money, flatter networks generally perform better because there are fewer layers of network components. By minimizing the number of tiers in the network, one reduces the potential for performance bottlenecks that can occur between network switching elements on their inter switch links (ISLs).

Consider the 3-tiered network depicted below in *Figures 1 and 2*. In *Figure 1*, we have a typical network that initially needs to scale to connect 512 device ports in an any-to-any configuration. Typical switching and routing products provide 32 ports for edge and aggregation switches, and a maximum of 512 ports for core switches. Using a typical ISL loading ratio of 4:1 (oversubscribed) would yield the following shown in *Figure 1*.

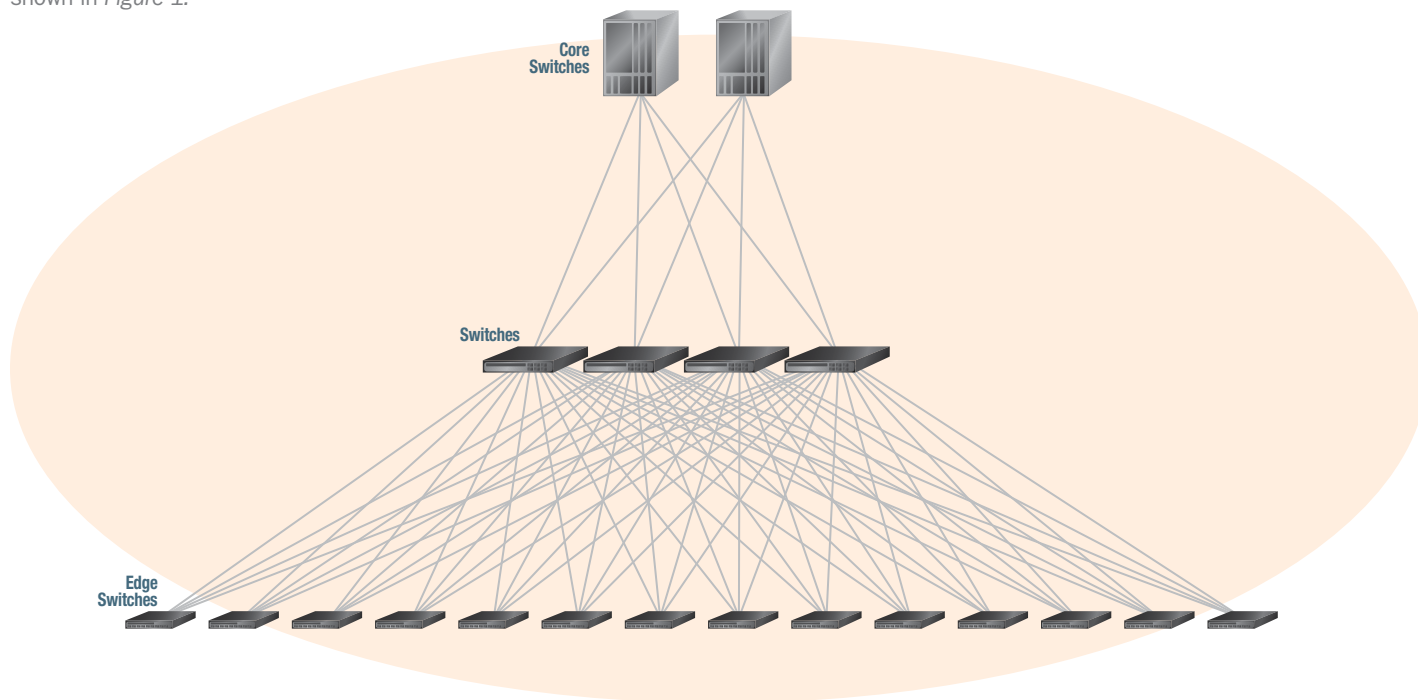


FIGURE 1 – THE TRADITIONAL 3-TIERED NETWORK

FIGURE 1. The Traditional 3-Tiered network

As server and storage resource requirements evolve, it is not uncommon to see connectivity requirements quadruple in a relatively short time. As these needs develop, the network shown in *Figure 1* could quickly evolve into that shown in *Figure 2*. Is this still a 3-tiered network, or has it evolved to grow even more layers? Using this many switching elements significantly increases the likelihood of performance bottlenecks over the ISLs that interconnect these boxes. Increased power, cooling, and management nightmares follow as well. Additionally, consider the cabling “rats nest” that accumulates over time in supporting this network. Instead of adding additional small switching elements to the middle of the network, a more desirable alternative would be the introduction of a scalable, low-power switching element that minimizes bottlenecks, latency, and management headaches.

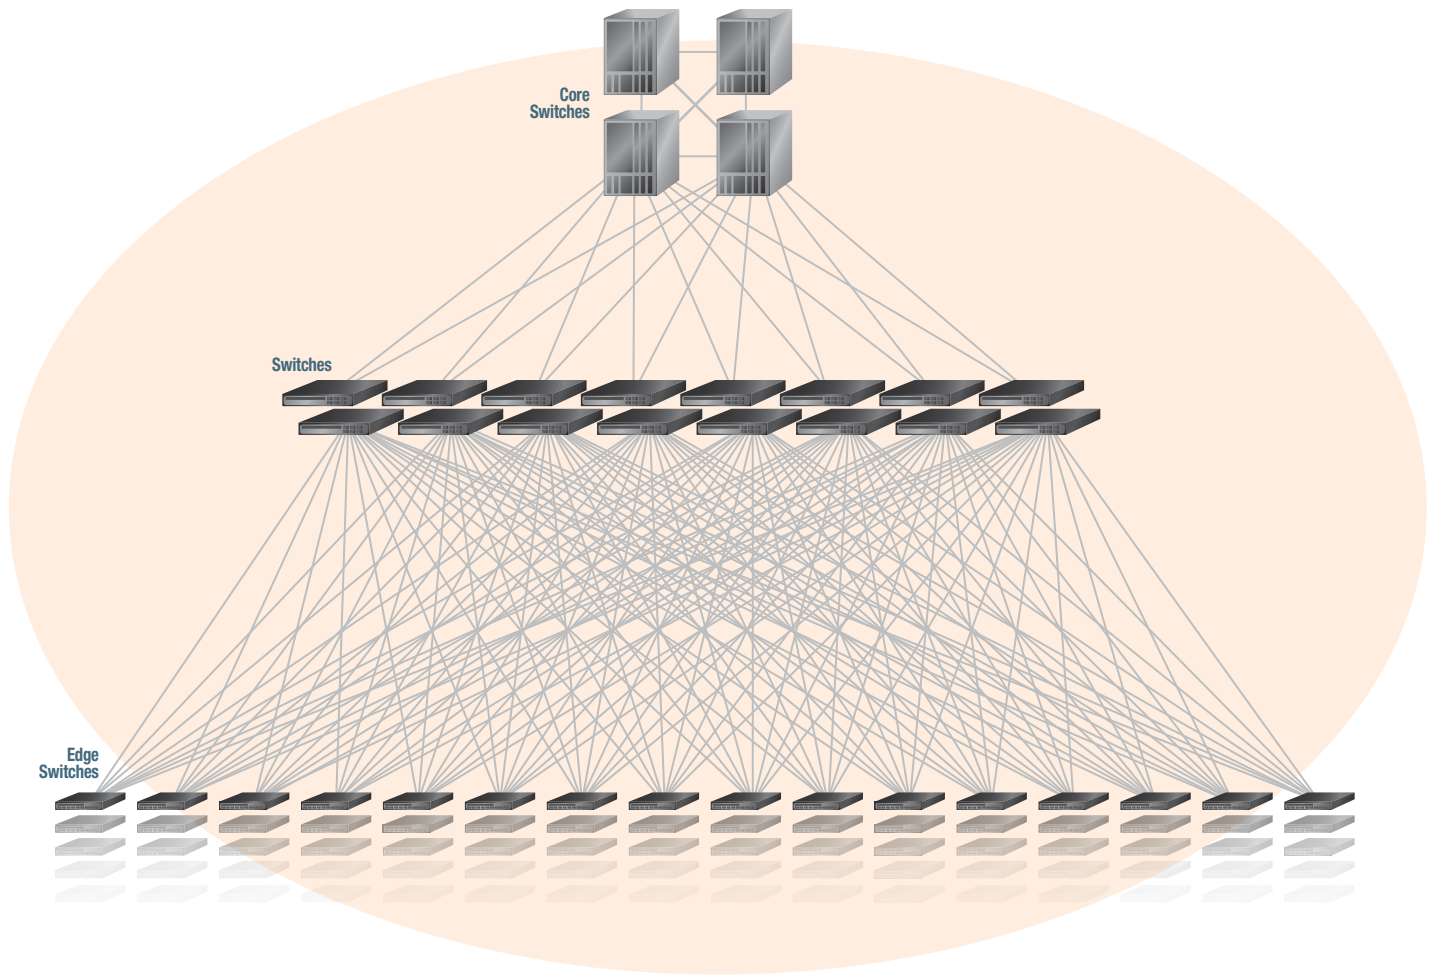


FIGURE 2. The Congested 3-Tiered network

FIGURE 2 - THE CONGESTED 3-TIERED NETWORK

Architecting your cloud network with minimized numbers of switching elements (boxes) can maximize its performance. For the time being such an ideal network model may seem hypothetical. However, there is an optimized solution that leverages the right combination of L1, L2, and L3 network switching and routing technologies available today. Specifically, this goal can be accomplished by the introduction of Layer 1 switching such as that available from ONPATH Technologies at the middle tier of the network, along with traditional Layer 2 and 3 switches placed at the edge and core. The ONPATH switch offers scalability to 4,000 non-blocking ports, enabling a flatter network. An additional benefit of the ONPATH solution is that it is significantly more cost effective and power efficient than traditional L2 and L3 switches. A redesign of the complex network in Figure 2 using a combination of Layer 1 and Layer 2/3 switching would yield the flattened network shown in Figure 3.

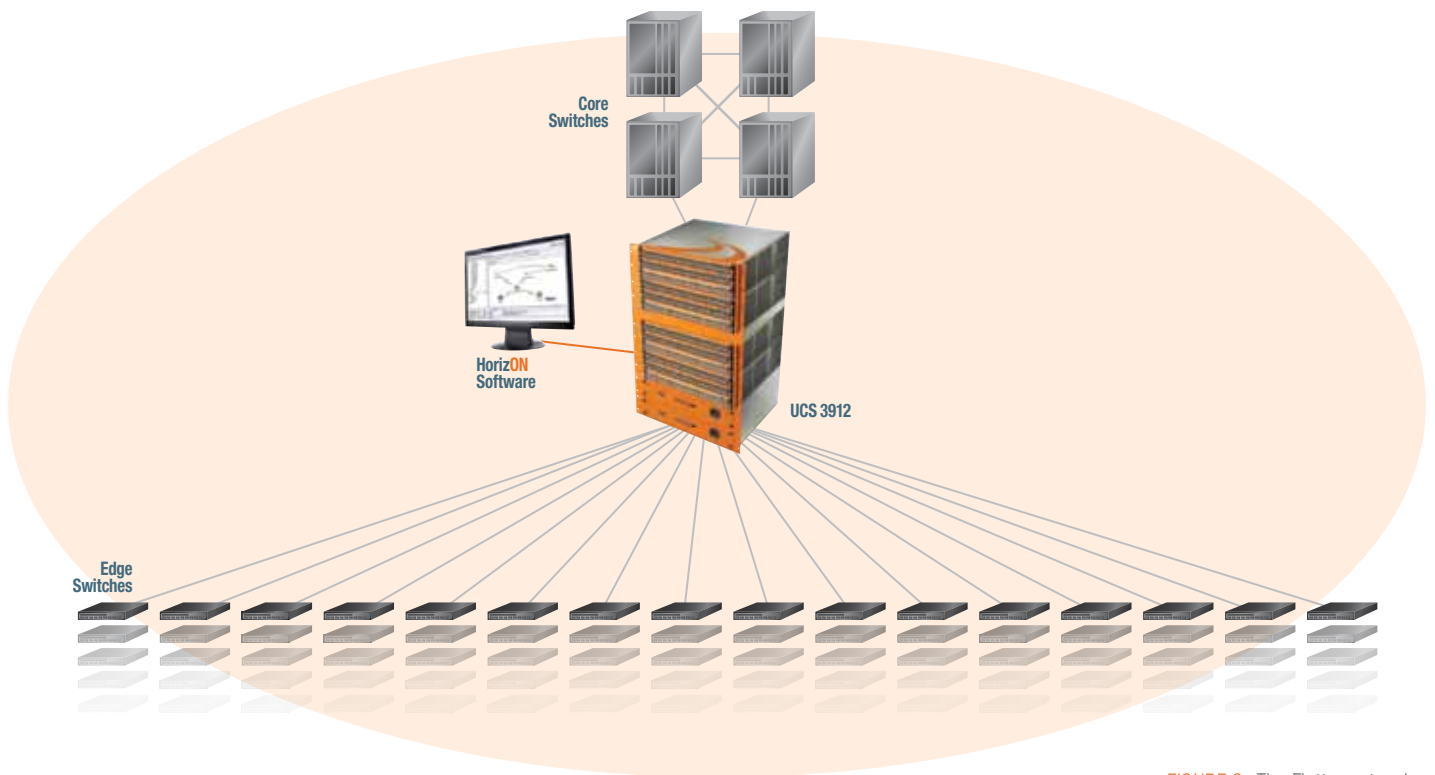


FIGURE 3. The Flatter network

FIGURE 3 – THE FLATTER NETWORK

In a flat network environment, there is no need to stitch a mass of boxes together into a network. This simplification not only eases traffic management, and results in greater reliability and improved uptime, but it also optimizes both CAPEX and OPEX. Let's explore each of the resulting benefits in a bit greater detail.

IMPROVED COST PER PORT

A flatter network will naturally result in a reduced overall cost per network port because the number of intermediary ports (inter switch links) will be reduced. Referring back to the example in *Figure 2*, it is obvious how price-per-port can soar depending on the number of switching elements and tiers within the network.

IMPROVED POWER AND COOLING

Power and cooling costs decline as parts are removed from the system, and with energy bills nearing historic highs this becomes increasingly relevant. The energy savings of a smaller floor plan with fewer operating pieces are dramatic - and this is exactly what a flatter network delivers. The rat's nest of cables typically hidden in floor tiles and crawl spaces beneath the data center further reduces cooling efficiency. Along the lines of reduced cost-per-port, a reduced power-per-switch-port can also be readily achieved. A flatter network can provide a significantly lower overall network power budget if architected correctly.

LOWER MANAGEMENT COSTS

Data Center and cloud network management costs today remain as high as seven times the cost of capital equipment. Often this complexity is unnecessarily driven by the fact that too many network switching elements are cobbled together over time in order to scale out the network. Managing the bandwidth, provisioning, partitioning, and features/functions for cloud networks that have evolved in this fashion can be very complex. By creating and maintaining a flatter network, management costs can be minimized.

IMPROVED SCALABILITY

Unforeseen growth, contraction, and consolidation are common in today's IT world. Data centers and networks can and must be designed with flexibility in mind in order to easily scale up or down as needs change. As illustrated, a flatter network will always be a more scalable network. Simply put, minimizing network elements significantly decreases the probability for bottlenecks and management complexity.

CONCLUSION

The evolution of cloud computing will drive exponential growth in connectivity requirements for next generation data center networks. The implications of this dramatic growth include network bottlenecks, increased cost, increased power requirements, and scalability limitations. Including Layer 1 switching as part of the data center architecture can assuage these demands and offer a flatter, more manageable cloud network.

Data center experts pondering whether the time has come to flatten their network should ask themselves several questions:

- Does our cloud network frequently experience interruptions? Would we like to see improved performance and reliability?
- Are we constantly purchasing new hardware and software in an effort to keep pace with expanding operations and bandwidth demand?
- Does a cluttered tangle of wires make it difficult to execute moves, adds, and changes?
- Is the power consumption of our data center facility a major expense? Would we like to reduce cooling bills? Do we have difficulty bringing new applications and services on-line because of insufficient data center power or cooling capabilities?
- Is our facility currently able to support continuous operations for higher speed cloud networks and higher density computing?
- Are we seeking improved scalability and sustainability in our network and data center?

If your answer to any or all of these questions is "yes," it may be time to explore how you can simplify and flatten your network using a new architectural approach that combines switching at the physical layer in concert with L2 and L3 switching.



ABOUT ONPATH TECHNOLOGIES

ONPATH Technologies is the leading provider of scalable connectivity and monitoring solutions for high-performance networks. With a history spanning over 25 years, ONPATH is a spin-out of Brocade Communications' physical layer switch business. ONPATH's Universal Connectivity System™ and HorizON™ Software deliver an advanced platform that automates and secures data center and test infrastructure to help network managers conserve time, increase utilization, and save money compared to manual patching or complex mesh switching architectures. Our patented switching technology and advanced software deliver the industry's most scalable, secure connectivity solution with the most in-depth view of your network. Scalable from 8 to 4,096 non-meshed ports, independent of your speed or protocol requirements, ONPATH delivers the flexibility and security required by today's data centers and the investment protection necessary for those of tomorrow. ONPATH currently has over one million installed ports throughout Fortune 1,000 and Government customers.

For more information on how ONPATH can help you with your initiatives, contact us today for an engineering application review or network consultation.

2000 Lincoln Drive East
Marlton, NJ 08053
609.518.4100
info.request@onpathtech.com
www.onpathtech.com