

Scalable Non-Blocking Infrastructure Management: A Comparative Study

Abstract

Today's enterprise environments impose demands for a large-scale, highly available, fault tolerant physical layer infrastructure. Typical storage area and data communication network environments, within a single data center or across geographically dispersed sites, can easily require several thousand ports of connectivity. The ability to plan for and build out network infrastructure in such a way that it is cost-effective, scalable, non-blocking, fault tolerant, and can be monitored and managed more effectively, is of great value to businesses.

OnPATH Technologies' Universal Connectivity System™ 2900 (UCS™ 2900) family of switch products provides a cost-effective solution to automating enterprise connectivity by scaling from a few hundred non-blocking ports to 4096 non-blocking ports and by simplifying overall infrastructure management. This technical paper discusses and compares approaches to designing and building large scale networks to meet today's demands at the physical layer.

Introduction

The Universal Connectivity System™ 2900 (UCS™ 2900) provides cost effective any-to-any non-blocking connectivity for enterprise environments. UCS 2900 systems simplify networks, automate changes and provide tools for monitoring and management of the infrastructure. The key value proposition of this offering is in the automation and virtualization of the network infrastructure layer. Having been deployed in Fortune 100 data centers and labs for more than two decades, UCS 2900 switches are the only solution designed for the enterprise that scales from hundreds of ports up to 4096 ports, to provide a completely non-blocking architecture in a single system.

Frequently, attempts are made to build a non-blocked, switched infrastructure layer over time using seemingly lower cost smaller port count switches. This approach tends to become very costly and has reliability and manageability problems that grow exponentially as configurations expand. The focus of this paper is to discuss and draw comparison between the use of a larger, more scalable product for automated switching infrastructure versus attempting to accomplish the same task by meshing together smaller switches.

The OnPATH Technologies Approach

The UCS 2900 product family cost effectively scales as infrastructure needs grow by offering a family of products, which make use of common hardware and software that can be added to over time.

The innovative approach used in the UCS 2900 product family leverages the widely accepted principles for non-blocking switching developed by Charles Clos at Bell Labs. These principles have been refined and augmented for modern Network and SAN environments over the past 2 decades. The basic premise is that an integrated, highly scalable, non-blocking switching platform can be built using switching elements arranged in a 3- stage configuration. OnPATH Technologies has perfected and patented many elements of this approach over the course of three product generations yielding a highly reliable, fault tolerant, cost effective solution.

By definition, non-blocking means that any IO device can make a connection through the switch to any other available IO device at any given moment in time, without regard to other connectivity already established in the environment. To understand the nature of a non-blocking switched infrastructure, consider a switching array with i inputs and o outputs. Typically, i and o would be equal to each other and constitute n number of full-duplex user ports (or paths). The middle switching stage must have $2n-1$ duplex switching paths in order

to provide full non-blocking any-to-any connectivity for the n - IO ports. Research into building networks with less than $2n-1$ core switching connections in a “rearrangably non-blocking” array has proven this approach to be non-viable. The ability to manage such a configuration is not practical in that it poses the need to randomly tear down established connections that are carrying live data.

The UCS 2900 system makes use of an innovative approach to providing all necessary $2n-1$ interconnections between the switching elements as an integrated part of the system. The integration of these elements highlights an obvious advantage in reliability, management and cost savings relative to the alternative approach of building an array out of smaller switching systems. The small system approach is burdened by the need to connect user ports together in a 3-stage array using external fiber or copper cable.

The UCS 2900 3-stage switching array is depicted in the following diagram (Figure 1). The Input and Output stage resides on the port cards. These are the same boards, only drawn separately in this diagram to help visualize the 3- stage architecture.

The middle stage is comprised of central switching cards. In the case of both port and middle stage cards, there is a non-blocking switch element resident on each module card. This fact, along with the innovative physical interconnect between the boards, allows customers to grow the port capacity as their needs increase. Configurations up to 4096 ports are accommodated simply by adding additional boards and chassis'. The patented interconnect scheme provides the final piece of the 3- stage non-blocking array. The middle stage of the array has at least $2n-1$ switching paths (where n is the number of I/O user ports) to provide a completely non-blocking switch.

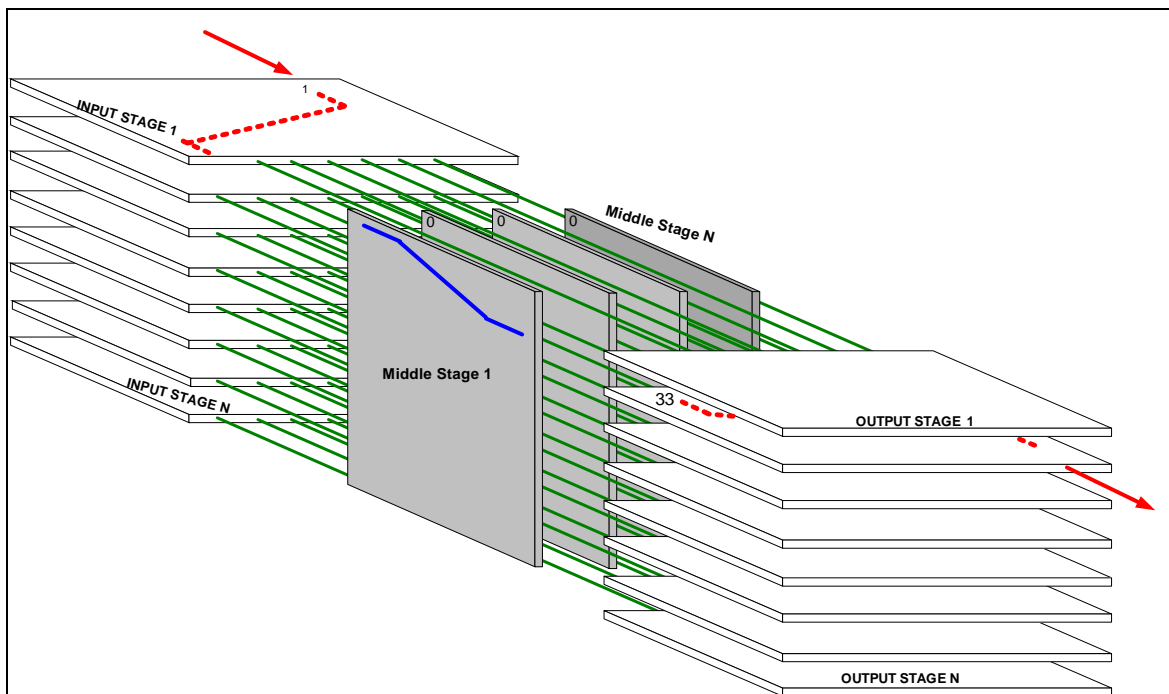


Figure 1: The Universal Connectivity System 2900 3-stage switching array.

Comparison to the “Small Switch Approach”

To draw a comparison, consider an alternative approach of building a non-blocking physical layer switching infrastructure using 288- port systems. First, consider an equivalent to the UCS 2910 model, which has 1024 fully non-blocked ports in a single chassis.

In order to build a functionally equivalent non-blocking switch out of 288- port switches, one would need to design and construct the arrangement as diagramed in Figure 2 (below).

The cost and complexity of building an arrangement as shown in this figure, along with the necessary cabling and optical interfaces factored into the design, makes this approach undesirable and very costly. In this case, eighteen 288- port switches, or nearly 5200 ports of connectivity, as well as over 2000 additional middle stage cables and optical interfaces, would be required to build the same non-blocking 1024 user port array that is housed within a single footprint in the UCS 2910.

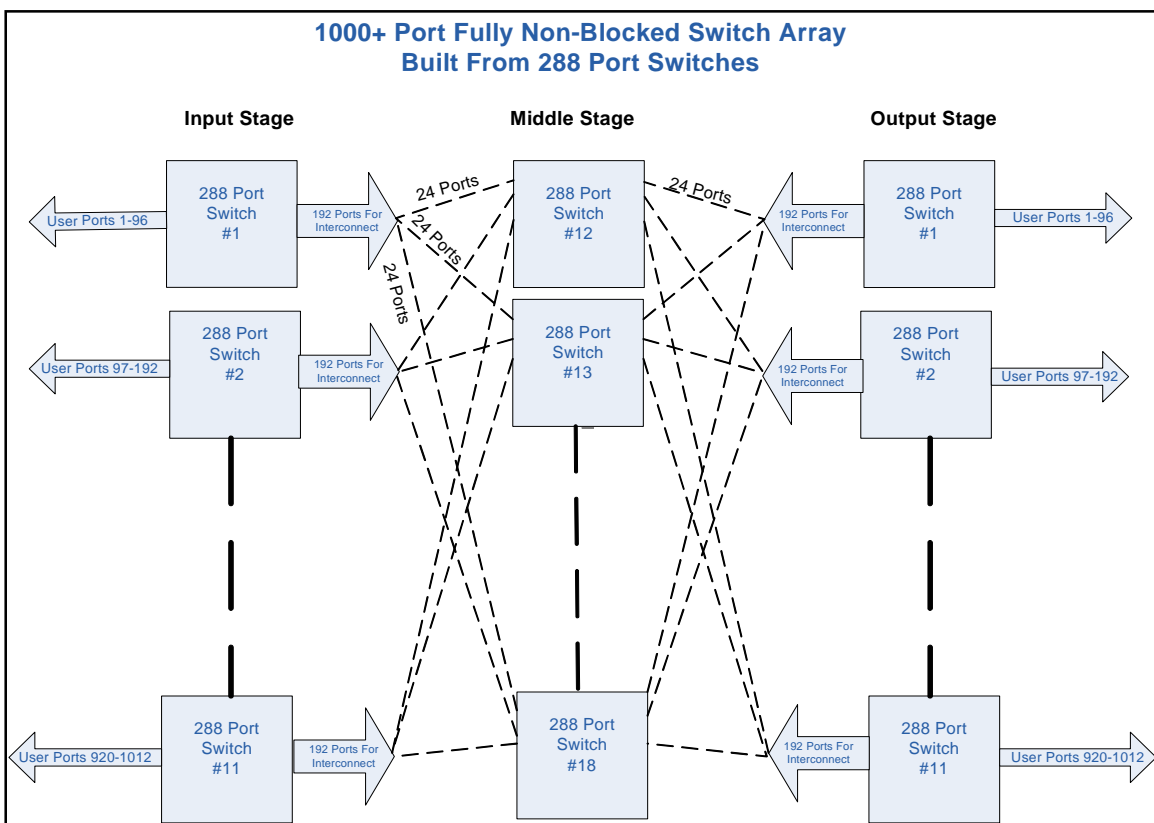


Figure 2: A total of eighteen (18) 288- port switches are required to build a true 1000+ port fully non-blocked switch array; Total required ports = 5184. Additional cables and optical interfaces are required as well.

Finally, draw another comparison between the UCS 2910 and building a 3-stage array, again using 288- port switches only this time allowing for significant probability of blocking in the 288- port array. In other words, allow the available user ports to outnumber ports going to the middle stage by a ratio of 2:1 on each 288- port switch.

While such a configuration is not recommended, as it can quickly lead to blocked circuit paths even with a light amount of connectivity, it is a useful exercise for this comparative analysis (see Figure 3).

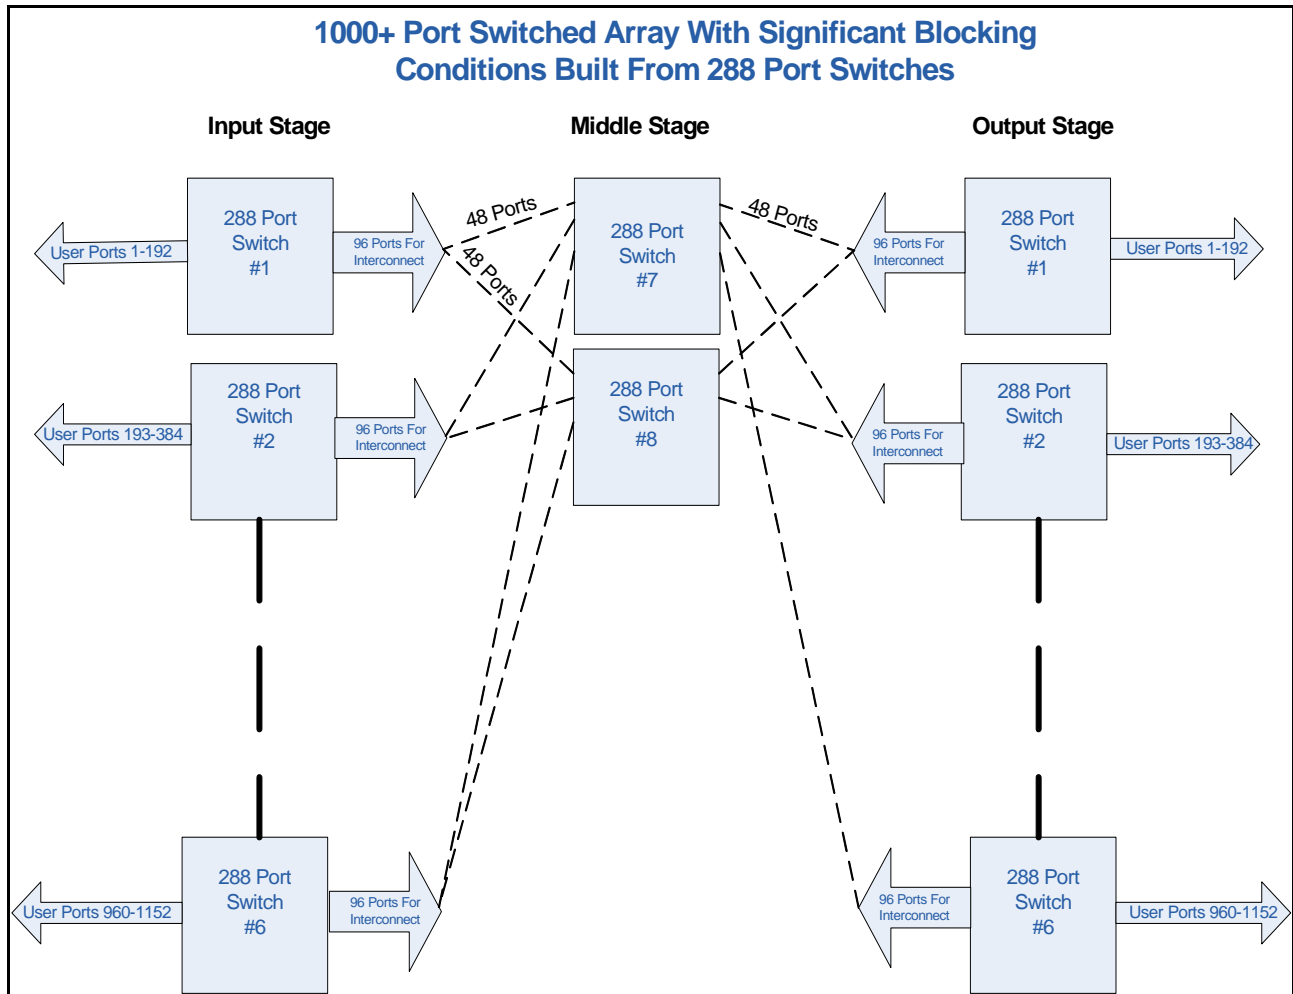


Figure 3: A total of eight (8) 288- port switches are required to build a 1000+ port switch array with significant blocking conditions; Total required ports = 2304. Additional cables and optical interfaces are required as well.

The result shown in this design is not practical in most enterprise environments where blocked circuits at the physical layer are not acceptable. However, even in comparison to this oversubscribed configuration, relative costs for the UCS 2900 would be significantly less expensive and the resulting solution much more reliable and scalable, as this configuration requires eight 288- port switches, or 2304 ports of connectivity, to build the above blocking array.

Conclusion

Today's enterprise customers demand highly available, fault tolerant, scalable infrastructures to support their mission-critical SAN and network environments. Implementing an automated physical layer solution that can grow cost effectively and provide any to any connectivity (non-blocking) and ease of management is critical. Based on decades of experience in the most

demanding network environments worldwide, the Universal Connectivity System 2900 product line, with offerings ranging from several hundred to 4096 non-blocking ports, provides a more cost effective and time proven solution to enterprise infrastructure connectivity than designing mission-critical networks with smaller port count switches.

About OnPATH Technologies

OnPATH Technologies provides automated connectivity solutions to help organizations simplify their physical layer infrastructure to improve the availability, manageability, and performance of their IT environments. Through the delivery of automated connectivity solutions, OnPATH Technologies simplifies management and reduces costs associated with this critical virtual connectivity infrastructure. The company's solutions meet the stringent requirements of enterprise organizations and government entities with complex, dynamic and mission-critical IT environments. OnPATH Technologies is located in Marlton, New Jersey. For more information, visit www.onpathtech.com, e-mail info.request@onpathtech.com, or phone 1.609.518.4100.