BRIGHTWAYS UNITEDSA-1776 DATASHEET

Cloud and AI Hyperscale Switch

- *U.S. Patent 10,630,606 "System, Method, And Architecture For Data Center Network Switching"*
- *U.S. Patent 11,206,225 "Hyperscale Switch And Method For Data Packet Network Switching"*
- *U.S. Patent 17,567,090 "Hyperscale Switch Element For Data Center Network Switching"*
- *U.S. Patent D906305 "Network Switch"*
- *Patent Cooperation Treaty (PCT) PCT/US2019/038238 Written Opinion Of The International Searching Authority*

TABLE OF CONTENTS

UNITEDSA-1776

In the rapidly evolving landscape of Cloud and AI technology, the importance of power efficiency in network infrastructure cannot be overstated. Enter the Brightways UnitedSA-1776 Cloud Hyperscale Switch—an innovative solution that sets a new standard for performance and energy consumption. This groundbreaking switch is meticulously designed to address the demanding needs of modern data centers,

effortlessly combining high throughput with remarkable power efficiency unseen before now. With its ability to support an extensive range of Ethernet modulations and facilitate over 2,000 connections per system, the SA-1776 is not just a switch; it is a pivotal element that empowers businesses to thrive in energy-conscious environments while meeting their ever-increasing connectivity demands.

HIGHLIGHTS

Connectivity Features

I/O Connections per Line Card:

- 180 x 10G Ports
- 132 x 25G Ports
- 45 x 100G Ports
- 18 x 400G Ports

Power Efficiency

Power Consumption:

- 1.9W per 10G port
- 2.1W per 25G port
- 6.0W per 100G port

Innovative Technology

Patented Switch Engine Technology:

Pspine™: Significantly reduces power usage and latency.

Interwatt Chip:

Future-Proof Capability

Field Upgradeability:

- Supports modular upgrades (10G, 25G, 50G, 100G, 400G) via simple license upgrade.

Reprogrammable FPGA

- Fully compatible with Pspine, allowing seamless scalability from 10G to 400G without new hardware.
- Easy license download for rapid deployment.

Performance Features

Buffer Capacity:

- 12GB deep buffers per line card for enhanced performance during high load transients.

Cloud Switching Technology:

- Requires only 1U of power supply, compared to competitors demanding up to 6U, ensuring efficiency and space savings.

Installation & Compliance

Operational Parameters:

- Voltage: 120/230V AC

- Frequency: 50/60Hz

- Dimensions: 17.5" (W) x 31.5" (H) x 26.5" (D)

- Weight: 450 lbs fully loaded

Safety and Compliance:

- Meets all relevant safety and EMI standards, including CSA, CE, CQC, RoHS, China RoHS, EAC, Nordic Swan, and BIS certifications.

Networking Features

Jumbo Frame Support:

- Packet sizes up to 9300 bytes.

Custom Linux-based Network Operating System (NOS):

LanBolt: Powered by OpenVswitch, enabling seamless L2 and L3 protocol integration with robust VLAN support.

Upgrade Path

L2 to L3 Transition:

- Begin with an L2 switch and upgrade to L3.

User Experience

GUI Customization:

- Intuitive and easily customizable graphical user interface.

Expert Support:

- Access to Brightways engineers for assistance.

Unlock the potential of your network with Brightways' innovative solutions, designed for performance, reliability, and future scalability.

MAJOR TECHNOLOGICAL INNOVATION

In both the automotive industry and the realm of cloud and AI switching, there is a relentless drive to innovate, largely fueled by the need for efficiency, performance, and sustainability.

Figure ES-4. Percent Change in Fuel Economy, Horsepower, and Weight Since 1975

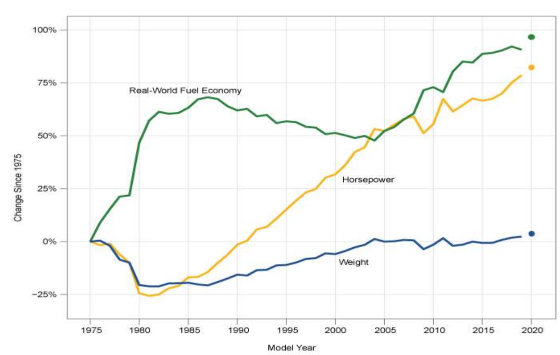

In the automotive sector, the push for innovation has led to breakthroughs like electric vehicles and advanced fuel technologies. The focus is not just on enhancing engine performance but also on improving fuel efficiency and reducing emissions. Innovations such as hybrid systems and autonomous driving technology exemplify this quest for smarter and greener solutions.

Similarly, in the cloud computing and AI switching domain, there is a continuous pursuit of advancements aimed at optimizing performance and reducing energy consumption. Companies like Brightways are pioneering technologies that enhance networking capabilities while minimizing electrical usage. Innovations like hyperscale switching and custom systems on a chip (SoC) demonstrate a commitment to achieving greater efficiency and lower latency, much like how the automotive industry seeks to maximize fuel efficiency and reduce carbon footprints.

Both industries are working towards creating systems that are not only powerful but also environmentally sustainable. The drive for innovation in cloud and AI switching mirrors the automotive industry's efforts, as both sectors strive to harness cutting-edge technology to deliver better performance while being mindful of resource consumption. This parallel pursuit underscores the universal goal of enhancing the capabilities of existing technologies, whether on the road or within the digital infrastructure.

THE HOP

Understanding the concept of "hops" in networking switches is fundamental for grasping power consumption dynamics in these systems. Each hop represents a point where data packets are processed as they travel through the network. The architecture of switches, particularly in high-speed environments like data centers, involves specialized chips and components, such as transceivers and MAC units, which handle this processing.

A critical aspect to note is that transceivers contribute significantly to overall power consumption—up to 70%. They are essential for transmitting data but do not directly participate in the routing process. This means that the more hops data takes, the more transceivers come into play, leading to higher energy consumption without necessarily improving data routing efficiency.

The Brightways network topology addresses this by minimizing the number of hops a packet encounters. By reducing the number of transceivers involved, it decreases both power consumption and latency. This understanding is essential because it highlights the interplay between network design, efficiency, and energy usage. Optimizing network topology not only enhances performance but also contributes to more sustainable operation in data centers and networking environments.

CURRENT HYPERSCALE NETWORKING

In the 1950s, the Clos Network was introduced as a solution to the persistent issue of call blocking within telephone switching offices, ostensibly a nod to innovation. Developed by Charles Clos, this multi-stage circuit-switching design features two layers of packet switches. While the lower layer connects to external devices like server network interface cards (NICs), the upper layer, known as uplink switches, is meant to provide alternative routes to alleviate congestion.

While the Clos Network enabled traditional plain old telephone service (POTS) to connect analog phone lines without busy signals, it can be argued that this approach is now outdated and hardly sufficient in today's fast-paced technological landscape. The initial success with POTS does little to address the significant pitfalls encountered in packet switching, specifically regarding the critical need for power-efficient solutions. This inefficiency raises doubts about the viability of the Clos Network in modern applications, emphasizing a disconcerting gap between historical innovation and current operational demands. As progress continues, reliance on outdated methodologies like the Clos Network may hinder advancement and adaptability in an increasingly competitive market.

Clos Loss

Impact of CLOS Network Topology on Power Consumption Performance

Introduction

The CLOS network topology, while widely used in data center networks for its scalability and call blocking tolerance, presents various challenges concerning power consumption performance. This document outlines key factors influencing the negative impact of CLOS configurations on power efficiency.

Network Structure

Multi-Tier Architecture: The CLOS topology typically consists of multiple tiers (e.g., spine-leaf), resulting in numerous switching nodes. Each additional tier increases the number of switches and interconnections, leading to greater power consumption.

Increased Hop Count: Each hop in a CLOS network requires processing by transceivers and networking logic elements. As data traverses multiple tiers, the cumulative effect of increased hops can lead to significantly higher power usage across the network.

Power Consumption Factors

Transceiver Overhead: In a CLOS configuration, the vast number of switches necessitates the deployment of numerous transceivers for data transmission. As transceivers can account for up to 70% of total power consumption, the high density of these components in CLOS networks poses a substantial energy burden without an accompanying improvement in routing efficiency.

Redundant Connections: CLOS networks may employ multiple paths for redundancy and load balancing. While this enhances reliability, it also results in wasted power due to underutilized links and active components that do not actively contribute to data transmission at any given time.

Inefficiencies in Packet Processing: The distributed nature of CLOS networks can introduce latency due to the need for packets to be processed at multiple intermediate switches. Each of these processing points consume additional power, compounding the inefficiencies associated with increased hop counts.

Impact on Performance Metrics

Latency: The inherent design of CLOS networks can lead to longer traversal times for packets as they navigate through more switches. Increased latency affects overall network performance and can necessitate additional resources to mitigate delays.

Thermal Management: Higher power consumption in CLOS networks contributes to significant heat generation. This requires enhanced cooling solutions in data centers, further amplifying energy costs associated with operating these infrastructures.

Conclusion

While the CLOS network topology offers valuable benefits such as scalability and call blocking tolerance, the design intricacies and associated power consumption challenges cannot be overlooked. Increased hop counts, transceiver density, redundant connections, and inefficient packet processing collectively undermine the potential advantages of CLOS architectures, leading to higher operational costs and lower overall energy efficiency in data center environments. For organizations aiming to optimize power consumption performance, considering alternative network topologies or improving design strategies may prove beneficial.

UNITEDSA-1776 PRODUCT OVERVIEW

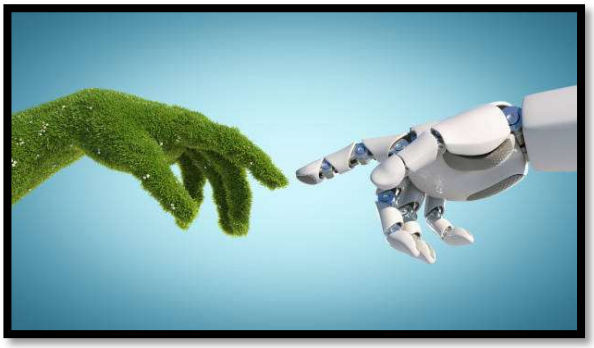

Introduction:

As data centers worldwide consume power equivalent to the output of dozens of nuclear reactors and growing, the need for efficient power solutions has never been more critical. With investments in internal power infrastructure often exceeding those for servers and switches, optimizing energy consumption is essential for the industry. Current estimates indicate that cloud networking switches account for approximately 20-40% of the overall electricity load in data centers, a figure that continues to rise due to increasing network virtualization and complexity.

Key Features:

Modular Design: The UnitedSA-1776 is a flexible, modular switch that accommodates future growth and customization.

High-Performance Line Cards: Includes sixteen line card slots that can be configured with a variety of line card options:

10G, 25G, 40G, 100G, or 400G: Tailor your switch configuration to match your data center's specific bandwidth requirements.

Power Efficiency:

Designed specifically to address the power consumption challenges in commercial data centers and AI centers, the UnitedSA-1776 incorporates advanced technologies that help reduce overall energy usage while maintaining high performance.

Summary:

The Brightways UnitedSA-1776 Cloud Hyperscale Switch is an innovative solution that meets the evolving needs of the data center industry, ensuring that power consumption is addressed without compromising on performance.

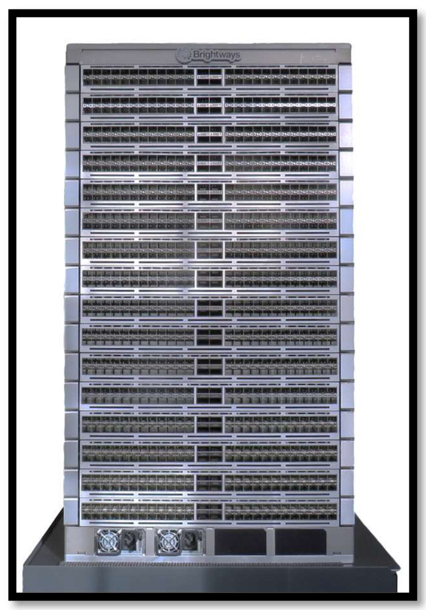

LINECARDS

The UnitedSA-1776 comes in four different Linecards – 1771, 1772, 1787, and 1788. The 1771 comes with a mix and match of 48 SFP+ and 4 QSFP+ ports and the 1772 comes with 36 QSFP+ ports. The SFP+ ports can operate as either 10GbE or 50GbE and The QSFP+ ports can operate as either 10, 25, 50, or 100GbE. The 1787 and 1788 offer 36 QSFP-DD ports and 32 OSFP ports respectively, and these can satisfy 10 – 400 GbE connections.

1771 LINECARD

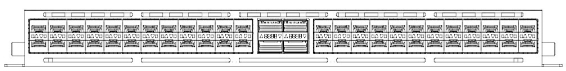

- 48 SFP+ ports, 4 QSFP ports
- 10, 25, 50 GbE configurations
- 64x 10, 25, 50 GbE

1772 LINECARD

- 36 QSFP+ ports
- 10, 25, 50, 100 GbE configurations
- 132x 10GbE, 96x 25GbE, 45x 100GbE

1787 LINECARD

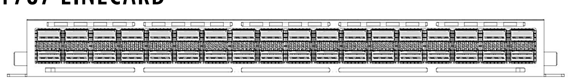

- 36 QSFP-DD ports
- 10, 25, 50, 100, 400 GbE configurations
- 192x 10GbE, 132x 25GbE, 45x 100GbE, 18x 400GbE

1788 LINE CARD

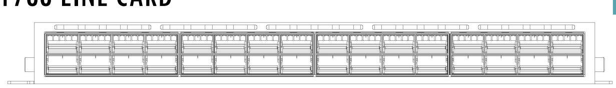

- 32 OSFP ports
- 10, 25, 50, 100, 400 GbE configurations
- 192x 10GbE, 132x 25GbE, 45x 100GbE, 18x 400GbE

PSPINE TM

The patented groundbreaking technology behind Brightways Pspine represents a significant leap forward in networking capabilities. Here are some of its key virtues:

Innovative Architecture: The two-hop, mesh-like structure fundamentally changes how networks are designed, allowing for more direct and efficient connections between devices. This innovation simplifies network management while enhancing performance.

Power Efficiency: One of the standout features of Pspine is its ability to drastically reduce energy consumption. Traditional Clos networks demand multiple hops, leading to higher power usage. In contrast, Brightways' technology minimizes these hops, resulting in lower operational costs and a smaller carbon footprint.

Reduced Latency: With a streamlined connection process, the Pspine technology boasts significantly lower latency. This means faster data transmission and improved user experiences, especially crucial for applications requiring real-time data processing.

Scalable Bandwidth: With a remarkable total bandwidth of 115 Tbps, the Pspine is designed to handle the demands of modern networking and Cloud LANs. It can efficiently support high data traffic volumes, making it a future-proof solution for growing businesses.

Dynamic Performance Optimization: The ability to shift from a mesh configuration to a two-layer network under heavy congestion showcases the system's intelligence. This adaptability ensures optimal performance during peak usage times, maintaining service quality even under stress.

Forward-Thinking Design: Brightways Pspine is not just a solution for today's challenges; it's a forward-looking technology that anticipates the future needs of networking, ensuring longevity and relevance in a rapidly evolving landscape.

In conclusion, Brightways Pspine's patented technology leads the charge in revolutionizing network design, marrying efficiency with cutting-edge performance and setting a new standard for the future of networking.

Here is a classic Leaf-Spine (Clos) Network below. Notice the two tiers of ASIC chips.

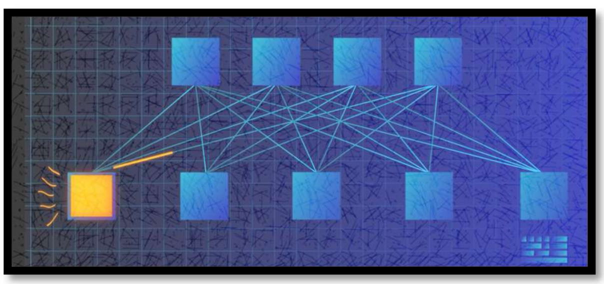

And,

Here is Brightways Pspine. Notice the one tier network meaning an entire layer of power consuming ASICs is eliminated and there is only one hop from one chip to another to go from input to output..

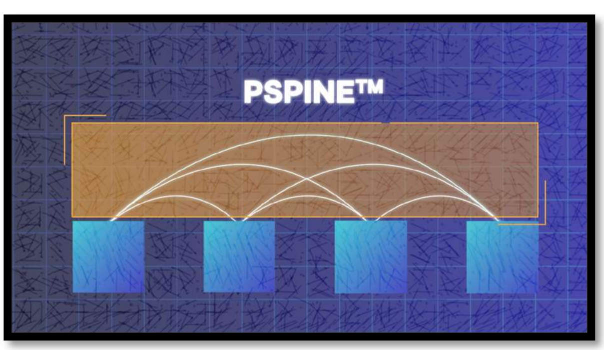

INTERWATT™ NETWORKING FPGA

The Brightways patented InterWatt FPGA networking chip represents a remarkable advancement in networking technology. Its world-class design and capabilities offer several key virtues that make it stand out in the industry:

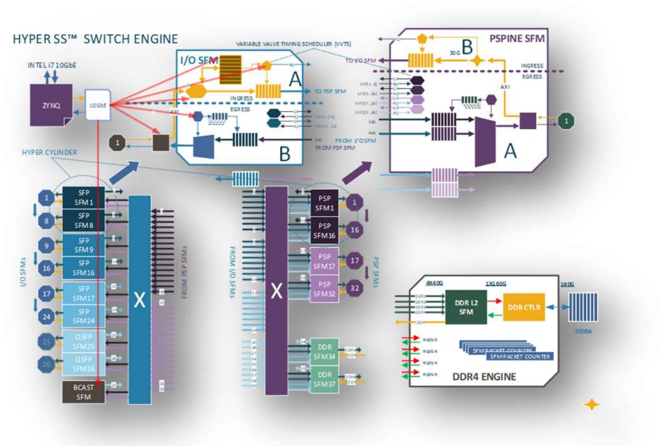

High Customizability: The InterWatt FPGA allows for extensive customization, enabling users to tailor the chip's functionalities specifically to their unique networking requirements. This adaptability ensures that it can evolve alongside changing demands and technologies.

Enhanced Processing Power: With its robust architecture, the InterWatt FPGA delivers exceptional processing capabilities. This enables faster data handling and improved throughput, allowing networks to efficiently manage large volumes of information without bottlenecks.

Low Latency Operations: The chip is designed to minimize latency, providing rapid data transmission that is critical for applications requiring immediate responses, such as real-time analytics and high-frequency trading.

Energy Efficiency: InterWatt FPGA is engineered to optimize power consumption while maintaining performance. This feature not only lowers operational costs but also contributes to sustainability by reducing the carbon footprint of networking operations.

Advanced Security Features: Security has been a priority in the development of the InterWatt FPGA. It integrates advanced protocols and encryption methods, safeguarding data transmission against potential threats and vulnerabilities.

Scalability: The design of the InterWatt FPGA accommodates growing data traffic needs. Users can scale their operations seamlessly without the need for complete system overhauls, making it a future-proof investment.

Seamless Integration: The chip's architecture allows it to easily integrate with existing network setups. This minimizes downtime during upgrades and enhances overall system performance without significant disruptions.

Real-time Data Processing: The capabilities of the InterWatt FPGA support real-time data processing, which is crucial for applications that rely on immediate data analysis, such as IoT devices and smart grid systems.

Versatile Applications: From telecommunications to data centers and beyond, the InterWatt FPGA is versatile enough to serve a wide range of industries, adapting to various use cases and performance requirements.

In summary, the Brightways InterWatt FPGA networking chip is a revolutionary component that combines versatility, efficiency, and performance. Its innovative design is set to redefine networking capabilities, placing it at the forefront of modern technology solutions.

Switch Flow Module (SFM)

The fundamental building block of the InterWatt FPGA Switch Flow Module (SFM) is designed to work in synergy with the Pspine architecture, addressing key challenges such as power consumption, reprogrammability, and future expansion.

Power Consumption: The SFM is engineered to optimize power usage by implementing efficient data handling and routing mechanisms. By minimizing the number of active components and hops that data must traverse, the SFM reduces overall energy expenditure. This approach allows data centers to maintain high performance without the excessive power demands typical of traditional architectures.

Reprogrammability: One of the standout features of the InterWatt FPGA SFM is its reprogrammability. This allows users to customize the functionality of the switch according to specific needs or changing requirements. As technologies evolve or new protocols emerge, the SFM can be updated or reconfigured without the need for physical hardware changes. This adaptability ensures that the networking infrastructure can remain relevant and efficient over time.

Future Expansion: The SFM is designed with scalability in mind, allowing for easy integration of additional modules or components as needed. Whether a data center experiences growth in user demand or technological advancements necessitate upgrades, the InterWatt FPGA SFM can accommodate such changes seamlessly. This capability not only extends the lifespan of the switch infrastructure but also helps organizations to future-proof their networking investments.

By combining these elements—enhanced power efficiency, the ability to be reprogrammed for diverse applications, and scalable expansion options—the InterWatt FPGA Switch Flow Module, in conjunction with the Pspine architecture, presents a robust solution for modern data center networking challenges.

Hyper Cylinder

Overview:

The Hyper Cylinder is a state-of-the-art architectural solution engineered to mitigate congestion in data flow systems. This innovative design strategically organizes Input/Output Switch Fabric Modules (I/O SFMs) and Packet Switch Fabric Modules (PSP SFMs) into smaller functional clusters, facilitating optimized competition for inter-module connections.

Key Features:

Modular Grouping:

- Reduces contention by grouping SFMs into smaller sets, enabling intra-group connection competition.

Crossbar Connectivity:

- Employs a crossbar switching mechanism, facilitating direct pathways between I/O and PSP SFMs.

Reduced Competition:

- Minimizes the competitive load on Pspine SFM links from 32 SFMs down to 8 SFMs, thus enhancing throughput.

Collision Mitigation:

- Decreases the probability of data collisions considerably, promoting higher reliability in connection establishment.

Improved Connection Success Rates:

- Enhances the likelihood of first-attempt successful connections, critical for achieving optimal system performance metrics.

Enhanced Switch Engine Performance:

- Facilitates seamless operation of the Interwatt switch engine, diminishing misfire occurrences and optimizing latency performance.

Performance Benefits:

Lower Latency:

- Efficient routing and reduced collision rates lead to improved response times in data transfer.

Scalability:

- The modular architecture scales effectively, accommodating future expansion with minimal congestion risks.

Robust Efficiency:

- Maximizes pipeline usage, ensuring high-performance throughput under varying load conditions.

Applications:

The Hyper Cylinder architecture is ideally suited for high-performance networking environments, including data centers, cloud computing platforms, and real-time data processing applications, where minimizing latency and maximizing throughput are paramount.

Transverse Virtual Output Queueing – TVOQ

Market Switches Virtual Output Queuing (VOQ) Overview

Architecture Design:

Traditional VOQ mechanisms are engineered to mitigate congestion by implementing individual ingress buffers dedicated to each engress port across the crossbar architecture.

Key Features:

Head-of-Line Blocking Elimination: By having separate buffers, traditional VOQ significantly reduces the bottlenecks usually encountered in a single-queue structure.

Drawbacks:

Increased Complexity: The duplication of resources for each egress increases the complexity of the system design and management.

Higher Power Consumption: The need for multiple buffering resources leads to elevated power requirements, impacting the overall efficiency of the system.

Brightways TVOQ Implementation

Streamlined Approach:

Brightways introduces a novel TVOQ architecture utilizing a single FIFO (First In, First Out) queue mechanism, designed to enhance performance while simplifying the architecture.

Advantages:

Simplification of Architecture: The use of a single queue reduces the complexity associated with resource management and buffer allocation.

Improved Efficiency: Lower power consumption is achieved due to minimized resource duplication.

Interwatt Chip VOQ Enhancements

Congestion Look-Ahead System:

The Interwatt Chip VOQ integrates an advanced congestion look-ahead algorithm that proactively assesses packets at the front of the queue.

Operational Benefits:

Anticipation of Blockages: By analyzing the state of packets, the system effectively forecasts potential congestions, enabling preemptive actions.

Packet Selection: This innovative approach allows the selection of packets that are likely to proceed without incurring a busy signal on the initial attempt, optimizing throughput and reducing latency all while keeping the proper order of packets.

Conclusion

The evolution from traditional VOQ to more advanced implementations like Brightways TVOQ and Interwatt Chip VOQ highlights a significant shift towards efficiency and simplicity in network architecture, addressing the challenges posed by congestion while minimizing resource usage.

Variable Valve Timing Scheduler – VVTS

On-Demand Connection Technology

Overview:

The On-Demand Connection Technology revolutionizes how networks establish connections, moving away from traditional scheduling systems that often create bottlenecks and inefficiencies. This innovative approach allows for immediate, real-time connections based on current network demand, ensuring optimal resource utilization and enhanced performance.

Key Features:

Dynamic Connection Establishment: Connections are created instantaneously as needed, eliminating the delays associated with pre-defined scheduling and ensuring a more responsive network environment.

Resource Efficiency: By reducing the need for complex scheduling algorithms and extensive packet reassembly, this technology minimizes resource consumption and cuts operational costs.

Real-Time Response: The network autonomously adjusts to incoming traffic, allowing multiple connections to be established simultaneously, thus enhancing throughput and reducing congestion.

Benefits:

Improved Network Performance: Experience reduced latency and increased bandwidth availability as connections are made on-the-fly, responding directly to network traffic needs.

Lower Power Consumption: Reduced reliance on complex scheduling mechanisms leads to significant power savings, contributing to a more environmentally friendly networking solution.

Simplicity and Reliability: Offers a straightforward approach to connection management, resulting in a more reliable network architecture that can quickly adapt to changing conditions.

High-Availability Buffering System – HABS

A switch must seamlessly navigate situations where multiple input ports are directing incoming packets to a single output port. This scenario can quickly overwhelm the output port, leading to a backlog of traffic as the combined data volume exceeds its throughput capacity. If left unaddressed, this congestion can cause severe packet loss of the switch, effectively freezing all internal pathways.

The HABS expertly alleviates this problem by redirecting incoming packets to high-speed buffers meant to cushion high volume bursts of traffic, with a substantial 12 GB capacity per linecard. This HABS buffering can hold up to 7.5 million packets, which equates to about 3.2 seconds of buffering at 10 Gbps. As conditions normalize, the HABS efficiently releases these packets in a controlled manner, restoring the Interwatt system to optimal performance. This proactive approach not only enhances reliability but also ensures uninterrupted service.

On-Demand Packet Diversion Algorithm (OPDA)

Overview

Introducing the innovative On-Demand Packet Diversion Algorithm, a dynamic solution designed to optimize network performance by intelligently routing packets in various directions, akin to a Clos network. This cutting-edge technology effectively circumvents congestion by utilizing available bandwidth on-the-fly, ensuring seamless data flow and enhanced network efficiency.

Key Features

Dynamic Packet Routing: The algorithm detects congestion in real-time and reroutes packets instantly, allowing them to "scoot around" bottlenecks and maintain a steady flow of data.

Congestion-Aware Operation: Leveraging advanced monitoring techniques, the algorithm identifies paths with free bandwidth and diversifies packet delivery to optimize resource utilization.

Efficient Bandwidth Usage: This approach ensures that available bandwidth is harnessed effectively, minimizing latency and maximizing throughput during peak traffic times.

Seamless Transition: Once congestion is alleviated, the algorithm gracefully reverts to its normal power conservation mode, minimizing energy consumption and reducing operational costs.

Flexibility and Adaptability: Capable of adapting to varying network conditions, the algorithm ensures that your system remains agile and responsive, providing a robust solution under fluctuating loads.

Benefits

Enhanced Network Performance: By actively managing congestion, the algorithm helps maintain high data rates and consistently optimal network performance.

Reduced Latency: The proactive routing of packets around congested pathways allows for faster data delivery, leading to improved user experiences.

Cost-Effective: With a focus on conserving power and resources, this algorithm not only boosts performance but also helps reduce operational expenses.

Scalability: Ideal for modern data centers and high-capacity networks, the On-Demand Packet Diversion Algorithm scales seamlessly with growing bandwidth requirements.

Conclusion

The On-Demand Packet Diversion Algorithm represents a significant advancement in network technology, providing an intelligent, responsive solution to congestion that enhances performance without compromising energy efficiency. Embrace the future of network management and ensure your systems operate at their peak potential.

LANBOLT NOS

Unlock the true potential of your network with the LanBolt Network Operating System (NOS), a powerful Linux-based platform driven by the cutting-edge Open vSwitch engine. Experience seamless integration of a full stack of Layer 2 (L2) and Layer 3 (L3) protocols, boasting exceptional VLAN support—all crafted on reliable, industry-standard software.

With LanBolt, you elevate your infrastructure to new heights, delivering unmatched performance through a blend of reliability and efficiency. This innovative solution features hot-pluggable design and dynamic orchestration, ensuring seamless failover from the main line card to a secondary linecard, so your network remains robust and operational even in critical situations. Embrace the future of networking with LanBolt!

The UnitedSA-1776 Cloud hyperscale switch represents a groundbreaking solution that can seamlessly integrate into a live network, automatically identifying routing information and mapping the LAN topology. Alternatively, it can be configured through an external OpenFlow agent like an SDN Controller. A primary and secondary line card is powered by an AI enabled multicore ARM microprocessor, effectively functioning as a single board computer that excels at managing essential tasks such as MAC Learning, MAC Propagation, MAC Aging, and telemetry. With Brightways, you are harnessing exceptional processing power to drive the control plane.

TELEMETRY “CONNECTIQ”

The SFM delivers a wealth of telemetry data essential for effective system monitoring and assessment, and the ConnectIQ GUI makes it available in real-time. This includes vital metrics like line card health, power state variables, Ethernet packet statistics, and crucial ingress and egress traffic statistics, as well as packet latency distribution. Interwatt methodically collects these important measurements, while the Control Plane compiles telemetry from each line card. This valuable information is readily available to network administrators.

Notably, FIFO depth statistics are key indicators for network performance, helping to pinpoint areas with heavy traffic and offering critical insights for proactively addressing potential network issues. By leveraging this data, network administrators can optimize performance and enhance overall network reliability.

ADVANCEMENTS AND CHALLENGES IN NETWORKING CHIP TECHNOLOGY

Prominent manufacturers of networking chips are advancing technology by integrating a greater number of transceivers onto single chips, achieving speeds of up to 400G. However, this high level of integration in cloud modular switches does not improve power efficiency or speed. While application-specific integrated circuits (FPGAs) can greatly enhance Top-of-Rack (ToR) switches, which can feature as many as 256 I/O ports on a single chip, modular switches—with port counts often exceeding 2,000—do not reap the same benefits.

In switch designs, it is common for the number of transceivers available for external connections to be reduced by half. This is because half of the transceivers are allocated to support external NIC (Network Interface Card) connections, while the other half are dedicated to accessing the internal switch fabric. This 2:1 reduction in the number of chip transceivers available for NIC connections is crucial, as the bandwidth of the switch fabric must closely match that of the I/O transceivers. This matching helps prevent congestion and maintains low latency.

Ultimately, increasing the number of transceivers on an FPGA does not lead to improved power efficiency or reduced latency for modular or hyperscale switches. The total number of transceivers and data hops remains constant, and transceivers are the primary consumers of power, accounting for about 70% of the total on-chip power. In summary, the number of transceivers significantly impacts power consumption, regardless of the semiconductor technology used. This highlights the need for a strategic approach in networking architecture when designing efficient networking solutions.

EMPOWER THE NETWORK AND UNLOCK EFFICIENCY

The Cloud Modular Switch can effectively replace the traditional Top-of-Rack (ToR) switch, like the UnitedSA-1776, by consolidating power-hungry Core Switches into a more efficient framework. Its higher integration leads to reduced power consumption and latency, making it a better choice for commercial data centers.

Similarly, the Brightways Pspine architecture can replace the standard Clos two-layer Leaf-Spine LAN due to significant reductions in power usage and latency. Since power infrastructure costs make up about 65% of data center expenses, solutions that address these factors will be more competitive.

The Pspine architecture supports up to 200,000 Network Interface Card (NIC) ports, similar to Leaf-Spine, but operates as a two-hop switch network instead of a four-hop network such as the Leaf_spine. The Pspine topology results in 35% less power usage, mathematically, and lower latency compared to Leaf-Spine.

In short, using the Brightways Cloud Switch enhances power efficiency and reduces latency while offering notable cost savings by switching from the Leaf-Spine structure to the Pspine topology. This change will improve operational speed and prepare the LAN for greater efficiency.

TECHNICAL SPECIFICATIONS

Line card	1771	1772
SFP+	48	-
QSFP+	4	36
10G Ports	64	126
25G Ports	64	96
50G Ports	TBD	TBD
100G Ports	-	45
Pkt Buffers	8GB	8GB
Max Power 10G	122W	239W
Max Power 25G	134W	201W
Max Power 50G	TBD	TBD
Max Power 100G	-	270W
Watts per port 10G	1.9	1.9
Watts per port 25G	2.1	2.1
Watts per port 50G	TBD	TBD
Watts per port 100G	-	6.0
MAC Addresses	437,760	437,760
Latency	1.2us	1.2us
Weight	16 Lbs.	16 Lbs.
Dims (in)	15.5 X 20 X 1.7	15.5 X 20 X 1.7

Notes:

1. Not including optics.
2. All Watts data are max power
3. Current pinnacle market switches require Fabric Modules practically doubling per port power in those switches
4. Brightways innovations eliminate the Fabric Module

Line card	1787
QSFP-DD	36
10G Ports	192
25G Ports	120
50G Ports	TBD
100G Ports	45
400G Ports	18
Pkt Buffers	8GB
Max Power 10G*	288W
Max Power 25G*	300W
Max Power 50G*	TBD
Max Power 100G*	TBD
Max Power 400G*	TBD
Watts per 10G port	1.5W
Watts per 25G port	2.5W
Watts per 50G port	TBD
Watts per 100G port	TBD
Watts per 400G port	TBD
MACs	294,912
Latency	1.2us
Weight	16 Lbs.
Dims (in)	15.5 X 20 X 1.7

Notes:

1. Same as above table

REFERENCE CONFIGURATIONS

Here are several sample configurations to expand the LAN beyond a single UnitedSA-1776 switch and support networks with over 100,000 connections. These configurations are ideal for L2 server access. For instance, in the configuration shown below, where two 1787 linecards each connect to a different switch, the total number of LAN connections (X), using the Pspine style of LAN, is calculated as follows:

192 (10G ports) x 14 (10G linecards) x 45 (100G ports) x 2 (100G linecards) = 241,920 connections.

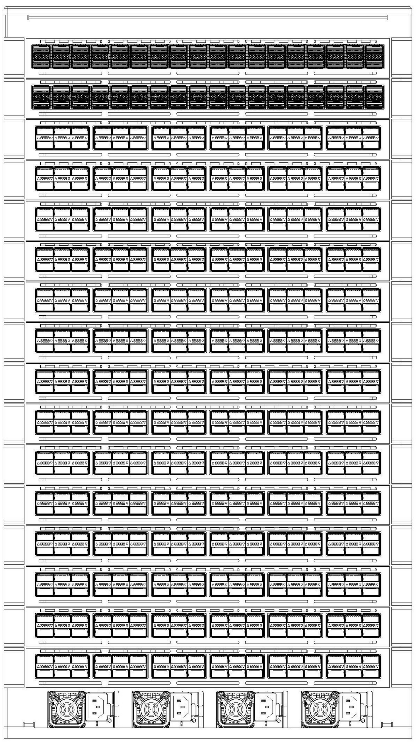

- Example described above
- 2x 1787 either 100, 400 GbE
- 14x 1772 either 10, 25, 50, 100 GbE

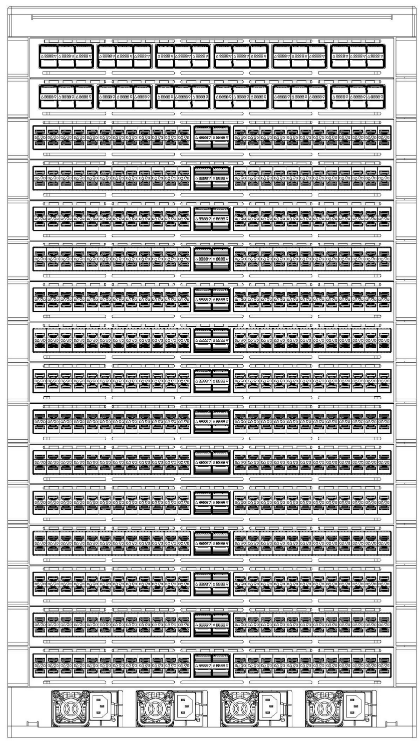

- 2x 1772 100 GbE
- 14x 1771 either 10, 25, 50 GbE

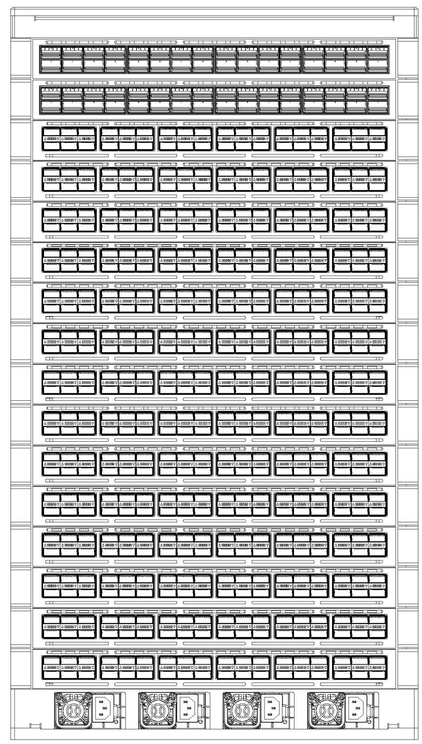

- 2x 1788 100, 400 GbE
- 14x 1772 either 10, 25, 50, 100 GbE

Chassis	
Height	18U (31.5”)
Depth	27.5”
Dimensions	17.5”x27.5”x31.5”
Weight of Chassis	150 lbs.
Weight with Line cards	425 lbs.
Power supply slots	4
Power supply bay height	1U
Power supplies	PET2000 Bel
Each Power supply Continuous	2000W
Power supply Efficiency	Platinum 94%
Total available power	6kW (N+1)
Fans	8
FRU*	Fan Door
Fan dimensions	120x120x38 mm

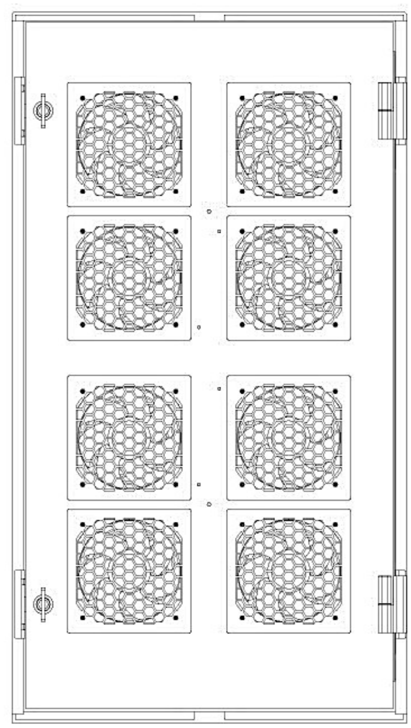

Notes

1. FRU - Field Replaceable Unit
2. Fan Door easily, quickly drops on mounting hinges and quick connect with two ribbon cables

LAYER 2 FEATURES

- Jumbo frames 9300 bytes (now)
- 802.1w Rapid Spanning Tree
- 802.1s Multiple Spanning Tree Protocol
- Rapid Per VLAN Spanning Tree (RPVST+)
- VLANs
- LACP
- MLAG
- Link Layer Discovery Protocol
- 802.3x Receive Flow Control
- IGMP

LAYER 3 FEATURES

- Routing Protocols: OSPF, OSPFv3, BGP, MP-BGP, IS-IS, and RIPv2
- Equal Cost Multipath Routing (ECMP)
- VRF
- Bi-Directional Forwarding Detection (BFD)
- Unicast Reverse Path Forwarding (uRPF)
- VRRP