



Pica8 White Paper

An Enterprise Approach to White Box Networking



Contents

Introduction	1
Powering Effective White Box Switches	1
Defining an “Open” Switch	2
The Myriad Benefits of Open, Disaggregated Networking	3
ELIMINATE VENDOR LOCK-IN	3
LOWER TCO	3
DELIMITED FLEXIBILITY AND HIGHER PERFORMANCE	3
IMPROVED RELIABILITY	3
EASE OF USE	3
THE PICA8 APPROACH: ENTERPRISE-READY	4
PICOS NETWORK OPERATING SYSTEM	4
PicaPilot Plus PICOS Makes White Box Networks Enterprise-ready	4
CrossFlow	5
Choice of Hardware Plus Superior Support	5
Realize the Full Benefits of Pica8	5

Introduction

White box switches running a modern, Linux-based network operating system (NOS) are disrupting the enterprise campus network market, providing a viable alternative to customers who have spent decades bending to the will of their entrenched switch and router equipment vendors.

Similar to the virtual server concept, where the server software is abstracted from underlying hardware, with white box switches the network operating system itself is abstracted – or disaggregated – from the underlying switch hardware. This creates the potential for an “open,” portable NOS that can run on a wide variety of switches from multiple vendors. No longer is an enterprise permanently wedded to a single vendor’s hardware and software development cycle.

The white box networking movement is now proving to be just as beneficial to enterprise networks as virtual servers were when they first came on the scene. It delivers dramatically simplified network deployment, operations and support, along with increased flexibility in terms of optimized hardware and software, all at lower costs and with simplified deployment and management.

But not all approaches to white box networking are the same. Pica8, for example, takes an approach that squarely addresses large enterprise requirements. Its PicaPilot™ technology removes almost all of the complexity currently associated with network configuration, operation and management by enabling hundreds or more switches to be managed as one logical switch using the company’s PICOS® Linux-based NOS. At the same time, PICOS’s Crossflow™ feature enables Layer 2, Layer 3, and software defined networking (SDN) control planes to run simultaneously and allow traffic to flow over the same switch port – all at the same time. This functionality paves the way for new levels of security and policy management. And the PICOS operating system is built on an unmodified Debian Linux kernel, making it easy for developers to work with using familiar tools.

Powering Effective White Box Switches

At the hardware level, a white box switch is really no different from the switches enterprises have been buying for years from major vendors such as Cisco and Juniper. A handful of well-established manufacturers – including Accton, Delta Networks, Foxconn and

ESCAPE FROM THE DATA CENTER:

PICA8 BRINGS SIMPLIFIED LEAF-SPINE ARCHITECTURE TO THE REST OF THE ENTERPRISE

Enterprise and data center networks traditionally used a three-tier model, with access switches closest to endusers feeding larger aggregation switch/routers which, in turn, connect to larger core routers that form the network backbone.

In this architecture, network architects must configure redundant pathways to ensure resiliency, typically using the Spanning Tree Protocol (STP). A potential drawback is that STP deactivates all but the primary network route. Should that route fail, it will bring up a backup path, which is then used until the primary comes back online. While that does provide redundancy, it can also lead to bandwidth constraints should the primary (or backup) network path become congested.

An alternative architecture known as leaf-spine addresses those limitations.

The traditional leaf-spine is a two-tier architecture, in which the leaf switches connect to end devices, such as servers and firewalls, and the spine switches connect to leaf switches. This approach was first used in data center networks, as it is a good fit for the “east-west” nature of data center traffic.

Consider a data center rack full of servers. At the top of the rack may be a pair of switches, known as Top of Rack (ToR) switches because of their physical location. Each server in the rack connects to the two ToR switches, for redundancy. These ToR switches are the equivalent of leaf switches in a leaf-spine topology.

Each leaf switch then connects to multiple spine switches; in a data center, they likely connect to every spine switch (see Figure 1). That means

Quanta Cloud Technology – supply their commodity hardware to both the white box market and the major switch vendors. Collectively, these manufacturers collect more than \$25 billion in annual switching revenue – a sure sign of a healthy, well-established market.

Typically, the hardware consists of standard 1U, 48-port Ethernet switches and can support speeds from 1G to 100G. They are based on a choice of ASICs, again from well-established vendors such as Broadcom, Cavium, Intel, Marvell and Mellanox, all with long histories in producing network chipsets. Alternatively, a proprietary switch could use a custom ASIC to implement some desired, non-standard feature. Service providers, for example, may choose to go the custom route, but enterprises typically would have no such specialized ASIC requirements.

Chip vendors also supply an application programming interface (API), which is used by the NOS to control interactions with the ASIC. These vendors may also include an SDK (Software Development Kit) to program the ASIC, such as to set up a VLAN or ACL entry. Broadcom, for example, recently released SDK Logical Table (SDKLT), which is an open source SDK that’s intended to enable a high level of programmability for its chipsets.

The final component of a white box switch is the NOS itself. Here’s where the white box switches differ significantly from those of legacy switch vendors. A vendor such as Cisco or Juniper takes the same commodity hardware and loads its own, proprietary NOS on top. The white box alternative is to instead load an open, Linux-based NOS. It’s a similar approach to using Linux on a commodity server instead of, say, Windows Server.

Defining an “Open” Switch

It’s the choice of the NOS that makes a switch open – or not.

Open networking software uses an open source Linux-based OS. The Open Compute Networking Project, for example, is an effort by the Open Compute Project (OCP) aimed at creating a set of disaggregated, open network technologies, including a Linux-based NOS and developer tools.

Similarly, the Open Network Install Environment (ONIE) is an open source initiative driven by a community of vendors to define an open “install environment” for white box switches. Also a project of the OCP, ONIE is intended to enable an ecosystem where end users can choose among different NOSs and install them on a common set of white box switches in the same manner that they provision servers. Open Network Linux, another OCP project, is one example of such an open-source NOS that uses the ONIE install environment to install into a white box switch’s flash memory.

there’s no need for spine switches to connect to one another. Rather, all ports on a spine switch are used to connect to leaf switches, using either Layer 2 switched connections or Layer 3 routed links. From a logical perspective, all switches are the same distance from one another – reachable in a single “hop.”

But the PicaPilot technology from Pica8 takes leaf-spine to another level for enterprise switch stack and chassis switch replacements. It extracts all the complexity from traditional two-tier leaf-spine networks and flattens the network by managing both the leaves and the spines as though they were all a single, logical switch with no hierarchy. Pica8 has paved the way for the data center-proven leaf-spine architecture to extend to the wider enterprise by turning it into the simplest and easiest to deploy network architecture of any kind – able to be managed by any SysAdmin (see figure 2). PICOS with PicaPilot maintains all the Layer 2/Layer 3 features of a traditional 3-tier design, including:

- Resilient multi-path fabric ensuring the applications and users get the highest availability
- A switching fabric capacity of 176 Gbps for 1 GbE and 1.28 Tbps for 10 GbE platforms
- Simple and flexible scale-out topologies:*
Up to 2,304 1 GbE access ports with 1.2:1 oversubscription ratio** Up to 1,536 10 GbE access ports with 2:1 oversubscription ratio*
- Support for SDN protocols and external programmability, helping you drive towards agile network orchestration

It’s a highly scalable and secure network that’s also easy to upgrade, because it’s based on cost-effective, plug-and-play infrastructure.

** Recommended configurations using Pica8 pre-loaded systems. Actual customer configurations may vary depending on platforms and oversubscription ratios.*

AT&T has also announced its own entry to the open networking community with its Disaggregated Networking Operating System (dNOS). The dNOS project provides a software framework to speed the adoption and use of white boxes in a service provider's infrastructure, AT&T says. Given that AT&T has more than 100,000 IP/MPLS routers in its network, that is a significant vote of confidence for the concept of disaggregated networking and open NOSs. With this project, AT&T is following the lead of web-scale companies, such as Facebook and Google, that have already standardized their internal networks on white box switches running a Linux NOS.

The Myriad Benefits of Open, Disaggregated Networking

The collective vote of confidence from AT&T and the web-scale giants makes sense, because open, white box networking has proven itself in large data centers as a new best practice. Now those same benefits are penetrating enterprise campus and branch office networks, freeing company after company from the restrictive, closed network technology cycle that they've been saddled with for decades.

ELIMINATE VENDOR LOCK-IN

The idea behind using an open NOS is that it can be ported from one white box hardware platform to another. That means there's no more vendor lock-in, in terms of hardware or software. As advances come along in software or hardware, enterprises are free to take advantage of them without the complexity involved in migrating from one proprietary switch vendor to another.

LOWER TCO

Pairing commodity hardware with open-standards-based disaggregated NOS software makes white box switches far less expensive than traditional legacy appliances in terms of capital costs. For example, the hardware/software list price for a fully configured white box version of a Cisco 6500-class chassis switch runs around 20 percent of the cost of the legacy Cisco switch – and the white box “replacement” provides significantly higher density and performance than lower price. Additionally, simplified management tools like ZTP and PicaPilot make ongoing operations far easier, significantly reducing ongoing operating costs. It all adds up to reduction in total cost of ownership that can easily exceed 50 percent.

DELIMITED FLEXIBILITY AND HIGHER PERFORMANCE

By using a switch-portable NOS, enterprises can select the switch that best matches each specific deployment in terms of redundancy, density, speeds, features, power profile, and port counts. Yet they will still have consistent and unified activation, programming, network management, monitoring, special features and behavior across their network infrastructure.

IMPROVED RELIABILITY

White box networks are also inherently more reliable than the legacy systems they replace, offering improved redundancy for a couple of reasons. One is the use of a higher-reliability chassis, with redundancy features built in. Such switches suddenly become more affordable given that enterprises can typically purchase two modern white box switches for automatic fail-over for the price of one (heavily discounted) switch from a legacy vendor.

The other factor is that the disaggregated network concept, when used with an appropriate NOS, enables the deployment of a leaf-spine network architecture throughout the enterprise, not just in the data center. The leaf-spine architecture provides improved redundancy and availability, because every leaf switch has a direct connection to every spine switch (see sidebar).

EASE OF USE

An open NOS based on Linux borrows technology from the server realm that promotes ease of use, including zero touch provisioning (ZTP) and automated licensing. Once a switch is physically connected to the network,

ZTP enables the automation of provisioning and configuration processes, typically using a Dynamic Host Control Protocol (DHCP) server.

ZTP routines can also take advantage of open source Linux tools such as Puppet and Chef, which began life as tools to automate server configuration tasks. These tools have now been adapted to provision switch configurations by the Open Source community. So, just as racks of servers and VMs are added to a cluster using Puppet or Chef, network switches and routers can be configured in the cluster by the same tools.

THE PICA8 APPROACH: ENTERPRISE-READY

While the white box network concept initially found favor in data centers, where its leaf/spine architecture is ideal for handling east-west traffic, Pica8 has unique technology that breaks the leaf-spine concept out of the data center and lets it flow across the entire enterprise.

PICOS NETWORK OPERATING SYSTEM

It starts with the PICOS NOS, which is built on a stable, unmodified Debian Linux stack. The fact that it's unmodified is important, because it means developers know exactly what they're dealing with and can use all their traditional, familiar development tools and orchestration platforms, including Chef, Puppet and Salt Open. Regular Debian updates are also simple to install, using one of the many available automation routines.

PICOS runs standards-based Layer-2 and Layer-3 protocols while also providing full support for the broad palette of key enterprise features that are not found in data center white box solutions. This enterprise-specific feature set includes wide ranging capabilities, such as support for voice VLANs, legacy Cisco phones, dumb VoIP phones, PVST, OSPF, VRF, NAC, secure remote access, QoS, 1G/2.5G, and more. All of this is controlled via a network Command Line Interface (CLI), which will be familiar to enterprise network administrators and is their preferred method for controlling enterprise network infrastructure.

PICOS also supports ZTP and is qualified to run on proven, high-performance hardware. It's ideal for converged server/storage/network installations because it provides an automation framework combining scripting with Linux, orchestration and workflow tools.

Two PICOS-only features – PicaPilot and Crossflow – combine to make Pica8's NOS the only white box solution that can address all of today's enterprise network requirements as well as any future network services due to its granular flexibility and ease of deployment.

PicaPilot Plus PICOS Makes White Box Networks Enterprise-ready

PicaPilot enables network managers to control dozens or hundreds of multivendor white box switches as if they were a single, logical switch with a single, consolidated IP address. This new Pica8-only capability has multiple ramifications for enterprise networks.

For one, there's no need to invest in large, expensive switch chassis. Rather, you can have multiple switches located either in the same rack or dispersed throughout a campus environment, all being managed as a single switch.

The ability to manage many switches as one also slashes operational overhead. All configuration, policy and security changes are now far easier to implement, with a single update effectively applying to hundreds of switches. It reduces the number of individual network elements you need to manage by a factor of 10 to 50 times or more. Management headaches associated with all these individual network elements melt away with PicaPilot.

PicaPilot is also what enables the leaf-spine configuration to be deployed across the enterprise, not just in the data center. Leaf-spine is not only simpler than the traditional three-tier network architecture (access, aggregation, core), it provides better performance and resiliency (see sidebar).

With PicaPilot, PICOS-powered Ethernet switches use Multi-Chassis Link Aggregation (MLAG) technology to connect devices, enabling each one to connect to a pair of Pica8 switches with all links running active/active to improve resiliency. There's no need to block certain links, as with the spanning tree protocol (STP), resulting in improved bandwidth utilization and performance. MLAG peer switches synchronize forwarding state between them, so if a leaf or spine switch fails, traffic is automatically rerouted for continuous uptime.

CrossFlow

Another Pica8 innovation, CrossFlow, likewise breaks new networking ground.

Previously, the state of the art in IP switching was hybrid Layer 2/Layer 3 switches, where some ports operated with the traditional L2/L3 stack software while others were under the control of a software defined network (SDN) controller. But each port supported just one or the other – L2/L3 or SDN, not both.

CrossFlow enables each switch port to support L2/L3 traffic as well as SDN traffic at the same time. This enables security updates, for example, to be sent to a switch without interrupting normal traffic. That can have important implications for large enterprise networks, by providing a dynamic way to insert policy into the network.

Say a large retailer notices suspicious activity at one of its stores coming from a certain set of MAC addresses, threatening to take down network service to the location. Using CrossFlow, its security operations team can push out a security policy change, instructing the switches at the location to either drop, mirror, or redirect traffic from the offending addresses to a security analytics platform – with no interruption to service anywhere in the network. Similarly, the capability can be used to conduct deep security monitoring on switch ports, again without interrupting traffic.

Choice of Hardware Plus Superior Support

All of these advanced capabilities work on your choice of hardware, from a simple 1G switch to a fully loaded 100G model, from a variety of vendors (see Pica8's complete hardware compatibility list: https://www.pica8.com/resources/?resourcelib_category=hardware-compatibility).

You can also take advantage of the type of concierge support from Pica8 that the legacy vendors used to provide 30 years ago, minus today's egregious legacy support contracts. Bug fixes come out quickly and enhancements are routine. If you're used to the cold shoulder and long wait times to address issues from the big-name switch makers, you'll find Pica8 to be a welcome change.

Realize the Full Benefits of Pica8

White box networking provides a welcome change to companies large and small that are looking for an alternative to the switching status quo: vendor lock-in with expensive legacy network providers.

Pica8 is the first, and only, white box networking solution specifically tailored for use in large, dispersed enterprise networks. Its PicaPilot technology turns even the most complex network into one that is both easy to deploy and manage, effectively reducing the number of networked devices by as much as 50 times.



Built on an unmodified Debian Linux kernel, the PICOS NOS gives enterprise DevOps teams total control over their network destiny. Innovative features such as CrossFlow make networking more flexible while also providing real-time security.

In addition to eliminating vendor lock-in, Pica8 enables enterprises to lower their network TCO by 50% or more thanks to drastically decreased Opex and Capex costs – even when compared to heavily discounted legacy solutions.

To learn more about white box networking and Pica8's groundbreaking enterprise software solutions, visit: www.pica8.com.