2025

All WhiteBoxes are NOT the SAME

Mark Harris

Edgecore Networks

12/21/2025

All WhiteBoxes are NOT the SAME

Experience & Technical Prowess Outperforms Low-Cost Offerings

Executive Summary

The networking industry increasingly faces competition from low-cost switching products that appear, on the surface, to match name-brand devices by incorporating the same commercial switching silicon. However, the switching chip is only one piece of a much larger engineering equation, which typically takes decades for any vendor to master.

Product reliability, electrical performance, thermal stability, component quality, and manufacturing maturity ultimately determine total cost of ownership (TCO) and long-term operational performance. Low-cost manufacturers frequently lack the engineering depth and production discipline required to ensure device longevity, component resilience, and consistent build standards.

This whitepaper explores the fit-and-function engineering details that separate high-quality, professionally engineered switches from cheaper alternatives, even when both use the same silicon. It also examines how supply chain, large-scale production capacity and global manufacturing maturity (including automation) provide name-brand vendors like Accton Technologies with performance consistency, product availability, and proven supply stability that smaller suppliers simply cannot match.

When purchasing business critical infrastructure, there are many factors to consider beyond technical specifications and list price. Long-term demonstrable value is key.

1.0 Why “Same Chip” Does Not Mean “Same Switch”

While low-cost switches may use the same switching ASIC as a premium enterprise device, the surrounding engineering accounts for most of the real-world performance and reliability. Simply put, the means to transform a raw switching chip (made by Broadcom or one of the other large ASIC suppliers) into an Enterprise-class switching solution, has far more to do with a wide range of engineering and manufacturing processes than it has to do with the schematics involved. For more than 25 years, Accton has proven that customer satisfaction, product reliability, support and even total cost of ownership are greatly affected by our leadership build quality rather than the technical performance of the embedded switching chip itself.

Critical engineering and manufacturing domains worth discussing with your short-list vendors include:

- Printed circuit board (PCB) layout
- Power design, delivery and electrical protection
- Mechanical build quality and tolerances
- Component selection and supply chain
- Thermal dynamics and management
- Reliability testing and verification
- Manufacturing consistency and sourcing discipline

It is essential to recognize that the capabilities of a high-performance chip alone cannot offset shortcomings in supporting engineering design and manufacturing processes. The overall system performance is determined by the quality of the switch's design, construction, protection, assembly, testing, and subsequent support.

2.0 Typical Cost-Cutting found in Low-Cost Designs

Manufacturers who focus on low-cost products often cut down on the essential development simulation and testing, quality of many of the ancillary components used across the bill-of-materials, as well as the mechanical components such as less robust connectors and cabling assemblies, thinner gauge and less sturdy metal frames, and even fewer screws holding everything in place. They may also exclude important quality components like RFI or EMF filters or reduce engineering work by making PCBs with fewer layers and thinner copper. These cost-saving measures aren't always listed in product datasheets but often become evident during extended use and can affect long-term customer satisfaction and system availability.

2.1 Power Supplies: The Leading Cause of Field Failure

A network switch's power system is one of the most critical—and most failure-prone—subsystems. While the exact failure rate depends heavily on the specific design, component quality, and operating environment, industry data consistently points to power supplies as a weak link. Reputable brands such as Accton who have delivered millions of devices over the years simply have the depth of experience to know how to specify and when to include better and higher quality power handling components to mitigate these risks and support the remainder of the system.

To reduce delivered device cost, low-end (and low-cost) vendors often:

- Specify generic power supplies that deliver the correct voltage and current as required by the design, but do so with a much lower efficiency. Low-cost power supplies may have a rated efficiency of 60% or less, while their higher-quality alternatives may deliver that same required power but with efficiencies at 90% or higher. The long-term cost savings between 60% and 90% efficiency is staggering.
- Use low-quality capacitors which may cost 75% less than the quality variants but exhibit pre-mature failures and emit electronic noise or other odd behaviors.
- Eliminate surge and transient protection which may cause unintended emissions which may be overlooked and affect other operations
- Reduce component over-spec margins, narrowing the environmental tolerances
- Purchase power modules from unknown suppliers with uncertain results which may be far different than the components which were specified in the design
- Use low-cost and poorly designed switching supplies without robust filtering nor mechanical stability, introducing unpredictable operation and early failures

The resulting scenarios include:

- Voltage fluctuation and instability as conditions change
- Premature component aging and degraded performance over time.
- Unexpected shutdowns under load as tolerances are exceeded
- Vulnerability to brownouts and line noise

Premium vendors invest heavily in power delivery reliability, including derated components and long-life capacitors rated for continuous operation across a wide range of thermal conditions. This is one of the key reasons premium network switches operate reliably for many years (see appendix) while bargain devices often fail early in production- sometimes catastrophically.

2.2 Thermal Design and Heat Dissipation

Today's network switches run far faster and hotter than anything previously seen. And while thermal design was not as critical when network speeds were lower and x86 technology ruled, in today's GPU-powered AI data centers, failure of even a single 100+ Tbps network switch can idle millions of dollars in GPU investments and affect the balance sheet.

Effective thermal engineering requires:

- Proper airflow modeling as devices are more compact in size and densely packed. Internal obstructions to air flow are significant.
- Heat sink sizing and material choice based on load and lifespan. And while heat sinks may be made out of steel or copper, heat transfer and costs vary widely
- Correct component spacing and placement based on throughput, emissions and signal timing with a keen awareness of airflow requirements at the same time
- Multi-layer PCB designs with thermal planes are essential but costly. Knowing how many layers and copper thickness can have dramatic impact on overall performance
- Sensors and fan control systems provide the subsystem to move heat effectively, but the fans themselves consume non-trivial power, so airflow, power consumption and cost must be balanced. Lower cost fans may use sleeve bearings rather than ball-bearings (which reduce friction) which often results in lower lifespans of the entire switch, and during operation require more power to move air due to their inefficient and higher friction designs.
- Simulated hot-zone analysis must be undertaken to spot thermal placement and design issues and correct prior to general production.

Low-cost products may:

- Use lower efficiency or lower-airflow fans, increasing the heat retained within the device, and therefore reduce component lifespans. Lower capacity and quality (e.g. sleeve versus bearing) fans cost less but reduce overall switch reliability and decrease its projected lifespan dramatically.
- Overcrowd components, a common mistake as all components have operating environments specified by their manufacturer. Crowding components reduce air flow and will also impact device lifespan.
- Skip thermal simulation and miss the assurance that their designs are properly cooled across the entire unit. This simulation is costly but helps spot thermal problems before general production and field-based product quality issues and reports.
- Choose uncoated PCBs or cheaper laminates. As many newer designs for high-speed switches include 32 or more layers, Accton understands the value of using the highest quality PCB resins and the appropriate copper layer thickness.
- Poorly layout internals making them too close to the margins which causes failure and reliability problems when environmental conditions vary.

Make no mistake, heat is the enemy of electronics. Poor thermal design shortens product life, increases component drift, and leads to unpredictable performance degradation- and ultimately device and infrastructure failures.

2.3 Mechanical Construction and Structural Durability

A network switch is a physical device which must be mounted in and connected to other physical components. This installation and de-installation requires a level of mechanical robustness to withstand these changes repeatedly over time. In specific, the switch must withstand:

- Continuous 24×7 operation and the thermal expansion cycles involved in maintenance operations and remediation processes.
- Rack vibration caused by colling fans and other components in the surrounding areas.
- Shipping stress, including the associated dead-drop cycles, crush and temperature/humidity cycling.
- Cable strain and connector stability. Common connector insertions may occur more than once, so all external connectors must be able to withstand those stresses.

Premium vendors typically use:

- Rigid chassis construction with thicker gauge sheet metal, reinforced corners, stainless steel mounting threads (e.g. PEMnuts)
- Reinforced PCB mounting and spacer assemblies to ensure the PCB is held in place but not stresses during movement
- High-quality connectors with locking mechanisms and plated surfaces.
- Chassis braces or stiffeners to reinforce all metal chassis surfaces

Low-cost switches may instead use:

- Thinner gauge stamped metal with tapped screw threads rather than thread inserts
- Non-locking or underspecified connectors. In some cases, these connectors do not bite into wire insulation ignoring the mechanical reinforcement available
- Mechanical assemblies without vibration tolerance, including carefully modelled and designed subassembly mounting systems, easily identified where liquid silicone or heated thermal plastics are used
- Weak front-panel port assembly retention, omitted brackets for securing port to front-panel plates
- Lower-cost fans which use sleeve or nylon bearings rather than ball bearings

This can result in:

- Ports failing under cable weight or repeated insertions/removals.
- Solder joint or PCB traces cracking and failing from stress due to repeated mechanical stress and vibration under stress
- Audible rattles or vibration, poorly assembled metal parts which ill-fit together or have tolerances which are too small for applications
- Louder fans with shorter usable life as they warp and wear
- Chassis warping over time, causing mounting and stability problems

While often invisible at purchase time, these problems present themselves over time and directly impact long-term field reliability.

2.4 PCB Quality and High-Speed Signal Integrity

When networking throughput was smaller (up to 100Gbps), PCB layout and optimization were less critical. Component placement was fairly arbitrary, and signals would move reliably. Once the 100Gbps threshold per SERDES was exceeded as seen with AI Fabric deployments, the physics associated with signal transfer became a major issue. It is no longer acceptable to design PCB for today's 100Gbps and above using designs with fewer layers, less copper and older component placement strategies. Today's AI-driven High-speed Ethernet design requires precise PCB engineering, including:

- Controlled impedance simulation and matching
- Shielded routing of all traves as every trace can essentially become an antenna
- Optimized trace lengths to propagate signal more quickly, given the speed of light nature of electricity.
- Dedicated ground and reference planes to reduce interference and crosstalk, extraneous signal and reduced usable throughput.
- Multi-layer board stack-ups are essential, with Accton's recent switches utilizing 20 layers or more, and multi-board header sandwiching techniques to reduce or eliminate flexible interconnect cables altogether.

Low-cost manufacturers may:

- Copy reference designs, modifying as needed, without simulation or total system design review
- Use thinner or fewer PCB layers to save costs, increasing signal contamination
- Reduce copper weight to reduce costs, missing much of the opportunity to transfer signals and power more reliably
- Compromise on shielding or ground planes, causing cross talk and signal integrity issues
- Skip extensive system testing, where all ports are utilized concurrently and throughput and error rates are studied

This can lead to:

- Higher packet loss at line rate, retransmission, congestion
- Link instability, bit error rates
- Increased EMI emissions, interference within the chassis and externally within the data center itself
- Shortened device lifespan due to dielectric breakdown, mechanical connections failing, header reliability issues, component failures.

Premium systems ensure electrical data signal integrity through extensive modeling, simulation, and intensive qualification— incurring costs that lower-end manufacturers typically avoid.

3.0 Manufacturing Scale and Production Maturity

Beyond engineering and component selection, manufacturing scale and quality discipline dramatically influence long-term product performance and consistency. It should be no surprise that the quality of the manufacturing of a complex system like today's high-performance switches demands designs that can be manufactured reliably. To It's not enough to have a hand-assembled prototype, but the proof is when multiple factories can build hundreds of units each week, each one identical to the next. This is the approach Accton pioneered decades ago to assure customers received solutions that perform for years to come.

3.1 At Scale, Manufacturing is Repeatable & Controlled

Name-brand network manufacturers like Accton typically have spent years refining manufacturing processes, deploying robotics and even automation to achieve the required long term delivered product quality. Key business attributes which drive the manufacturing choices:

- The ability to operate multiple high-volume factories spread across multiple geographies
- Maintain a strong global supply chain and the sourcing agreements which minimize risks to production capacity
- Implement standardized production processes which can be deployed in multiple locations with identical results
- Use statistical process controls (SPC) to assure delivered product quality
- Track serial-level quality and return data, to create an environment focused on continuous improvement
- Perform multi-stage automated testing and burn-in, spotting and correcting problems before the customer is impacted
- Maintain incoming material inspection programs to ensure suppliers are also delivering components consistently

With millions of units produced annually and one of the lowest returns rate, companies like Accton benefit from:

- Process learning curves which reduce manufacturing costs while increasing product quality and reliability
- Supplier relationship maturity with clear understanding about capacity and delivery
- Continuous improvement across all aspects of manufacturing, across all new and existing sites
- Predictable component procurement, based on product forecasting and run-rate business lines
- Tight lifecycle planning to assure components are available throughout their entire lifespan, at the quantity needed across all channels.

3.2 Small or Off-Brand Producers Face Challenges

Scale is a major factor for complex products which directly and negatively impacts build quality. Low-cost manufacturers often:

- Produce in small or inconsistent batches where each device or batch can be considered bespoke in nature.
- Change in suppliers based on price at the time of production. Referred to as the “spot” market, everything from sheet metal panels, power subassemblies and capacitors drives a high level of product inconsistency.
- Lack automated testing capabilities due to the high costs to implement such programs at scale.
- Have limited ability to trace failures back to process issues. These makers ignore root-cause failure analysis, which eliminates continuous improvement opportunities
- Operate without predictive component lifecycle planning and ignore component refinement planning to assure reliable componentry is always a core tenant.

This ill-conceived cost cutting leads to:

- Wide variability between production runs, changing expected performance and long-term product reliability
- Difficulty obtaining identical replacements since the behavior of any new unit may be different than the previous unit, even when model numbers are the same.
- Unpredictable spare part availability which may affect support service level agreements and downtime expectations and commitments.
- Fragile deployments, such as AI, are very demanding and operate at higher throughputs, temperature and power consumption than previous devices.
- Rapid design changes that make products incompatible with earlier units.

3.3 Volume Matters in Reliability

With millions of devices shipped, Accton has proven that large-scale manufacturing when designed from the ground up with quality in mind can easily improve product reliability because:

- Production detail and re-work / troubleshooting data is gathered across thousands of devices to identify root causes to enable resolutions to be made at the source. Formal processes are put in place to determine root cause and resolutions.
- Field return trends can be corrected quickly at the source
- Quality problems surface early and can be designed out in the manufacturing processes, tooling or testing processes.
- Designs reach maturity faster due to broad deployment

In contrast, small production volumes seen with low-cost manufacturers cannot generate comparable return analytics, leaving issues undiscovered until devices fail at customers' sites.

4.0 Support, Serviceability, and Risk

While precision engineering, scheduled supply chain and mature manufacturing discipline minimizes risk, it cannot eliminate the laws of physics—hardware failures are an inevitability. The true measure of a supplier is not just the raw performance and durability of their device by design, but the velocity and depth of their response when the unexpected occurs. This is especially true now that every business is a connected business and outages become catastrophic when their duration lengthens.

4.1 Mature Vendors Provide:

- A means to handle technical support issues globally with cost-effective and agreed upon ticket response times. And while every mature vendor has different support options, setting expectations and working with service partners is critical.
- Spare units delivered rapidly to replace in-service devices that fail. Many mature providers can work with end-user logistic teams for spares depot and advanced replacement options.
- Historical RMA analysis to identify trends or anomalies that might otherwise be overlooked to correct in further production runs. This concept of continuous improvement is typically out of reach for smaller providers with limited resources.
- Predictable replacement availability as part of an end-user's operational planning processes.

4.2 Low-Cost Vendors Often Provide:

- Limited or outsourced support with little access to the specific engineering and production teams that actually created the products.
- Simple “refund or replacement only” policies which ignores the root problem, and simply assumes that a new device will perform better
- No formal failure analysis to allow long term reliability to be studied and included in vendor scorecard exercises
- Inconsistent spare unit availability as repair stock levels vary when higher cost unit prices drive lower repair stock quantities.

Once production ends or individual components become unavailable, low-end switch providers may not maintain suitable stock levels for repair processes, forcing premature rip-and-replace cycles. Mature vendors calculate reasonable levels of stock for returns/RMA requirements.

5.0 Total Cost of Ownership: The Real Economics

While lower-cost networking products may be 30–60% cheaper up front (even when they use the same core switching silicon chipset), customers and operators often pay more over the lifecycle due to:

- Increased early failures and the larger number of remediation cycles associated with troubleshooting, and the logistics of higher rates of repair and replacement.
- Labor cost from field replacements when service providers must be engaged to provide on-site physical device remediation. Onsite visits are costly, so switches that increase these visits significantly increase the real TCO.
- Unplanned network downtime due to product failures can incur costs in the range of thousands of dollars per hour to millions. Lower cost switches typically cause more unplanned downtime, which directly translates to increased bottom-line TCO.
- Shorter product lifespan due to heat, thermal and component aging as discussed above. Designs which cut corners typically fail much earlier, with lifespans that can be half of those seen with highly engineered mature alternatives.
- Unpredictable supply and spot pricing due to smaller batches of devices being manufactured. In fact each batch of a low-cost switch may have unique pricing and for devices that are based on the most popular silicon, these prices typically rise over time. Mature providers enter longer-term contracts which manages cost of goods in a much more predictable fashion.
- Higher replacement frequency for lower cost solutions directly affects the TCO on any given solution.

Premium devices, designed with deeper engineering rigor and manufactured at scale, deliver:

- A longer life, as the higher quality and better spec components are used. The specs are well within their designs limits and are engineered to perform comprehensively beyond what is drawn on a schematic.
- Higher mean-time-between-failure (MTBF) simulation and testing during the engineering phase assures product longevity.
- Consistent manufacturing quality across fabrication centers worldwide assure product uniformity regardless of manufacturing site or date.
- Predictable parts availability due to volume purchase commitments assures that the components that were designed up front continue to be exactly as built rather than relying on substitutions which may have different operating specs than the original.
- Lower total operational cost over any reasonable period of time for IT equipment (typically 3 years or longer).

Over a multi-year deployment, the difference in TCO becomes significant and measurable. And when used in a scaled deployment (like a data center), the added hard costs and higher operational risks become very apparent.

6.0 Conclusion

The industry myth that "all Whitebox switches are created equal" is a dangerous oversimplification. While two devices may share the exact same switching silicon, they are by no means equivalent in performance, reliability, or Total Cost of Ownership (TCO). The silicon is merely the engine; the true differentiation lies in the chassis that houses it—specifically the power architecture, PCB signal integrity, thermal dynamics, and mechanical resilience.

Low-cost alternatives often achieve their price points through invisible compromises—cutting corners on component quality and manufacturing tolerances. These deficiencies lie dormant until deployment, where they manifest as costly outages and degradation. Conversely, vendors with deep engineering maturity and industrial-scale rigor deliver devices that verify quality at the architectural level. By prioritizing holistic design over short-term savings, high-quality switches:

- **Maximize Operational Resilience:** Function reliably across wider thermal and electrical envelopes, drastically reducing remediation cycles and OpEx.
- **Ensure Deterministic Reliability:** Leverage superior component vetting and rigorous stress-testing to minimize Mean Time Between Failures (MTBF).
- **Preserve Signal Integrity:** Maintain consistent performance over an extended lifecycle, preventing the degradation of internal connectors and power systems that plague cheaper units.
- **Secure Long-Term ROI:** Transform the network from a consumable expense into a durable asset, delivering superior economic value in scalable environments.

In business-critical networks, silicon provides the capability, but engineering discipline and manufacturing maturity guarantee quality and execution.

At-A-Glance Comparison

Category	Name-Brand Engineering Manufacturer	Low-Cost / Off-Brand Manufacturer
Typical MTBF (Mean Time Between Failures)	**150,000–300,000 hours** (17–34 years theoretical under controlled conditions)	**30,000–70,000 hours** (3–8 years theoretical, often significantly lower in field conditions)
Power Supply Quality	High-grade capacitors, derated components, protected power rails, surge handling, long-life PSU design	Low-cost components, minimal surge protection, shorter lifespan, greater failure risk under voltage fluctuation
Thermal Engineering	Full thermal simulation, high airflow design, quality heat sinks, optimized layout, internal temperature monitoring	Small fans, less simulation, higher heat concentration, limited airflow control, higher component aging rate
PCB Design	Multi-layer boards, controlled impedance routing, heavy copper planes, EMC suppression, signal integrity engineering	Minimal-layer boards, thinner copper, copied layouts, reduced shielding, increased risk of EMI and packet loss
Mechanical Structure	Rigid chassis, reinforced mounting points, strong port retention, resistant to vibration/shipping stress	Thin sheet metal, weak mechanical support, ports easily flex or fail, increased chance of solder joint fatigue
Component Sourcing	Stable supply chain, audited component vendors, long-term lifecycle commitments	Supplier changes based on cost, inconsistent BOM, limited sourcing transparency, spot buys
Manufacturing Capacity	High-volume production across multiple factories; strong statistical process controls (SPC)	Small batch runs, minimal statistical tracking, less consistency between production lots
Quality Testing	Full lifecycle testing including HALT/HASS, thermal chambers, shock/vibration, burn-in testing	Basic functional test only; little or no stress or accelerated aging testing, may not be verifiable
Product Consistency Across Builds	High repeatability; serial-level traceability; extensive return analytics, automated processes	Lower consistency: units from different lots may behave differently or use different component mixes, manual steps
Spare Parts Availability	Predictable multi-year spare inventory and forward lifecycle planning	Limited or unpredictable; product discontinuations may make spares unavailable quickly
Serviceability and Support	Structured RMA programs, replacements within days, detailed hardware failure analytics	Refund or exchange only; minimal analysis; lack of structured replacement flow

Category	Name-Brand Engineering Manufacturer	Low-Cost / Off-Brand Manufacturer
Expected Service Life in Real Deployments	7–12 years in continuous operation	2–6 years; often shorter in demanding environments
Total Cost of Ownership (TCO)	Lower over the product life due to reduced failures, stable supply, and long lifecycle	Higher over the product life due to replacements, downtime, and unpredictable support

About Edgecore

Edgecore Networks is a division of Accton Technology, the market leader in providing open standards-based disaggregated network and computing solutions to the hyperscale, large enterprise, service provider and campus communities. Edgecore Networks provides off the shelf high performance switches which enable architectural flexibility, unlimited scalability and a total cost of ownership which is half of competitive approaches. With more than 35 years of experience and millions of devices shipped, our business model is centered on building long-term partnerships that foster trust and innovation.