




R.O.C.S.™ 2

Going Green: Optical Switches for Sustainable Data Center Operations

Modern data centers are driving up energy consumption

Data centers are essential for modern society as they enable cloud computing and big data storage, vital for business and personal applications. However, data centers require significant amounts of energy to function, and their energy consumption is a growing concern for sustainability.

One way to address this issue is by reducing the energy consumption of the network switches in the data center, which can account for up to 30% of the total energy consumption. This white paper explores using an optical switch to reduce energy consumption in a data center.



Artificial intelligence is driving up energy consumption

Energy consumption explosion


The rise of cloud computing, big data, and artificial intelligence has led to an explosion in the energy consumption of data centers. As a result, the costs associated with powering and cooling data centers have become a significant challenge for data center operators.

Environmental impacts

The exponential growth in data center energy consumption has also contributed to increased greenhouse gas emissions, leading to concerns about the impact of data centers on the environment.


Innovation challenges

The challenge for data center operators is to find innovative solutions to reduce energy consumption and associated costs without compromising performance or availability. The use of optical switches is one such solution that can help data centers address these challenges.



The Real Amount of Energy A Data Center Uses – AKCP.com, 2022

Optical circuit switching delivers energy savings and great performance



Efficient use of the network infrastructure

Optical switches being bit-rate agnostic, they allow efficient use of transceivers from different generations and data rates in the same infrastructure. This efficiency reduces overall energy consumption as the network switches can route more data in less time.

Reduce cooling requirements

An optical circuit switch consumes less power than an electronic switch, lowering operating costs. Additionally, the lower heat output from an optical switch reduces the energy needed for cooling the data center, further lowering the operating costs.

Deliver energy savings in AI and HPC

In AI and HPC, where systems run 24/7 at peak capacity and at a high price per kWh, this translates into potential annual savings of millions of dollars.

Eco-Friendly

Therefore, by implementing optical circuit switches, data centers can operate more sustainably and become environmentally friendly.

Mission Apollo

Landing Optical Circuit Switching at Datacenter Scale

Optical circuit switches have reduced electrical usage by over 30%

The world's first [large-scale deployment](#) of optical circuit switches for data center networking was introduced by Google.

The Apollo layer replaces the Spine blocks for [large cost and power savings](#) by eliminating the electrical switches and the optical interfaces.

This also substantially improves their infrastructure's evolution, reliability, cost efficiency, power consumption, and application performance.



A game-changing approach to tackle the biggest failure points

Lower the number of networking equipment

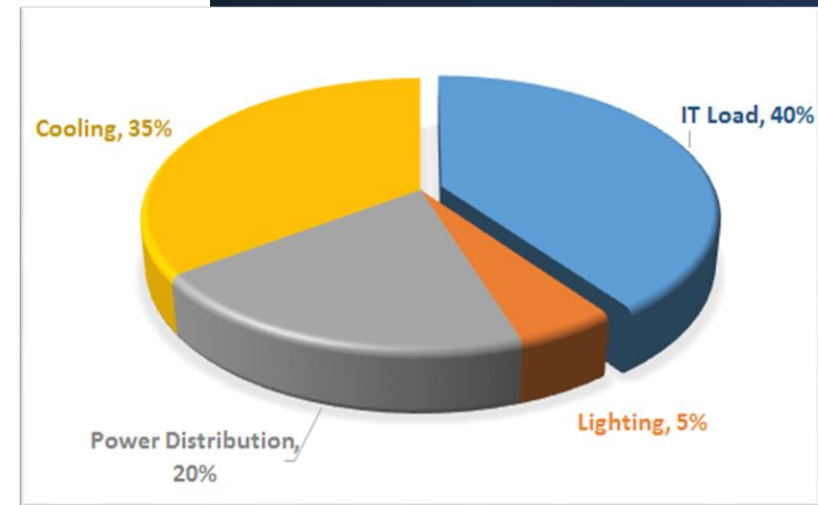
With each generation of higher-speed transceivers comes higher power distribution requirements. Considering a typical 400G transceiver and a 64-port 400G-enabled network switch operating at maximum capacity, the power footprint, including cooling, is well over 1 kW. Minimizing the use of such power-hungry equipment is one solution to achieving sustainability in AI and HPC data centers.

Prevent transceiver failures

The increasing density of network switch ports has created a significant cooling challenge. With dozens of power-hungry high-speed transceivers in a single unit, successfully maintaining the ideal operating temperature is no easy task. Transceiver failures due to poor thermal management are one of the most important factors.

Leverage optical circuit switching

Introducing an optical switch into the rack helps tackle these failures. Optically switched communication links reduce both the number of transceivers involved in the communication path and their failure rate. It also leads to a reduction in the workload on the packet switches in the network, avoiding the deployment of additional packet switches.



Breakdown of a Data Center Power Consumption - Chhachhi, 2015

R.O.C.S.™ 2

Optical switching and fail-over inside the datacenter

Leverage compact and low port-count optical switches to deploy a decentralized architecture

R.O.C.S.™ 2, a 2x2 dual optical circuit switch, provides a solution that can tackle the biggest issues in data centers; Power consumption and cooling. The compact size and low port-count eliminate the need to provision fixed-location rack units. On one hand, given this flexibility, each optical switch can be installed locally if needed, minimizing the impact on cable management. On the other hand, optically switching the link with a low-density optical circuit switch avoids the challenge of maintaining proper operating temperature for the transceivers as seen with a high-density packet switch.

Designed for the most demanding datacenter computing workloads

Datacenters specializing in artificial intelligence, machine learning, and high-performance computing face significant challenges in increasing their computing capacities while keeping cooling systems within budget. R.O.C.S. 2 devices have been designed to meet the needs of these markets. A R.O.C.S. 2 device is scalable, fault-tolerant and maintenance-free and can operate 24/7. This high availability and failover capability without the need to provision additional power and cooling resources allows these data centers to reduce their environmental footprint.

Provide an effective solution for transceivers failures

The signal monitoring capabilities provided by R.O.C.S.™ 2 help deploy an effective solution in the event of transceiver failure. In the event of a link failure, the device can automatically switch to a secondary link without the need to provision more network ports. This approach provides significant energy savings compared to deploying network redundancy to address these outages and helps keep operational expenses low by providing quality of service and a maintenance-free solution.

Compatibility with mixed generations transceivers maintaining deployment flexibility

The independence of the R.O.C.S. 2 switch to data rate also allows multiple generations of transceivers to be used in the network architecture. Similar to Google's Apollo fabric, this means that all the different aggregation blocks (AB) can use the same device providing good flexibility and scalability. This smooths out the power and cooling demand curve.

R.O.C.S.™ 2

Dual 2x2 Optical Circuit Switch



Power Consumption: < 0.5 W

Optical losses: < 1.5 dB

Bandwidth: 1260 nm to 1620 nm

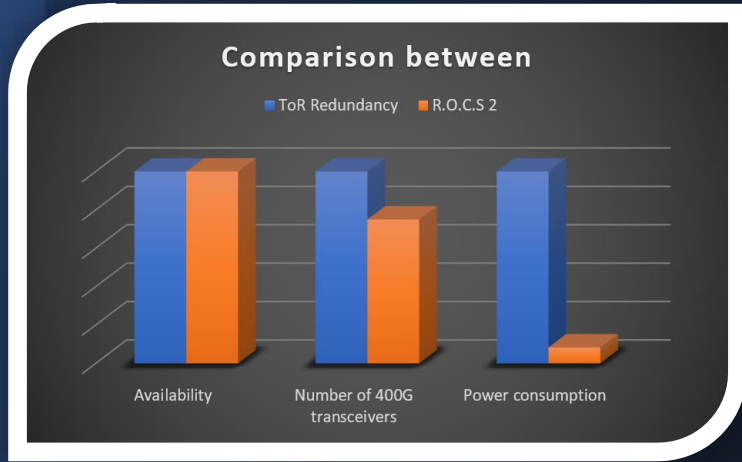
PDL: < 0.1 dB

Crosstalk : < -60 dB

Switching time: < 10 ms

- ✓ Automatic switching (Fail-over)
- ✓ Integrated optical power monitoring
- ✓ Adjustable power threshold
- ✓ Fail-safe on power loss

Lowering the number of equipment while unlocking high availability and power savings



High availability

In this use case, the R.O.C.S.TM 2 device is installed in a configuration that restores the server's connectivity when a downward transceiver fails by re-purposing an upward transceiver as a downward transceiver.


Lowering the number of equipment

Such configuration avoids the deployment of a second ToR and two additional 400G transceivers compared to deploying ToR redundancy leading to significant power savings while offering high availability.

Allocating room for computing resources.

R.O.C.S.TM 2 devices are compact and scalable, allowing better use of space, thus decreasing cooling requirements and granting freedom for the growth of computing resources.

Use-Case Switch-side Auto-failover




Complement your existing infrastructure

Top-of-rack (ToR) switch redundancy and hybrid optical circuit switching offer different approaches to enhancing reliability and performance within an AI/HPC datacenter.

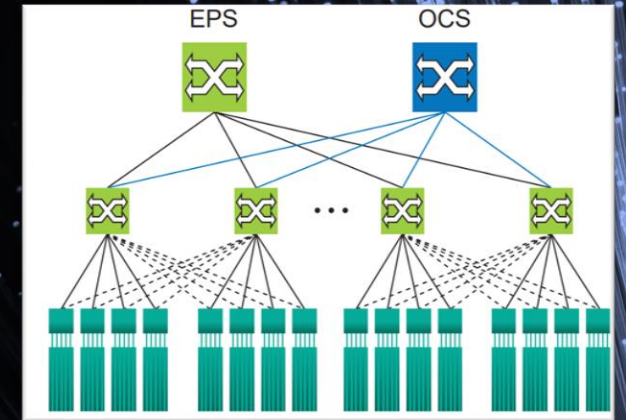
ToR switch redundancy involves deploying multiple switches to ensure network resilience and mitigate single points of failure.

On the other hand, hybrid switching leverages optical communication technologies to establish dedicated connections between compute resources. This reduces latency and congestion, enabling efficient data transfers for AI and HPC workloads.

Both approaches can complement each other to create a reliable, sustainable, high-performance datacenter infrastructure.



Popular ToR and ToR Switch in Data Center Architectures – FS.com, 2022



EPS/OCS Hybrid Networking – NTT Technical Review, Vol. 14, 2016



Let's talk

www.aeponyx.com