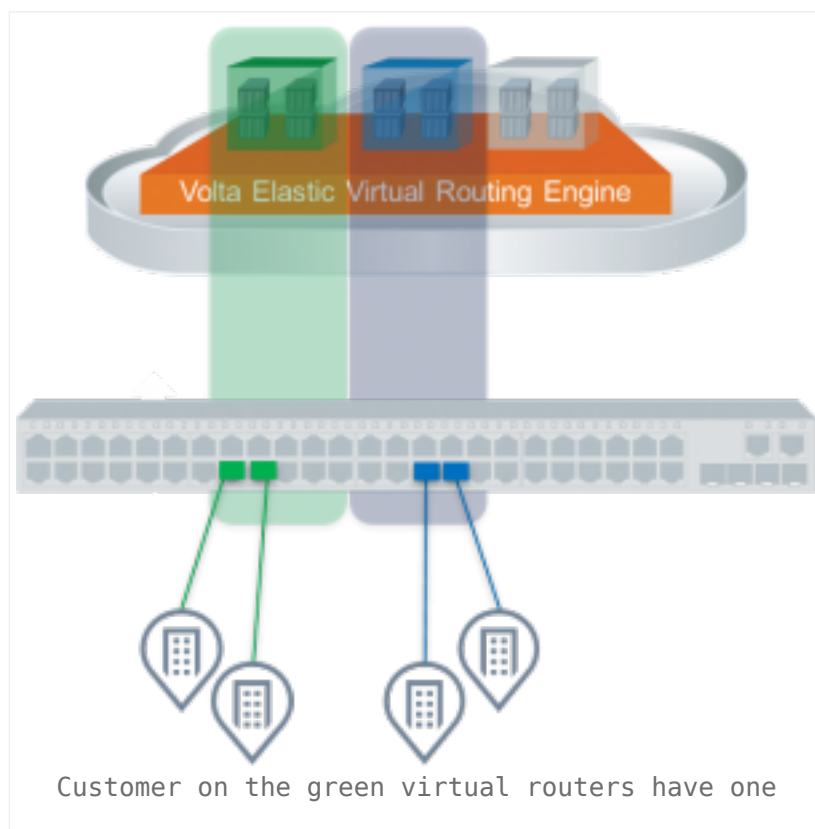


Provider Edge Router

Service provider revenue depends on the provider edge router to deliver enterprise-class offerings like MPLS VPNs. The provider edge is under pressure from two interrelated trends. First, enterprise customers are requesting more advanced services with more frequent changes. Second, these customers also need more bandwidth which requires an investment in higher port speeds. Thus, services are becoming more complex and often must be chained together from several different physical and virtual network vendor devices. This leads to large, complicated configuration files, often tens of megabytes in size, where multiple service configurations are mixed together without a clear separation between customers. Over time, the provider edge has become complicated, operationally expensive and difficult to scale which leaves little room for innovation.

Legacy routers are constrained by how they deliver device management and the control plane. These proprietary chassis-based routers run the management and control plane software on a device controller (e.g. Juniper Routing Engine or Cisco Supervisor Engine) that is essentially a single blade x86 based server. This software is both memory and processor intensive. As routers have to deliver more services, over more ports, with larger configurations, the device controller's memory and processing do not scale with it, creating a bottleneck in service delivery. Providers need both a way to both lower the per-port cost of the data plane and scale out the control plane.

This is particularly true of MPLS edge routing. MPLS will remain essential to enterprise-class services for the foreseeable future, but providers must reduce the cost at the service edge to maintain profitability.



service while customers on the blue VR have another.

The port cost can be addressed by using white box switches. This is a very cost-effective approach to add ports to the provider edge. By providing cloud-native networking software running on a public or private cloud infrastructure, Volta enables network operators to scale out the control plane using the most cost-effective means of adding processing and memory. Each virtual router runs the bulk of the control plane in the cloud so processing can scale cost-effectively. This approach has a separate instance of the routing protocols for each virtual router, enhancing scalability, reliability, and customization of each virtual router.

As a result, Volta supports up to 255 separate virtual routers on a single white box switch. In this multi-tenant scenario, each virtual router is a separate set of processes in the cloud as its own administrative domain that is configured to meet a given customer's exact requirements. It significantly simplifies service creation and provisioning by having a separate configuration file for every customer. This all leads to faster service delivery, as well as reducing the errors in the process. Volta's networking service library is delivered as YANG (RFC 6020) data model or as a Python library. A single API into the cloud simplifies integration into a network operator's OSS and BSS.

This simplifies the service edge which enables faster service delivery and more robust solutions for the customers while reducing CAPEX and OPEX. Volta's solution is fully standards-based which ensures interoperability with legacy routers. This allows providers to gradually add white box switches. In cost comparisons, Volta found that adding white box switches using our cloud software saved 90% of the cost of adding ports to existing legacy routers.