CloudBackbone does the heavy lifting

By 2020, 85 percent of applications will be deployed on the cloud. Serving as network traffic hubs, data centers (DCs) will bear cloud services and backbone networks will create large-scale DC interconnections (DCI) by connecting WAN. But, are traditional backbone networks up to the task?

By Zhang Liqiang, Zhou Fei, and Gao Bo

Cloud services have three major requirements that networks need to deal with: fast TTM; hard-to-predict, randomly changing traffic; and different bearing requirements for different services.

As an agile solution for the ultra-broadband age, CloudBackbone comprises two parts: one, the network physical layer, which includes backbone network routers and transport devices, and two, the network control layer – the layer's brain, which is implemented by Network Cloud Engine (NCE).

Two industry wins

Increases in DCI information access and data backup have caused inter-DC traffic to increase, boosting CAGR to more than 30 percent. To meet future bandwidth requirements for cloud services, Huawei has launched NE9000, the industry's most fully integrated backbone router. It's also rolled out the industry's first 4T routing line cards, yielding a total system capacity of up to 80 Tbps – four times higher than competitors' devices.

Huawei's transport devices provide a 320–640 Tbps all-optical switching platform that matches large-capacity OTNs with OXC all-optical cross-connections, specs that can keep up with the next 5 to 10 years of service development.

With its 4T line cards, one NE9000 backbone router can provide 800 x 100GE ports, meeting the required quantity of interconnections between 100GE ports for connecting DCs into a full-mesh, flat backbone network. The NE9000 provides industry-leading 400GE ports that work with the transport devices to support E2E 400 Gbps transmission links. The router improves port efficiency fourfold, simplifies network topology, and requires fewer optical fiber links. The Solar 5.0 network processing chip cuts energy use to 0.4 W/G, half the industry average, which slashes OPEX and solves equipment room issues.

Service packages can increase new subscribers to more than 10,000 per month, but also drive up service bandwidth requirements, with single-node capacity likely to reach 25.6 Tbps over the next five years. Due to the space and power supply limitations of equipment rooms, operators prefer a single product to multiple devices. Due to the new 4T line cards, a single NE9000 can deliver 80 Tbps in capacity, equivalent to the capacity of three competitor 2+4 clusters.

The NE9000 meets five-year service requirements, uses far less space in equipment rooms, and saves 1 million kWh over five years.

DCI and DCN

Different departments usually manage DC network (DCN) service deployment and DCI deployment, manually transferring and coordinating tenant requirements. Weeks or months can elapse between issuing requirements and provisioning services, which is far too long for fast-moving cloud services. Through the NCE, CloudBackbone can coordinate DCN and DCI deployment and automatically transfer intra- and inter-DC tenant and service requirements. It takes just minutes to deploy services after user requirements are issued, massively shortening TTM and service innovation.

A large-scale DC generally bears tens of thousands of users with a wide array of requirements on network delays. For example, financial transactions have strict latency requirements, while data backup is hard on bandwidth. Traditional models use the same link to bear different users' services, so they compete for bandwidth and can't ensure service quality. The NCE automatically maps user requirements and services in a DC to VPN tunnels with different capabilities. One VPN tunnel is mapped to each tenant, so E2E tenant-level SLAs can be flexibly set. Operators can then charge subscribers according to network capabilities, adding value to pipes.

Keeping it real-time

The random, unpredictable nature of cloud services like inter-DC VM migration and network hotspot events can result in traffic bursts, congested links, and low traffic utilization. Because traditional traffic engineering (TE) is based on predicted network conditions, it cannot solve this issue. These conditions are where NCE proves its worth.

Like an urban transport control center, NCE has a global network view. It can collect network-wide usage data on link bandwidth and recalculate routes based on bandwidth balancing policies. It can then redistribute specified traffic in real time from a congested link to one with a lighter traffic load, improving link bandwidth utilization by up to 50 percent.

To ensure user QoE, CloudBackbone can adjust traffic based on factors such as delay, using lower-latency links for VIP services and diverting common services. China Telecom and Tencent have already deployed Huawei's traffic optimization technology, with great success.

CloudBackbone calendars bandwidth so operators can customize bandwidth assurance on-demand, based on service characteristics, for example, setting a period where 200 Mbps

CloudBackbone's innovative award-winning IP+Optical solution for collaborative planning uses a Super Controller for planning, protection, and management.

of bandwidth is assured and another where only 100 Mbps is assured.

InfoVision 2016 award winner IP+Optical

The IP and optical layers are independently planned and maintained on traditional backbone networks and links are manually deployed. Provisioning a single link not only takes weeks, but backup link planning is done separately, which wastes resources on re-planning redundancy.

CloudBackbone's innovative award-winning IP+Optical solution for collaborative planning uses a Super Controller for planning, protection, and management. It also coordinated the IP-layer and optical-layer controllers, improving network resource utilization and maintenance efficiency. For example, planning inter-layer links can simplify changing a redundant link configuration from N:N to N:1, thus reducing construction costs. Additionally, collaborative protection enables the IP layer to reach peers through optical-layer links to improve network reliability, preventing a multipoint link failure from interrupting services if the IP layer can't find or restore the failed links.

The Super Controller automatically sets up IP-layer and optical-layer links, reducing provisioning time from weeks to minutes, greatly simplifying O&M and minimizing maintenance costs. Huawei's IP+Optical solution has cut TCO by 40 percent for the Latin American operator America Movil, and slashed automated service provisioning and deployment to minutes.

Multi-service integration

Due to increasing requirements, backbone networks need to shift from bearing single services to multiple services. This is necessary to maximize network resource utilization and avoid overlapping investment.

CloudBackbone can virtualize a single physical backbone router into multiple independent virtual systems (VS). Each VS is allocated an independent slicing unit for its exclusive use, enabling hardware resources such as backplane and power supply to be shared while isolating the control and forwarding planes. Control and forwarding resources can be flexibly configured according to service requirements and each VS is isolated, preventing interference between them.

The VS technology in CloudBackbone executes network slicing so that different network planes virtualized from one physical network can bear different services. Slicing helps operators improve resource utilization, use fewer physical nodes, and cut network construction costs. Control, management, and the physical planes used by different services are isolated, enabling each plane to be independently managed and upgraded – using a dedicated network for each service guarantees network security, reliability, and user experience.

Adapting backbone networks to the requirements of cloud services is vital for operators to digitize. The Huawei CloudBackbone solution helps them build ultra-broadband, agile, reliable, and competitive next-gen backbone networks.