Scale Geographical Multicloud Demand

Problem

Exponential traffic volume growth between cloud-based applications adversely affects user experience because it congests traditional centralized MPLS network topologies. Offloading traffic to the public internet is not secure, and dedicated connections to individual clouds are financially and operationally unsustainable.

Solution

Adapt to changing business needs and support new business models across time zones, with shifting traffic patterns and new regulations. Leverage edge-based ecosystems of service providers (cloud, network, SaaS) that can tailor resources and services based on business-driven policy decisions invoked via real-time configuration control (SDN/NFV). Drive inter-cloud traffic through the edge node via virtual interconnections to clouds. Minimize inter-cloud data transfer and store service quality policies in local data repositories at the edge node; drive traffic demand to the interconnected edge node mesh where digital ecosystems route traffic to other nodes most efficiently. Dynamically rewire services and connectivity, continually adapting to planned and unplanned business and technical environment shifts. Adjust cloud utilization across and within edge nodes to reflect usage and cost policies.

Constraints

1. Multicloud resource management is complex if each connection is separately managed with each cloud-based application having different needs over time.
2. Dynamic volume, bandwidth and infrastructure resource policy management wasn't considered feasible in traditional IT experiences.
3. Fear of inadequate cloud resource infrastructure causes firms to over-provision in fixed-price arrangements; yet, growing digital economy-based demand is dangerously unpredictable.
4. Distributed multicloud application proliferation requires resiliency tactics other than centralized failovers.
5. Multicloud workloads must interact in real time with resilient connections.
6. Diverse workload interactions (e.g., collaboration, transactions and analytics) require sophisticated QoS transmission policies ensuring delays don't degrade user experience.
7. Long distances impede high throughput and responsiveness in MPLS networks despite the use of SD-WAN.

Steps

1. Employ edge-based cloud ecosystems. Virtually cross connect dedicated circuits to clouds on demand.
2. Use SDN/NFV to segment traffic with different bandwidths based on application QoS needs.
3. Install cloud balancing with cloud bursting capabilities, ensuring effective resource utilization.
4. Install cloud resource monitoring and accounting to align and inform cloud-balancing policies.
5. Solve for exponential growth of demand across the distributed enterprise, deploying applications in geographically dispersed clouds through a mesh of interconnected digital edge nodes, leveraging partner ecosystems where private, one-to-many connections can be established.
6. Tailor each node for local services, enabling control of performance and cloud utilization scaling on demand.
7. Leverage predictive analytics and inform resource allocation policies across the distributed enterprise.

Reference View

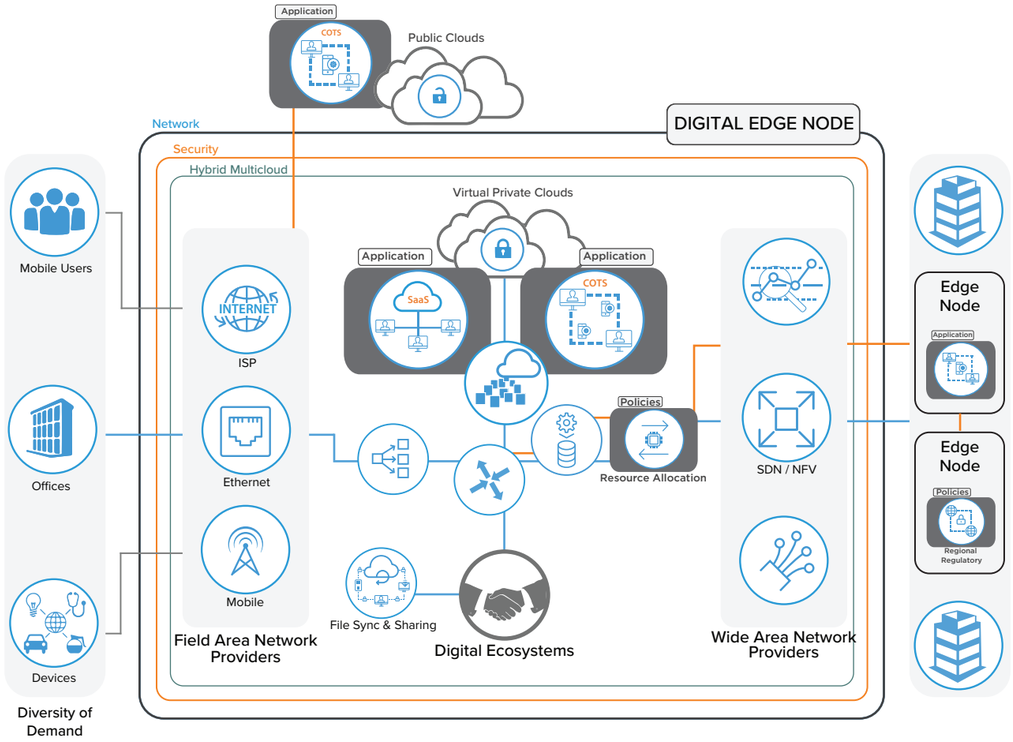

Forces

- Time to market combined with a compelling need to minimize the IT debt of aging infrastructure is driving new business applications to be delivered via cloud platforms to remain competitive in the digital economy.
- Expansion into new geographic markets with new partners cannot wait months for standard IT deployments.
- Providing relevant, timely value for cloud-based systems of insight and engagement across the proliferation of devices and their expanding capabilities to a growing mobile workforce is unsustainable using traditional centralized network architectures.
- Demand fluctuates sufficiently to challenge assumptions about capacity management in traditional architectures.
- Centrally planned capacity management cannot meet demand.

Results

Technical

- Most traffic is localized using optimal paths via edge nodes to users of cloud-based workloads. Backhaul to corporate data centers is minimized.
- Cloud resource usage is dynamically tied to demand.
- Dynamic routing enables self-healing in the case of local bottlenecks.
- Adding more bandwidth is programmatic, increasing flexibility and agility.

Business

- Dynamic, real-time enterprise able to respond to changing demand by flexibly adjusting infrastructure supply—saving resources and reducing waste without re-architecting.
- Real-time and predictive analytics inform digital economy business strategies.
- Partnerships enable regionally tailored business models on a global scale.
- Planning for periodic shifts in demand is a policy versus engineering issue.