



Microsoft's Cloud Networks



Microsoft's Cloud Networks

Microsoft's customers depend on fast and reliable connectivity to our cloud services. To ensure superior connectivity, Microsoft combines globally distributed datacenters and edge computing nodes with one of the largest fiber backbones, providing multiple terabits per second of capacity to over 70 points of presence around the world.

This strategy brief provides an overview of the global network Microsoft has built to bring cloud services to enterprise customers, small businesses, governments, and consumers, and how we are optimizing our network for performance, cost, and reliability.

Hyper-scale

Microsoft's cloud supports over 200 online services including Bing, Microsoft Azure, Office 365, OneDrive, Skype, and Xbox Live. We have invested over \$15 billion in building a highly scalable, reliable, secure, and efficient cloud infrastructure. Today we interconnect over 1600 unique networks with multiple redundant points, allowing us to move more than 110 million petabytes through our global Microsoft datacenters, over more than 1.4 million route miles of fiber.

To support cloud services at this scale, we have built a globally distributed networking platform to enable high performance, low-latency, and low transactional-cost services.

Our cloud network is built on the following principles:

Agile – driving standardization and automation to reduce capacity cycle time and increase speed and performance.

Scalable – investing in a network infrastructure today that will support future growth, while reducing supply side risk, demand variability, and improving capacity capabilities.

Reliable – reducing the service dependency on network state enabling fast upgradability and higher performance.

Visible – maintaining advanced telemetry to continuously monitor current operational state and model planned state for greater availability.

Efficient – driving to reduce functional cost per megabit, while increasing throughput and speed.



Global network overview

We think about our network in three major components—inside the datacenter, between our datacenters and edge nodes, and our geographic reach throughout the internet ecosystem.

Inside our datacenters, we connect more than 1 million servers to the network fabric, which contains the routers, load balancing, firewalls, Domain Name Service (DNS) servers, and many other services. The datacenter fabric then connects to the core backbone and our inter-datacenter fabric.

At the physical layer, Microsoft invests globally in fiber assets to insure our ability to continuously scale our bandwidth. Our network consists of thousands of 10 and 100 Gb/s links allowing customer traffic to ingress and egress Microsoft's datacenters, as well as supporting inter-datacenter traffic for applications such as data replication for disaster recovery.

Outside the datacenter, we utilize metro optical links that allow us to communicate between multiple datacenters and connect to the long haul, which allows us to transit transcontinental or sub-sea distances.

We also connect to edge sites where we attach to the Internet through third party colocation facilities, service providers, and multi-system operators.

Azure ExpressRoute enables customers to create private connections between on-premises or colocation infrastructures and Microsoft's datacenters. ExpressRoute connections do not go over the public Internet, and offer more reliability, faster speeds, lower latencies, and higher security than typical connections over the Internet. In some cases, using ExpressRoute connections to transfer data between on-premises and Azure can also yield significant cost benefits.

Network architecture

The growing demand for our cloud services requires a nimble network architecture that is able to efficiently scale out without undue complexity. We are migrating from a classic network design with a point-to-point, simple tree configuration, to a CLOS-based design that more seamlessly moves data throughout the network. Specific devices are not limited to performing specific functions; rather, those functions are now performed in the servers as applications.

This approach provides the following advantages:

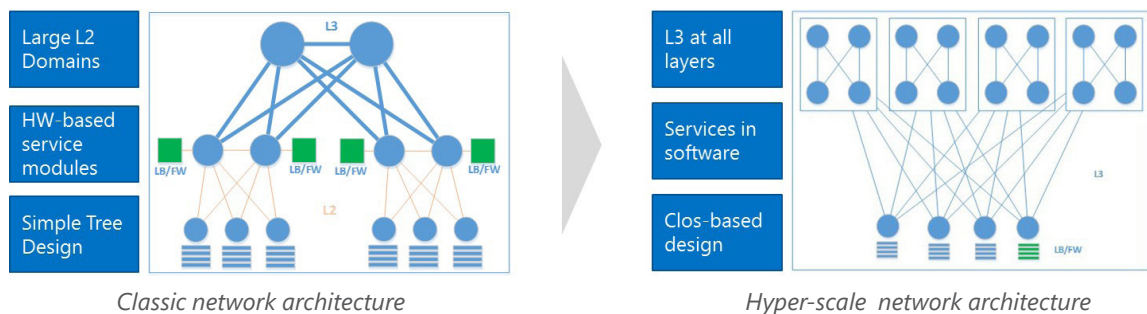
- Employing a more resilient design, automated monitoring and remediation, and minimal human involvement results in **higher availability**.
- Moving away from manual provisioning processes and high diversity to greater

automation for network provisioning and integrated processes delivers **greater agility**.

- Shifting from complex hardware configurations to an environment with simplified requirements and a more unified infrastructure provides **improved efficiency**.

Network function virtualization

The transition of functionality from custom network hardware to virtualized software functions, within our standard compute nodes, is known in the industry as Network Function Virtualization (NFV). Microsoft has adopted NFV in several components of our architecture from ExpressRoute to firewalls. This provides both added flexibility for how we utilize our CPU cores, but also enhanced horizontal scalability.



Network infrastructure convergence

To stay ahead of the sprawling complexity that comes at the massive, hyper-scale that our cloud operates in, we have converged our infrastructure. The design of the datacenter, servers, storage, network, and applications are optimized on an end-to-end basis with common fungible components in a uniform design.

The fabric allows data to move seamlessly around the network and the design approach means that if we lose a certain section of a datacenter, we have the

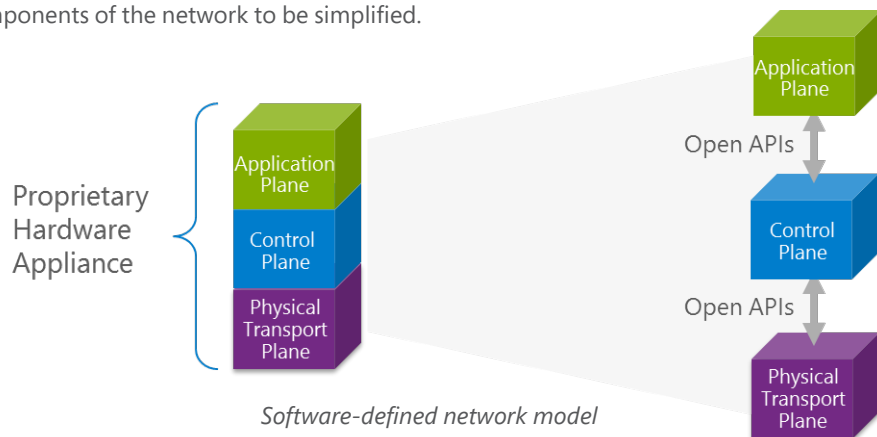
ability to have that application run elsewhere without interrupting the service. Standardized commodity servers can operate disparate workloads, and common blob and structured storage solutions allow us to move data around quickly while maintaining necessary protection. This convergence extends to the Domain Name Server (DNS), Content Distribution Network (CDN), and Edge nodes to connect in closer proximity to customers and improve their overall experience.

Software-defined networking

To improve flexibility and accelerate the adoption of advanced technologies into our network, we have broadly adopted software-defined networking (SDN). SDN is the architectural concept of separating the software control plane from the data plane. Traditional high-end network hardware is very computational-intensive, as it is required to perform table look-ups and determine the handling priority for every packet passing through the network. In an SDN environment, the control plane traffic, which guides packets to their intended destination, runs on a different set of servers, allowing the switches and other components of the network to be simplified.

Utilizing SDN, we are able to extract and separate the application, the control plane, and the transport of the data. This allows us to insert our own APIs to gain visibility of how the data flows, gain automated control over these data flows, and allows us to optimize network performance outside of the hardware refresh cycle.

As our cloud infrastructure continues its rapid build-out, SDN provides us with the ability to manage and optimize the complexity and scale of the systems from an end-to-end operational perspective, gaining greater agility and reliability.



Managing the technology lifecycle

As with most areas in IT, the technologies employed in the datacenter are continually advancing in terms of performance, efficiency, and acquisition cost. While we strive to employ the latest technologies across our cloud infrastructure, we also need to rationally and methodically manage the update process.

The approach we take can be compared to a "crop rotation," where we maintain and deprecate deployed architectural instances as a unit, instead of employing piecemeal device upgrades. Network and server infrastructure typically follows a

three to four year service lifecycle, while new technologies tend to offer compelling advantages roughly every 18 months. In the initial build-out of a datacenter, we bring in new servers, network equipment, and supporting infrastructure. Over a period of time, this equipment becomes dated, as new technologies are developed that offer performance and efficiency improvements. The crop rotation approach allows us to continually onboard and deploy new equipment as older technologies are deprecated and retired, while maintaining continuous operations.

Looking forward

As we continue to build out our hyper-scale network infrastructure to meet the accelerating demand for our cloud services, we are focused on scale, performance, and cost effectiveness. To remain competitive and offer superior services to our cloud customers, each infrastructure component must be holistically optimized on an on-going basis to ensure the lowest transactional cost is maintained while ensuring scalability, availability, and operational efficiency.

Reducing acquisition costs

At any given time, there are several efforts underway to drive down the capital costs for the components of our network infrastructure. These include assessing silicon photonics optics technologies becoming available in the industry, integrating optical subsystems into the switching infrastructure, designing network architectures that reduce dependence on fiber cabling, and integrating network technologies into the servers and racks themselves. These efforts help us remove unnecessary components from the end-to-end infrastructure and more optimally locate functions where they can be performed most cost-effectively.

Reducing operating costs

Using our deployed resources at the highest efficiency possible allows us to build and manage less of it. Any infrastructure we deploy that is not substantially utilized, is a wasted resource. To this end, we are focused on collocating related compute and storage workloads—known as affinization—to improve application performance and reduce the resources required to support the workload, and optimize the placement of traffic onto the network through the SDN techniques (outlined above) to maximize the utilization of existing network resources.

As we pursue these potential technological innovations to deliver better performance at lower cost, we remain committed to the four priorities that guide all of our efforts:

Customer-driven growth

- Respond quickly to support the services our customers are consuming.
- Continuously expand our dark fiber footprint.
- Expand capacity through strategic acquisitions that augment our network capacity.

Speed-orientation

- Expand our edge-node presence to bring direct network connections closer to our customers to improve the overall experience.
- Reduce the deployment and activation time for edge nodes, core and optical capacity, and the datacenter fabric.

Persistent availability

- Ensure that compute and storage infrastructure is connected and available.
- Keep processes in place for remediation in the case of disruption and maintain strict procedures for changes to the network.

Data-driven decisions

- Analyze current state and determine the roadmap for migration.
- Drive the adoption and migration towards a converged network.

Without a high quality network, there is no cloud. Microsoft is continually innovating to improve performance, increase availability, and reduce overall costs. The investments we make help ensure our customers gain the advantages of our cloud services while gaining the best possible experience.

Microsoft has extensive experience operating a cloud services infrastructure since 1995. As Microsoft's cloud services portfolio and infrastructure continues to grow we are making thoughtful investments to answer customer needs for greater availability, improved performance, increased security, and lower costs.

Contributors:

Brad Booth

Jeff Cox

Monica Drake

Vijay Gill

Darrell Porcher

Theresa Wescott

David Whipple

For more information, please visit www.microsoft.com/datacenters

© 2015 Microsoft Corporation. All rights reserved.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.