

IT@Intel: Building a Multi-Cloud-Ready Enterprise Network

Intel IT has adopted a new network paradigm that supports evolving business needs while continuing to meet Intel's stringent information security and privacy standards

Authors

Sanjay Rungta

Senior Principal Engineer
Intel IT

Chandra Chitneni

Principal Engineer
Intel IT

Sridhar Mahankali

Principal Engineer
Intel Information Security

Table of Contents

Business Challenge	2
Solution	3
Scaling the Private Cloud Network... ..	3
Upgrading Regional Cloud Interconnects and Scaling the WAN ..	3
Maintaining and Enhancing Network Security Posture	4
Automating Network Services	4
Results	5
Conclusion	5

Executive Summary

Intel's network requirements change constantly, and it is Intel IT's job to keep up. Increasing adoption of public cloud services, enormous data and network bandwidth demand growth, an increasingly complex cybersecurity landscape, and new business models and product lines are shaping a new approach to both our data center networks and Intel's WAN.

At the heart of the new network paradigm is software-defined networking that helps automate networking workflows and consistently maintain a healthy information security posture. Other aspects of our network transformation include the following:

- Upgrading data center networks from 10 and 40 gigabit Ethernet to 100 gigabit Ethernet and increasing network port utilization
- Taking advantage of regional co-location facilities to improve interconnects and scale the WAN
- Establishing flexible but security-conscious controls that support both on-premises and public cloud workloads

While we are in the midst of our network transformation journey, we are already experiencing cost and operational benefits related to scalability, support for multi-cloud, automation, and information security. We believe our multi-cloud network strategy can accommodate Intel's large on-premises private cloud as well as public cloud services so we can seamlessly meet Intel business needs and deliver maximum business value.

Contributor**Mario Villalta**, Industry Engagement Manager, Intel IT**Acronyms**

CaaS	container as a service
GbE	gigabit Ethernet
IaaS	infrastructure as a service
IoT	Internet of Things
PaaS	platform as a service
PoP	point of presence
SaaS	software as a service
SDN	software-defined networking
VRF	virtual routing and forwarding

Business Challenge

Intel's on-premises private cloud is extensive, consisting of 23 data centers worldwide with 56 modules and 297,500 servers with 437 PB of storage.¹ It is designed to meet most of the critical compute, storage, and network needs of Intel's design and manufacturing activities. But our network is in need of transformation due to several internal and external forces (see Figure 1).

The IT industry is rapidly adopting public cloud services, which can often deliver new capabilities more quickly than a private cloud—especially in the areas of container as a service (CaaS), software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). Similarly, Intel is increasingly adopting public cloud services for certain use cases, such as productivity and business applications. As a result, Intel IT has developed a multi-cloud strategy that is flexible enough to encompass both Intel's large on-premises private cloud and the ability to integrate public cloud services. Our goal is to seamlessly serve the needs of our internal customers' (Intel employees and business units) and deliver optimal business value.

¹ IT environment end-of-year 2019.




Figure 1. Internal and external forces are driving network transformation at Intel.

A robust and flexible multi-cloud network strategy must provide the following:

- **Ability to handle today's data explosion.** Intel's global design and manufacturing activity generates enormous amounts of data, which is increasing every year. Data traffic on the network, along with storage needs, are growing at least 30 percent annually. Our corporate network must be able to support the global development life cycle from chip design to manufacturing. At the same time, more generic workloads—such as productivity and office services—are moving to the public cloud. A multi-cloud network has to transport data across the world on Intel's private network and in/out of the public cloud for integration.
- **Cost efficiency and flexibility.** Our on-premises private cloud cost effectively manages Intel's increasing demand for critical compute needs. In fact, Intel saved more than USD 2.8 billion between 2010 and 2018 through the use of private cloud.² However, public cloud service providers can generally provide access to new capabilities more rapidly compared to a private cloud. A multi-cloud network must support the agility to use the right cloud for the right use case, keeping Intel's network costs down while still meeting business needs.
- **Scalability and manageability.** The number of devices at Intel is constantly growing—they range from servers in the data center to multiple employee devices (BYOD). Employees expect pervasive connectivity across all devices, anywhere, anytime. Plus, there are new types of devices, including a multitude of sensors in Intel's factories and a growing number of newly designed smart buildings that use public SaaS for solution delivery. A multi-cloud network must support a wide range of devices and easily integrate new devices—without driving up manageability costs.
- **Versatility and seamless connectivity.** Intel's "network" is not simple. It consists of globally dispersed site networks, data centers, a WAN, and internet connectivity. In addition, Intel's work model has evolved over the years. Instead of a stationary PC and cubicle, more employees use a variety of devices while moving around. Employees need access to the network from all their devices and must be able to access enterprise applications from anywhere (on-premises or off-premises). A multi-cloud network must be able to seamlessly integrate all these devices and be prepared for additional growth.
- **Security.** Protecting users and data is a top priority for Intel IT, and it becomes even more important—and challenging—in a multi-cloud environment. Whether data is on the corporate network, in the data center, or on its way to and from the public cloud, a multi-cloud network must incorporate state-of-the-art security features.

Over the last year, we have developed a multi-cloud network strategy that considers all of these priorities.

² See the Intel IT white paper, "Data Center Strategy Leading Intel's Business Transformation."

Solution

As our network needs evolve, we must continue to optimize our private network so that it remains cost-competitive with public cloud offerings. At the same time, we need to be flexible enough to integrate public cloud services where they make business sense for Intel. The following sections describe how we are scaling both the private cloud network and the WAN to accommodate Intel's multi-cloud world, and using software-defined networking (SDN) to automate network processes.

Scaling the Private Cloud Network

Our vision for the private cloud network is summed up as “a best-in-class network that powers Intel's digital transformation.” Two characteristics are imperative to achieve this goal:

- **Data-driven.** We need a data-aware network that can optimize traffic flow and maintain network quality based on software-defined capabilities (see [Automating Network Services](#) for more information).
- **Intelligent edge.** We want to empower new user experiences from diverse edge and Internet of Things (IoT) devices through scalable connectivity solutions.

As network traffic grows—as much as 15 to 30 percent year-over-year—we see the need to transition the data center network from 10 and 40 GbE to 100 GbE. Data center architecture is transitioning from a large chassis-based model to fixed small-form-factor switches to realize the benefits of faster innovation and cost reduction. To meet the aggregate traffic growth, we are beginning to use 4- to 16-way equal-cost multi-path routing (ECMP) to scale out. In parallel, to support cross-site data transfer, we are transitioning regional networks from 1 or 10 Gbps to 40 or 100 Gbps.

In addition to increasing the network capacity, we are working to increase the effective utilization of network ports. Over the last nine years, port utilization has climbed from 40 percent to 70 percent (a 1.75x increase). To improve network stability and scalability, we are migrating from a predominantly

Layer 2-based to a Layer 3-based network. This new architecture enables us to simultaneously use all available bandwidth on primary and secondary paths, improving capacity utilization. A Layer 3 architecture also eliminates the need for the Spanning Tree Protocol within our data centers, which does not scale well for large networks. Other technologies in use during our network transformation include network overlay, multi-chassis link aggregation, and tunneling to extend Layer 2 across data centers, over the Layer 3 topology.

Upgrading Regional Cloud Interconnects and Scaling the WAN

Due to diverse business and IT productivity use cases, we expect to use an optimized set of IaaS and SaaS providers. Most, if not all, of these cloud service providers have a point of presence (PoP) at key co-location hosting providers that also happen to be a hub for several WAN and internet providers.

To enable high-speed and reliable connectivity between Intel and these cloud service providers, we are establishing regional interconnects—many at 2x 10 Gbps—at key co-location provider locations (see Figure 2). This will allow us to cross-connect to multiple cloud service providers directly in the co-location facility. This approach optimizes cost, time to market, and operational efficiency.

This hub/spoke approach enables faster integration with service providers, while the co-location routers take advantage of virtualization (L2TPv3) to extend the connectivity from cloud service providers to Intel. Dynamic routing is handled using the Border Gateway Protocol (BGP).

We are now enabling the regional multi-cloud capability in the US, where we have a significant concentration of application environments and dependencies. We intend to extend this capability to additional regions by establishing similar hubs of multi-cloud connectivity. In each region, we use a strict vetting process that considers application performance and dependencies relating to co-location of applications and data.




Figure 2. Using carrier-neutral co-location facilities enables us to cost-effectively scale Intel's WAN.

Designing for Multi-Cloud Connectivity

Key features of a multi-cloud-ready network

- Hub/spoke approach
- Scalable to a large number of discrete public cloud accounts
- Comprehends future multi-cloud connectivity
- Enables foundational capabilities that allow enterprise services to gradually migrate from on-premises to public cloud as needed
- Control points for security as Intel Top Secret and external-facing application use cases are accommodated

Several business drivers make SDN a logical choice for the Intel network:

- **Cost efficiency.** Intel is growing, but IT budget and headcount are not—we must do more with fewer people. SDN and automation can help with these challenges.
- **Consistency.** Automated network configuration helps keep network deployment consistent from one data center or site to another around the world.
- **Versatility.** Automation can be deployed across the entire network environment, from office and lab to the data center to the WAN (although technology maturity is not equal across all these environments).

Our vertically integrated SDN platform (see Figure 3) will provide holistic analytics that can help further network transformation. These analytics capabilities include event management, service management, and access management—resulting in cross-domain, data-driven analytics.

We are not strangers to network automation—we have been deploying SDN-ready equipment since 2018. However, to increase efficiency, reduce technical debt, and free up staff for other activities, we are performing less customized automation and using more out-of-the box solutions.

We are pursuing an aggressive four-year adoption schedule (compared to a typical eight-year process), due to significant projected financial benefits. Our focus areas in each year are guided by domain-based technology maturity; for example, SDN technologies for the data center are currently more mature than SDN technologies for the office, lab, and WAN environments. Therefore, the schedule is as follows:

- Year 1 (2019) - Validation and maturing year (waiting for software stack to mature), with a focus on data centers.
- Year 2 (2020) - Full deployment in the data center and further maturation of SDN in the office, lab, and WAN environments.
- Year 3 (2021) - Begin full deployment in the office and lab environments, and a portion of the WAN.
- Year 4 (2022) - Finalize the deployment of the office and lab environments and the WAN.

Maintaining and Enhancing Network Security Posture

In the initial phase of our multi-cloud journey, we made a conscious choice to secure multi-cloud communication flows using on-premises firewalls and security controls. Service isolation is maintained and extended from the co-location-based WAN routers to on-premises routers using capabilities like virtual routing and forwarding (VRF).

Network security controls within a public cloud IaaS environment use a mixture of native network security capabilities provided by the cloud service provider and third-party capabilities, particularly where the cloud service provider capabilities may not meet Intel's security requirements. Our approach is to implement a similar set of security domains in the public cloud IaaS environments as we have on-premises. Any communication flows that cross security domains will be secured by centrally managed firewalls and intrusion detection technologies.

Automating Network Services

SDN technology allows the development of highly programmable networks with the help of workflow automation. SDN is a key enabler for our multi-cloud network strategy; it is our expectation that by aggressively investing in SDN capabilities throughout Intel's network, we can simplify infrastructure management, increase scalability, reduce operational costs, and improve the overall customer service life cycle. The hyperscale cloud service providers have realized these same results through SDN.




Figure 3. We are developing vertically integrated SDN capabilities to optimize the networking environment.

Ongoing Evolution of Network Services

As Intel's business needs evolve, so do network services

Intel's network has changed substantially over the years, as business needs changed and technology matured.

- **Gen 1.** As demand for corporate connectivity increased, we began transitioning from a bridged to a routed network in the early 2000s, with a focus on connectivity. Network speed increased from 100 Mb/s to 1 gigabit Ethernet (GbE).
- **Gen 2.** With the increasing adoption of internet connectivity at Intel, connectivity became even more important in 2010—and so did information security. Network speed increased to 10 GbE.
- **Gen 3 and beyond.** We are creating a data-driven, multi-cloud network with an intelligent edge and network speeds up to 100 GbE. We are now migrating to multiple 100 GbE inter-switch links to keep up with the traffic growth.

This evolutionary path illustrates how we continually meet the needs of our customers. In this way, our network resembles that of a communications service provider. In fact, Intel's multi-cloud network operates at a carrier-grade level.

Automated and templated configuration helps ensure that security configurations can be consistently and quickly enforced. But to reap the full benefit of automated configuration, we need to maintain tight control of changes in the templates' security configuration elements. This involves a formalized change and review process so we can consistently maintain a strong security posture. Additionally, our data center SDN solution enables a centralized and dynamically updated view of network device vulnerabilities and provides a quicker path to remediation. This significantly improves our visibility into data center network OS vulnerabilities as they are discovered.

Results

- **Scaling.** While most of our scaling efforts to date have been focused on the data center, we are also making strides in scaling the WAN. We have already established a US West Coast PoP and anticipate building similar hubs elsewhere.
- **Multi-cloud.** When a business use case arises where the public cloud provides value, such as for SaaS or IaaS, we help the Intel business unit decide which provider is best suited for the use case. Currently, we have defined two or three CaaS/IaaS/PaaS providers and consume services from several SaaS providers for enabling richer collaboration, employee productivity, sales and marketing framework, personnel benefits, and several more business capabilities

- **Automation.** SDN deployment is underway in Intel's Design data centers, and we are already enabling the configuration consistency and dynamic centralized visibility into the SDN-enabled part of the network. We expect to extend this framework to Enterprise data centers in 2020.
- **Security.** We are establishing demilitarized zones in the public cloud networks that we use, similar to what we have done for the corporate network. We are establishing automated guardrails like "no internet-facing instantiation in internal-facing security domains" to further secure our environment while enabling business flexibility.

Conclusion

As Intel begins another 50 years of technology advancement, it relies on Intel IT to provide the necessary networking capabilities to support the business. Our network strategy accommodates exploding amounts of design, testing, and manufacturing data on our private cloud network along with migration of certain enterprise workloads to a multi-cloud environment. Flexible yet comprehensive information security controls help protect users and data, while SDN automates networking workflows and increases operational efficiency. We are excited about our network transformation efforts, which enable Intel to continually transform the world through innovation and collaboration.

Related Content

If you liked this paper, you may also be interested in these related papers:

- Intel IT's Multi-Cloud Strategy: Focused on the Business
- Taking Enterprise Security Beyond the Edge
- Securing the Cloud for Enterprise Workloads: The Journey Continues
- Data Center Strategy Leading Intel's Business Transformation
- Security Architecture Enables Intel's Digital Transformation

For more information on Intel IT best practices, visit intel.com/IT.

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: Improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation on [Twitter](#) or [LinkedIn](#).



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

The information provided in this paper is intended to be general in nature and is not specific guidance. Recommendations (including potential cost savings) are based upon intel's experience and are estimates only. Intel does not guarantee or warrant others will obtain similar results.

Information in this document is provided in connection with intel products and services. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in intel's terms and conditions of sale for such products, intel assumes no liability whatsoever and intel disclaims any express or implied warranty, relating to sale and/or use of intel products and services including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

© Intel Corporation. All rights reserved. 1221/WWES/KC/PDF 341841-002