Gartner.

Hype Cycle for Enterprise Networking, 2020

Published 8 July 2020 - ID G00441509 - 105 min read

By Analysts Andrew Lerner, Danellie Young

While 5G and SD-WAN remain highly hyped, COVID-19 instantly shifted I&O's priorities, creating the need to dynamically, rapidly and temporarily scale to support a large remote workforce. This accelerated interest in cloud-based and consumption-based networking and security technologies.

Analysis

What You Need to Know

The COVID-19 pandemic instantly shuffled enterprise networking priorities. This accelerated hype and interest levels surrounding cloud-delivered networking technologies that enable rapid scale and offer consumption-driven pricing, including zero trust network access (ZTNA) and secure access service edge (SASE). Further, this demonstrated that I&O leaders cannot simply build out large network infrastructures and wait for demand. Instead, to remain competitive, enterprises must deliver just the right amount of network functionality and services, at the right time, and at the right cost. For more information about how peer I&O leaders view the technologies aligned with this Hype Cycle, please see "2020-2022 Emerging Technology Roadmap for Large Enterprises."

The Hype Cycle

Networking technologies are being driven heavily by digitization and cloud computing, and this research describes the 30 most hyped technologies/innovations in networking. For each, we define it and identify (1) the value to enterprises, (2) the current level of adoption and (3) anticipated rate of growth. I&O leaders should use this research to determine if and/or when to invest in these technologies. Overall, the technologies covered in this Hype Cycle changed approximately 40% compared to 2019. This is the second year in a row in which the networking Hype Cycle has changed by approximately 40%, after a decade of mostly incremental change.

Six technologies new to this Hype Cycle include telemetry-based routing, Wi-Fi 7, SD-branch, function accelerator cards (FACs), ZTNA and standard private LTE. Thirty-percent of technologies are on the rise,

climbing toward peak hype, including service mesh, SD-branch and standard private LTE. Approximately 25% are at or very near the peak including SASE, 5G and multicloud networking. Another 25% are at, or sliding into, the Trough of Disillusionment, including intent-based networking and NFV/uCPE. The remaining 20% of the technologies are on the plateau including SD-WAN, and Ethernet switching fabrics.

Several technologies are moving through the Hype Cycle rapidly including SASE and Wi-Fi 6, as more vendors offer and market these capabilities. Intent-based networking is in the Trough of Disillusionment as tangible real-world benefits and adoption are limited; while SD-WAN continues rapid movement as a mainstream technology. Software-defined networking (SDN) is obsolete as true SDN technologies (not just technologies marketed as SDN) have not achieved any significant market traction.

Figure 1. Hype Cycle for Enterprise Networking, 2020

The Priority Matrix

The Priority Matrix for this Hype Cycle shows networking technologies, by impact, against a timeline for the number of years until mainstream adoption. It is useful for ranking which technologies an organization should examine first based on maturity and business impact. The 2020 enterprise networking priority matrix indicates that networking technologies with potentially transformational benefits will not reach mainstream adoption in the near term (less than two years). However, FACs (NextGen SmartNICs), SASE, low-earth-orbit satellite systems and named data networking remain transformational technologies worth noting.

Therefore, I&O leaders adopting these elements should be either in more mature organizations or explicitly willing to assume the risks of earlier adoption, perhaps starting by piloting them in parts of their organization. Less mature and risk-averse organizations should focus on the business impact to network transformation and consider more proven technologies such as Ethernet switching fabric and SD-WAN.

Gartner continues to observe the intersection of old and new innovations, especially in the areas of agility and accessibility, with mainstream adoption expected over the next two to five years. This is apparent on this Hype Cycle, with technologies such as 5G services, network automation and Wi-Fi 6 (802.11ax) showing as highly beneficial in the same time frame. Similarly, networking will see moderate impact from identity-based segmentation (microsegmentation), Kubernetes, SD-branch, and zero trust network access among others, although IPv6 is of lower benefit.

Figure 2. Priority Matrix for Enterprise Networking, 2020

Priority Matrix for Enterprise Networking, 2020

benefit	years to mainstream adoption			
	less than two years	two to five years	five to 10 years	more than 10 years
transformational			FACs (NextGen SmartNICs) Low Earth Orbit Satellite Systems Secure Access Service Edge (SASE)	Named Data Networking
high	SD-WAN	Network Automation Wi-Fi 6 (802.11ax)	NVMe-oF Standard Private LTE	
moderate	Ethernet Switching Fabric	400 Gbps Ethernet 5G Services Cloud-Managed Networks Enhanced Internet Services Identity-Based Segmentation (Microsegmentation) Kubernetes Networking Multicloud Networking Network On-Demand Services SD-Branch Service Mesh Software-Defined Cloud Interconnect Zero Trust Network Access	Edge Networking Intent-Based Networking NFV and uCPE Open Networking/OCP Networking Telemetry-Based Routing Wi-Fi 7 (802.11be)	
low			IPv6	

As of July 2020

Source: Gartner
ID: 441509

Off the Hype Cycle

Technologies omitted from this year's Hype Cycle include:

- Hybrid WAN — A valuable and mature technology that is widely deployed in the enterprise, thus is no longer one of the most hyped technologies in networking.

- 802.11ac Wave 2 — A valuable and mature technology that is widely deployed in the enterprise, thus no longer one of the most hyped technologies in networking.

- 2.5/5 Gbps Ethernet — A valuable and mature technology that is widely deployed in the enterprise, thus no longer one of the most hyped technologies in networking.

- 25/50 Gbps Ethernet — A valuable and mature technology that is widely deployed in the enterprise, thus no longer one of the most hyped technologies in networking.

- FWaaS — This is covered in security-related Hype Cycles.

- SRE — A valuable and hyped technology that is covered in other Hype Cycles.

- uCPE — This has been consolidated and is now covered under uCPE and NFV.

- NPMD — A valuable but mature technology that is covered in other Hype Cycles.

- Brite-Box Switching — A valuable technology that is no longer one of the most hyped technologies in networking.

Technologies that were renamed in this year's Hype Cycle include:

- Enhanced Internet Services — Renamed from "enhanced internet delivery" to "enhanced internet services" to better align with client terminology.

- OCP — Renamed to "open networking/OCP networking" to cover a broader set of technologies relevant to clients including brite-box switching and SONiC.

- NFV — Renamed to "uCPE and NFV" to capture the most hyped aspects of the technology.

- Microsegmentaton — Renamed to "identify-based segmentation" to align with market trends toward identity-based zero trust network security postures and to more accurately reflect the need to define segmentation strategies for modern applications based on identities of the workloads, applications, services and microservices instead of network segment location or IP addresses.

On the Rise

FACs (NextGen SmartNICs)

Analysis By: Anushree Verma

Definition: Function accelerator cards (FACs) are a class of network interface hardware that help improve and accelerate server availability, bandwidth performance and data transport efficiency in a network, besides enabling connectivity to a network. While all FACs are essentially NICs, not all NICs/SmartNICs are FACs. It comes with an in-built processor, onboard memory and peripheral interfaces, and is deployed either as an ASIC, an FPGA or an SoC and, hence, should be programmable in most cases.

Position and Adoption Speed Justification: FACs have started growing in popularity since end of 2019, and these have just started shipping in volumes from early 2020. They accelerate a variety of network-specific capabilities including network services, security and storage functions. FACs will play a critical role in modern cloud architecture and software-defined networking .Traditional NICs typically depend on servers to run complex networking stacks, such as encryption for security or load balancing, etc., besides the basic network processing tasks. This takes away the processing capacity from VMs as the CPU time required to process each packet adds to the latency per packet before transferring it to a NIC. This results in latency issues. By offloading a whole host of network functions, services including security and storage from the server, FACs, improve server availability, bandwidth performance, and data transport efficiency besides enabling connectivity to a network.

However, these are still being adopted by hyperscale or other cloud service providers for various use cases now such as Amazon Web Services (AWS) Nitro System. They can be very useful at the edge where there are varying workload scenarios. Enterprises have minimal adoption until now. Given the benefits, FACs is carving out a separate market for itself, which will cannibalize the existing NIC market as well. By 2024, we estimate that one in three 10G+ and higher NICs shipped will be a FAC. FACs are estimated to grow at about 119% compound annual growth rate (CAGR; 2024-2019) while enterprises who have just started adopting it is estimated to grow at 115% CAGR (2024-2019).

User Advice: Discuss FAC options with network vendors or procure it directly and customize if you operate at a massive scale in order to optimize your server availability. This can be done through techniques such as load balancing to split workload across multiple physical or virtual servers and offloading of different server functions such as security (SSL), compression and virtualization.

It can be deployed either as a stand-alone device or integrated onto the motherboard as well.

However, not all SmartNICs available in the market are FACs and lead to server optimization so careful evaluation according to the use case is imperative.

Business Impact: FACs help improve:

1. Server availability and thus optimize resource utilization (CPU cores, memory, storage and network) with lower latency and jitter.

2. Flexibility due to programmability of FACs for a nondisruptive upgrade or adaptability for varying workloads whenever required

This will thus provide new, more agile systems which lead to overall network scalability and efficiency.

Benefit Rating: Transformational

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Broadcom; Ethernity Networks; Mellanox Technologies; Pensando Systems

Recommended Reading: "Market Trends: Function Accelerator Cards Disrupting Traditional Ethernet Adapter Market"

Named Data Networking

Analysis By: Sylvain Fabre

Definition: Named data networking is an architecture for the future internet that associates unique names (similar to URLs) to blocks of data that can be stored, digitally signed and transmitted across nodes. The current system transports data containers (packets) between two endpoints using an IP address.

Position and Adoption Speed Justification: The National Science Foundation (NSF) began development of named data networking (NDN) in 2010 and it evolved into a consortium of academic and industry members.

Named data networking is an architecture for the future internet. NDN follows the information-centric networking (ICN) approach and is in a development stage. The key pillars of the architecture and basic protocols have been defined. A global research community is testing them in a test bed of 39 nodes (as per WUSTL, May 2020) installed at universities across North America, Europe and Asia. This research community and leading vendors are developing additional architectural elements and more efficient implementations. NDN will also provide an architecture with the ability to effectively and securely support broadcast/multicast of content (with native caching) and the Internet of Things.

Being a clean-slate design, NDN has the opportunity to introduce leapfrog innovation in networking, but this might also hinder its adoption. As IPv6 has demonstrated, introducing large-scale changes in the network is not easy. There are no commercial implementations of NDN yet and we have not seen interest from

enterprises. NDN has been considered as one of the technologies for the next generation of mobile networks (5G), but so far no sign of adoption in the current 5G commercial implementations (which started back in 2018). Networking vendors and service providers should evaluate how NDN could impact their portfolio and strategy in the long term. They should evaluate the pros and cons of playing an early innovator role in NDN development.

Verizon and Cisco tested using named data as opposed to location identifiers like IP addresses to simplify network architecture. The use case was video delivery optimization. Expected NDN benefits include dynamic adaptive streaming, and load balancing of media. Also, forwarding and caching of content in the network can increase traffic localization and free up backhaul and core via bandwidth saving leveraging multicasting. As NDN is agnostic of the access layer, it also enables a more efficient handoff between different access methods, such as 5G and Wi-Fi — a coexistence which will continue as 5G gets rolled out.

Gartner expects the use cases, as well as the benefits of NDN to complement that of edge computing, at least regarding caching, video optimization, application acceleration and bandwidth saving, but adoption remains uncertain even within a 6G time frame. However, the Named Data Networking Consortium has several members from the telecom industry (see Sample Vendors section below) which does indicate interest from the vendor community.

User Advice: Although enterprises do not need to get involved at this stage, NDN has the potential to drive the first real technology change in the set of protocols that underpin the internet, which were designed in the 1970s. This advancement will overcome intrinsic limitations in areas like scalability, security, efficient content distribution and mobility support. Until recently, the industry has elected an incremental approach to address shortcomings, remediating specific problems as they arose. But this also introduced increasing complexity, making networks hard to design and operate. This strategy of patching the architecture to fix each problem independently will not work forever and will not support long-term (greater than 10 years) expansion for the IT and telco industries.

However, some concerns remain as to how likely adoption of NDN will be:

- There is a risk of NDN vanishing before maturing because it does not prove value in terms of enabling capabilities that can't already be done — in addition, introducing an internetwide replacement for IP could take decades even if it was useful.

- Anything involving a naming scheme would require agreement on naming conventions and standardization, which could lead to a scenario where NDN and TCP/IP coexist for 30 years.

NDN needs to confirm that it can add value to 5G — or at least 6G — for example with respect to video and other content delivery, as well as helping tackle issues created by mass densification associated with mmWave deployments as well as heterogenous networks across different access technologies.

Business Impact: NDN addresses some of the fundamental weaknesses of the internet, for example, those around security and efficient content distribution. It will enable completely new services and applications, creating new markets and companies like the internet did more than 20 years ago. But it will also need to support current services and applications. To succeed, NDN will need strong support not only from academia, but also from the IT and telco industries — something not to be taken for granted considering that such a high level of innovation could reshape the competitive landscape. If NDN succeeds, it will have a major impact on the networking industry (including all IT vendors) because NDN changes everything, not just routers and switches.

Benefit Rating: Transformational

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Sample Vendors: Cisco; Fujitsu; Huawei; Intel; Juniper Networks; Panasonic; Viasat

Wi-Fi 7 (802.11be)

Analysis By: Tim Zimmerman; Mike Toussaint; Bill Menezes

Definition: IEEE 802.11be is the next amendment proposal to IEEE 802.11 for advances in WLAN. The amendment builds upon the existing IEEE 802.11ax standard and potentially will be called Wi-Fi 7.

Position and Adoption Speed Justification: IEEE, the standards body of 802.11-based technologies, started the Extremely High Throughput (EHT) Working Group in September 2018. The new amendment will use the unlicensed spectrum newly allocated at 6 GHz to create more efficient use of noncontiguous spectrum and allow 320 MHz bandwidth and 16 spatial streams. This will boost theoretical performance with higher speeds upward of 30 Gbps, increase WLAN capacities, and lower latency while maintaining backward compatibility with existing 802.11 standards. Since the amendment is using existing technical functionality but allowing for more spatial streams to coexist and wider bandwidth for higher performance, we expect to move through the Hype Cycle much faster than other amendments.

User Advice:

- Clients must review their use cases as 802.11ac Wave 2 or 802.11ax capabilities can support the vast majority of enterprise use cases.

- Advanced use cases for immersive technologies such as virtual reality and 8K video streams are expected to proliferate over the next three to five years.

- The amendment is only in the IEEE Working Group.

- End users are advised to follow the progress of the standard, but products will not be available for another 12 to 18 months.

Business Impact: Wi-Fi 7 (802.11be) will focus on the deployment of video traffic and the throughput requirements for 4K and 8K video streams, though many applications may be addressed using Wi-Fi 6. New high-throughput, low-latency immersive applications such as virtual or augmented reality will continue to be developed for the enterprise market.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Sample Vendors: Broadcom; Cisco; Extreme; Hewlett Packard Enterprise (Aruba); Huawei; Intel; Juniper Networks (Mist Systems)

Recommended Reading: "Magic Quadrant for Wired and Wireless LAN Access Infrastructure"

"Critical Capabilities for Wired and Wireless LAN Access Infrastructure"

Enhanced Internet Services

Analysis By: Mark Fabbi

Definition: Enhanced internet services are a collection of over-the-top (OTT) internet-based service offerings to improve the reliability and performance of internet-based traffic. Enhanced internet services include features such as optimized, dynamic path determination, performance optimization and secure, automated tunnel creation. Enhanced internet services may be sourced by enterprise customers or built into SaaS services, and are increasingly relevant due to growing cloud adoption.

Position and Adoption Speed Justification: Enhanced internet services are still relatively nascent, though similar security services such as security web gateway as a service (SWGaaS), popularized by Zscaler, have been around for a number of years. We are observing increasing interest in enhanced internet services (such as with Mode networks, Teridion and others), as a growing proportion of enterprise applications are deployed in the cloud and rely on internet services for access. As the internet becomes an increasingly integral component of a corporate WAN strategy, the need for a more assured, consistent internet experience also increases. Architecturally, it makes sense to deliver network functions OTT, as it aligns with emerging traffic patterns. With the rapid adoption of software-defined WAN (SD-WAN) allowing

direct internet access for cloud-delivered corporate applications, enterprises are starting to explore ways of achieving better application performance and reliability across internet connections. Enhanced internet services address the internet performance and reliability gap between SD-WAN and cloud-delivered applications especially for geographically distributed organizations. Other approaches to provide more consistent performance include network resident gateways offered by many major carriers, and SD-WAN service-based solutions such as from Aryaka and Cato Networks.

It should be noted that enhanced internet services do not refer to security services such as Zscaler, Cisco Umbrella or Palo Alto Networks' Prisma Access. While similar in concept (internet traffic directed to a cloud resident service offering) the services provided are different. Enhanced internet services focus on enhanced routing, optimization and control, while OTT security services enforce corporate security policies on internet traffic, most typically for non-mission-critical or unknown internet destinations. It should also be noted that some of the security services may use enhanced internet techniques across their own backbone networks. Enhanced internet services also exclude cloud-based management of on-premises network infrastructure. Several functions that have traditionally been delivered as dedicated markets or features in other markets will make sense to deliver OTT, such as IPv6 gateways, WAF, DNS and global redirection services. In the long term, it is possible to see more service consolidation across related OTT service offerings such as enhanced internet delivery, cloud-based security services and CDNs. Consolidation may occur due to acquisition or internal development.

User Advice: Enterprises with global or extended regional footprints that are increasing usage of public cloud-based applications should consider OTT enhanced internet services.

Business Impact: Enhanced internet services enable enterprises to improve reliability and performance for cloud resident applications. It is particularly relevant for organizations using multiple cloud services with a distributed user footprint (global or panregion). These services may also be of value to midsize enterprises moving to an all-internet-based WAN transport.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Anapaya; Apcela; Expereo; Fastly; Instart Logic; Mode; Teridion

SD-Branch

Analysis By: Andrew Lerner

Definition: SD-branch products allow multiple branch network functions to be managed as a single

construct. SD-branch products must support:

1. Four network functions; WAN gateways, wired switching, WLAN and firewalls.

1. Unified configuration, policy, reporting, visibility and automation across the four functions, via a single console.

2. Zero-touch configuration for initial provisioning, and automated operational tasks such as troubleshooting, reporting and typical moves/adds/changes.

3. A fully supported, documented and published API.

Position and Adoption Speed Justification: SD-branch technology enables enterprises to deploy a single policy in a central manner that is automatically deployed to multiple devices at a location, and to multiple locations. The notion of unified management of branch office network infrastructure is not new; in prior years, this has been referred to as BOB (branch office in a box). However, we are seeing increasing interest in this concept being referred to now as SD-branch, both by end-user clients and vendors who market the term extensively.

The enterprise need for SD-branch generally stems from the need for greater simplicity, agility and increased security via consistency of policies across branch devices. SD-branch solutions can be single vendor or multivendor; however, in practice, we anticipate that most enterprises will ultimately aim to have SD-branch deployments be single vendor. SD-branch solutions can be sourced directly via enterprises in a DIY fashion and/or procured from service providers who offer it as a turnkey managed service.

SD-branch is complementary to SD-WAN as SD-branch solutions can manage SD-WAN products.

User Advice:

- Use SD-branch products if operational changes to branch networking equipment such as gateways, WLAN or switches are a substantial challenge or drain on personnel.

- Shortlist multiple suppliers and run a real-world pilot to validate functionality within your environment before making a long-term or strategic purchase from any specific vendor.

- Don't assume that all products marketed as SD-branch actually provide central policy, configuration and troubleshooting across multiple functions in a simplified manner. In other words, be wary of overmarketing of the term.

- When evaluating SD-branch solutions, prioritize the specific network function that is of highest importance for your environment, and focus on that function. Not all vendors have level feature capability

across all functions.

- Prefer SD-branch products that allow for management of multivendor environments to aid with transition and/or to protect existing investments.

- Include SD-branch solutions when refreshing branch equipment, but do not require that all functions are replaced at once.

Business Impact: With SD-branch, IT staff can manage branch infrastructure in a simplified and more efficient way. This can reduce opex costs for branch networking/security equipment, and increase uptime/availability in remote branches due to simplified and integrated policy and UI. This also helps to add new locations quickly and adding new apps/user groups easily. This will be most compelling for smaller, distributed branch locations where it is challenging to get network personnel on-site.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Cisco; Fortinet; Hewlett Packard Enterprise; Huawei; Juniper Networks; Versa Networks

Recommended Reading: "Magic Quadrant for WAN Edge Infrastructure"

"Critical Capabilities for WAN Edge Infrastructure"

Software-Defined Cloud Interconnect

Analysis By: Lisa Pierce

Definition: Software-defined cloud interconnect (SDCI) is a form of structured network connectivity between the enterprise and public cloud service provider (CSP). It employs a WAN connection to SDCI hubs, who aggregate and intermediate connectivity between the enterprise site and CSP. SDCIs preprovision connectivity from their hubs to public CSPs (IaaS, PaaS, SaaS providers). SDCI's highest value is in the breadth and rapidity of connectivity available to many (dozens to hundreds) of CSPs across a broad range of global locations.

Position and Adoption Speed Justification: Although less than 10% of enterprises have adopted SDCI, we anticipate adoption will rise concurrent with their proactive efforts to rationalize multicloud connectivity to 25% of enterprises by YE23. We believe this because selecting the best method of connecting to public CSPs is essential to optimizing the end-user experience, performance and cost. Many clients find the range

of potential choices to connect their sites to CSPs to be overwhelming. In addition to SDCI, clients can employ: (1) ad hoc internet connectivity, (2) carrier-based cloud interconnect (like AT&T NetBond), (3) colocation hubs like Equinix, or (4) hyperscale cloud provider native connectivity offers (like AWS Direct Connect). However, SDCIs are one of the few solutions Gartner recommends to connect to multiple CSPs. Today, most enterprises connect to 20+ public CSPs — use of multiple CSPs will continue: by YE23, Gartner forecasts that almost 25% of all application services spending will be spent on SaaS, and 41% of all computing and storage spending will be spent on IaaS. (See "Forecast: Public Cloud Services, Worldwide, 2017-2023, 4Q19 Update" and "Forecast Alert: IT Spending, Worldwide, 4Q19 Update.")

User Advice: Today, four critical considerations all enterprises must address include:

- Getting there: Today, enterprises must separately order WAN connectivity to the SDCI hub from carriers and ISPs; providers who offer SDCI functionality don't include this in their service. But an enterprise provisions CSP connectivity once per site: upon connecting to the SDCI hub, the SDCI provider thereafter provides flexible connectivity to multiple CSPs. With increasing frequency, colocation providers offer or host SDCI functionality, which simplifies WAN connectivity architecture.

- Hybrid complications: Assure that end-to-end reliability, security and performance meet your requirements in hybrid environments like (1) mixed computing and applications environments, and (2) multicloud applications. Clients who use carrier-based SDCI will often find that mixing and matching CSP is supported; mixing and matching carriers often will not.

- Governance and compliance: In addition to proving compliance with prominent security and industry requirements, depending upon regulatory circumstances, the SDCI provider must also meet data residency and applications accessibility requirements.

- Alternate solutions: A growing number of providers are offering connectivity to multiple cloud providers, including enhanced internet backbone providers and IaaS providers. Ultimately SDCI is an enabling technology used by many types of providers who want to support flexible connectivity to multiple cloud providers, it is not a stand-alone category of service provider.

Business Impact: SDCI's flexibility allows enterprises to connect the CSPs on an ad hoc basis, all while providing superior performance and compliance. SDCI allows for rapid provisioning via programmable controls and a centralized dashboard to respond to changing demands of organizations and their business. In addition, many SDCI providers offer value-added services, such as security features and inter-SDCI hub WAN transport on an intraregional and interregional basis.

Key benefits SDCI solutions provide include: (1) a wider array of public CSP choice than carrier-based cloud interconnect services, which typically interconnect to fewer than a dozen public CSPs; (2) offering more rapid and flexible interconnectivity choices than just cross-connection-based colocation offers, allowing clients to connect to/disconnect from CSPs in minutes, rather than days; and (3) compared to solutions from

CSPs themselves (like AWS Direct Connect), offering a uniform method to connect to multiple CSPs (versus one). In contrast, using the public internet cannot assure consistent end-to-end performance, and can impair and complicate efforts to assure adequate reliability and governance requirements are consistently met.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Aryaka; Cato Networks; CoreSite; Equinix; InterCloud; Megaport; PacketFabric; PCCW Global; Teridion; Unitas Global

Recommended Reading: "Five Key Factors to Prepare Your WAN for Multicloud Connectivity"

"Comparing Network Architectures for Interconnecting Data Centers and Clouds"

At the Peak

Low Earth Orbit Satellite Systems

Analysis By: Bill Menezes

Definition: Low earth orbit satellites currently provide global or regional narrowband voice and data network services, including to areas with little or no existing terrestrial network coverage. LEO systems operate at lower altitudes (1,200 miles or less) than predominant geostationary systems (about 22,000 miles). Planned LEO broadband systems of up to thousands of satellites could support significantly lower latency and, depending on system technology, broadband data speeds that are greater than throughput of current GEO and LEO systems.

Position and Adoption Speed Justification: Demand for low earth orbit (LEO) services is well-defined. A further growth driver is the lower cost of launching smaller LEO satellites, which can cost around $1 million compared with the $50 to $100 million cost of a geostationary (GEO) satellite, and can be clustered on a single launch rocket. But newer and planned systems will move slowly through the Hype Cycle given the lengthy time frames to plan and deploy them and to develop inexpensive directional antennas. Only about 53% of households globally had internet access at YE19, according to the International Telecommunication Union. Lack of broadband access hinders economic growth, thereby limiting enterprise business potential in underserved regions. However, not all of the next-generation LEO systems planned for extending broadband to at least 60% of the world's population in the coming decade will come to fruition. One of the systems closest to commercial deployment — OneWeb's 648-satellite, LEO constellation — halted with

OneWeb's March 27 bankruptcy filing.

User Advice: Enterprises with current or planned business interests in remote or underserved global regions should closely follow LEO system development to align narrowband and broadband connectivity requirements for IoT and general data networking with technology capabilities and service availability. Among planned systems:

- OneWeb had launched 74 of its satellites by the time it declared bankruptcy. During 2Q19 the company began the process of trying to sell its radio spectrum, potentially enabling a successor to take over its plan. The system targets downlink speeds of multiple Gbps at round-trip latency of 10 ms to 30 ms. Terminals for fixed and mobile applications would provide a broadband satellite connection plus 2G, 3G and 4G LTE device connections.

- SpaceX venture Starlink as of 2Q19 had launched about 350 of its satellites, with plans to operate a 4,425-satcom LEO constellation providing global broadband internet access. The company scheduled launches of 180 additional satellites in 2020, targeting possible initial commercial service to North America by year's end. Global commercial broadband coverage is targeted for 2024. The full constellation will require more than 100 successful launches.

- Telesat, which already operates a number of GEO satellites, plans a constellation of 300 LEO satellites to provide global broadband connectivity. The Canadian company launched a test satellite in January 2018, targeting full commercial service in 2023.

- Amazon's "Project Kuiper" is seeking regulatory permission for a LEO constellation of 3,236 satellites to provide broadband internet access to underserved areas as well as to Amazon data-dependent services such as AWS. As of 2Q20, there was no announced service date.

Other planned or existing LEO services include those supporting narrowband data for IoT and digital imaging, or satellite phone and messaging services.

Business Impact: Planned LEO satellites will enable broadband connectivity for all remote or underserved geographies for consumer or enterprise use cases. These services also will be able to provide low-latency backhaul for terrestrial technologies such as remote cellular towers and Wi-Fi hotspots. It will possibly spur new development of those networks in areas where high costs have prevented fiber or other wired WAN backhaul connections. These systems will require customer infrastructure such as directional or phased-array antennas at manufacturing volumes large enough to make them cost-effective.

Just as significant is the use case for narrowband IoT, which requires simpler, less expensive ground antennas and only intermittent connectivity to endpoints or gateways that may be served with a small number of satellites rather than global constellations. Enterprise also may benefit for LEO backhaul for private networking in use cases such as manufacturing, healthcare or natural resources.

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Amazon; Iridium; ORBCOMM; SpaceX; Telesat

Recommended Reading: "Market Insight: New Satellite Constellations Enable Revenue Opportunity for CSPs to Complement IoT Connectivity Services"

"Satellite Communications Strengthen Resilience Planning"

"Market Trends: New Satellite Constellations Will Provide Revolutionary Opportunities for Connecting the IoT"

Standard Private LTE

Analysis By: Sylvain Fabre; Tim Zimmerman

Definition: Private LTE is defined as a private wireless network based on Long Term Evolution (LTE) technology infrastructure to interconnect people/things in an enterprise business. Private LTE offering can include voice, video, messaging, broadband data and IoT/M2M features. Access is only available to the enterprise own devices with authorized SIM cards, increasing security, with better and predictable performance as the infrastructure is not publicly shared. .

Position and Adoption Speed Justification: Private LTE is the 16th most-searched term among networking technologies considered for Gartner Hype Cycle, which is a composite metric based on Gartner search, Gartner inquiry, and Google trends.

There is an increasing interest in deploying private LTE networks by enterprise business customers, in various vertical industries such as manufacturing, mining, oil, utility and railroad companies; IoT services, university campuses, stadiums/arenas etc., to support their growing business, and emergency services.

Private LTE networks are designed and deployed with a focus on high capacity, high speed and high security capabilities to increase operational efficiencies of business customers. Enterprise customers are using private LTE networks to either complement or replace Wi-Fi or their traditional private radio technologies such as PMR, VHF/UHF and others. Private LTE is also being deployed to replace traditional wireline technologies such as Ethernet and Multiprotocol Label Switching (MPLS) for remote office connectivity. Unlike the traditional private radio networks, private LTE can also support growing M2M/IoT

communications. CSPs are increasingly focused on investing in private LTE networks. CSP suppliers are also focused in developing private LTE solutions to offer directly to enterprise customers or offer through CSPs. With the introduction of LTE-U and Citizens Broadband Radio Service (CBRS, a spectrum band from 3.5 GHz to 3.7 GHz in the USA) — shared public spectrum, private LTE is shaping up as a major scenario for enterprise mobility. CBRS testing started in a number of U.S. cities across 36 states; unlike other unlicensed spectrum options that can get crowded, access is dynamically managed by the SASs to is minimize interference.

U.K. and Germany have already allocated spectrum in the same band as CBRS, and that other bands are expected (such as 6GHz).

From an IoT/M2M perspective, the roadmap from 4G to 5G promised capabilities requires a much higher density of endpoints, up to 1 million per sq. km.

User Advice: Private LTE is one of the opportunities for CSPs to grow their revenue. CSPs should prepare a strategic plan to invest in private LTE infrastructure and grow revenue over the next several years. A major advantage for CSPs is that private LTE networks can be deployed using their existing LTE infrastructure with very minimal additional investments. Private LTE service should have the following features: On-demand high-speed LTE access, on-demand high capacity LTE access, high-security features such as data encryption, eliminating internet-based traffic when needed, multifactor authentication (MFA), and secure/encrypted devices. Private LTE offering should also have features like mobile worker management, push-to-talk, push-to-video, and other business applications targeting various vertical enterprise customers. Private LTE offering should also have dedicated networks such as LTE-M, etc., to support IoT/M2M business requirements.

With the introduction of LTE-U and Citizens Broadband Radio Service (CBRS) — shared public spectrum, private LTE market sees CSPs and their suppliers competing directly. CSPs used to have the upper hand in this market thanks to licensed ownership of spectrum resources and they should use those assets to offer a differentiated private LTE service.

IO enterprise leaders should:

- Take advantage of the widening choice of suppliers by including CSPs, Vendors, Integrators in your RFQ for Private LTE.

- Consider the need for off-site mobility and national roaming.

For CSPs delivering Private LTE, we recommend:

- Also include VoLTE, VoWiFi, video, messaging, RCS, fixed-mobile convergence, etc. in their private LTE

offering.

- Use licensed ownership of spectrum resources to offer a differentiated private LTE service

- Differentiate offering by bundling other value-added services such as agile and scalable network with features such as SIM management, self-healing, closed loop automation, national roaming, etc.

Business Impact: The need for private connectivity in industry verticals is growing, as seen is the volume of Gartner inquiries on this topic. Verticals such as mining, oil, railroad, health, logistics, land development, banking, and several other verticals are seeking private connectivity on demand to support their premises. Such connectivity is not just human-to-human, it also includes human-to-things, things-to-things, etc. Private LTE networks is perceived as one aspect of increasing productivity and cost reduction. CSPs should invest in private LTE networks and develop customized offerings for various vertical industries with an ecosystem of partners.

For enterprise IT leaders. we recommend:

- Consider deploying a LTE PMN if your premises don't have access to cellular wireless coverage or wish an alternative supplier to CSPs.

- Include CSP suppliers in your RFQs as they are now a competitors to operators for private LTE solutions.

- Consider CBRS from non-cellular vendors to provide local coverage or additional network capacity for indoor and outdoor wireless services for both rural and urban environments.

- Get vendors to clarify their roadmap from 4G to 5G, such that disruption is minimal and interference managed during coexistence of these technologies.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Athonet; Cambium Networks; Cisco; CommScope; Druid Software; Ericsson; Huawei; Mavenir; NEC; Nokia

Recommended Reading: "Autonomous Things Ecosystems Open Opportunities for IT Services Providers in Telecom"

"Emerging Technologies: Venture Capital Growth Insights for mmWave 5G"

"Magic Quadrant for LTE Network Infrastructure"

"Industry Focus Can Expand Private LTE Footprint"

"Market Insight: Public Safety, First Responder and Mission-Critical Networks Migrate Toward LTE"

Kubernetes Networking

Analysis By: Simon Richard

Definition: Kubernetes (K8s) networking software enables pods and services within a Kubernetes cluster to communicate with each other and with the outside world. The software enables network policies and multitenancy to be implemented without requiring developers to concern themselves with IP address management. Networking products handle communications both inside the K8s pod via a container network interface (CNI) plug-in, and from the outside world via ingress controllers. There is also increased interest in communication between K8s clusters.

Position and Adoption Speed Justification: Kubernetes networking solutions are directly related to the adoption of containers, a practice that is growing rapidly in the enterprise. Kubernetes has become the de facto standard system for container orchestration, and networking Kubernetes requires both architectural and operational planning.

Public cloud providers offer their own natively integrated CNI plug-ins and ingress controllers that integrate with their cloud platform and load balancers, but customers can also deploy third-party plug-ins. For on-premises Kubernetes deployments, third-party CNI plug-in and ingress controllers are necessary.

The technology landscape associated with container networking solutions is fragmented, with many commercial vendor and open-source CNI plug-ins, and we see multiple viable solutions such as those listed on the Kubernetes Documentation. Network virtualization vendors and data center fabric have CNI plug-ins that extend their policy model to Kubernetes applications and that are supported by commercial Kubernetes platforms. On the ingress-controller side, most load balancing as well as some API gateway solutions extend their offerings to provide ingress controller functionality.

We anticipate the increasing use of commercially supported software CNI plug-ins during the next three years, as Kubernetes becomes more prevalent in on-premises enterprise production workloads. Today, CNI solutions implement Kubernetes network policies. Kubernetes network policies also play a role in service mesh technology, which controls service connectivity from Layer 4 to 7.

User Advice: Enable pod-to-pod, service-to-service and service-to-external-world communications.

Invest in data center network switching solutions that provide turnkey, RESTful API-based integration with Kubernetes, because network provisioning is driven outside of dedicated network tooling.

Avoid making Kubernetes-networking decisions in isolation. When sourcing, designing and deploying container networking solutions, ensure that the networking, cloud and development teams are involved as part of a cross-functional effort.

Eliminate manual network provisioning in Kubernetes-based environments. The automation and self-service provisioning of network resources will be required.

Avoid making long-term, proprietary or expensive investments in container networking solutions during the next 12 months, due to the nascent and rapidly evolving state of the market. Most Kubernetes platforms support multiple Kubernetes network solutions, enabling customers to change their CNI plug-ins.

When deploying Kubernetes, look first to extend your deployed network virtualization or data center fabric vendors/solutions from the cloud or data center networking.

Prefer ingress controllers that extend your existing load balancing or API gateway vendors.

Include requirements for turnkey Kubernetes network integration in network RFPs.

Business Impact: Kubernetes networking solutions provide connectivity for applications and services that are deployed in K8s clusters.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Sample Vendors: Cisco; Citrix; F5; Juniper Networks; Nuage Networks; Rancher Labs; Red Hat; Tigera; VMware

Recommended Reading: "Assessing Kubernetes-Based Hybrid Container Management Platforms"

"Containers Will Change Your Data Center Infrastructure and Operations Strategy"

"Innovation Insight for Service Mesh"

"Achieving Kubernetes Operational Readiness to Run Containers in Production"

Edge Networking

Analysis By: Tim Zimmerman; Joe Skorupa

Definition: Edge networking is a diverse set of communication technologies and architectures that connect end-user devices and sensors to the edge of the communication infrastructure. For WAN connected devices, this includes cellular, private cellular, LPWAN and satellite. For on-premises, connectivity includes traditional Ethernet and Wi-Fi, but also over 40 industrial automation protocols and building automation protocols.

Position and Adoption Speed Justification: Edge networking has been traditionally siloed. Cellular, private cellular and LPWAN connectivity are addressed by different vendors as compared to on-premises solutions, which, for many industrial automation protocols, were vendor-specific. The visibility of IoT for security risk and discovery as well as for network convergence and business benefits has also shone a light on the IoT solution architectures.

The focus on end-to-end IoT architecture as part of the overall network strategy, the advances in edge computing and communication technology, and the convergence of networking require integration of edge networking decisions into the overall architecture. Hence, this technology is moving slowly through the Hype Cycle. The flexibility of private cellular (such as CBRS), higher speeds of cellular 5G as well as the device integration of NB-IoT are indicating the convergence of solutions that require WAN coverage. The ability to use private cellular (CBRS or 5G) for on-premises applications, such as large manufacturing facility applications, shows the overlap of WAN with on-premises communication technologies. In the same vein of convergence, the more widely used industrial protocols are moving to an IP-based definition, which allows for coexistence with traditional enterprise communication infrastructures. Additionally, the availability of 802.11ax for local-area networking and the announcement of 802.11be (6 GHz) provide additional control of device connectivity options and higher performance. High-performance on-premises connectivity will allow device communications to flow upstream to edge computing platforms or core networking resources while addressing congestion and higher density. Implementation of the standards associated with Time-Sensitive Networking (TSN) for wired Ethernet connectivity will allow best-effort enterprise infrastructure to provide services that address latency-sensitive applications and promote additional convergence.

User Advice: End users should:

- Be aware that WAN and WLAN/LAN solutions today require different and separate communication infrastructures.

- Invest in shared spectrum communications (such as CBRS, 5G, NB-IoT) for large, open environments that could be indoors (utilities or manufacturing plants) or outdoors. 5G infrastructure is currently being rolled out across different geographies but private LTE (CBRS) can address location-specific requirements. It is imperative to understand whether 5G-specific attributes for throughput, latency and

device density are required by your applications, and if so, whether 5G service is available. Also identify whether NB-IoT or CBRS is available, and whether new devices (which will be required) support the required functionality.

- Invest in WLAN for indoor enterprise applications or defined outdoor environments such as stadiums or parks. Monitor the technology (both 802.11ax and 802.11be); do not pay any premiums for the new functionality and ensure that any short-term purchase is upgradable to the final standard.

- Mandate support for 802.11u to allow migration of applications connectivity from WAN to WLAN infrastructures.

Business Impact: The short-term value proposition will be the consolidation of many different, siloed communication types into one or two with the ability to move between two technologies. This will simplify procurement, deployment and operations. The longer-term value will be commoditization of the hardware and lower prices across all vertical markets globally to achieve the desired business outcome.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Belden (Hirschmann); Cambium Networks; Cisco; Federated Wireless; Hewlett Packard Enterprise; Hewlett Packard Enterprise (Aruba); Orange Business Services; Siemens; Verizon; Vodafone Group

Recommended Reading: "Exploring the Edge: 12 Frontiers of Edge Computing"

Multicloud Networking

Analysis By: Jonathan Forest

Definition: Multicloud networking products address networking requirements across multiple private and public cloud environments. They provide consistent networking policy, network security, governance and network visibility across multiple clouds from a single console without the administrator having to login to all those environments individually. Multicloud networking products are typically delivered as software and can be managed by the enterprise and/or delivered as a service.

Position and Adoption Speed Justification: Adoption of multicloud computing is farther along than multicloud networking. While Gartner clients aren't asking for multicloud networking in great numbers, overall market interest is increasing, as validated by recent Gartner social media conversation analysis that shows that

multicloud networking is a top data center network topic. With varying vendor functionality in terms of depth and breadth of configuration, visibility, and troubleshooting support across cloud service providers, multicloud networking solutions are currently seen as "nice to have" rather than "must have," which has limited its adoption.

Gartner believes that multicloud networking will become increasingly important since the market is evolving from consuming applications from a hybrid of data center and a single cloud provider to consuming from the data center and multiple cloud providers. This increases the complexity to offer the right and consistent level of performance, security, reliability, management, and cost, So, as these environments get more distributed and complex, multicloud networking solutions will become increasingly relevant as vendors expand their capabilities in the next two years.

Having said that, customers who use management consoles from their data center networking vendors to manage networking configurations inside public cloud network infrastructures will remain less than 1% through 2023.

User Advice: Multicloud implementations will need coordination and strategy across the enterprise to identify the type of services needed and deliver the benefits of a cloud environment. Organizations should also take into account security and policy access functionality across all their environments.

While this market is very new and it may be too early to make strategic investment decisions, organizations can implement multicloud computing without multicloud networking. Typically, after moving workloads to the cloud, adding an additional cloud service provider is incremental. At a certain point, the management of individual environments can become too onerous resulting in the need of multicloud networking solutions. Having said that, organizations looking to implement multicloud networking solutions should:

1. Prefer lighter-weight functionality by focusing on solutions that use native cloud-provider APIs versus solutions that require installation of heavy agents or overlays.

2. Confirm which public cloud native networking services are fully supported by validating visibility, troubleshooting and configuration management for VLANs, ACLs, and routing.

3. When cloud-native capabilities do not suffice, fill gaps or address mission-critical business functionality by using multicloud networking solutions.

4. Include your private and public clouds as part of your overall network strategy by defining management requirements and needs for a consistent management framework.

5. Choose the appropriate multicloud network connectivity option by evaluating cost, performance, and security.

Business Impact: Multicloud networking solutions can be data center networking solutions that extend to public cloud, public cloud solutions that extend to on-premises, or from vendors that don't have an installed base in either. There are multiple mechanisms to enable a multicloud networking capability, such as overlay networks, agents, virtual routers, and API-level integrations. Multicloud networking is similar conceptually to data center switching fabrics in that multiple components are managed as a single construct and policy is created centrally, and external programmability is enabled via API.

The ability to manage various cloud environments from a single console ensures consistency, reliability, security and simplicity. Consistent policy enforcement reduces complexity and ultimately better control and governance. This lowers operational support costs and makes it easier to pass audits.

While multicloud compute is more widely deployed, multicloud networking is still nascent. We estimate that multicloud networking solutions are currently used by less than 1% of the market.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Alkira; AppViewX; Arista Networks; Aviatrix; Cisco; Extreme; H3C; Huawei; Juniper Networks; VMware

Recommended Reading: "Best Networking Practices Inside the Public Cloud"

"How to Interconnect With Azure, AWS and Google Backbones"

"Five Key Factors to Prepare Your WAN for Multicloud Connectivity"

NVMe-oF

Analysis By: Julia Palmer; Joseph Unsworth; Joe Skorupa

Definition: Nonvolatile memory express over fabrics (NVMe-oF) is a network protocols that is taking advantage of the parallel-access and low-latency features of NVMe PCIe devices. NVMe-oF enables tunneling the NVMe command set over additional transports beyond PCIe over various networked interfaces to the remote subsystems across a data center network. The specification defines a common protocol interface and is designed to work with high-performance fabric technology including RDMA over Fibre Channel, InfiniBand or Ethernet with RoCEv2, iWARP or TCP.

Position and Adoption Speed Justification: NVMe is a storage protocol that is being used within solid-state arrays and servers. It takes advantage of the latest nonvolatile memory to address the needs of extreme-low-latency workloads and is now broadly deployed. However, NVMe-oF, which requires a storage network, is still emerging and developing at different rates depending on the network encapsulation method. Today, many NVMe-oF offerings that use fifth generation and/or sixth generation Fibre Channel (FC-NVMe) are available, but adoption of NVMe-oF within 25/50/100 Gigabit Ethernet is slower. In November 2018, the NVMe standards body ratified NVMe/TCP as a new transport mechanism. In the future, it's likely that TCP/IP will evolve to be an important data center transport for NVMe-oF. The NVMe-oF protocol can take advantage of high-speed networks and will accelerate the adoption of next-generation storage architectures, such as disaggregated compute, scale-out software-defined storage and hyperconverged and composable infrastructures, bringing super low latency application access to the mainstream enterprise. Unlike server-attached flash storage, shared accelerated NVMe and NVMe-oF- can scale out to high capacity with high-availability features and be managed from a central location, serving dozens of compute clients. Most storage array vendors have already debuted at least one NVMe-oF capable product with nearly all vendors expected to do so during 2021.

User Advice: Buyers should clearly identify workload where the scalability and performance of NVMe-based solutions and NVMe-oF justify the premium cost of an end-to-end NVMe deployment. There are a select variety of highly performant workloads that can utilize the technology, such as AI/ML, high-performance computing (HPC), in-memory databases or transaction processing. Next, identify appropriate potential storage platform, NIC/HBA and network fabric suppliers to verify that interoperability testing has been performed and that reference customers are available. Should buyers not desire the immediate performance gains and associated costs, then they should investigate how simply and nondisruptively their existing products migration path to ensure investment protection for the future.

Most storage vendors already offer solid-state arrays with internal NVMe storage, and during the next 12 months an increasing number of infrastructure vendors will offer support of NVMe-oF connectivity to the compute hosts. HCI vendors will deliver NVMe storage in an integrated offering during the next 12 to 18 months, but customers need to verify the availability of NVMe-oF networks between HCI nodes to see significant performance improvement. Similarly, when customers require NVMe-oF storage networks that encompass switches, host bus adapters (HBAs), and OS kernel drivers, IT infrastructure modernization will be required.

This potential requirements of making major changes to the data center storage networking and servers are slowing down the adoption of NVMe-oF solutions in mainstream enterprises. However, due to the better interoperability and availability of NVMe-oF over Fibre Channel within the next two years, I&O leaders implementing NVMe-oF will likely choose to deploy it within an existing Fibre Channel SAN infrastructure. Investment protection for customers with existing fifth-generation or sixth-generation FC SANs is compelling because customers can implement new fast NVMe storage arrays and connect via NVMe-oF to servers while using the same media. Therefore, old and new storage, network switches and host bus adaptors can

run together in the same FC-based storage network (SAN), with SCSI and NVMe storage separated by zones, as long as compatible fifth- or sixth-generation FC equipment is used. Furthermore, as support for NVMe-oF expands and matures, I&O leaders will have an additional choice of either deploying NVMe-oF with RDMA RoCEv2 or NVMe-oF over TCP/IP-based products to leverage latest Ethernet deployment, thereby easing the transition and providing investment protection.

Business Impact: Today, NVMe SSDs and NVMe-oF offerings can have a dramatic impact on business use cases where low-latency requirements are critical to the bottom line. Though requiring potential infrastructure enhancements, the clear benefits these technologies can provide will immediately attract high-performance computing customers who can quickly show a positive ROI. Designed for all low-latency workloads where performance is a business differentiator, NVMe SSD with NVME-oF will deliver architectures that extend and enhance the capabilities of modern general-purpose solid-state arrays. Most workloads will not need the multimillion IOPS performance that these new technologies offer, but most customers are demanding the lower, consistent response times provided by NVMe-based systems.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Dell EMC; Excelero; Hitachi Vantara; IBM; Kaminario; Lightbits; NetApp; Pavilion Data Systems; Pure Storage; StorCentric

Recommended Reading: "Top 10 Technologies That Will Drive the Future of Infrastructure and Operations"

"2019 Strategic Roadmap for Storage"

"Prepare Your Storage and Data Management Strategy for the Impact of Artificial Intelligence Workloads"

"Critical Capabilities for Solid-State Arrays"

Service Mesh

Analysis By: Anne Thomas

Definition: Service mesh is a distributed computing middleware that optimizes communications between application services. It provides proxy and/or lightweight mediation for service-to-service communications, and supports functions such as authentication, authorization, encryption, service discovery, request routing, load balancing, self-healing recovery and service instrumentation.

Position and Adoption Speed Justification: A service mesh addresses the lightweight middleware requirements of service-to-service communications (east-west), especially among microservices running in managed container systems. These technologies are evolving rapidly. Many commercial and open-source solutions are now generally available.

Hype surrounding service mesh technology accelerated in early 2017 when Google, IBM and Lyft launched the Istio open-source project to produce a service mesh for Kubernetes. Numerous vendors now contribute to the project and provide commercial Istio-based products, and many people associate the service mesh market exclusively with Istio, even though it isn't the most mature product in the market.

Many clients have expressed confusion about the relationship between service meshes and other API mediation technologies, such as API gateways and application delivery controllers (ADCs.) A service mesh is lighter weight, and therefore doesn't replace traditional API mediators. (See "How a Service Mesh Fits Into Your API Mediation Strategy.") Unfortunately, management, federation and interoperability between the various API mediators and service meshes haven't been addressed by the vendor community, yet.

User Advice: Application leaders responsible for API management and microservices middleware should:

- Deliver secure and resilient miniservices and microservices operations by adopting a service mesh.

- Limit code dependence on any particular service mesh technology by favoring approaches that reduce vendor lock-in, such as sidecar proxies (over library-based implementations).

- Reduce cultural issues and turf wars by assigning service mesh ownership to a cross-functional PlatformOps team that solicits input and collaborates with networking, security and development teams.

- Accelerate knowledge transfer and consistent application of security policies by collaborating with I&O and security teams that manage existing API gateways and application delivery controllers.

Business Impact: A service mesh is a powerful piece of middleware that improves development and operations of microservice-based applications. It ensures reliable, resilient and secure service-to-service communications. It provides deep visibility into the services, enabling proactive operations and faster diagnostics. It automates complex communication concerns, thereby improving developer productivity and ensuring that certain standards and policies are enforced consistently across applications.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Amazon Web Services; Buoyant; F5; Google; HashiCorp; Istio; Kong; Microsoft; Netflix; VMware

Recommended Reading: "How a Service Mesh Fits Into Your API Mediation Strategy"

Secure Access Service Edge (SASE)

Analysis By: Joe Skorupa; Neil MacDonald

Definition: Secure access service edge (SASE, pronounced "sassy") delivers multiple capabilities such as SD-WAN, SWG, CASB, NGFW and zero trust network access (ZTNA).

SASE supports branch office and remote worker access. SASE is delivered as a service, and based upon the identity of the device/entity, combined with real-time context and security/compliance policies. Identities can be associated with people, devices, IoT or edge computing locations.

Position and Adoption Speed Justification: SASE is driven by enterprise digital business transformation: the adoption of cloud-based services by distributed and mobile workforces; edge computing and business continuity plans that must include flexible, anywhere, anytime, secure remote access. While the term originated in 2019, the architecture has been deployed by early adopters as early as 2017. By 2024, at least 40% of enterprises will have explicit strategies to adopt SASE, up from less than 1% at year-end 2018.

By 2023, 20% of enterprises will have adopted SWG, CASB, ZTNA and branch FWaaS capabilities from the same vendor, up from less than 5% in 2019. However, today most implementations involve two vendors (SD-WAN + Network Security), although single vendor solutions are appearing. Dual-vendor deployments that have deep cross-vendor integration are highly functional and largely eliminate the need to deploy anything more than a L4 stateful firewall in the branch office. This will drive a new wave of consolidation as vendors struggle to invest to compete in this highly disruptive, rapidly evolving landscape.

SASE is in the early stages of market development but is being actively marketed and developed by the vendor community. Although the term is relatively new, the architectural approach (cloud if you can, on-premises if you must) has been deployed for at least two years. The inversion of networking and network security patterns as users, devices and services leave the traditional enterprise perimeter will transform the competitive landscape for network and network security as a service over the next decade, although the winners and losers will be apparent by 2022. True SASE services are cloud-native — dynamically scalable, globally accessible, typically microservices-based and multitenant. The breadth of services required to fulfill the broad use cases means very few vendors will offer a complete solution in 2020, although many already deliver a broad set of capabilities. Multiple incumbent networking and network security vendors are developing new or enhancing existing cloud-delivery-based capabilities.

User Advice: There have been more than a dozen SASE announcements over the past 12 months by

vendors seeking to stake out their position in this extremely competitive market. There will be a great deal of slideware and marketecture, especially from incumbents that are ill-prepared for the cloud-based delivery as a service model and the investments required for distributed PoPs. This is a case where software architecture and implementation matters

When evaluating SASE offering, be sure to:

- Involve your CISO and lead network architect when evaluating offerings and roadmaps from incumbent and emerging vendors as SASE cuts across traditional technology boundaries.

- Leverage a WAN refresh, firewall refresh, VPN refresh or SD-WAN deployment to drive the redesign of your network and network security architectures.

- Strive for not more than two vendors to deliver all core services.

- Use cost-cutting initiatives in 2020 from MPLS offload to fund branch office and workforce transformation via adoption of SASE.

- Understand what capabilities you require in terms of networking and security, including latency, throughput, geographic coverage and endpoint types.

- Combine branch office and secure remote access in a single implementation, even if the transition will occur over an extended period.

- Avoid vendors that propose to deliver the broad set of services by linking a large number of products via virtual machine service chaining.

- Prioritize use cases where SASE drives measurable business value. Mobile workforce, contractor access and edge computing applications that are latency sensitive are three likely opportunities.

Some buyers will implement a well-integrated dual vendor best-of-breed strategy while others will select a single vendor approach. Expect resistance from team members that are wedded to appliance-based deployments.

Business Impact: SASE will enable I&O and security teams to deliver the rich set of secure networking and security services in a consistent and integrated manner to support the needs of digital business transformation, edge computing and workforce mobility. This will enable new digital business use cases (such as digital ecosystem and mobile workforce enablement) with increased ease of use, while at the same time reducing costs and complexity via vendor consolidation and dedicated circuit offload.

COVID-19 has highlighted the need for business continuity plans that include flexible, anywhere, anytime, secure remote access, at scale, even from untrusted devices. SASE's cloud-delivered set of services,

including zero trust network access, is driving rapid adoption of SASE.

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Akamai; Cato Networks; Cisco; Citrix; iboss; Netskope; Open Systems; Palo Alto Networks; VMware; Zscaler

Recommended Reading: "The Future of Network Security Is in the Cloud"

"Magic Quadrant for Cloud Access Security Brokers"

"Market Guide for Zero Trust Network Access"

"Market Trends: How to Win as WAN Edge and Security Converge Into the Secure Access Service Edge"

"Quick Answer: Cost-Effectively Scaling Secure Access While Preparing for a Remote Workforce"

5G Services

Analysis By: Bill Menezes

Definition: 5G services comprise local or wide-area cellular data connectivity based on 3GPP Release 15 or later, providing the next generation of cellular communications networking to follow 4G LTE. Providers will base services, such as new or enhanced end-user or IoT applications, on key 5G performance requirements of up to multigigabit mobile data throughput. Services will also be based on low-latency data transmission and support for massive deployment of machine-to-machine communications that are supported by Release 16 and later.

Position and Adoption Speed Justification: 5G is the most-searched term among networking technologies considered for Gartner Hype Cycles, based on a composite metric comprising Gartner search, Gartner inquiry, and Google trends. As of 2Q20, more than 70 service providers globally had launched commercial fixed or mobile wireless services using 3GPP-compliant 5G technology, according to the Global mobile Suppliers Association (GSA). Expanding network availability will keep 5G services moving rapidly through the emerging state. However, coverage will grow slowly in some regions and service providers have identified few use cases requiring 5G performance capabilities instead of widespread alternatives such as Wi-Fi or 4G LTE. Ratification of the next two 5G technology specifications, Releases 16 and 17, expected in 2020 and 2022, respectively, will bring significant performance improvements beyond data throughput.

These will support ultra-low latency, massive coverage density for low-powered IoT endpoints and network slicing. Gartner still expects that by 2025, less than 45% of communications service providers (CSPs) globally will have launched 5G.

Early 5G launches, such the U.S. service by AT&T and T-Mobile U.S., focused on wide-area 5G coverage for general use cases using low- or midband spectrum in "non-stand-alone" architectures complemented by 4G LTE. Verizon Wireless initially focused on clusters of higher performance 5G coverage using millimeter wave spectrum requiring significantly more transmitters. SK Telecom's initial launch supported a 5G smartphone for mobile service. Other CSPs, such as NTT DOCOMO, KT, Telstra, China Mobile, EE and Swisscom were among carriers that began providing similar services.

User Advice: 5G continues to suffer from immaturity, substantial hype and unrealistic expectations about features and availability sets, caused by infrastructure and service provider marketing. Enterprises must incorporate realistic networking assumptions for business plans by working with business leaders and network service providers to identify the availability of 5G at specific locations. Optimal 5G performance will require use of low- and midband cellular spectrum plus millimeter wave spectrum. We do not expect 5G using millimeter wave spectrum to become readily available outside of dense urban areas. Further, identify locations where available millimeter wave signals may encounter signal propagation challenges from obstacles such as building walls, window glass and heavy foliage.

Identify opportunities to pilot 5G networks and services to deliver innovation, based on current 5G capabilities, such as high-speed data. At the same time, identify applications where currently available service provider technologies (such as LTE-A) will support usage scenarios requiring up to 1 Gbps data speeds and latency up to 30 milliseconds.

For specialized local area use cases, identify situations in which private networks and future 5G features can be combined with edge computing to enable new business capabilities and integrate with innovation plans.

Service provider "5G service" actually will be based on an ecosystem of licensed, unlicensed and shared spectrum bands. Technologies will include not only proprietary prestandard specifications (such as Verizon's V5GTF), but also standards-based 5G New Radio, legacy 2G, 3G, 4G LTE, LTE Advanced (LTE-A), and Wi-Fi, providing mobile and fixed access services to a common core WAN service.

Business Impact: Gartner expects 5G services to be a key enabler of fixed and mobile enterprise IoT strategies (such as in manufacturing) that will require very low latency connections for large numbers of endpoints. 5G services will move beyond throughput improvements to address various new technical requirements and market drivers/expectations that take advantage of capabilities the 5G standards will define. These eventually will include narrowband IoT and LTE Cat-M1 capabilities for ultra-low-powered devices. Service providers also plan to include 5G as an element of SD-WAN offerings. New service

providers also may emerge to utilize 5G features such as network slicing or "network as a service" for edge computing. Specific performance and cost impacts will remain unclear until CSPs refine their plan structures and prices beyond the initial commercial launches.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: AT&T; BT; China Mobile; EE; NTT DOCOMO; SK Telecom; Telstra; Verizon

Recommended Reading: "Innovation Opportunities Will Be Enabled As 5G Evolves Through 2025"

"Don't Expect 5G to Replace Wired Access WANs Anytime Soon"

"Four Key Ways Enterprises Should Plan for 5G"

400 Gbps Ethernet

Analysis By: Naresh Singh

Definition: It is an IEEE 802.3bs Ethernet standard, that transmits Ethernet frames at 400 Gbps. Although it is achieved by combining multiple 25 Gbps lanes or multiple cable strands running at 100 Gbps, now offerings are available based on four-level pulse amplitude modulation (PAM4) technology that enables the transmission of frames at a speed of 400 Gbps through eight lanes of 50 Gbps.

Position and Adoption Speed Justification: 400 Gbps was ratified as an IEEE Standards Association standard in December 2017 (802.3bs). Products supporting it became commercially available in mid-2018, from multiple vendors including mainstream and niche providers. Although relatively new, it has been able to achieve swift and early but limited adoption among mega cloud providers like Amazon, Apple, Facebook, Google and Microsoft. The primary use case is data center backbones and interconnects, and application is driven by AI/ML fabrics and multimedia content.

The use of PAM4 in 400 Gbps serves a critical need. Higher Ethernet bandwidths, when applied traditionally with parallel cables, can lead to heavy and thick cable cores that can lead to operational challenges in the data center and thus can be impractical in many cases. With PAM4, this challenge is addressed by obtaining double the rate of bits per signaling unit to get to an effective bit rate of 50 Gbps per lane using a single pair of optic fibers. This leads to savings in both the cables and lasers used, thus giving it an edge over existing approaches to achieve higher Ethernet bandwidth.

In January 2020, IEEE 802.3cm that enables 400 Gbps over multimode fiber for shorter distance use cases has also been ratified. OIF, an industry group, has also been developing the 400ZR coherent WDM-based 400G modules for distances up to 80 km for usage in data center interconnects and CSP metro access and interconnects. Offerings based on this have started appearing in the market for sampling purposes. These options for a wide range of use cases position 400G suitably for broader adoption by both service providers and enterprises.

One of the biggest challenges with this technology in the recent past has been the lack of general availability of 400 Gbps capable optical transceivers, especially transceivers capable of going beyond 2 km and those serving shorter-range usage below 100 meters. However, due to multiple initiatives by IEEE and efforts of industry bodies and multiple multisource agreements (MSAs), a number of optical modules and transceivers including short-range direct attach copper cables with splitters and active optical cables are now appearing in the market.

It is estimated that over 90% of the demand for 400 Gbps comes from hyperscale cloud providers and to an extent communications service providers. For enterprises though, it is forecast to remain niche through the next three years, with overall shipment levels reaching 1 million ports only by 2024. As cost-effective 400 Gbps optics that are capable for transmissions over several kilometers become available, we expect the technology to also become more relevant for CSPs and large global enterprises.

User Advice: Some best practices that users can adopt include the following:

- Enterprises with large greenfield data centers should potentially look at 400 Gbps for their data center core and plan their network architecture accordingly with an aim to build a scalable and optimized data center network. As the technology is backward-compatible with 200 Gbps, 100 Gbps, 50 Gbps and 25 Gbps, it can be a good investment protection when used with a backward-compatible transceiver like QSFP-DD or OSFP.

- Enterprises with a need to reduce growing optical cable footprint due to increased use of parallel single-mode technologies for 40 Gbps and 100 Gbps should evaluate 400 Gbps PAM4 offerings as an option.

- As the technology is relatively new, a number of optical interconnect options are there and more will emerge. But initially prices are likely to remain high. Explore third-party optical transceiver providers as a means to bring down costs as you look into testing out this technology.

- Data center clusters and disaggregated systems that need consistently lower latency to perform well will benefit from 400G PAM4 systems.

Business Impact: 400 Gbps Ethernet will remain a niche technology for enterprises through the next three years; however, for large data centers it will grow into a good option for top of rack (ToR) aggregation and the core layers of the network. Those looking at implementing I/O-intensive workloads such as AI/ML and

already have a significantly high-density container workloads could find it justifiable even in the short term. Also, this technology has high potential to replace 40G and 100G for high-speed data center interconnects.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Accton Technology (Edgecore Networks); Arista Networks; Cisco Systems; Juniper Networks; NVIDIA (Mellanox Technologies)

Recommended Reading: "40G Is Dead — Embrace 100G in Your Data Center!"

"How to Avoid the Biggest Rip-Off in Networking"

Network On-Demand Services

Analysis By: Danellie Young

Definition: Network on-demand services offer real-time changes to port-related functions via a web portal or API. These functions include changes to access/port/path for both transport and overlay, transport service features (QoS), WOC on demand, add/delete/adjust WAN services across individual and multiple WAN services, and creating and adjusting capacity of connections to specific destinations (i.e., cloud services or extranet).

Position and Adoption Speed Justification: Network on-demand has been available for some time, though it was initially limited to a single service (usually MPLS) and did not allow connections to new endpoints. To support the adoption of cloud and IoT, enterprise networks are under pressure to transform and become more agile. In response to this, a new generation of network on-demand services are much more flexible, typically supporting multiple services in addition to connections to new endpoints such as internet and cloud services. They can also support a migration of network traffic from MPLS to internet services, and the addition of cloud and extranet endpoints.

However, growing enterprise interest and adoption of SD-WAN and network function virtualization (NFV) services require providers to offer methods to support self-service and improve port agility for better optimization and orchestration across WAN links. Capabilities enabled by enhanced portals or APIs will allow the enterprise to make the necessary adjustments quickly and easily in support of newer offerings and solutions. Gartner estimates approximately 25% of enterprises use NOD services today and that number will grow to 50% by 2025. Over the next three to five years, we expect increased network on-demand adoption, as features and control over dynamic bandwidth (NOD) services continue to evolve and add

greater value to enterprises in terms of speed and flexibility.

User Advice: Organizations requiring high levels of agility and better cost optimization and control from their carrier services should select network service providers that make dynamic network services available or soon plan to offer them. Such services should ideally include voice and data services and must enable enterprises to achieve lower costs and tight alignment to usage levels. Enterprises should seek network service providers focused on greater self-service and that offer services either via a portal or an API to enable enterprise network managers to add/remove services and adjust their capacity on demand. Service providers are not typically able to offer such on-demand transport capabilities over their partners' infrastructure due to lack of operational integration.

Business Impact: These services are typically delivered over carrier-owned, WAN connections (often via Ethernet access) or bundled into offers from system integrators or third-party managed service providers. They are at their most valuable when connecting to major enterprise hub sites and data centers. Such services extend beyond simply a provisioning mechanism for existing services and enable new self-service approach/type to be added to the network environment (such as QoS modifications). This allows for optimal agility to support network growth and evolving demands, including growing use of cloud applications and temporary demands for bandwidth, such as during scheduled data-replication activities (e.g., backups).

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Sample Vendors: AT&T; CenturyLink; GTT; Masergy; Telstra; Verizon

Recommended Reading: "Network On-Demand Services Bring Unprecedented Agility to Enterprise WANs"

Network Automation

Analysis By: Josh Chessman; Andrew Lerner

Definition: Network automation tools automate the configuration, visibility, troubleshooting, reporting and maintenance of physical and virtual network devices, and services. Network automation tools can increase agility and operational efficiency, lower costs, reduce human error, and improve compliance with configuration policies.

Position and Adoption Speed Justification: Based on interactions with clients and confirmed via conference polling, network automation often ranks as a top No. 2 priority with networking teams in terms of strategic investment. "Network automation" is the sixth most-searched term among networking technologies

considered for Gartner Hype Cycle, which is a composite metric based on Gartner search, Gartner Inquiry, and Google Trends. Further, network vendors are now embedding automation into nearly all management suites and it is becoming a baseline capability in markets such as SD-WAN. There are independent third-party tools that work across vendors that are also growing in popularity. Most organizations are using some form of network automation in production today, however, the depth of usage is limited. We estimate that roughly 30% of network changes are automated today with that number increasing to 50% in 2023. The overall market growth remains positive as the technology continues to move past the Peak, with Gartner estimating growth across all aspects of the market.

The discipline is held back partly by a lack of organizational process maturity pushing teams toward taking a pragmatic approach to resolving their organizations' specific requirements. This has frequently resulted in a cultural reluctance to adopt DevOps tools and principles as these promote agility while NetOps is focused on reliability and uptime. While a lack of scripting skills has been a barrier in the past, many tools have moved beyond basic scripting requirements which will lead to broader adoption of network orchestration tools.

User Advice: We recommend:

- Measure, reward and socialize network automation within your organization by targeting business outcomes, such as faster service delivery, improved application availability or reduced operational expenses.

- Implement network automation in a gradual fashion by focusing on new network build-outs (both on-premises and cloud based) and non-change activities first, such as reporting and troubleshooting, and then establishing a baseline and evolving toward more critical activities.

- Invest in personnel by shifting the hiring and training focus away from hard network skills and toward those with more flexible skill sets, including coding (in a language such as Python) and familiarity with deployed or identified automation solutions. Additionally, cross-pollinating networking teams with adjacent DevOps personnel will be useful due to the significant overlap in skills and processes.

Business Impact: Network automation tools can lower cost and improve operational agility. Consider network automation tools as components of a broader automation strategy. This demands participation in strategic, companywide deployment and configuration automation strategies (which are usually implemented as part of an IT service support management toolset), and integration with configuration management tools for other technologies, such as servers and storage.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Apstra; BMC; Cisco; CloudGenix; NetBrain; Red Hat; SaltStack; SolarWinds; VMware

Recommended Reading: "Jump-Start Network Automation to Scale Digital initiatives"

"Cool Vendors in Enterprise Networking"

"Market Guide for Network Automation"

"3 Ways to Improve Network Automation"

"NetOps 2.0: Embrace Network Automation and Analytics to Win in the Era of ContinuousNext"

Sliding Into the Trough

Telemetry-Based Routing

Analysis By: Bob Gill; Andrew Lerner

Definition: Telemetry-based routing dynamically steers and/or influences network traffic based on a broader set of information vs. traditional IP routing. Intelligence is derived from real-time network measurements, such as path latency, server availability, or policies such as circuit cost or data sovereignty. This is different from traditional IP routing which uses a hop-by-hop method based on destination IP address. TBR is typically packaged within a broader offering versus sold as a stand-alone fashion.

Position and Adoption Speed Justification: The concept behind telemetry-based routing (TBR) is not new, and other terms have been used previously to describe similar or partial functionality including, application-based steering, and performance routing.

As a stand-alone technology, packaged by itself, adoption of TBR in the enterprise is extremely limited. Instead, it is embedded into other products and services including SD-WAN, CDN and DNS. Further, it is increasingly being packaged with noninternet backbones to provide alternative transport to internet and MPLS. It may never achieve widespread adoption directly from enterprises as a stand-alone technology (that they acquire, manage and support themselves in an isolated fashion).

SD-WAN, CDN, DNS and "middle-mile" acceleration products and services embed TBR functionality to improve performance, cost and/or security in support of enterprise digital and cloud initiatives. This is because while the internet is foundational to modern business communications, the constantly changing nature of it makes predictability, resilience, and performance guarantees difficult to achieve. Routes across the internet are often guided by proximate-to-the-next-node information, rather than knowledge about the

end-to-end path, or network as a whole. Real-time network intelligence, which we describe as telemetry, is maturing as the ability to collect and analyze vast amounts of real time traffic measurements from across the internet becomes more powerful and accessible. Like a GPS may reroute a motorist due to changing traffic conditions, TMR maps the relevant network intelligence to specific routing requests, such that the message is routed by the lowest-latency route, most cost-effective route, or bypasses certain countries, etc.

User Advice: We recommend the following:

- IT Leaders should prefer network services and products that embed telemetry-based routing concepts in SD-WAN, CDN, DNS and carrier-based services.

- IT Leaders should use TBR features when delivering internet-based, mission-critical applications that require high-performance. However, they should not expect to acquire TBR as a stand-alone product offering.

- SaaS and content providers should demand telemetry-based routing techniques from the network services and products they utilize, to maximize potential performance and costs while delivering content.

Business Impact: TBR can improve performance, security and optimize costs for applications and services accessed over the internet, which can improve customer experience.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Sample Vendors: Akamai; Anapaya Systems; Cloudflare; NS1; Teridion

Recommended Reading: "Market Guide for CDN Services"

"WAN Practices to Optimize Application Performance"

"Magic Quadrant for WAN Edge Infrastructure"

Wi-Fi 6 (802.11ax)

Analysis By: Tim Zimmerman

Definition: Wi-Fi 6 (802.11ax) is the latest iteration of the IEEE 802.11 WLAN family. Its main enhancements are allowing the network to control device connectivity for the first time and to improve the efficiency of

existing 2.4GHz and 5GHz spectrum, thereby increasing throughput in densely populated areas. As such, its goal is to support a larger number of devices including IoT that are properly connected to the network.

Position and Adoption Speed Justification: IEEE, the standards body of 802.11-based technologies, started the High Efficiency WLAN (HEW) Study Group in May 2013 to examine the most pressing needs for the next-generation Wi-Fi technology. Ratification of the new Wi-Fi 6 (802.11ax) standard was expected in late 2019 but was delayed until 2020 and perhaps longer due to COVID-19. Prestandards-based networking equipment has already been released by a majority of the key/leading vendors. Successive Wi-Fi technologies have increased per-device throughput at an impressive rate through the years (from 802.11b's 11 Mbps to 10 Gbps expected from the new standard).

The number of IoT devices and the convergence of building automation and line of business devices onto the enterprise communication infrastructure continues to contribute to congestion. 802.11ax will be able to more intelligently use network resources instead of letting the device make connectivity decision as previous versions of the standard allowed. Advancements allow differing streams to communicate to multiple devices simultaneously as well as cutting latency by as much as 75%.

User Advice:

- We advise clients not to pay a premium for any adoption of Wi-Fi 6 (802.11ax) unless existing wireless solution do not provide the performance and functionality needed to meet defined end-user requirements.

- For IT leaders, Wi-Fi 6 (802.11ax) represents an opportunity to significantly improve performance for special use cases such as dense device deployments, they are advised to monitor the standardization timeline and product availability to find the right entry point for future infrastructure upgrades as well as availability of client devices that will support the new standard.

- We advise clients that "Wi-Fi 6 certification" currently does not mean 802.11ax compliant since the standard is not ratified. Any organization that purchases prestandard products should have the ability to upgrade or update their product at no cost to be standards compliant.

Business Impact: Wi-Fi 6 (802.11ax) is the first Wi-Fi technology that will be able to support large-scale IoT usage scenarios and other high density environments using existing 2.4GHz and 5GHz spectrum via its use of OFDMA.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Broadcom; Cisco; Extreme Networks; Hewlett Packard Enterprise (Aruba); Huawei; Intel; Juniper Networks (Mist Systems); Qualcomm

Recommended Reading: "Magic Quadrant for Wired and Wireless LAN Access Infrastructure"

"Critical Capabilities for Wired and Wireless LAN Access Infrastructure"

Open Networking/OCP Networking

Analysis By: Mark Fabbi

Definition: Open networking concepts originated in the data center to support the disaggregation of hardware and software networking platforms. This leveraged Open Compute Project (OCP) efforts promoting open-source network hardware and software designs and the commercial, brite-box, offerings of vendors such as Dell EMC and Mellanox Technologies. Open software solutions are based on Linux network operating systems and tools. More recent work expands the focus beyond data center networking to campus switching and wireless access points.

Position and Adoption Speed Justification: OCP-based networking is still in its early adoption phase for enterprise deployments, but is well established in cloud and service provider networks. However, open, brite-box networking solutions, which bundle disaggregated hardware and software into a single supported package represent a more mainstream approach to leveraging the open networking movement. Brite-box solutions from vendors such as Dell EMC and Mellanox have over 4000 production customers. On the other hand, Hewlett Packard Enterprise (HPE) abandoned their Altoline open switch portfolio which contributes to the uncertainty in the market. Recent acquisitions of Mellanox and Cumulus Networks (both by NVIDIA) contribute both hope and uncertainty in the market. In the broad data center market, original design manufacturer (ODM) (or white-box) switches (of which OCP designs represent a significant portion) accounted for 28.4% of unit shipments of data center switches in 2019 (up from 19.8% in 2018). This is with an average selling price 74% less than traditional OEM switches. OCP-compliant hardware is available from many white-box switch providers, either as stand-alone switches or coupled with commercial software, often as part of branded (brite-box) offerings. Brite-box solutions will be much more attractive to enterprise buyers as the complexity of integrating hardware and software as well as support is dealt with by the brite-box provider. Software components such as the Open Network Install Environment (ONIE), SAI and Open Network Linux (ONL) promote further abstraction of the underlying hardware, making switch OS functions more transportable across switches using different underlying switch silicon.

OCP-compliant switches are available to support a wide range of 10/25/40/50/100G applications for both leaf and spine requirements, including chassis options. Standard network components should become available in a number of Linux distributions to further the efforts of the open-networking movement. Such contributions lower the barrier to entry for new network innovation. However, the growing number of open-source software options (Open Network Linux, SONiC, OS10 Open Edition, FBOSS, OPX, etc.) could

potentially fragment efforts and delay broader adoption, though during the past year we have seen an increased focus on the Microsoft contributed SONiC switch OS software.

Most of the initial work under OCP networking has focused on data center products with a very broad set of possible commercial applications. However we have observed expansion of OCP networking into new areas, such as WLAN access points and campus (with designs for a stacking, PoE capable access layer switch now available), branch and optical networking platforms.

User Advice: We recommend the following:

- Large data center and service provider operators should demand support for OCP networking for future purchases to help reduce operating costs, as they offer competitive and more flexible alternatives to OEM hardware and software.

- Organizations considering open-source hardware options can gain further leverage by assessing SAI-based software options as they become available, as these will further lower their dependency on hardware components.

- In addition, support for specific APIs that are aligned with chosen orchestration and operational environments is key.

- Any organization pursuing a disaggregated open-source approach must have advanced technical expertise to deal with the support and integration requirements.

- Enterprises wanting to lower their network expenditure or create a customized operational environment should evaluate the commercial offerings of vendors targeting open networking and brite-box switching, rather than focusing on the open-source foundational technologies.

- Include brite-box switching on shortlists in specific usage scenarios, such as new application deployments — e.g., analytics cluster and virtual desktop infrastructure (VDI) pods — and to drive new operational initiatives, such as automated development/test and DevOps.

Business Impact: Open networking and OCP initiatives will help reduce related capital expenditure costs by as much as 70%, increase operational flexibility when compared with traditional network approaches while reducing vendor lock-in. Decoupling of hardware and software can allow for changing deployments as use cases or demand shifts as changes can be made in software rather than hardware replacements which lowers the investment risk. OCP software should lower the barriers to entry for innovation in hardware and software, which will benefit the broader networking market and allow for more customization as required.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Alpha Networks; Arista Networks (Mojo Networks); Cumulus Networks; Dell EMC; Edgecore Networks; Mellanox Technologies; Microsoft; Nephos Technologies; Quanta Computer

Recommended Reading: "Magic Quadrant for Data Center and Cloud Networking"

Zero Trust Network Access

Analysis By: Steve Riley

Definition: Zero trust network access (ZTNA) creates an identity- and context-based, logical-access boundary around an application or set of applications. The applications are hidden from discovery, and access is restricted via a trust broker to a set of named entities. The broker verifies the identity, context and policy adherence of the specified participants before allowing access, and prohibits lateral movement elsewhere in the network. This removes the application assets from public visibility and significantly reduces the surface area for attack.

Position and Adoption Speed Justification: ZTNA is a synthesis of concepts promulgated by the Cloud Security Alliance's software-defined perimeters (SDP) project, by Google's BeyondCorp vision, and in O'Reilly's *Zero Trust Networks* book. Early products on the market tended to focus on use cases involving access to web applications. Newer, more complete products work with a wider range of applications and protocols.

As more organizations suddenly find themselves transitioning to much more remote work, hardware-based VPNs exhibit limitations. ZTNA has piqued the interest of those seeking a more flexible alternative to VPNs and those seeking more precise access and session control to applications located on-premises and in the cloud. ZTNA vendors continue to attract venture capital funding. This, in turn, encourages new startups to enter an increasingly crowded market and seek ways to differentiate. Merger and acquisition (M&A) activity in this market is underway, with several startup vendors now having been acquired by larger networking, telecommunications and security vendors.

User Advice: Organizations should evaluate ZTNA for any of these use cases:

- Opening up applications and services to collaborative ecosystem applications, such as distribution channels, suppliers, contractors or retail outlets without requiring the use of a VPN or DMZ.

- Normalizing the user experience for application access — ZTNA eliminates the distinction between being on and off the corporate network.

- Application-specific access for IT contractors and remote or mobile employees as an alternative to VPN-based access.

- Extending access to an acquired organization during M&A activities, without having to configure site-to-site VPN and firewall rules. The merged companies can quickly and easily share applications without requiring the underlying networks and/or identity systems to be integrated.

- Enabling users on personal devices — ZTNA can improve security and simplify bring your own device (BYOD) programs by reducing full management requirements and enabling more-secure direct application access.

- Cloaking systems on hostile networks, such as systems facing the public internet used for collaboration.

- Carrying encryption all the way to the endpoints for scenarios where you don't trust the carrier or cloud provider.

- Permitting users in potentially dangerous areas of the world to interact with applications and data in ways that reduce or eliminate risk prone to originate in those areas.

- Securing access to enclaves of IoT devices if the device can support lightweight SDP agent or a virtual-appliance-based connector on the IoT network segment for connection.

Business Impact: The benefits of ZTNA are immediate. Similar to a traditional VPN, services brought within the ZTNA environment are no longer visible on the public internet and, thus, are shielded from attackers. In addition, ZTNA brings significant benefits in user experience, agility, adaptability and ease of policy management. For cloud-based ZTNA offerings, scalability and ease of adoption are additional benefits. ZTNA enables digital business transformation scenarios that are ill-suited to legacy access approaches. As a result of digital transformation efforts, most enterprises will have more applications, services and data outside their enterprises than inside. Cloud-based ZTNA services place the security controls where the users and applications are — in the cloud. Some of the larger ZTNA vendors have invested in dozens of points of presence worldwide for low-latency access.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Akamai; AppGate; Cato Networks; Cisco; Netskope; Perimeter 81; Proofpoint; Pulse Secure; SAIFE; Zscaler

Recommended Reading: "Market Guide for Zero Trust Network Access"

"Zero Trust Is an Initial Step on the Roadmap to CARTA"

"Solving the Challenges of Modern Remote Access"

"Quick Answer: Cost-Effectively Scaling Secure Access While Preparing for a Remote Workforce"

"The Future of Network Security Is in the Cloud"

Intent-Based Networking

Analysis By: Andrew Lerner

Definition: An IBN system is a closed-loop system to help design, provision, and operate a network, based on business policies. IBNS is typically packaged as software and a full IBNS includes four key subcomponents:

- Can translate a higher-level business policy to a network configuration.

- Can automate network activities across network components.

- Has awareness of network state/health.

- Provides continuous assurance and enables dynamic optimization.

Position and Adoption Speed Justification: IBNS is being hyped by networking vendors because of the promise to improve network availability and agility simultaneously, enabling digital business transformation. However, as of early 2020 real-world enterprise adoption of full IBNSs is nascent. We estimate fewer than 100 full deployments that meet all phases of the definition. The detailed description of the subcomponents of IBNS includes:

- Translation and validation: System can take a higher-level business policy (what) as input from end users and convert it to the required network configuration (how).

- Automation: System can configure appropriate network changes (how) across existing network infrastructure.

- State awareness: System ingests real-time network status for systems under its control.

- Assurance and dynamic optimization: System continuously validates that business intent is being met; can take corrective actions when it is not.

Each of the four subcapabilities of intent can be deployed independently. There is much heavier adoption of certain individual components, including automation. These components delivered individually add value, but are not full intent when they aren't delivered together as part of a closed loop.

In the last 18 months, the terminology around intent (including intent-based, intent-driven, IBN, etc.) has largely been taken over by vendor marketers, as there is rampant overmarketing of intent by networking vendors. Things such as automation, abstraction and programmability are described as intent. Unfortunately, many products are marketed by vendors as intent that falls short of the full capabilities of intent.

We anticipate that the number of full closed-loop IBNS commercial enterprise deployments to remain below 200 through 2020, and increase moderately by the end of 2021. We expect adoption to be pragmatic — associated with new build-outs and/or network refresh initiatives. Through 2020, early rollouts are likely to be in larger-scale environments for well-defined and specific use cases, such as spine/leaf data center networks. We expect that data center networking vendors will increasingly incorporate assurance, validation and dynamic remediation into their networking suites, getting closer to a full IBN.

We recommend and anticipate that adoption will be phased, with the different IBNS subcomponents enabled gradually over months and years. Technologically, IBNS requires the ability to abstract and model network behavior, which has proved difficult historically, particularly in multivendor environments. Furthermore, IBNS represents a substantial cultural shift in how networks are designed and operated, which will create barriers to adoption in many risk-averse enterprises.

User Advice:

- Deploy IBNS pragmatically in phases, because each of the four individual subcomponents adds value. For example, organizations can deploy an IBNS in "notify mode," whereby a skilled engineer must approve a proposed automated change suggested by the IBNS.

- Mandate support for open, RESTful APIs when purchasing new networking infrastructures to support integration within an IBNS moving forward.

- Choose an IBNS that supports multivendor network infrastructures and extends into public cloud environments to support a broader range of use cases and avoid vendor lock-in.

- Tune-out vendor marketing regarding products that are listed as intent.

Business Impact: Gartner sees the biggest benefits from IBNS as improving network agility and availability, and supporting unified intent and policy across multiple infrastructures. When the technology matures, a full IBNS implementation can reduce the time to deliver network infrastructure services to business leaders by 50% to 90%, while simultaneously reducing the number and duration of outages by at least 50%. Although

agility and availability are the high-level benefits, IBNS provides several other specific benefits, including:

- Reduced operating expenditure (opex) — Reduce opex associated with managing networks and free up senior-level network resources to focus on more important strategic tasks.

- Performance optimization — Intent-based algorithms provide better traffic engineering versus traditional approaches, such as routing protocols. This can improve application performance.

- Reduction in dedicated tooling costs — Intent tools may circumvent the costs of other related network infrastructure tooling, because automation and orchestration are embedded in IBNS.

- Better documentation — IBNS provides real-time self-documentation, which also includes the rationale (intent) behind design/configuration decisions.

- Improved compliance — IBNSs simplify auditing, due to the algorithmic correctness of configurations, direct mapping to business intent and ongoing, dynamic, real-time validation.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Apstra; Cisco; Forward Networks; Gluware; Huawei; Intentionet; Juniper Networks; NetYCE; VMware

Recommended Reading: "Innovation Insight: Intent-Based Networking Systems"

"Cool Vendors in Enterprise Networking"

NFV and uCPE

Analysis By: Bjarne Munch; Mike Toussaint

Definition: Universal customer premises equipment and network function virtualizations entail the virtual instantiation of network functions such as firewall, WAN optimization, SD-WAN, routing, etc. running on uCPE hosts which use standard CPUs (such as x86). NFV is not specific to the WAN, and can entail virtualization of any network function. However, NFV in conjunction with uCPE is specific to WAN deployments.

Position and Adoption Speed Justification: Universal customer premises equipment (uCPE) and network function virtualization (NFV) have grown out of the requirement to support and consolidate multiple network

functions at the WAN edge, while increasing flexibility and agility of deploying these functions for both network providers and end users alike. uCPE and NFV technologies address this requirement by enabling multiple dedicated network functions, from multiple vendors, to run concurrently on a single server platform. Thus, it enables increased deployment agility, additional service chaining capabilities, automation of device provisioning and orchestration, and reduced power requirements at branch sites.

While there are advantages to deploying uCPE with NFV technologies, obstacles to adoption are moderate to high. For instance, we have observed that while carriers have favored uCPE and NFV form factors, adopters have found performance issues have arisen with one or more NFVs running on a single uCPE platform. Additionally, the premise of "any" NFV running on uCPE platforms has become problematic due to integration and heterogeneity, thereby causing vendors to certify functionality on a small subset of uCPEs — which patently negates the very concept of a universally supported host infrastructure. We have also observed challenges with support for network interface modules in uCPE platforms because they are vendor-specific. Therefore, adding network interface modules to the uCPE from different vendors is not supported.

We commonly observe clients deploying at least two of the following NFVs at branch locations, including SD-WAN, and firewalls. Several network service providers (NSPs) with NFVs running on uCPEs, which they provide.

While end-user-managed uCPE-based NFV services are still nascent, and still facing substantive challenges, there is growing adoption of managed network-based NFV services. This is especially true of network service providers deploying network nodes across the geographies which they cover because uCPE and NFV technologies enable a broadening set of functions and vendors, combined with increased availability of web-based customer portals for self-service and service-chaining capabilities.

This Hype Cycle covers uCPE and NFV together, as it relates to WAN and branch scenarios. NFV has broader applicability outside of WAN which is not included in this analysis.

User Advice: When refreshing WAN edge equipment from a managed network service provider, prefer uCPE to custom-built appliances, but recognize that, as of early 2020, the technology is still nascent and consequently may present price/performance/compatibility obstacles that offset the agility advantages.

Ensure that any proof of concept or pilot is performed with all functionality required on the uCPE to ensure these functions operate acceptably together on the chosen hardware.

For network-based NFV, ensure that virtual functionality is located close to the enterprise locations (ideally not more than a 10 ms to 20 ms round-trip delay between the site and the point of presence, typically a distance of a few hundred miles).

When managed services are a preference, choose providers that support the range of functions required by the enterprise and, if applicable, any preferred vendors, and conduct a pilot before signing any contract to ensure that the provider can deliver and manage the service as promised.

Conduct a thorough total cost of ownership analysis to validate that the promised agility justifies the additional expense.

Business Impact: For the enterprise, there are three main value propositions of NFV:

- Network-based NFV can improve enterprise network agility because it enables enterprises to rapidly deploy new functionality where needed.

- Network-based NFV can facilitate the deployment of network functionality in several locations where it can offer an optimized architecture, such as an SD-WAN endpoint close to cloud services or a virtual firewall deployed in providers' facilities instead of a large number of small branch offices.

- For branch office uCPE-based NFV, business impacts are mainly from the ability to consolidate more functionality on fewer appliances, as there is still limited enhanced agility availability.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: AT&T; NTT Communications; Orange Business Services; Versa Networks; VMware

Recommended Reading: "Pump the Brakes on Network Function Virtualization Services"

"NFV/uCPE Is a Deployment Option — Shift Focus and Resources to SASE and Edge Computing Opportunities"

"Magic Quadrant for Network Services, Global"

"Critical Capabilities for Network Services, Global"

"5 Options to Secure SD-WAN-Based Internet Access"

Climbing the Slope

Software-Defined Networking (SDN)

Analysis By: Mark Fabbi; Joe Skorupa

Definition: Software-defined networking (SDN) is an architectural approach to designing, building and operating networks that promised increased agility and extensibility by abstracting the network topology and control plane. However, SDN products never made it to mainstream enterprise adoption. Rather, SDN spawned innovations in automation, orchestration, segmentation and the disaggregation of network hardware and software.

Position and Adoption Speed Justification: SDN remains a major topic of discussion across multiple network markets and in many vendor marketing efforts. However, true SDN technologies have not achieved any significant enterprise market traction and should not be considered by any enterprise networking organization. The hope that SDN would allow for the decoupling of the control plane from network hardware and foster independent software innovation never came to fruition and there are effectively no SDN technologies available in the mainstream marketplace today.

While true SDN solutions have not had any significant market adoption, the development of SDN and the threat to established market players had a profound, positive effect on subsequent market developments. SDN clearly influenced the increasing use of white-box switches, the open-source hardware and software movement (supported by the Open Compute Project) and the development of independent network switch software providers. More important to the enterprise market was a shift in focus of traditional network vendors innovation around operations and management. This has led to improvements in agility and automation, simplified operational requirements and a general adoption of function-wide configuration (i.e., management that spans devices in a data center, campus or WAN environments). Without the threat of SDN it is unlikely the operational advances from traditional vendors would have occurred.

User Advice: While SDN is clearly obsolete in the enterprise market there are still many organizations that site SDN as a cornerstone of their future strategy and architecture. What is critical for enterprises is to understand what they are trying to accomplish when they think "SDN." We recommend the following:

- Don't get caught up in the hype and vendor claims that commercial products are SDN or engage in any discussions or planning to deploy SDN. SDN is not the answer to any enterprise networking challenge today.

- Focus on the desired outcomes you are trying to achieve such as increased automation, virtual segmentation, external orchestration, control and programmability of the network or decoupling physical hardware from software switch operating systems.

- Select an operational/automation framework first — then decide on networking vendors and products. Decoupled hardware and software and independent network overlays provide a way of establishing long-term operational models that are independent of underlying hardware if that is a desired outcome.

- Evaluate both hardware infrastructure and software overlays approaches. Evaluate not only vendor promises but the operational requirements to actually achieve a stated benefit. Hold any considered or deployed vendor accountable to deliver your desired outcomes and don't believe marketing hype surrounding SDN-related claims. Evaluate reference accounts and pay particular attention to implementation and ongoing operational costs and investments to ensure that benefits can be realistically achieved.

- Develop cross-functional collaboration and investigate methodologies to better integrate server, virtualization, network, security and application teams. These teams can help identify key use cases — both short-term, such as self-service development environments and microsegmentation; and long-term, integrating networking more broadly into data center orchestration.

- Allocate time and resources to evaluate technologies and a shortlist of relevant vendors — both incumbent and nonincumbent — in order to arrive at a solution that best meets the needs of the cross-functional organization.

Business Impact: There is no direct commercial impact of full SDN solutions. However, downstream innovation can increase network agility, simplify management, improve security and lead to reductions in operational and capital costs while fostering cross-functional collaboration. New network solutions should focus on reducing or eliminating the "human middleware" problem that has plagued traditional network solutions for the past two decades. By bringing network operations into more streamlined and automated operational process that have been architected in the virtual compute environment, user organizations can bring application deployments in line with the increasing speed of business. Available network overlay technology can create new competitive environments shifting the focus from physical infrastructure to software and operational features.

Benefit Rating: Low

Market Penetration: Less than 1% of target audience

Maturity: Obsolete

Sample Vendors: NEC

Recommended Reading: "State of SDN: If You Think SDN Is the Answer, You're Asking the Wrong Question"

Identity-Based Segmentation (Microsegmentation)

Analysis By: Adam Hils; Neil MacDonald; Jeremy D'Hoinne

Definition: Identity-based segmentation (also referred to as microsegmentation, zero trust network segmentation or logical segmentation) uses policy- and workload-identity-driven firewalling (typically software-based) or differentially encrypted network communications to isolate workloads, applications and processes in data centers, public cloud IaaS and containers. This includes workloads that span on-premises and multiple public cloud IaaS providers.

Position and Adoption Speed Justification: With more servers being virtualized or moving to infrastructure as a service, traditional firewall, intrusion prevention, and antivirus rarely follow the fast pace of deployment for new assets. This leaves the enterprise vulnerable to attackers gaining a foothold and then moving laterally within enterprise networks. This has created increased interest in visibility and further segmentation and zero trust networking based approaches for east-west traffic between applications, servers and services in modern data centers. The increasingly dynamic nature of data center workloads makes traditional network-centric segmentation strategies complex, if not impossible, to apply. Further, the shift to microservices container architectures for applications has also increased the amount of east-west traffic and further complicated the ability of network-centric firewalls to provide this segmentation. The extension of data centers into public cloud also has placed a focus on software-based approaches for segmentation, in many cases, using the built-in segmentation capabilities of the cloud providers. Growing interest in zero trust networking approaches has also increased interest in using application/service identities as the foundation for adaptive application segmentation policies. This is critical to enforce segmentation policies in the dynamic networking environments used within container-based environments.

User Advice: Security and risk management leaders should use the following guidelines when implementing identity-based segmentation:

- Don't oversegment. Oversegmentation is the leading cause of failure and an unnecessary expense for segmentation projects.

- Don't use IP addresses or network location as the foundation for segmentation policies. Use the identities of applications, workloads and services, either via logical tags, labels, fingerprints or stronger identity mechanisms such as certificates.

- Start with a network flow mapping project to understand application and server flows before undertaking the segmentation project. Leading microsegmentation vendors provide this capability to help enterprises on their journey to microsegmentation.

- Apply continuous adaptive segmentation. Start with new assets, then close existing gaps. Identify quick wins, and mix zoning governing principles when needed.

- Adopt a risk-based approach and look beyond technical considerations when segmenting. Consider the business processes and the value information being protected.

- Consider products with established security expertise, such as those from security vendors targeting this

market. Isolation alone isn't segmentation: If mediated communication is needed between zones, this requires different functionality than merely keeping them apart.

- Architect for consistent segmentation policies across on-premises and public cloud IaaS; using approaches such as host-based controls or using the native APIs of the underlying cloud fabric. Alternatively, several vendors use virtual appliance or container-based approaches to provide this capability.

- Ensure that your segmentation strategy extends into containers and container networking environments.

- Plan for coexistence of traditional firewalls and microsegmentation approaches for at least the next five years and seek products that can support using both.

Business Impact: Identity-based segmentation is a form of zero trust networking and is used to reduce the "blast radius" if and when an attacker breaches the enterprise network by reducing the ability of the attacker to spread laterally. It also enables enterprises to enforce consistent segmentation policies across on-premises and cloud-based workloads, including workloads that host containers meeting compliance requirements. In addition, several solutions provide extensive visibility and visualization of flows for baselining and anomaly detection. For some specific scenarios, like PCI reduction of scope, microsegmentation can be used to avoid costly network reconfiguration.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Amazon Web Services; Cisco; Edgewise; GuardiCore; Illumio; Microsoft; Palo Alto Networks; ShieldX; vArmour; VMware

Recommended Reading: "Market Guide for Cloud Workload Protection Platforms"

"Zero Trust Is an Initial Step on the Roadmap to CARTA"

"Control Network Security Complexity, Inefficiencies and Security Failures by Minimizing Firewall Diversity"

"Solution Comparison for Microsegmentation Products"

SD-WAN

Analysis By: Andrew Lerner

Definition: Software-defined wide-area network (SD-WAN) products replace traditional branch routers. They provide several features: dynamic path selection, based on business or application policy; centralized policy and management of WAN edge devices; and zero-touch configuration. SD-WAN products are WAN transport/carrier-agnostic, and can create secure paths across multiple WAN connections. SD-WAN products can be hardware- or software-based, and managed directly by enterprises or embedded in a managed service offering.

Position and Adoption Speed Justification: Rampant client interest in SD-WAN products continues, and we estimate that more than 25,000 customers have deployed SD-WAN products in production networks, which is over 600,000 branch locations. We expect continued rapid growth of SD-WAN deployments, and forecast vendor revenue to grow at a more than 23% compound annual growth rate (CAGR) for the next three years. In conjunction with a hybrid WAN topology, SD-WAN improves availability, cost and performance for enterprise WANs. Organizations moving to hybrid or internet-only WAN transport are driven toward SD-WAN products, because of their improved path selection functionality and manageability. Large numbers of vendors (several dozen) are competing in the market, including incumbent network and security vendors, startup vendors and smaller vendors with regional or vertical focus.

User Advice: Networking leaders should refresh their branch WAN equipment by implementing SD-WAN when they're migrating apps to the public cloud, building hybrid WANs, equipment is at end of life, or managed network service/MPLS contracts are up for renewal. Follow a comprehensive SD-WAN selection process by evaluating a diverse set of vendors and running a pilot. This is particularly important now, because not all offerings on the market are stable and scalable. Include network security teams in the design, planning and implementation, because SD-WAN-enabled hybrid WANs directly affect placement of security controls, such as firewalls and secure web gateways (SWGs).

Business Impact: The main purpose of emerging SD-WAN products is to create simpler and more cost-effective branch office WANs that map to modern application and cloud architectures. These products are significantly faster, easier to deploy and more manageable than traditional, router-based solutions. The benefits of an SD-WAN approach are substantial, compared with traditional, router-based WAN architectures, including reduced capital and operational expenditures (capex/opex) at the WAN edge, improved provisioning times, and the potential for enhanced branch availability.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Sample Vendors: Aryaka; Cato Networks; Cisco; Fortinet; Palo Alto Networks; Silver Peak; Versa Networks; VMware

Recommended Reading: "Technology Insight for SD-WAN"

"Magic Quadrant for WAN Edge Infrastructure"

"Solution Comparison for SD-WAN"

"Assessing the Strengths and Weaknesses of SD-WAN Technology"

IPv6

Analysis By: Neil Rickard

Definition: Internet Protocol version 6 (IPv6) is the next version of Internet Protocol (IP). It's designed to overcome several key limitations of Internet Protocol version 4 (IPv4) — in particular, the exhaustion of public IPv4 addresses.

Position and Adoption Speed Justification: Public IPv4 address exhaustion means new public IP addresses will be IPv6 only. The continued growth of the internet, including consumer internet access in emerging markets, mobile devices and the Internet of Things (IoT), will increasingly be accommodated using IPv6. Google IPv6 statistics show nearly 30% of users accessing Google do so using IPv6. However, the installed base of IPv4 infrastructure is huge and it will not migrate soon as migration costs are high. Also, there is an active market for IPv4 address transfers between organizations. Hence through at least 2025, the public internet will carry both IPv4 and IPv6 traffic.

Many network operators have been early adopters of IPv6 to facilitate continued growth. However, for enterprises, the actual depletion date for public addresses isn't critical because most enterprises use private IP addressing inside their organizations. Additionally, they typically already have the few public IP addresses they need for external services. Although most enterprises have little need to transition their internal systems to IPv6, they do need to enable IPv6 support on their public internet presence to ensure successful communication with IPv6-based consumers. Many organizations have already done this.

Enterprises are increasingly discovering pockets of IPv6 appearing within their network environments, especially in IoT use cases. These include sensor networks, process automation systems and building environmental control networks, because many of these systems will only support IPv6.

User Advice: Selective deployment of IPv6 is now advisable, especially for public-facing services, but a large-scale replacement of existing IPv4 services is still not recommended as migration costs and risks are substantial. Network planners should separate their IPv6 strategies into public internet presence, user LANs, data center LANs, and specialty networks such as IoT.

Enabling IPv6 on the enterprise public internet presence is now becoming critical. Because an expanding part of the internet's population will be natively IPv6 — especially in Asia, the developing world and 3G/4G/5G mobile networks — public-facing services (such as websites) should be migrated to IPv6 as soon as possible. Nearly 30% of all websites (by traffic volumes) have already enabled IPv6. Supporting external IPv6 connectivity will improve performance and visibility for enterprises' IPv6-only customers, partners and remote employees connecting via an IPv6-only internet provider or IPv6-only mobile devices. Organizations that do not enable IPv6 for these services will be forced to connect to such users via operators' network address translation gateways, which could easily become overwhelmed and may be unable to cope with complex translations.

There are far fewer benefits to deploying IPv6 internally to the enterprise. The migration costs are high for established IP networks, and attempts to transition even midsize networks have revealed many unexpected problems and hidden costs. As IPv6 is a new protocol and tends to be deployed with public addressing, rather than network address translation (NAT), network security must be completely rethought. It is generally easier to continue with private IPv4 addressing internally to the enterprise; if additional IPv4 addresses are required for external connections, these can be obtained through address trading. Gartner estimates that fewer than 5% of all intranets are IPv6-enabled.

Although we do not currently recommend migrating to IPv6 for all internal systems, enterprises should proceed with their IPv6 planning, including determining which systems and equipment are IPv6-ready and which are not, as well as establishing a transition roadmap. Larger network service companies, such as NTT Communications, offer professional services to assist in this activity. As businesses replace network and IT infrastructure, every new item of hardware or software with network capability should have IPv6 support so that the eventual transition will be easier and less expensive. Leading network services and cloud services support both IPv4 and IPv6 but it is important for enterprises to determine the options, limitations and cost of such support.

New greenfield networks should be built to support both protocols from the outset. In addition, organizations that undertake a large-scale IPv6 deployment are likely to need DDI platforms: a DNS, a Dynamic Host Configuration Protocol (DHCP) and an IP address management platform.

Expect to support both IPv4 and IPv6 through at least 2025. In some specialty networks, IPv6 provides a specific advantage when, for example, scale is critical, intercompany/interagency addressing is required, or networks are frequently set up and torn down.

Business Impact: The key for most businesses is to avoid connectivity disruptions that new users running the IPv6 protocol could create. The adoption of IPv6 beyond this requirement will typically be time-consuming and costly with few, if any, benefits.

Benefit Rating: Low

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: BlueCat; Cisco Systems; Citrix; Infoblox; NTT Communications; Tata Communications; Verizon

Ethernet Switching Fabric

Analysis By: Evan Zeng

Definition: Ethernet fabric is a type of Ethernet switching construct (Layer 2 and/or Layer 3) in which multiple physical and virtual devices are interconnected to build an integrated network system. The system is managed like a single logical entity and provides end-to-end network connectivity with consistent latency and availability. Programmable fabrics are advanced type of Ethernet fabrics, which are mostly built with advanced functions like network programmability, orchestration, network analytic and automation to enhance day 2 operation experience.

Position and Adoption Speed Justification: Gartner estimates there were more than 50,000 deployment cases worldwide by end of 2019 with increasingly high penetration of programmable fabrics. Enterprises' demand on DevOps and network automation generates solid, ever-increasing demands on Ethernet fabrics. As growing AI/ML based network analytics and intent-based technologies being adopted, end users' skill requirements on Ethernet fabrics for day 2 operation are getting lower, reducing the current skills gap. This will help further increase its adoption at much faster speed. Meanwhile, as enterprises have pervasively started adopting the hybrid, multicloud architecture, Ethernet fabrics have different form factors at the cloud and even edge, expanding its adoption.

User Advice: When there are data center network refresh projects, I&O leaders should:

- Adopt Ethernet fabrics when you have data center network modernization projects. Further, adopt programmable fabrics, if programmability, network analytics, automation, and orchestration across physical and/or virtual network infrastructure are key for your vision.

- Select Ethernet fabric as an easier-to-deploy and managed underlay network platform for your overlay network solutions (such as VMware NSX), or as an integrated network fabric across physical and virtual network to support your Hyperconverged infrastructure deployment.

- When selecting Ethernet fabrics, prefer those that can support network analytics and intent-based networking features to advance network automation to the next level.

Business Impact: Ethernet fabric has become popular in data centers as adopters claim these end-user

benefits:

- Automation — Enables end-to-end network automation by providing centralized management and policy control across physical and/or virtual network, significantly reducing human efforts and network opex.

- Availability — Enhanced network availability by dramatically enhancing network automation to reduce human errors in day 2 operation. The added-on network analytics and intent-based offerings can even further such improvement.

- Programmability — The northbound APIs provide programmable access to network resources, enabling infrastructure-as-a-code management on network. Some Ethernet fabrics even support data plane programming at switch devices, furthering the programmability benefits.

- Orchestration — Programmable control of network elements interfacing with CLI and/or open protocols and/or APIs, enabling DevOps philosophy on network.

- Network Analytics — Modern Ethernet fabrics use network analytic, AI/ML and telemetry technologies to collect and analyze network traffic, find out anomalies, and self-conduct close-loop network response, reducing human efforts and network downtime.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Arista Networks; Cisco; Dell EMC; Extreme Networks; Hewlett Packard Enterprise (Aruba); Huawei; Juniper Networks; Mellanox Technologies; New H3C Group; VMware

Recommended Reading: "Innovation Insight for Ethernet Switching Fabric"

"Market Guide for Network Traffic Analysis"

"Magic Quadrant for Data Center Networking"

"Market Trends: Evolve Ethernet Fabrics to Seize Opportunities in the Edge Networking Market"

"Building Data Center Networks in the Digital Business Era"

Cloud-Managed Networks

Analysis By: Bill Menezes

Definition: Cloud-managed networks (CMNs) enable centralized configuration and management of secure, remotely deployed enterprise network devices via a cloud-based portal. This eliminates on-premises physical or virtualized management hardware, or installation software on enterprise data center servers. CMNs are an integrated solution with a portfolio of network devices and a required, subscription-based cloud management software platform.

Position and Adoption Speed Justification: Cloud-managed network solutions have expanded beyond initial adopters with limited IT staff (such as SMBs) that need to support campus switches or WLAN access points. Such use cases typically entailed distributed/geographically diverse locations (such as schools, retail stores and branch offices). Functionality has expanded to include wired connectivity for management, provisioning, guest access, policy enforcement and other typical location network service applications. Managed software-defined WAN solutions and WAN security appliance management also now are part of the cloud portfolio. Gartner forecasts that by 2023, 70% of new WLAN deployments will used cloud-based controllers, up from 35% in 2019, keeping the technology moving to the Plateau of Productivity. Adoption will continue increasing with deployment of CMN solutions in the data center. Gartner expects 10,000 data center customers using CMN by 2023, up from about 250 in 2020.

User Advice: Organizations with limited IT staff should deploy cloud-managed networks for remote or distributed locations where ease of provisioning and ongoing administration are essential for networking equipment.

CMNs in this context are different from traditional managed service offerings, in which a service provider's staff may use a multitenant, cloud-based platform to manage multiple customer networks, and which are developed to be "cloud native."

Assess all branch office requirements before adopting the model and selecting a vendor, especially if you typically have used an on-premises application suite. Although CMNs often are chosen initially for WLAN, the model can apply to other network devices. Gartner increasingly sees vendors building a common orchestration between the LAN, WLAN and WAN, and sometimes security, with some solutions based on a cloud management platform, for example.

Remember the CMN model requires an ongoing subscription in order to maintain management functionality. Calculate total cost of ownership (TCO) for the expected lifetime of the deployment by assessing the migration path from cloud-managed to non-cloud-managed in terms of support and portfolio integration to account for potentially changing needs. Some vendors' cloud-managed hardware may not work with their on-premises management and control infrastructure; generally, vendors are unable to manage other vendors' hardware, and may have limited or nonexistent integration with on-premises applications.

Calculate from a technical and cost perspective the impact of integration with existing infrastructure,

scalability, support processes, and compliance with security policies and WAN requirements.

Identify potential issues raised by the location (country), uptime and security standards of the vendor's data center.

Business Impact: CMNs can benefit organizations that want to streamline and optimize network operations, or need to activate remote branches and manage them remotely with limited availability of on-site technical staff. Examples include retail store chains, coffee shops and restaurants, small hotels, waiting rooms in healthcare facilities, schools, small businesses, and remote offices. Although Gartner often sees vendors propose CMN as a solution for campus LANs, large enterprise sites with complex LAN and WLAN infrastructure are a less suitable use case than branch offices. Adoption by large enterprises with more than 100 branch offices remains limited. Those cases require a deeper assessment of technical compatibility with preexisting infrastructure, compliance with security policies and a more thorough ROI analysis.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Arista Networks; Cisco Meraki; Cumulus Networks; Extreme Networks; Hewlett Packard Enterprise (HPE); Huawei; Juniper Networks

Recommended Reading: "Technology Insight for Cloud-Managed Networks"

"Critical Capabilities for Wired and Wireless LAN Access Infrastructure"

"Magic Quadrant for Data Center and Cloud Networking"

Appendixes

Figure 3. Hype Cycle for Enterprise Networking, 2019

Hype Cycle for Enterprise Networking, 2019
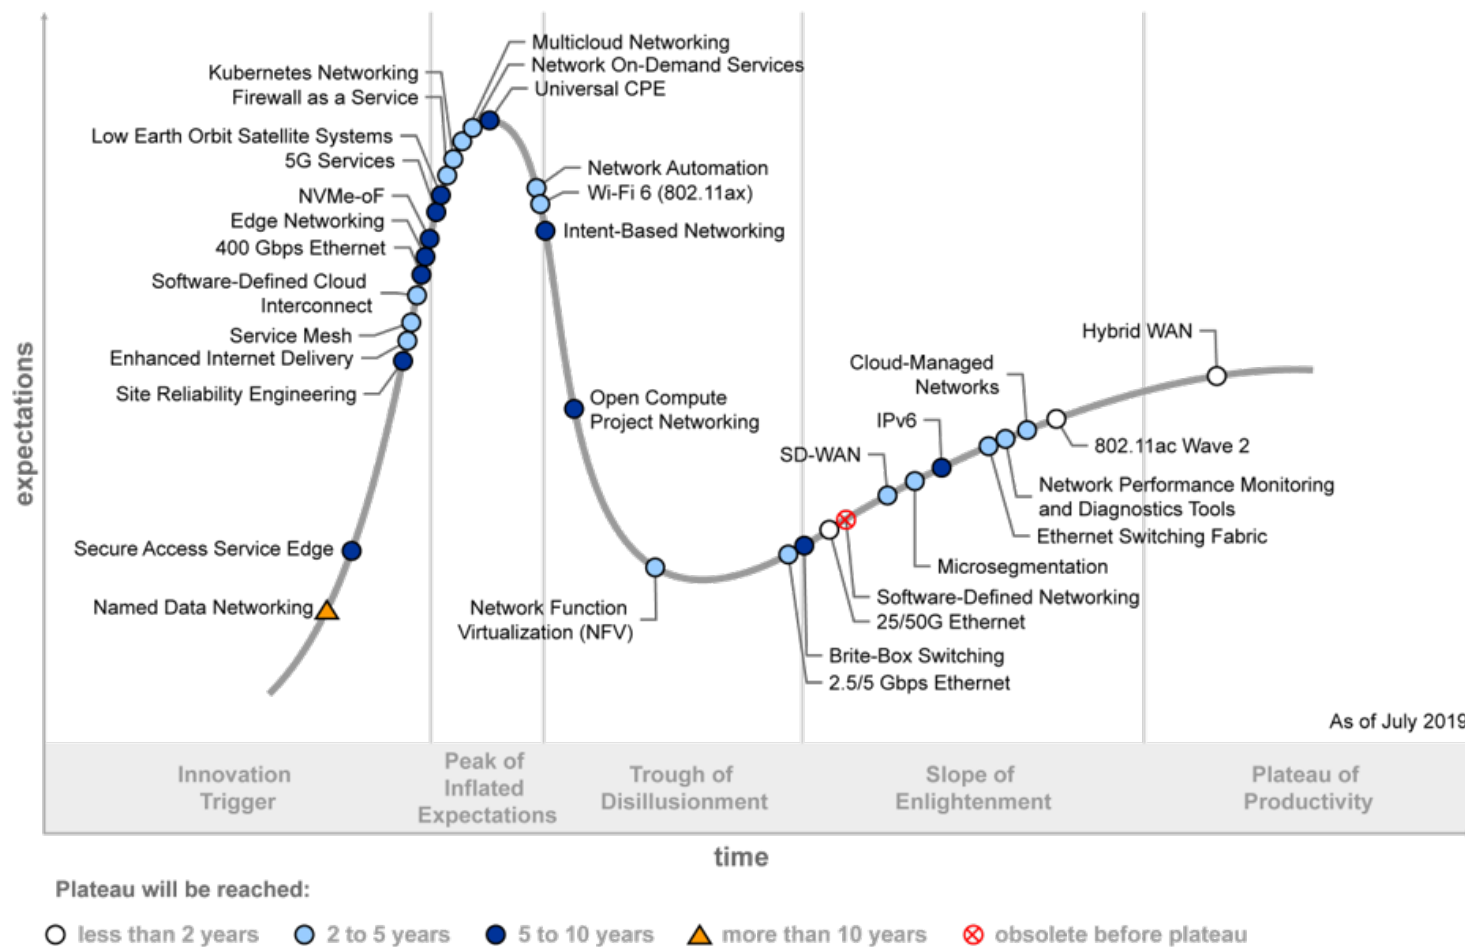
Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 1: Hype Cycle Phases

Enlarge Table

Phase	Definition
Innovation Trigger	A breakthrough, public demonstration, product launch or other event generates significant press and industry interest.
Peak of Inflated Expectations	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the technology is pushed to its limits. The only enterprises making money are conference organizers and magazine publishers.

Trough of Disillusionment	Because the technology does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
Slope of Enlightenment	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the technology's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
Plateau of Productivity	The real-world benefits of the technology are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
Years to Mainstream	The time required for the technology to reach the Plateau of Productivity.

Source: Gartner (June 2020)

Table 2: Benefit Ratings

Enlarge Table

Benefit Rating	Definition
Transformational	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
High	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
Moderate	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
Low	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner (June 2020)

Table 3: Maturity Levels

Enlarge Table

Maturity Level	Status	Products/Vendors
Embryonic	▪ In labs	▪ None
Emerging	▪ Commercialization by vendors ▪ Pilots and deployments by industry leaders	▪ First generation ▪ High price ▪ Much customization
Adolescent	▪ Maturing technology capabilities and process understanding ▪ Uptake beyond early adopters	▪ Second generation ▪ Less customization
Early mainstream	▪ Proven technology ▪ Vendors, technology and adoption rapidly evolving	▪ Third generation ▪ More out-of-box methodologies
Mature mainstream	▪ Robust technology Not much evolution in vendors or technology	▪ Several dominant vendors

Source: Gartner (June 2020)