ONUG

www.onug.net

Building the New Network Cloud

A Playbook for the Journey to Agile Network Infrastructure

Authored by ONUG Collaborative Network Cloud Working Group Members:

Steven Wood - Chair of ONUG Network Cloud Working Group, and Principal Engineer, Cisco
Mark Fishburn - President, MarketWord Inc.
Nuno Ferreira - Vice President, Technology, Volterra/F5
William Collins - Principal Cloud Architect, Alkira
Hammad Alam - Vice President, Solution Architecture, Aviatrix
John Gonsalves - Principal Field Evangelist, Aviatrix
Sal Rannazzisi - Associate Director, Merck
Zaheer Aziz - Solutions Architect, Cisco

TABLE OF CONTENTS

This playbook intends to give guidance for organizations to make these informed choices that reflect not just their business requirements but allows them to migrate to exploit the constant evolution of Secure Cloud ecosystems. The term Cloud Ecosystem is used to raise awareness that there are many types of ecosystem providers with different characteristics - including, but not exclusively, public Cloud and SaaS providers.

01 Introduction

This work describes an exciting new approach for the fulfillment of evolving business requirements: The Network Cloud. It brings a set of tools and structure for the journey to an infrastructure that has both the necessary elasticity and agility required for today's enterprises.

It transforms the entire Cloud Ecosystem into a dynamic business resource empowering enterprises and businesses in general to choose an infrastructure that supports their unique business requirements - reversing "**the way it's always been**" where businesses were forced to fit what was being provided.

This playbook intends to give guidance for organizations to make these informed choices that reflect not just their business requirements but allows them to migrate to exploit the constant evolution of Secure Cloud ecosystems. The term Cloud Ecosystem is used to raise awareness that there are many types of ecosystem providers with different characteristics - including, but not exclusively, public Cloud and SaaS providers.

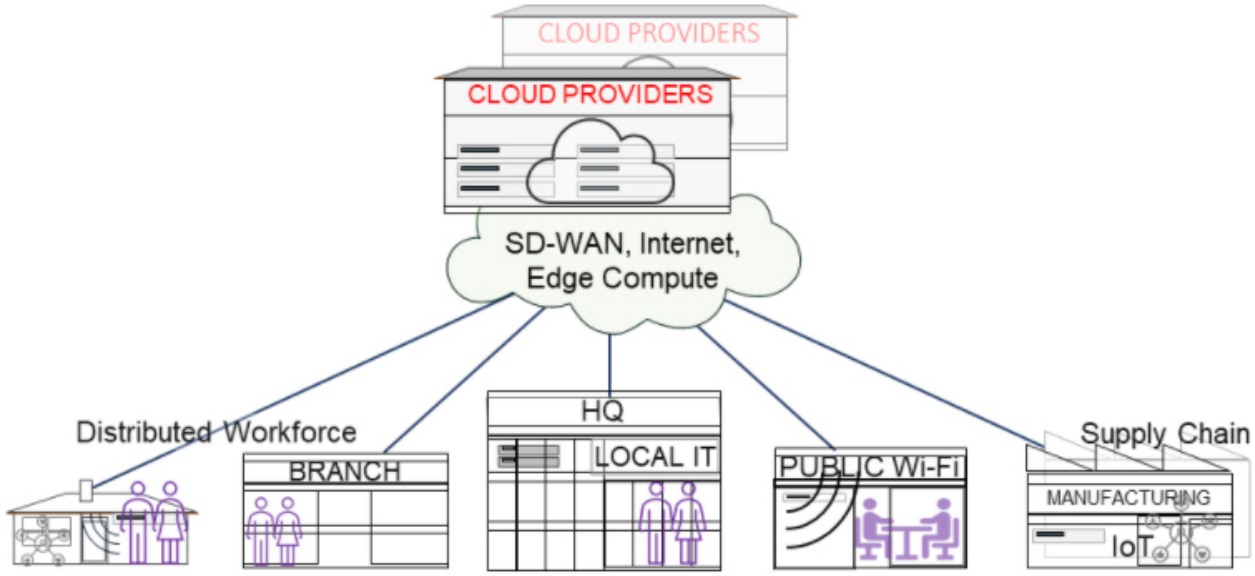

Abstracted Secure Multi-cloud Ecosystem

What makes it "**elastic and agile?**"

Its ability to adapt/scale - expand, contract and deploy seamlessly - over time as businesses grow, distribute their work forces, and change business models and policies. It also applies to the properties of the underlying network and security infrastructure to adapt and change as the business needs change, with the need for an infrastructure refresh.

The Network Cloud relies on the infrastructure's ability to automate securely identifying, controlling, connecting and monitoring users, applications and devices.

Purpose

This report describes the key attributes that enterprises, government agencies should look for when selecting cloud providers, integrators, Managed Service Providers (MSPs), or in-house management as prime contractors.

It defines a new approach to layer and structure the thinking needed to deploy an agile, cloud-like network solution for infrastructure that is almost always in a state of migration:

Network Cloud Business and Digitalized Architecture

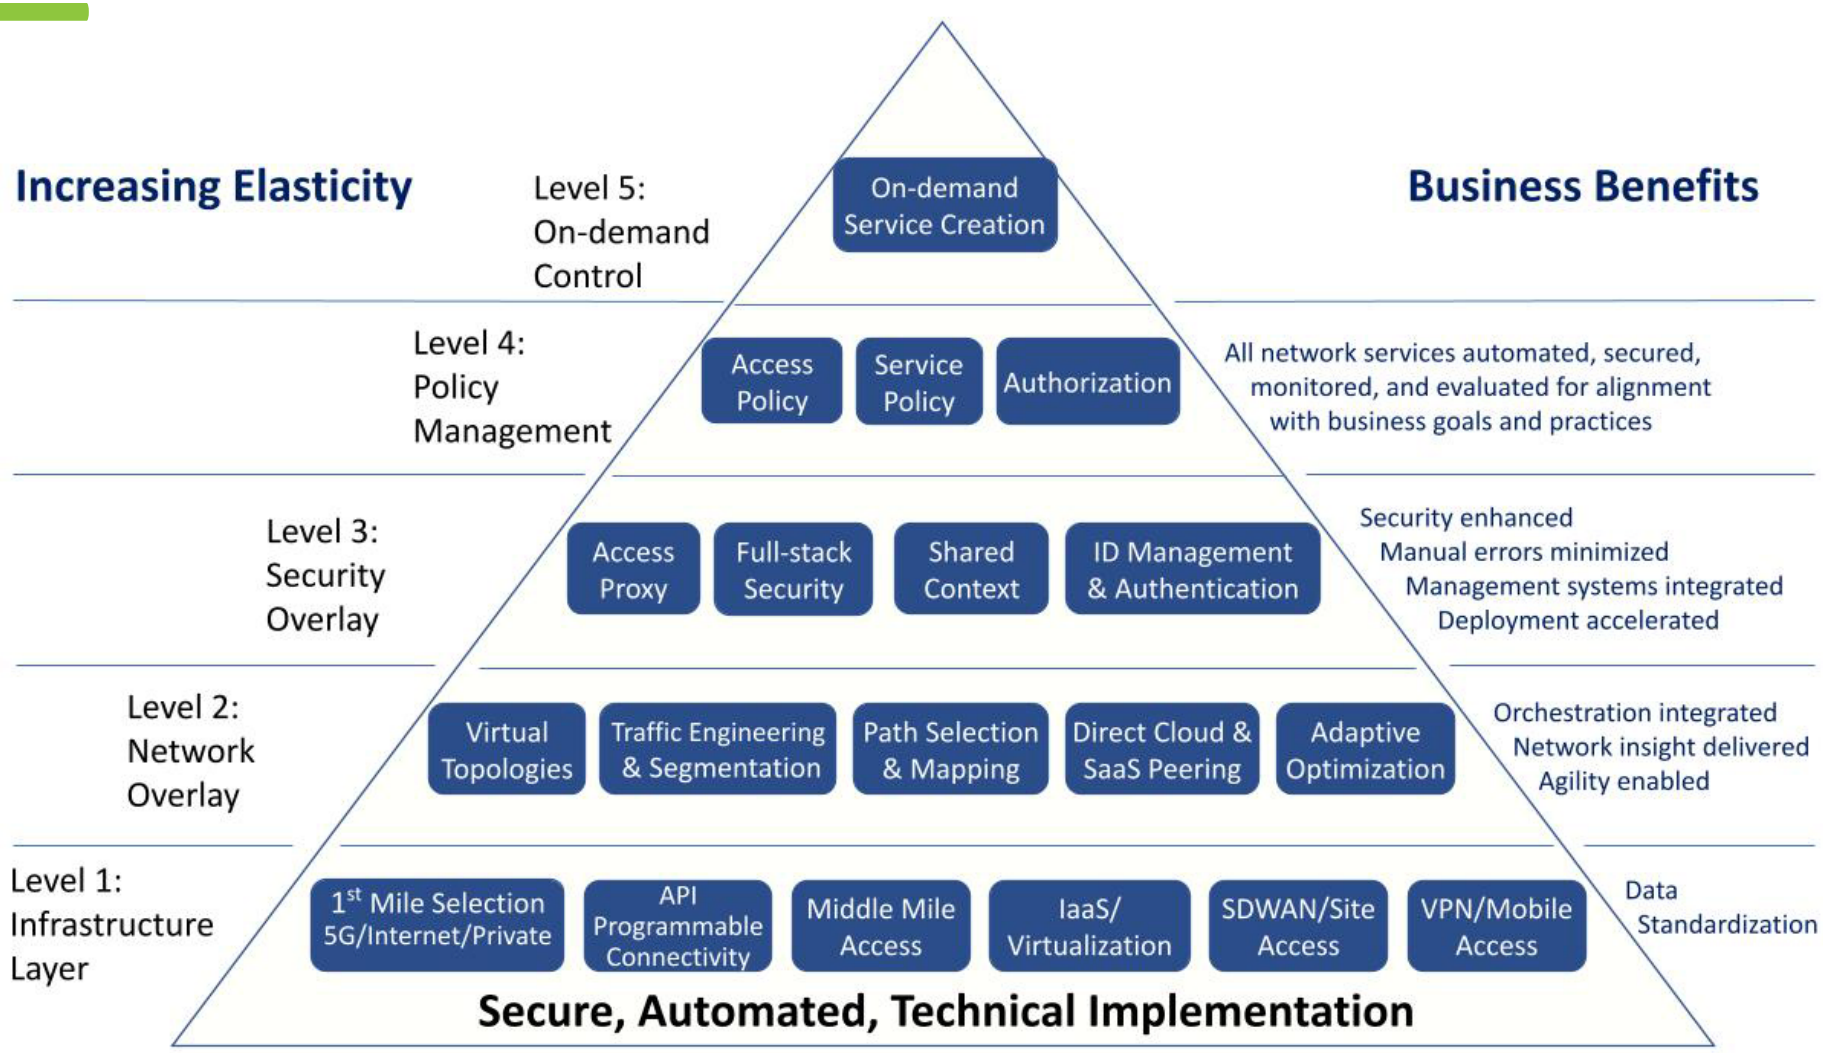

The intention is to provide simplicity, lower cost of migration and to avoid abdicating key business responsibilities.

How does it all work?

It abstracts the underlying secure services and networks, application platforms and multi-cloud connectivity. It also brings the agility to adopt new technologies such as Secure Containers, Zero Trust, 5G access, SD-WAN, multi-cloud, connected workloads and more. The intention is to facilitate smooth response to constantly evolving business conditions, policies and application requirements.

Network Cloud relies on some key transformations in the enterprise technology stack:

01 **Network-as-a-service:** deployment of WAN and cloud connectivity as a service deployed at colocation facilities allows for the properties of elasticity to be realized at the network layer. It enables more advanced services like SD-WAN and VPN over private networks, not just public.

02 **Zero trust security:** use of cloud-based security, with rich identity and posture for policy-controlled access to applications and other enterprise resources. In addition, securing and protecting traffic flows needed for fast on-boarding of users on any network access.

03 **Layered architecture stack:** abstraction of connectivity, networking, security and application infrastructure, with each layer consumed via API-exposed services. Expert knowledge of how services are rendered are no longer necessary to consume or deploy a full-service stack.

It focuses on choices that leverage and raise awareness of these developments so that we may delegate implementation to partners without abdicating the understanding of their services. Critical to migration and delegation is the use of agreed, secure APIs with an accredited code base. These services must scale in terms of capacity and performance to support enterprise demand. It also aims to simplify the complexities of the evolving multi-cloud ecosystem.

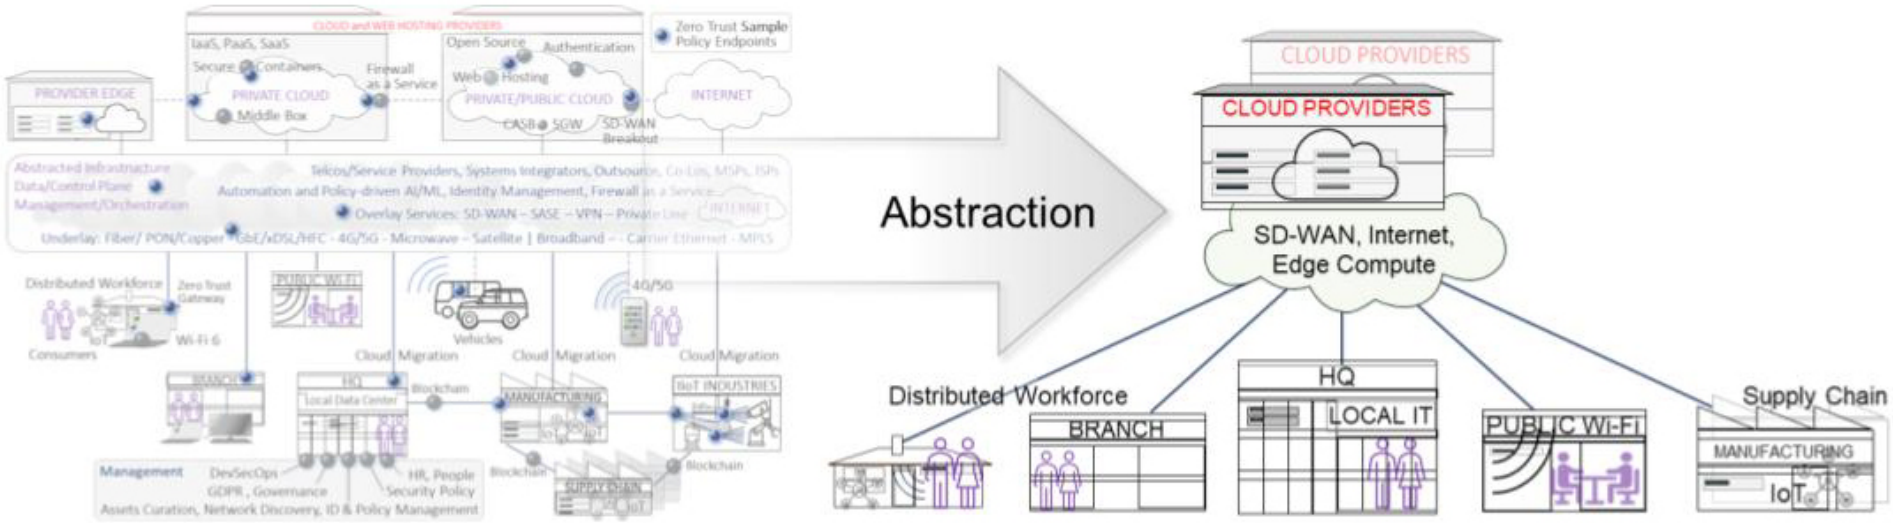

Scope

The paper covers business benefits and overview, architectures, and how to get started in your enterprise.

It is not the intention of the paper to describe a new competing or accommodating architecture or even a specific use case of SASE plus Zero Trust, Workloads or SD-WAN. The purpose is to identify the technology disruptions, available approaches and how to leverage them to best fit the end-user business goals, roadmap and policies.

02 Benefits

Introduction of Network Cloud's elastic infrastructure services has measurable benefits in terms of business agility, cost and time to market. There are many benefits to your business, network and security operations.

2.1 Business Benefits

Agility: Rapidly adjust to changes in your business, whether they be moving to a hybrid-work model, enabling secure work environments in non-traditional, uncontrolled locations and deploying new applications for your employees and customers, wherever they may be.

Faster Time-to-Market: Respond rapidly to and improve customer experience and employee and partner productivity. Deploy new access where your business needs it. Onboard new users and applications rapidly without network dependencies and connect to cloud and SaaS infrastructure without major infrastructure deployment projects. Scale and manage capacity to meet long-term and seasonal demand.

Future of Work: Meet the needs of the new distributed workforce. The Network Cloud's elastic infrastructure allows you to deploy access anywhere so you can work from anywhere. Rapidly onboard new applications and users over any network access in a controlled and secure way. Integrate remote and private access to enterprise resources in Cloud/SaaS or private DC into a single access and onboarding method

Reduced OpEx: The elastic infrastructure is built around a service-driven, competitive, cloud-like consumption model - pay-as-you-go, consumption-based pricing with capacity reservation, integrating the use of subscription-based services to enable on-demand, seamless ability to migrate to new services while simplifying or outsourcing the management and complexity of connectivity infrastructure as you digitize your business and organization.

Increased Productivity and Faster Deployments: Optimized application performance to enhance employees' productivity and provide capacity for enhanced customer experiences. This being achieved by automation and API-driven service abstractions eases deployment of new infrastructure to support your business. Faster time-to-service for IT projects via separation of concerns: networking, cloud and security

Maintain Control: The elastic infrastructure puts control of critical services in your hands via automation and service abstraction. Security, network and cloud services are in the hands of the enterprise to enable changes on your schedule, not the vendor or service provider.

2.1 Network Cloud Security Benefits

Automated, Secure Access to Resources Following Zero Trust Practices and Principles

Motivation

Rapid, secure access to data, and workflows, fast onboarding of new services that scales has always been the goal of automated systems. In a perimeter-less multi-cloud ecosystem where every element is subject to attack, protection has become the overriding imperative to protect the viability of the organization. Zero Trust is most likely the only approach that can meet the requirement.

Applying Zero Trust Principles	Zero Trust is a systematic strategy, desirable service attributes and a set of principles addressing the demise of the traditional network perimeter and focused on the un-authorised exfiltration of confidential data. The principles are applicable to network data and control plane, to data and code in containers, and governs the access across workflows in multi-cloud applications. There are no limitations on how and where it can be applied other than those of limited physical devices.
Implementing Zero Trust	The common elements to all implementations that adopt Zero Trust are Policy Management (The overall control and decision-making process), Identity, Access Control, Policy enforcement and Monitoring.
Identity Management	The identification of the requestor including allowable roles such as level of privilege, administrative, managing of resources or just simply access to a resource reflects the business and or customer role. Authentication is commonly handled via the Identity Management to ensure the requesting user, device or application proves who or what they are. Delegated authentication is often a necessary user function.
Posture and Access Control	Defines the policy associated with access to the requested resource. It includes user, device and network posture, the context of which can be used in a decision to grant resource access.
Policy Management	Evaluates the requests matching request with access policy to approve the request at the designated enforcement point or take actions where attempts were made that were not in policy or where not normal. Monitoring would ensure that out-of-character actions were guaranteed, blocked or permitted.
Business-driven Policy	Users are authenticated and authorization to specific areas is based on a defined business logic that is applied to the infrastructure - end to end. Once established as a policy – for example managing a firewall or turning up a service – then a requesting application can be automatically given access to and manage the resource.

Transport Security

Environmental Segmentation: Segmentation between routed domains will enforce network boundaries.

Encrypted Transport: Transport between user and application should be encrypted all the way through to guarantee authenticity. This begins with use of wire-speed MACsec encryption of layer 2 network traffic to help prevent intrusion and control plane attacks such as denial of service attacks on routers and firewalls.

Transport Firewall: Network policy rules for Layers 3 and 4 require enforcement on the routing processes/devices. These policies also follow the logic explained on the previous topic where the policy itself is derived from a business logic in conjunction with the transport construct / segments.

Policy-driven Networking Rate Limiting: The ability to rate limit specific applications, users or devices helps protect against denial of service attacks.

Consistency: The same security policy, and implementation framework from any source or destination. Access network independence.

- **Elastic Security**
  - Support for scale up and down and maintain security constructs to the new members or upon deleted members of the enforcement point/cluster
  - Scaling up and down should be instructed via threshold mechanisms - ex. if bandwidth consumption increases, build another node/member
  - New nodes/members adopt the same configurations as the previous ones - ex. BGP neighbor configurations should also be established by the new node (for load balancing and as such security to be maintained over this new node).
- **Policy mobility**

  Important in the serverless context where services can be spun up and down and be dealt with by different "enforcers" in a matter of minutes.

Application Security

Web Application Firewall	For signature based control. The advantage of having signatures is that it is easy for the WAF engine operator to create and update them when new attacks are discovered.
Monitoring and Reporting Suspicious Behavior	Suspicious and user behavior surveillance is critical to detect malicious users, applications and devices vs. authorized interactions. This includes rate limitations, out-of-bounds service rate, time use or other out-of the ordinary access detection, scoring and monitoring.
Time-Series Anomaly Detection	Similar to the above, time series focused on application access patterns can avoid potential issues. Correlating application errors with potential traffic anomalies can prevent adverse business impact.
Service Policies	Layer 7 service policies for granular level security or segmentation. In a world where traffic is all the same, this is key.
Network Segmentation	In the past, segmentation was used as a grouping construct to establish network identity for access enforcement at the network layer. As application access evolves towards Zero trust and user–identity based access control, its use in the context will diminish. However, this will take a long time, and the two methods will coexist. Network segmentation will continue to be used for traffic separation, access control to resources not ready for ZT migration and for asset access containment (minimizing blast radius). As such, automated segmentation, such as SD-WAN or VPN, will continue to be a necessary service for any Network Cloud deployment. Ideally, user-context should be shared between network and application layer to allow for full policy classification to be written by the enterprise application security team.

03 Infrastructure Transformation

A number of critical disruptions in IT over the past few years have contributed to the ability to realize elastic networks:

01 The movement of applications away from private data centers toward SaaS and public cloud and a similar movement of IT users off-prem to remote and mobile environments has changed the traffic patterns in the network greatly, requiring optimizations for secure and performant access.

02 The broad adoption and innovation around the software defined public data center (i.e., the public cloud and SaaS). This revolution is driving the expected consumption model for application delivery, security and networking itself towards an on-demand model.

03 Network innovation targeted at reducing latency to cloud applications as well as simplified, on-demand connection to clouds with sophisticated transport methods and protocols expected to be abstracted and delivered as a service.

03 Movement towards identity- based security policy and posture, for users and applications. Evaluation of user identity, access posture (e.g., public or private) and centralized control of classification policy, with centralized or distributed enforcement is unifying the security domain as its own user-to-application "overlay" across public and private access networks. This includes also the emerging trend of Zero trust end-to-end application layer security.

These disruptions have resulted in new technologies, services and consumption models that make Elastic Network Architecture possible. Analyzing these disruptions and the transformations is the base for understanding the properties of the Elastic Network as well as the technology architecture needed for an implementation.

3.1 The Network Transformation

The WAN Disruption

- The Internet is evolving from a network-of-networks to a network of data centers
- SDCI* & Multiple Provider Backbones
- Large Private POP and Colo footprints
- Short-term contracts, Cloud-like Service agility
- Trending toward single ISP first-mile access

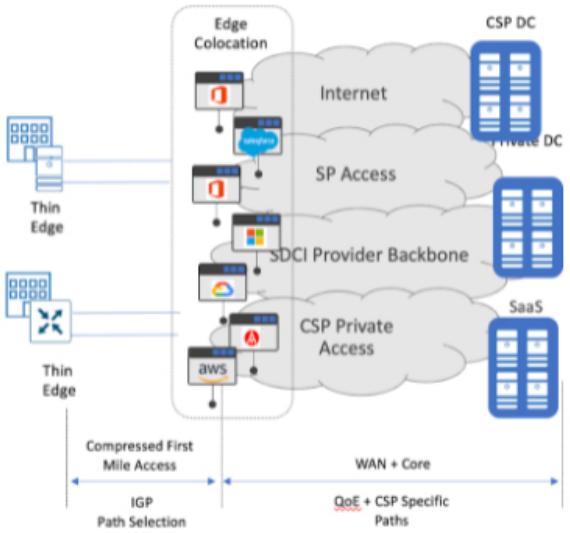

*SDCI – Software Defined Cloud Interconnect

Over the past few years since SD-WAN first disrupted the WAN networking space, the emergence and adoption of cloud-based services have meant that applications are no longer running in a relatively small number of privately owned and maintained data centers. Delivery of private applications has shifted to public cloud IaaS and PaaS providers, and consumption of third-party applications has shifted dramatically to SaaS delivery.

As a result, application traffic patterns over the WAN have shifted from private-data-center focused hub-and-spoke to "**one-to-many,**" with many new and different cloud peerings becoming the dominant pattern in the enterprise network. Many of these services were, at first, made available via the internet as it is the easiest way to deliver ubiquitous connectivity and is thus driving the rising importance of internet-based WAN traffic. However the performance, reliability and security of public internet service is not always sufficient for enterprise business digitization. As a result, the WAN architecture changed again.

These changes resulted in a restructuring of internet traffic patterns:

01 The internet is changing from a network of networks to a network of public (and private) data centers

02 Those public data centers (IaaS and SaaS) providers are peering locally to ISPs, moving the IP service edge much closer to the enterprise locations and users.

03 Traffic is being groomed from the first mile directly onto provider-specific private networks at colocation facilities, avoiding the IXP transit peering which is optimized for cost and not performance.

04 WAN backhaul is becoming a star-topology with colocation facilities at the hub, becoming the new peering point in the WAN network, replacing the private data center in this role.

05 More and more traffic is coming from non-corporate-owned sites: home offices, public WiFi access and mobile users.

06 First mile access is trending toward a single ISP for the first hop with all the cloud providers directly peering with ISP data centers. The goal is to eliminate IXP paths and hot-potato routing policies that congest access via legacy public internet.

WAN network delivery is evolving to meet these changes with many new entrants bringing high performance networks with new API driven, cloud-like business and consumption models. One goal is to attract traffic to optimized network infrastructure, and get your traffic off the public internet as soon as possible. Another is to transform the process of WAN and Cloud connectivity to an on-demand service.

Software-Defined Cloud Interconnect or "**SDCI**" providers are deploying global high-capacity networks that are programmable and controllable via APIs. This allows on-demand creation of network connections to any cloud datacenter region or region-to-region connections in a matter of minutes. Service agility, and short-term monthly contracts, allow you to think about your WAN deployment in a much more business agile fashion, turning up or down WANs to modulate your infrastructure costs in a nearly on-demand fashion.

SD-WAN and VPN services are being integrated with SDCI to provide "**full-stack**" SD-WAN service. Integrated APIs between the SD-WAN and SDCI WAN layer allow both overlay and underlay connectivity to be deployed from the same SD-WAN service automation platform. Through the power of APIs, enterprises can build an SD-WAN aggregation network on-demand in minutes, with a global footprint of aggregation sites worldwide. Enterprise sites and remote clients can connect to the nearest SD-WAN aggregation site for private cloud, public cloud and SaaS access.

The bottom line: Cloud, SaaS and networking vendors are deploying SD-WAN and Cloud service points-of-presence at hundreds of colocation facilities worldwide. These edge PoPs are within most major metros and thus close to likely very close to your enterprise sites, and mobile workers customers and employees, who need access to applications from both controlled (**e.g., on-premise private networks**) and uncontrolled networks (**e.g., public Wifi and Internet**).

3.2 Security Transformation

Traditional network security struggles to keep up with rich and evolving applications

- Constant application innovation/changes
- Proprietary data structures
- Protocol evolution (web sockets, TLS 1.3, QUIC, cert pinning)
- Protected (encrypted) data
- Extending Zero Trust for SaaS with new dimensions: DESTINATION BASED DIFFERENTIATION

Traditional security, therefore, fell behind the application evolution and still operates as an infrastructure component with little or no integration with the application itself. It is based on 5-tuple rules that allow or block traffic but are not aware that applications move and potentially are reachable on other interfaces or data centers. This lack of discovery services makes traditional security very rigid with lack of elasticity. **Furthermore, creating traffic rules based on 5-tuples is no longer effective because over 80% of your traffic is now either a tunnel or an API**. This change is very important because if the traffic is potentially all the same (same or similar 5-tuple) then traditional security and microsegmentation will be ineffective.

Key to elastic infrastructure security is the need for end-to-end visibility and control. Today, this is not possible due to the native separation of devices or services and especially the teams who operate them.

There are multiple devices – well, let us call them services – needed to accommodate two simple things: users to talk to apps and apps to talk to other apps, which include services like:

- User Access
- DDoS Protection
- Routing
- API Gateway
- Firewalling
- Application Delivery Network or Controller (Load Balancer)
- WAF
- Ingress / Egress Gateway (Proxy)

Just to mention a few ...

Across these services, there are multiple teams that operate them, and these are normally isolated and not often communicating between them. It is, therefore, necessary to create a collaborative environment that allows these teams to share the same view of the services and infrastructure.

Ideally, this environment should be capable of correlating events between these services so that the teams can quickly identify issues and resolve them instead of praying for the problem to disappear because by then the damage has already been caused to the business.

These services, called "**the plumbing,**" need, more and more, to be integrated with the applications they are serving. Consider an example where an application that has three or four instances (replicas or just copies) or even a cluster that is highly available and therefore has multiple copies of the application. When these three or four instances are all working together, we are all good – the clients / users are happy, and the application admin is also happy. When there is a problem on the infrastructure (on any layer) that causes the application to start behaving differently, there is probably a need to get back to an optimum scenario, and this might require changes. For these changes to occur, there needs to be a good level of symbiosis between the "plumbing" and the application so that, for example, the application creates some more replicas of itself to try to accommodate the changes or the extra number of requests / traffic.

When this application creates more replicas, the infrastructure must also accommodate and adapt to discover these new replicas or copies and start sending traffic to them as well.This requires the infrastructure to be elastic and not relying on Operations teams to manually go through a change control process and execute them.

When application components scale up they need to be automatically discovered and able to reach and serve each other and/or clients/users that need access to them, they also need to accommodate all the security principles installed and deployed for the previous replicas. This is where identity comes into play.

When a new replica of an existing application is created, there needs to be an identity service to bootstrap and identify it (and possibly check if the replica is indeed a genuine replica or has been changed / compromised) so that the application receiving traffic or sending traffic, can identify itself to the infrastructure, client or user and maintain the same level of security in place. For example, the application could be requesting encryption by presenting its client certificate.

Regarding the infrastructure itself, if the services are aware of each other and information is correlated, then when new replicas are created, they will maintain the security services that were defined for the service. For example, if the operator decides that for this application the traffic must go through a firewall, then to a load balancer, a proxy and then through an API gateway, the replicas that are created as part of a scaling-up construct, will automatically flow through these services too.

This scaling behavior is very important because it is how your infrastructure will also evolve and scale (up, down, dynamically ...). Any infrastructure service should be able to accommodate a similar behavior as described above for this application.

04 Architecture: An Overlay Approach

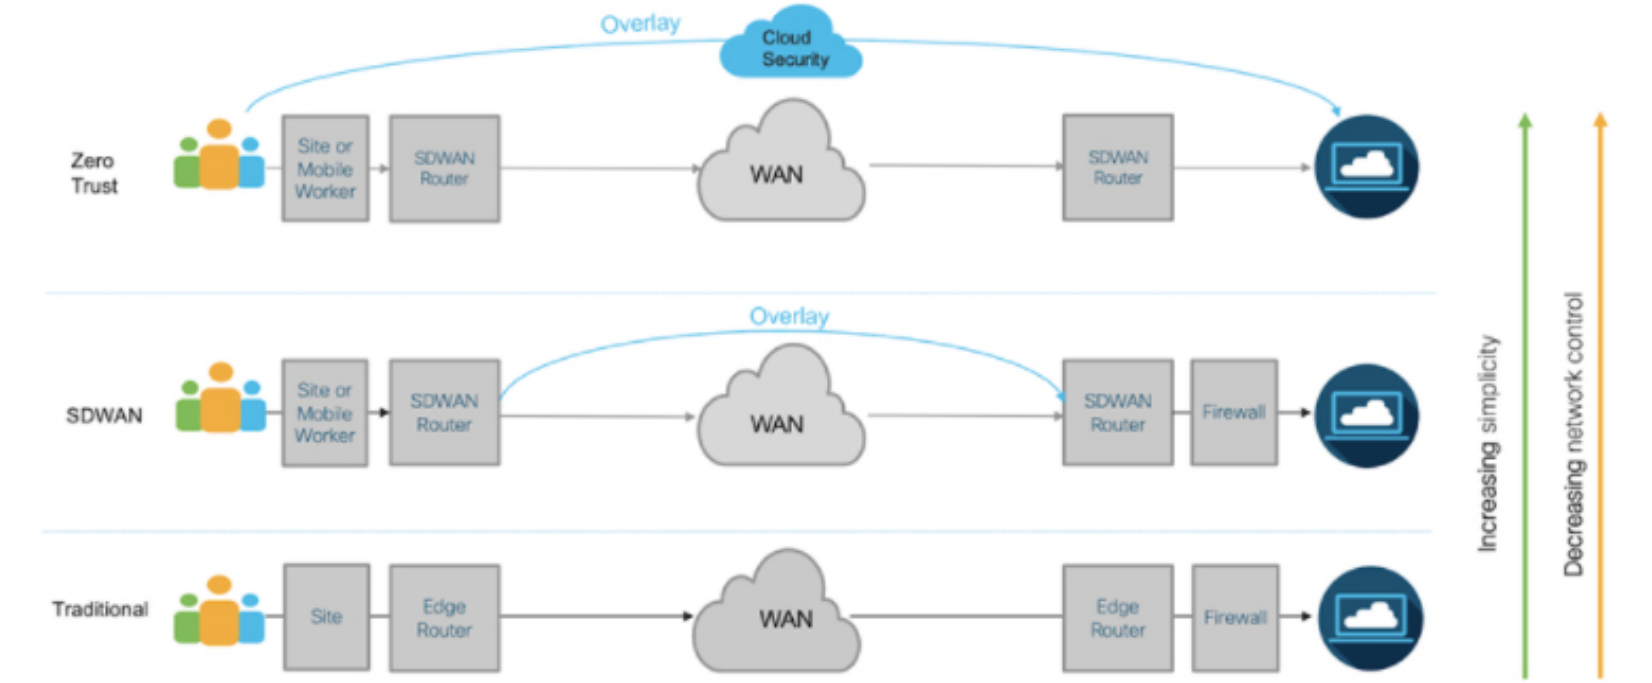

A good way to model the transformation of application access is to think in terms of overlays. Overlays have the property of simplifying the overall problem by the power of abstraction. SD-WAN overlays abstracted the IP network. One could simply use the any-to-any connectivity of the IP network to establish a network of virtual tunnels in a highly automated way, without being concerned about complex routing, multipathing and resilience needed in the underlay network. It allows the IT professional to focus more on connecting to SaaS and IaaS endpoints and delivering application SLA policy. The underlying IP network properties are simply **"consumed."**

This is how we approach the deployment of elastic networks as well. In this case, we treat the security layer as an overlay as well since security applies to all layers of control and management plane.

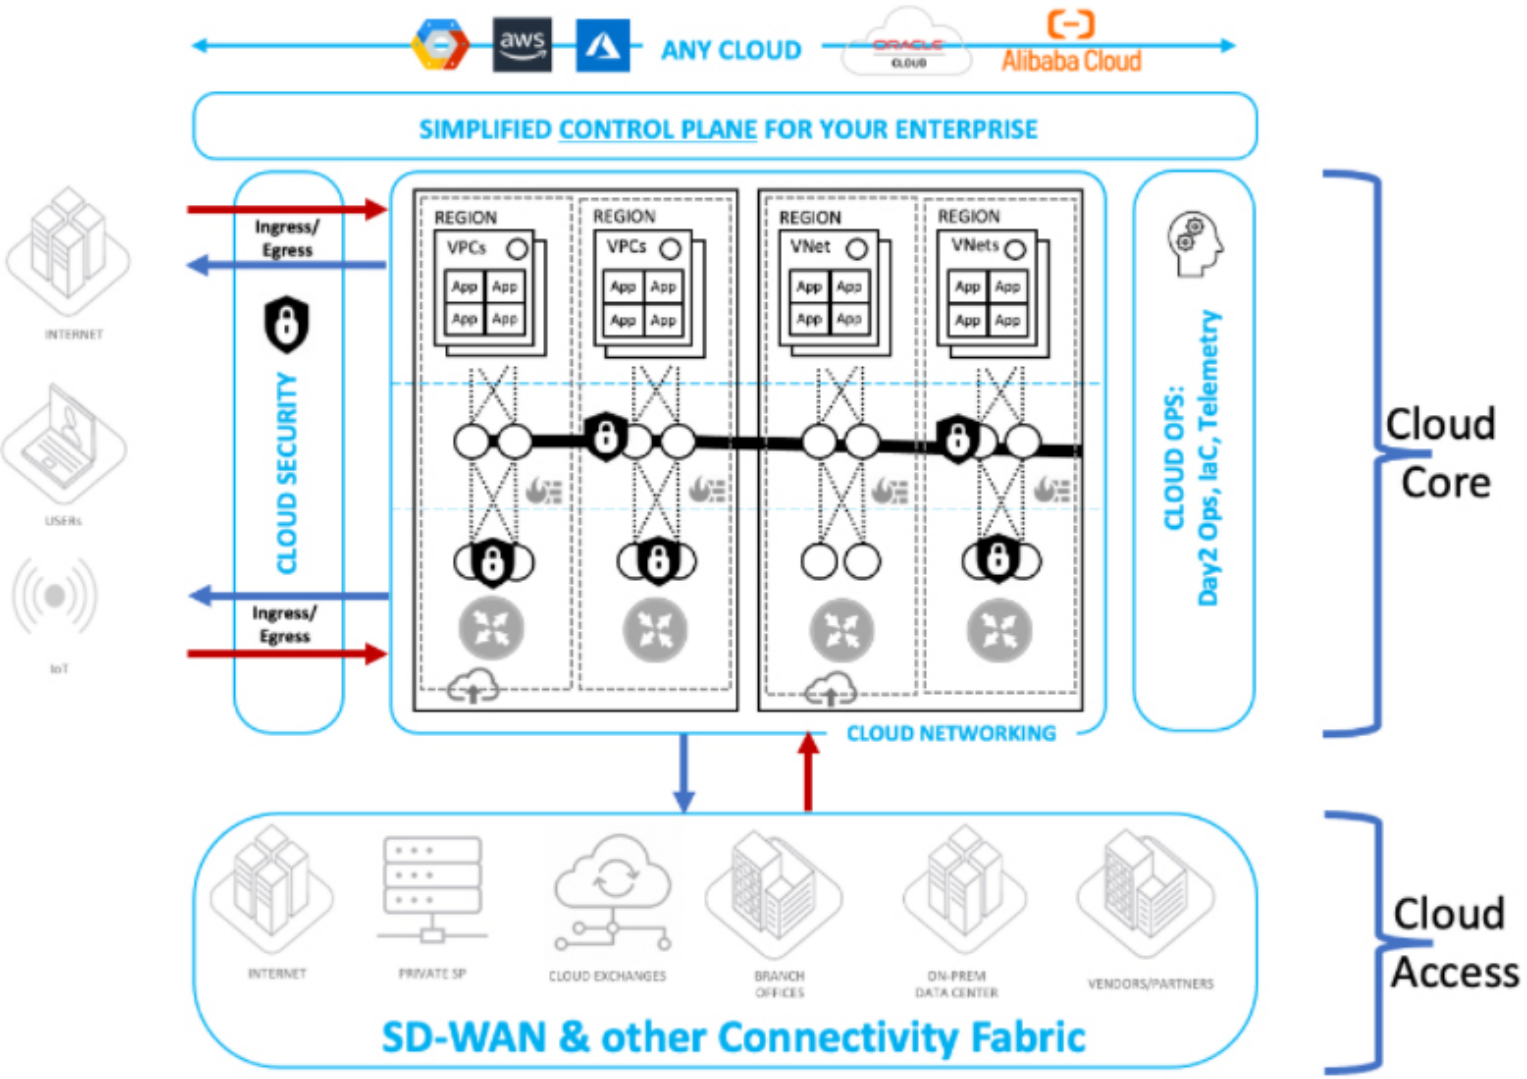

With this in mind, elastic connectivity must integrate with an architecture which is well-aware of connectivity inside the cloud, with "**cloud"** not being limited to a single region of the first cloud. Instead, the architecture must give multi-region and multi-cloud readiness without requiring a redesign when business demands infrastructure teams to support more providers.

The network transformation in the cloud should not wait till you are in the cloud. Its planning happens at the same time when you are planning your migration to the cloud architecture. **This architecture must take into account a reference Multi-Cloud Network Architecture (MCNA) approach which provides:**

- Secure, encrypted and high throughput connectivity to the cloud
- A repeatable architecture that helps you build in the first region of the first cloud efficiently
- Ensure Multi-Region and Multi-Cloud readiness
- Integration with SD-WAN for branch connectivity
- Integration with private path connectivity from cloud providers
- Integration with branch and other external connectivity over internet
- Ability to leverage internet for choice based internet egress and ingress services

4.1 Cloud Core

Cloud Core is providing the core features that are required to run the Enterprise applications in the cloud. It binds application workloads with network and security services. The Cloud Core layer can be viewed as a service bus. There is no concept of tiered DC or leaf-spine Architecture in the native Cloud hence one should not directly compare it with On-Prem DC Architectures. This is a key aspect of elasticity provided in the cloud - abstraction of service consumption from infrastructure and the on-demand horizontal scaling that it provides.

Cloud Core acts as a robust foundation that can span multiple public clouds and provides a common networking and security layer where application workloads are attached. In addition, it provides a service extension framework for advanced services like encryption, next-gen security such as deep packet inspection, traffic engineering, network correctness intelligence, etc. The workload attachment or connectivity within the cloud core layer is encrypted. In essence networking and security services are provided by this unified Cloud Core entity. The Architecture is flexible enough to extend these Cloud Core services all the way close to workloads if required based on specific Enterprise requirement(s). Cloud Core includes the following aspects :

- **Workloads:** Dev / Prod / etc. workload/instances/VMs deployed in VPC/vNET / etc. in collapsed or distributed manner
- **Management:** Shared-Services VPC for management instances/VMs
- **Operational:** Logging tools
- **Security:** applies at different layers. Security is a shared responsibility between the Cloud provider and Enterprise hosting their workloads. It can be provided by different options and means for example
    - L2-L4 Firewalls, NGFW, WAF, Security Groups, ACLs, etc.
    - Ingress and Egress Security for Internet traffic

Cloud Core also provides a **Service Extension Framework (SEF)** to insert new services on-demand or based on certain policies.

4.2 Cloud Access

The first area of focus for enterprises on their digital transformation, migrating to the public cloud is often getting to the cloud. Challenges include using public internet versus getting private connectivity circuits (like Direct Connect, ExpressRoute, Cloud Interconnect, etc.), integrating with existing MPLS provider contracts and getting the public cloud added as another leg to an existing network that connects data centers and branches.

Another major task to take into account in the first phase of connectivity to the public cloud is to understand how the networking would look inside the cloud. Networking-in-the-cloud and networking-to-the-cloud needs to work hand-in-hand for a smooth end-to-end experience for the applications residing in the public cloud.

Cloud Access components cover the technology and components required to connect Cloud to the On-Prem resources. The reality is that almost all enterprises are moving to Cloud but most are not 100% there yet and have a need to connect to on-prem applications: it will take some time for them to migrate entire workloads to the Cloud. This could be a lengthy process spanning months or even years as some workload might not be suitable to run in the Cloud due to compliance, legacy architecture, application complexities or other legal reasons.

With the shift to work-from-home and hybrid work becoming a dominant trend in the global workforce, any elastic network architecture must consider not only traditional branches, partner locations and data centers, but also mobile and work-from-home modes - which have traditionally been considered separate access domains operationally.

Creating an elastic network architecture requires abstraction of the plethora of different connectivity types, ideally unifying them with a single common overlay service like SD-WAN or IPSec VPN. **Some of the most common connectivity types are:**

- Internet - Public Access
- MPLS, Broadband, 4G/5G wireless
- Cloud Direct Connects - Private Links (Hosted or Enterprise-managed)
- Branch, DC and Remote Users connectivity through SDP SaaS solutions

While IPSec VPN is a good short-term access method for adding "another leg" on the existing private WAN, it is not the optimal access method for fully distributed elastic network access to cloud and SaaS. Managing and scaling hundreds of IPSec tunnels one-by-one quickly becomes a limiting factor for elastic automation.

Thus technologies such as SD-WAN - the dominant on-premise access method are expanding to cover client VPN and home gateway access with the goal of unifying both on-premise and off-premise access under one single access architecture.

4.2.1 SD-WAN for Network Cloud Access

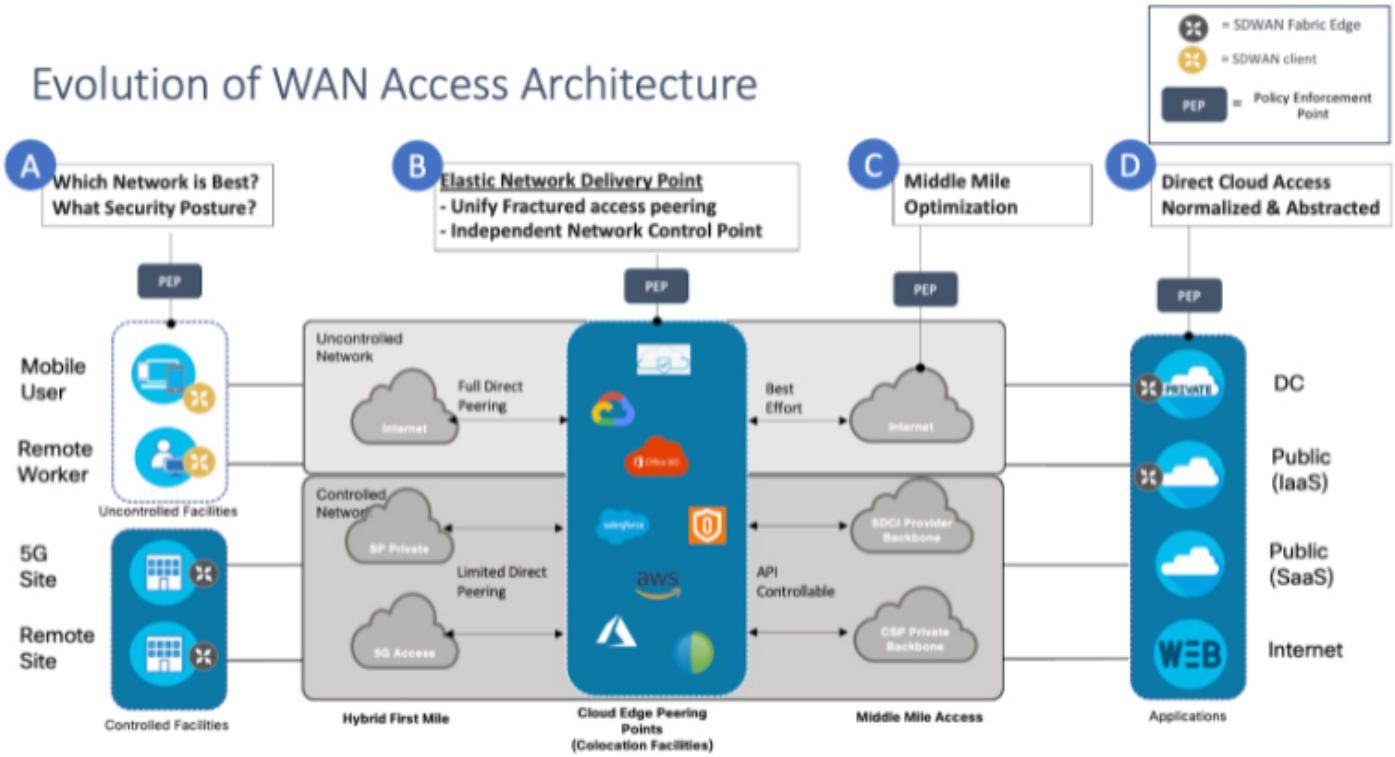

The changes in application consumption, traffic patterns and WAN structure, coupled with the emergence of remote work, lead to a new network structure as shown in the figure above.

What does this mean for WAN design and Network Cloud access? The key shifts are that Policy and service consumption is evolving to three major points in the network: the client devices themselves (private site-routers or endpoint clients/devices), the colocation edge, and the cloud datacenters. The colocation facility is a confluence point for all this change and is a key point in the network for consuming deploying elastic services for both corporate branch locations as well as teleworkers and mobile workers.

Looking at the network and the key policy enforcement points, the network pattern is breaking into three separate areas of interest:

01 The first mile access. A Single ISP domain, avoiding inter-exchange routing

02 **The Colocation edge** – an aggregation point where services can be delivered and a unified peering point for connection to services (both cloud and on-prem)

03 **The Middle Mile access:** where high capacity, high speed private access networks built by cloud and SaaS providers carry traffic at low-latency to application endpoints. For example, public cloud DCs and SaaS.

What does this mean for SD-WAN and WAN networking in general?

Recall that SD-WAN is an overlay networking solution that goes endpoint to endpoint – so policy control is primarily at these endpoints. Fundamentally, we have a new policy enforcement point in the network – the colocation edge, which we won't be able to see or control in a traditional SD-WAN if we go over the top of it. So let's look at how this affects deployment.

The first key policy enforcement point is at the client router or device: what device is being used, what app is being accessed, and is the network controlled or uncontrolled? These are all decisions that need to be made at the client and with access control determined globally by enterprise policy. Typically for SD-WAN this is network selection based on performance, load balancing or other traffic engineering policies. For elastic networks, client capabilities need to be unified for endpoint clients like phones/laptops as well network clients like routers so that security and access policies can extend across controlled and uncontrolled networks as well as all types of endpoint devices.

An emerging policy control point is the Colocation facility. These "**colos**" are dispersed across metro regions globally, with very close in proximity to sites and users. This globally distributed network edge is an independent control point that can provide direct peering access to cloud and SaaS provider networks across high capacity, low latency networks as well as global WAN networking. Today, this peering edge is fractured – many different providers providing different access methods to peer to services one-by-one.

The co-location facility is a key point in the network to achieve elasticity and provide direct peering to IaaS and SaaS services over performance optimized networks. It's a key control point for cloud service access, network interexchange and policy-based application routing. It is here that SD-WAN services can be used to unify fractured cloud peerings into a single access service method - API or GUI driven to simplify and speed deployment of new cloud and SaaS access. In order to achieve elasticity, new **"network cloud"** providers are deploying infrastructure and automation to enable on-demand consumption of scalable SD-WAN services that can be added into your existing SD-WAN overlay as and when needed, without the need for the enterprise to deploy any network infrastructure.

This new approach amounts to a "**Cloud-based**" aggregation site with full, preprovisioned public cloud access that can be deployed, connected and scaled programmatically via automation APIs and GUIs thus bringing the on-demand consumption properties of Cloud to the network.

The next control point is middle mile access selection. Here you can control traffic routing to Public cloud, SaaS and on-demand transit networks for creating WAN core networking. SD-WAN is fast becoming accepted as the unified access method here to do intent-based network selection and traffic mapping to SaaS, Cloud and Site-to-site direct peering over high speed, high capacity networks.

The final policy control point is the Cloud edge itself. The goal here is to deliver a unified access and policy controls right to the edge of the VPC infrastructure. SD-WAN can be extended here to bring your cloud data centers directly onto the SD-WAN, unify the access method and policy control between Cloud and Branch and indeed even be extended across the Cloud provider backbone to create a unified SD-WAN Core solution.

4.2.2 Role of Edge (IoT, 5G) in Cloud Access

With the amount of data that is generated these days by IOT and 5G type devices, it is important to provide computational resources close to where data lives. This gives rise to edge computing. The Cloud Edge components in architecture caters the need for these new connectivity points to the cloud.

Cloud Edge allows for significant improvement when it comes to data processing closer to the device generating or consuming that data. While both Cloud Edge becomes part of Cloud Access placed in between the on-prem devices and the Cloud Core, newer initiatives by cloud providers such as AWS Outpost, AWS Wavelength, Azure Stack, etc., are adding new dimensions where Cloud Core functionality needs to be extended to the edge.

4.3 APIs and Automation

To obtain the benefits of Elastic Network Architecture, automation plays a critical role. Network automation from the customer perspective has been hampered in many ways by variety. By adding many vendor devices, platforms, orchestration tools and unique protocols, and transports over which communication occurs, the opportunity for consistent, repeatable and reliable automation steadily decreases.

Runtime Automation Approach (Mutable)

Handling automation in this type of scenario, a Site Reliability Engineer would first provision a network construct. Once the construct is up and running, the Engineer would then use additional tooling to handle run-time configuration, e.g., packages, policies, rules and integrations. This creates a distinct separation between operational and desired state of the network. As a result, obstacles are presented when horizontally scaling. **Attributes of runtime automation include:**

Tight Coupling	Components have a heavy dependence on each other, causing cascading impact. The blast radius is large in outage events, causing a decrease in overall changes completed.
Mutable Infrastructure	Physical or virtual devices are configured, updated or modified in place at runtime. This introduces configuration drift over time leading to inconsistency.
Vertical Scaling	Network infrastructure scales by increasing the capacity of individual nodes on the network maximizing the power of individual devices or pairs of devices.
Inconsistent Interaction Surfaces	The combination of vendors, hardware models and software versions is generally present. CLI and API interaction surfaces vary in functionality and interoperability, leading to a combination of manual intervention, scripts and disjointed toolings.

Buildtime Automation Approach (Immutable)

Handling automation in this type of scenario, a Site Reliability Engineer would handle the configuration component, e.g., packages, policies, rules and integrations, and then provision the construct completely intact. Public cloud has had a direct impact on how automation is thought through. The opposite of the attributes listed above is embraced with immutability in Infrastructure as Code (IaC):

Loose Coupling	Components are detached, enabling them to work independently of each other as part of a larger group of systems.
Immutable Infrastructure	Infrastructure is configured at build time. Reducing the number of moving pieces at run time increases environment consistency, reliability and grounded DR capabilities.
Horizontal Scaling	Overall capacity is increased by adding additional nodes, usually of equivalent capacity.
Consistent Interaction Surfaces	Infrastructure is provisioned via proprietary provisioning tools or cloud-agnostic provisioning tools. In AWS, Terraform would use a provider plugin that leverages the AWS GO SDK to build infrastructure. This provides the same interaction surface across all AWS infrastructure products.

Advantages in Practice

Taking a **"shift-left"** approach by moving these components into the build process forces the appropriate teams to get involved earlier in the process. Instead of cloud security practitioners being brought in to remediate an environment that is QA getting ready to go Prod, this approach would integrate security policy as early as initial development pull requests to version control when standing up the initial infrastructure. With this methodology, risky configuration is also less likely to make it to production.

Shifting the way infrastructure is deployed can also remove barriers between the application and the infrastructure that underpins it. Imagine a scenario where product teams are rapidly updating and releasing new versions of their application. When the corresponding infrastructure gets updated in-place, outside of the whole process of the application, there is a higher likelihood that undesired behavior will occur. By handling all the necessary infrastructure required for an application, in the same lifecycle as the application, outcomes can become more predictable and reliable.

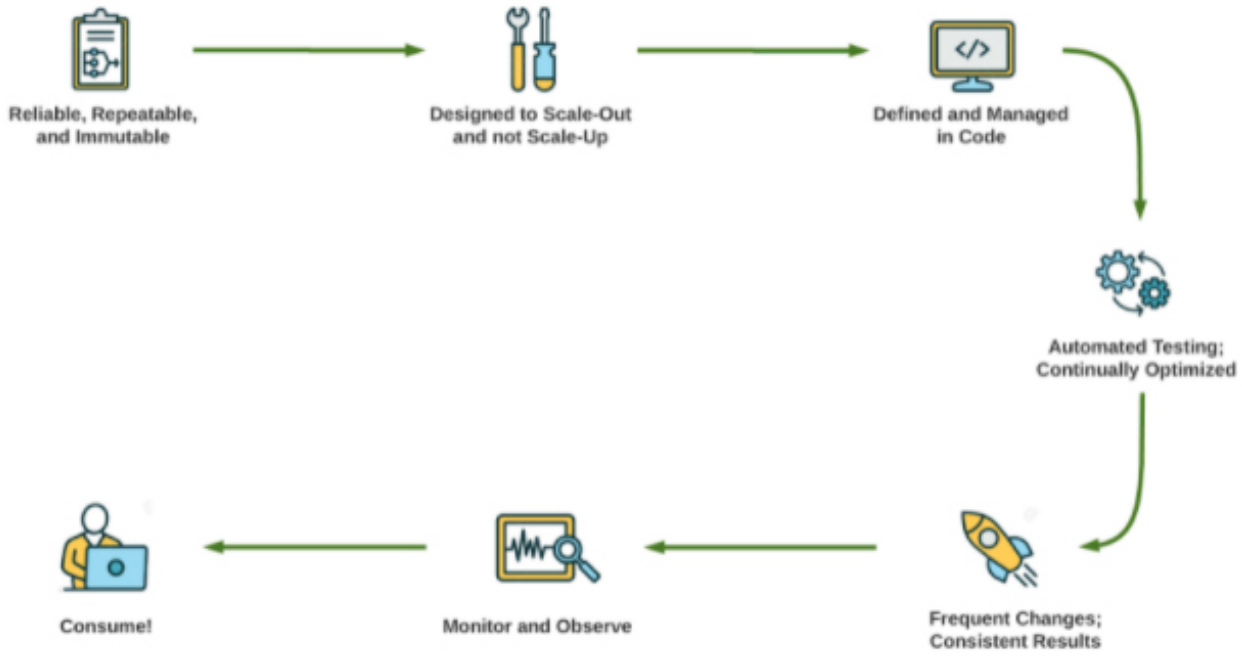

05 Complimentary Architectures

5.1 SASE

SASE is intended as a simplified WAN and security solution for a mobile, global workplace that relies on cloud applications and data. The formerly common solution of backhauling all WAN traffic over long distances to one or a few corporate data centers for security functions adds network latency when users and their cloud applications are globally dispersed, rather than on-premises. By targeting services to the edge at the connection source, SASE eliminates the latency caused by backhauling.

SASE brings together SD-WAN and a number security functions, usually including Cloud Access Security Brokers (CASB), Secure Web Gateways (SWG), antivirus/malware inspection, virtual private networking (VPN), firewall as a service (FWaaS) and data loss prevention (DLP), all delivered by a single cloud service at the network edge. The goal is to combine the cloud security functions with the transport agnostic overlay capabilities of SD-WAN stack to reduce deployment complexity.

SASE SD-WAN service enhancements may include traffic prioritization, WAN optimization and converged backbones to enhance reliability and maximize performance.

WAN and security functions are typically delivered as a single service at globally dispersed SASE points of presence (PoPs) located as close as possible to dispersed users, branch offices and cloud services. To access SASE services, edge locations or users connect to the closest available PoP. SASE vendors may contract with several backbone providers and peering partners to offer customers fast, low-latency WAN performance for long-distance PoP-to-PoP connections.

5.2 SDCI

Software-Defined Cloud Interconnect (SDCI) technology is a new type of interconnection service to connect an enterprise to a large variety of cloud, network, SaaS and internet service providers. SDCI uses an API driven approach to create a Cloud Exchange service at public colocation facilities, that Enterprises can use without the need to deploy and maintain their own network equipment. The SDCI provider deploys the infrastructure and provides a virtual service on top to each Enterprise customer. This service can support the use of multiple cloud providers and to support multi-cloud applications, interconnecting two or more Cloud Service Providers without traversing the Internet. This may sound like an MSP WAN service, but the critical difference is the consumption model: portal- or API-driven service creation, consumption-based billing, and full-stack service automation transforms the private WAN to an on-demand cloud service allowing you to deploy a private WAN in minutes. SDCI vendors operate high performance programmable private networks that are pre-plumbed to multiple cloud and SaaS providers to offer customers fast, jitter-free and low-latency WAN performance for PoP-to-PoP and PoP-to-service connections.

Software-Defined Cloud Interconnect is complementary to SASE. It provides the WAN access layer delivered as a single service at globally dispersed points of presence (PoPs) located as close as possible to dispersed users, branch offices and cloud services. To access cloud security services for SASE, edge locations or users connect via SD-WAN to the closest available PoP and then cross-connect from there to the cloud security service of choice. One of the key differentiators between SDCI and SASE is that SDCI automates the provisioning of the overlay (SD-WAN) network, the underlay networks (especially new programmable middle-mile networks that are pre-plumbed with multi-cloud connections) and the cloud-resident connectivity.

5.3 SDP

SDP, or Software-Defined Perimeter, is a security framework that controls access to resources based on identity. By establishing a perimeter via software versus hardware, an SDP hides an organization's infrastructure — regardless of where it is located — from outsiders, while enabling authorized users to access it.

The framework is based on the U.S. Department of Defense's (DOD) Defense Information Systems Agency's (DISA) "**need to know**" model from 2007, in which all endpoints attempting to access a given infrastructure must be authenticated and authorized prior to entrance. In 2013, the Cloud Security Alliance (CSA) released its SDP working group guidance, which incorporated elements of DISA's work with security standards from the National Institute of Standards and Technology (NIST) and other organizations.

SDPs provide secure access to network-based services, applications and systems deployed in public and/or private clouds and on premises. The SDP approach is sometimes said to create a "**black cloud**" because it obscures systems by cloaking them within the perimeter so outsiders can't observe them.

SDP software is purpose-built to give medium and large organizations the perimeter security model needed for zero trust applications and workload-centric network connectivity. In addition to reducing the attack surface, SDP's virtual boundary around the network layer also eliminates vendor chaos by allowing for installation on any host, without network reconfiguration or appliance lock-in.

5.4 Network as a Service (NasS)

Network as a Service brings Software-Defined Networking, programmability and API driven operation to the WAN, hybrid cloud, multi-cloud and interconnection services. NaaS models are inclusive of integrated software, licenses and support of services that are delivered via the NaaS platform in a flexible consumption or subscription-based model. An enterprise that consumes NaaS should simply have the ability to consume and scale products and services as required based on their business scale.

06 Network Cloud and your Enterprise: Journey Map

As enterprises are going through their digital transformation focused around public cloud adoption, they find themselves with different kinds of challenges depending on their cloud maturity and cloud journey.

The chart below ties together the business and technical elements with an intent-based approach driven by the business policy. The intention is a visual checklist to make sure that necessary tactical elements are addressed in a structured order. The idea is that the more decisions are made in a top-down manner, the more it can follow changing business requirements and innovation or experimenting locally has a framework into which it can fit.

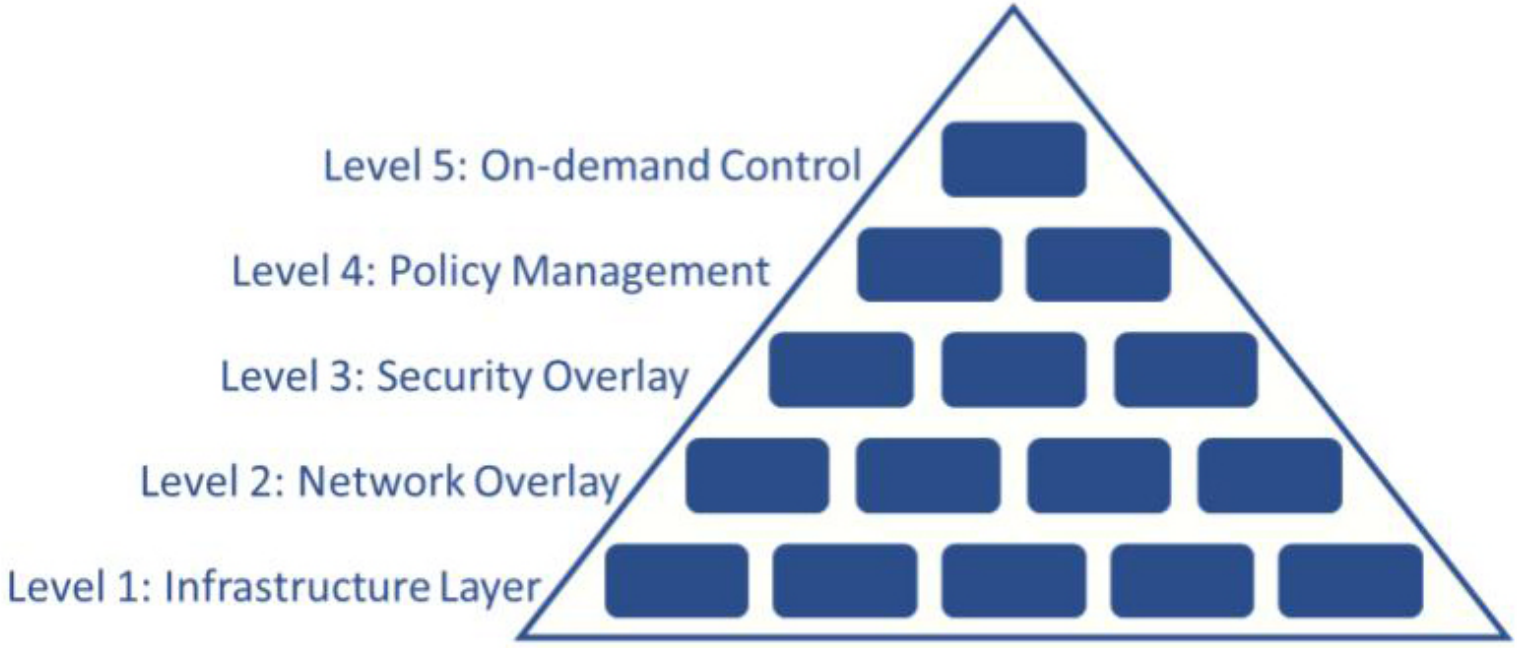

Network Cloud Business and Digialized Architecture

This allows on-demand, rapid adoption of new technologies without risk to existing stable mission critical operations, then rolled out as required as your living checklist for your Network Cloud implementation.

The journey map segments the various technologies needed to realize elastic networking into levels or layers of concern. We represent the overall system in a layered model to partition and separate areas of concern to help organize the problem and think about it in a structured way. Each layer has a set of services or capabilities that are desired to be elastic in nature.

Within each layer of concern, services are defined and implemented. A layered approach allows key services at each layer to be exposed and consumed via service APIs, yet abstracts the sophisticated nature of underlying service configuration and allows separation of concerns within the enterprise. The goal is to make it possible for an Enterprise operator to work at any layer of the pyramid and consume the lower layers as third-party delivered services, ideally on demand via API integrations. These integrations can be custom IaC or packaged software GUI automation, or both.

Infrastructure Layer: the devices, endpoints and connection services that make up physical connectivity to connect any client to any service required with the enterprise business model.

Network Overlay Layer: The network overlay layer is the enterprise network "canvas" upon which the business requirements of your network can be realized. SD-WAN is emerging as the unified method for overlay networking in both WAN and cloud network virtualization layer that is key to abstraction and elasticity of the network. The goal of the network overlay is to provide the abstraction layer for the devices and services that make up the infrastructure layer.

Security Layer: manages the identity management, the roles and authentication of authorized users, application APIs, and devices, implementing the policies of the Policy Management Layer. This includes software that controls and manages Infrastructure and Network Overlay layers.

Policy Management: Defines and automates which access and service policies are to be applied. It governs the access control authorization applied by the Security Layer. It also automates the monitoring of the implementation of the layers below to provide alerts of Policy infringement.

On-demand Service Control: This layer defines the current implementation of automated on-demand services in alignment with current and planned business requirements.

6.1 Deployment Considerations

Like any large transition, planning is key. Developing a north star strategy of where you want to get to is necessary to lay the groundwork for any such transition. The north star should consider what major deployment, operational and security changes are required and what they mean in terms of changes for your organization. One possible way to structure this top down is to look at the key underlying transitions.

Evolving towards an elastic infrastructure is multifaceted, but there are some key transitions that are generally applicable:

- Transition to modern, identity-based user onboarding and access methods
- Use of IaaS-based networking and security to the maximum extent possible
- Developing a colocation based WAN aggregation and access strategy
- Access transformation: unifying private site and public site access
- Cloud and SaaS networking and access control
- Controller and API strategy
- Security compliance

When defining a north star transition plan, it is helpful to work in a layered approach as laid out in the previous section. Consider the following layers:

- Infrastructure Layer - basic IP network connectivity to devices and service access
- Network Overlay - your virtual network including controllers and APIs; all your enterprise network service elasticity is here
- Security Overlay - SASE and Zero touch deployment layer
- Policy
- Service Control

6.1.1 Infrastructure Layer

Generally speaking, the transition to EI starts at the infrastructure layer by evaluating required site access types: requirements for private site access, types of access (wireline, 5G), VPN and mobile access. There are a few key considerations here. The infrastructure layer should provide raw connectivity to these services in a similar consumption model to the internet. For example, most cloud and SaaS services are accessible via any public internet connection. The infrastructure layer should provide a similar model for connectivity, with options for a low latency, private consumption model.

Parts of the infrastructure layer itself are being virtualized and delivered as an on-demand service. As discussed earlier, WAN network service is also being virtualized and delivered as an on demand service - beyond traditional WAN managed services. SDCI (Software Defined Cloud Interconnect) services allow WAN and cloud connections to be deployed programmatically via API. These APIs can be integrated into overlay layer automation, either via network controller software or IaC software orchestration, to provision these infrastructure connections on-demand, without the need to deploy any physical infrastructure, except at the branch. Automation and programmability such as this is key to elastic infrastructure deployment.

6.1.2 Network Overlay

The overlay layer is responsible for abstraction of infrastructure, taking care of configuration and management of devices and orchestrating the base connectivity properties (raw connection, reachability, QoS, etc.) needed so that creation of the overlay can happen regardless of device and connectivity differences at the infra layer. These properties are leveraged from SD-WAN investments already made as SD-WAN serves as a unified connectivity virtualization service over any type of network connectivity, internet or private.

Increasingly, Edge, Aggregation, Private DC, Cloud and SaaS connectivity is being virtualized and delivered as an on-demand service. Here, the trend is towards simplification via overlays like SD-WAN and IPSec. SD-WAN is fast becoming the dominant option for most WAN networks that have many endpoints, including when site-to-site communication is needed. SD-WAN provides a unified access method for all these. IPSec can also be used for smaller sites, DC-to-Cloud direct connections or as a first step to offload to cloud services. However, large IPSec networks quickly become very difficult to manage in point-to-multipoint networks, have less-automated policy controls and thus are not recommended for a long-term solution.

Today, SD-WAN virtualization is being extended to the cloud to provide unified public cloud provider direct connects (AWS DX, Azure ExpressRoute, Google Partner Interconnect) via this overlay layer. SD-WAN automation abstracts the complexity of the different cloud providers and coupled with SDCI, can also provide on-demand enablement for rapid deployment and lifecycle management without the need to deploy any device infrastructure.

SaaS direct connects that are SD-WAN compatible are also emerging, allowing SaaS traffic to be integrated directly onto the SD-WAN fabric allowing enterprises to eliminate the separate "**Direct Internet Access**" leg needed from each site or client to access publicly exposed SaaS applications. In the future, SaaS applications will be able to be onboarded with the same connectivity and security as other enterprise sites. This will be key to fast deployment of SaaS and cloud connections: the overlay automatically provides the network layer confidentiality required by Enterprise policy.

The bottom line is that overlay network virtualization is a key requirement to build the elastic network "**stack.**"

Depending on how you use SD-WAN networking (i.e., as a true WAN networking service or simply just client-to-application access) there are two options to consider: a cloud security service with integrated SD-WAN and IPSec access which can be simpler to consume or standalone SD-WAN network-as-a-service with service insertion for security and direct cloud peering to create a layered approach, where services are attached to a unified SD-WAN fabric.

The SASE "**full-stack**" access integration approach currently splits the WAN into SD-WAN access network and a "**middle-mile**" peering network which may not be SD-WAN and may be a separate network domain completely, requiring its own management and configuration. The "full stack" approach is certainly a good one for providing scalability and operations required for implementing elastic networks, especially if an automated direct-peering approach to cloud services is all you require.

If your enterprise has a global WAN network and you anticipate the need for multicloud, SaaS and private DC access with security, then a unified SD-WAN fabric deployment where all sites: Cloud, SaaS and Security are all endpoint services on the fabric may be a more scalable, allowing your to achieve a single domain for networking that is operationally separate and abstracted from security. In this case, service insertion will be a preferred approach to adopting cloud security. With SD-WAN becoming increasingly available as a service in colocation data centers, along with Cloud Security services, such an approach is viable and can be deployed and scaled on-demand using centralized SD-WAN orchestration.

6.1.3 Security Architecture

To define your security architecture, first you need to identify what security elements you need to adopt for your infrastructure. This is dependent on two main points:

Company Policies	These policies are either mandated by regulation, compliance or business decisions
Application Needs	Yes, applications will drive the requirements for you to adopt security components. For example, if your application encrypts data, then you might not need to worry about transport encryption and you might not need to proxy the access with a TLS termination capable proxy, if your application contains sensitive data that can be highly attractive but should remain totally confidential, then you might need to add a DLP component

Most architects tend to take the approach of combining legacy and do a little market research to see how this legacy stuff has evolved, and then go with it, often forgetting that there will be a team who will need to operate this whole thing and they are probably not happy.

In terms of security for new applications (some call it modern applications or cloud native applications), these are applications that are mainly API driven and can be hosted pretty much anywhere. We believe that the minimum set of security functions needed are the ones below:

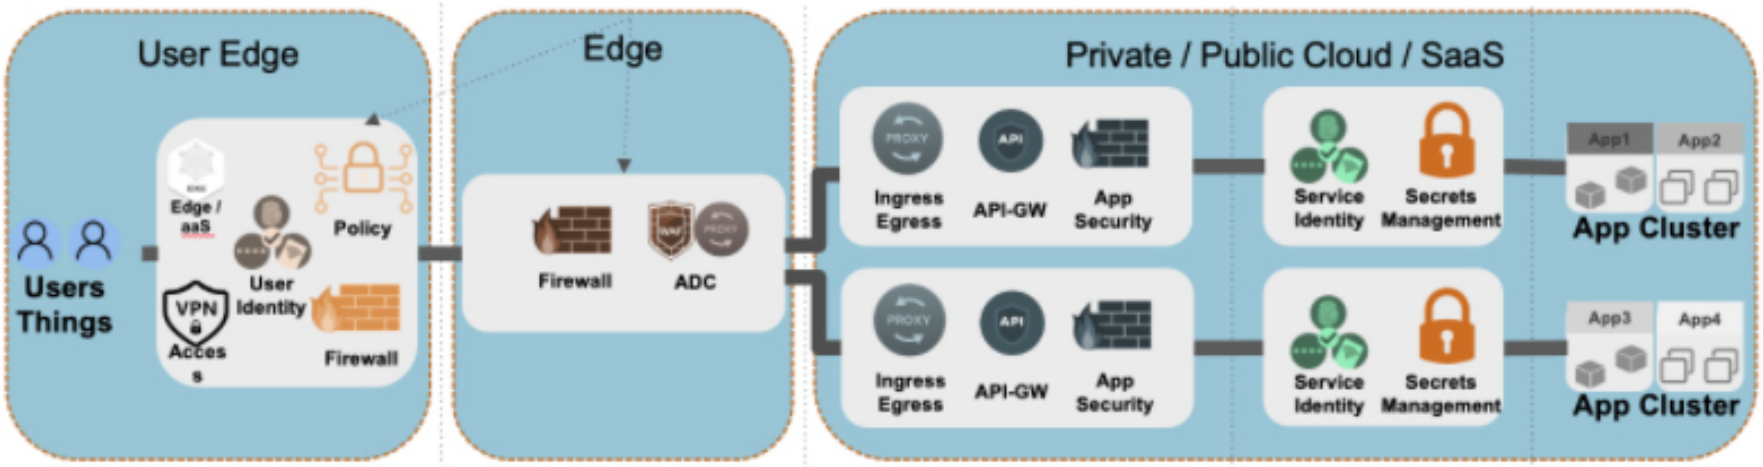

As can be seen, there are three main areas of specific security focus and each of these areas carries its own security functions:

User Edge

This area is where the users are and where they initiate their requests. It has been mentioned above that these users can be in several locations such as:

Corporate Office	(Where they are in fact connected already to the CORP Network)
Home Office	(Where the connectivity might be there - similar to a Corporate Office - or the more typical case where the user is just using home internet broadband and wants to access corporate resources
Mobile Location	(Public WiFi, shared office, friend's house, anywhere else really)

If we take an approach to address each of these situations individually, it could lead to several different architectures (to address the different types of users) and different types of final result. This would lead to a proliferation of devices, solutions, tools, etc. It also means that someone in that organization will need to operate in this environment and operations considerations will be key to successful deployment.

The solution needs to be elastic so that it is able to accommodate any user in any location accessing the resources they need to access. If the user changes their location, their access doesn't necessarily need to change (it can do, though, but for the majority it won't).

So based on the above, the functions we believe are crucial in this layer are:

01 **Identity** - The user needs to be identified. He needs to be able to have in hand some form of proof that he is who he claims he is (being this a token, certificate or any other item that the infrastructure will understand and able to correlate to who this person is)

02 **Policy enforcer** - A policy enforcer is some method of enforcing a policy that comes from Corporate. This can be a L3/4 policy or a L7 policy.

Edge

The edge is where users (internal or/and external) "hit" the company's infrastructure. This was, historically, seen as a DMZ where you would deploy a set of security functions and isolate a set of services to allow or deny certain types of traffic. Now the reality is that the DMZ area can no longer be confined to specific locations and users tromboned to them, in fact the Edge nowadays is widely spread because of the needs of the users and needs of applications themselves.

The edge functions as a place where data needs to be accessed or processed and as such it contains services and/or compute resources to accommodate these functions. The edge locations and can be:

- **CoLo:** (Colocation facilities where the Organization purchased some space)
- **Public Cloud:** (Cloud Edge where the Organization has edge services exposed to the internet or simply exposed to other corporate locations - including user locations)
- **Private Cloud:** (This will be your typical DMZ location- but with Edge functions)
- **Customer Remote Locations:** (Branch offices, ATM locations, basically any location where data needs to be either processed or accessed - or both)

The edge needs to contain a set of services that allow those two key functions to be executed (process/discover data and/or expose data to be accessed), and as such there are some key functions that need to exist such as:

- **Firewalling:** (Allow or deny certain traffic types, being this at L3/4 or L7)
- **Proxy:** (Expose services or make them appear where operators need them to appear or as a forward proxy for your users to access with content restriction defined by the operator)
- **WAF:** (Inspect content based on signatures and/or create new signatures to restrict specific content)
- **DDoS:** (Increasingly important to be able to stop Volumetric attacks and ones generated by BoTs)
- **API Gateway:** (API manipulation and authentication - ability to authenticate requests to specific APIs, discover and map them and apply controls to their relationships)

These services can be built inline or service chained via the data-plane.

Private / Public Clouds / SaaS

This is where the applications live. They can be Monolith servers that run one application, Virtualized servers that run multiple virtual machines with the same, or different applications, micro-services running on clusters like Kubernetes, pieces of code running on serverless environments or even SaaS applications offered on the cloud in the form of API requests (or others)

By now, it is well known that for each client (north-south) request there are multiple other requests generated between applications (east-west). We are talking about an exponential difference between client requests and app-to-app requests

It is therefore very important to have very well-defined east-west controls in place to ensure these applications are only really talking with the ones they should. A least privileged or a zero trust environment is desirable in this environment. Some services that is important to have here are:

- **Ingress / Egress Gateway:** To control who enters which "realm"
- **API Gateway:** To authenticate and authorize access between apps / APIs
- **App Firewall:** A firewall to govern traffic to, from or by an application or service

These services are infrastructure-based services but from an application perspective. The goal is to ensure that the applications or services are only exchanging information with other applications or services they really should. As an example, if we have an application with a service that processes payments and another service that processes emails (like an email service), then these two services probably shouldn't be allowed to communicate with each other.

Other services that are also important in this layer are tied to the application control itself:

- **Identity Service:** to identify the application when it is created or scaled
- **Secrets Management:** To manage secrets and keys to be used or served by the application

When we look at all these services, it might sound overwhelming but if we read above, this needs to be considered in a framework that adapts itself, is able to scale up and down, and is able to identify each service to allow them to talk to the right endpoints.

Overwhelming or not, sometimes it could be the difference between services running without any problems or services' confidential information appearing published on the web.

6.1.4 Orchestration and Service Control

Orchestration is simply the provisioning, modification and management of resources which deliver an application or service. Services and applications are now, more than ever, delivered from multiple sources making orchestration a critical piece of the workflow, especially when dealing with multi-cloud networking. As applications become more distributed, the approach shifts from interacting with devices to interacting with all of the infrastructure required to run a given application or service.

Automation vs. Orchestration

Orchestration is the next phase of automation for networking. The outcome transitions from single tasks with limited or no human intervention to completely policy-based or event-driven outcomes with no human intervention. These two approaches can be compared as follows:

AUTOMATION	ORCHESTRATION
Executes single low-level tasks	Executes multiple, sequential, high-level operations
Interacts with device CLI or scripting frameworks	Interacts via orchestration tooling which uses APIs that enable automation across vendors and services
Does not account for state of operational environment	Accounts for state, status and configuration of existing environment
Little or no governance or policy management	Extensive governance and policy management

Accounting for API Sprawl

As enterprises transform, they are increasingly transitioning from monolithic applications to more modern, cloud-native applications. This increases API consumption in the enterprise environment exponentially.

Provisioning Tools

To maintain consistency for cross-platform infrastructure, a provisioning tool is often used to act as an interpreter between the consumer and the infrastructure end-point, minimizing the exposure of the consumer having to directly interact with a vendor API. Provisioning tools operate in a declarative nature, describing the end state of the system without defining the individual steps to reach that end state.

07 Conclusions

The evolution toward Cloud, IaaS & PaaS, as well as SaaS has led to a revolution in how applications are delivered and how data center services are deployed and consumed. Service level automation, Infrastructure abstraction and SaaS lead to on-demand, elastic and agile deployment for enterprise applications.

This has resulted in changing network traffic patterns and new abstracted network delivery models. Network consumption is moving up the stack, trending toward abstracted, service level consumption, rather than the IP layer service building that has traditionally been done. This is an evolution beyond the IP layer that was brought about by SDWAN, where overlays and automation helped speed network deployment.

A similar pattern is occuring in security with the trend toward zero trust architectures, identity- and posture-based access control with flexible enforcement and unified security policies that can be enforced over multiple network access types, public or private. Security is becoming an overlay itself, one that can be deployed on top of any network type.

The movement toward IaaS delivery and abstraction of network and security services is the beginning of a new "**Network Cloud**" that brings the consumption flexibility of the software defined cloud data center to the network itself. This transformation will bring unprecedented speed, agility and on-demand control to the enterprise digitization journey.

Part of the ONUG Collaborative, The Network Cloud Working Group is comprised of Enterprise Cloud Consumers and Suppliers. The Working Group focuses on the use cases for leveraging the Network Cloud's elasticity and agility features to fulfill the business requirements of today's Multi-Cloud enterprise. Interested in joining the team? Contact us at Sponsors@ONUG.net