

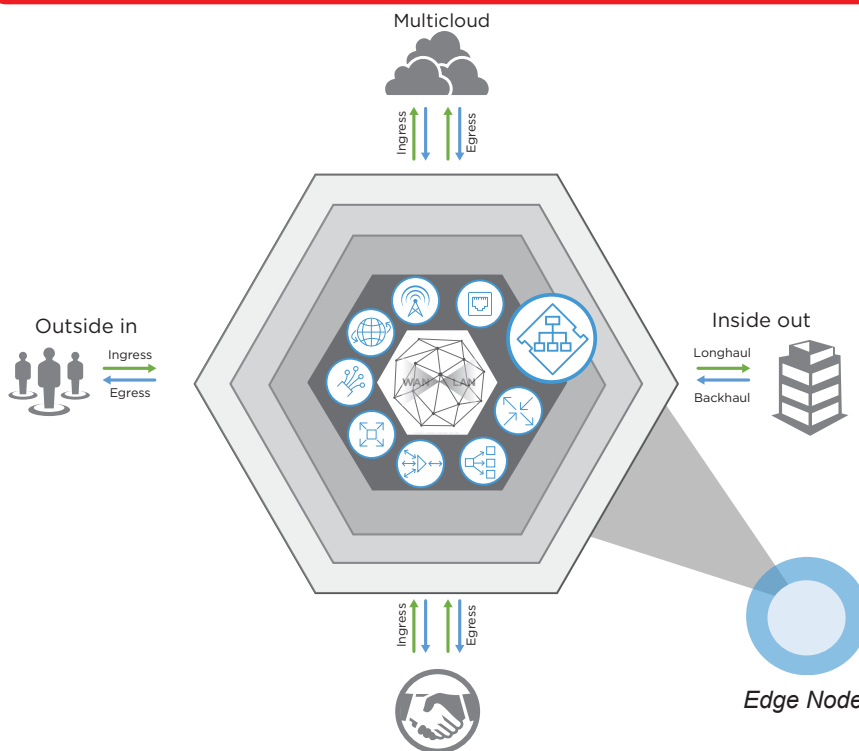
## Design Principles

- Improve performance (quality of user experience).
- Increase resilience (amid frequent change).
- Deliver scalability (cost-effective throughput).
- Ensure sustainability (design for growth).
- Minimize geographical distance.
- Place intersection points in areas of density (users, traffic and data).
- Change topology to adapt to business change.
- Provide vendor neutrality (maximize competitive choice and innovation).

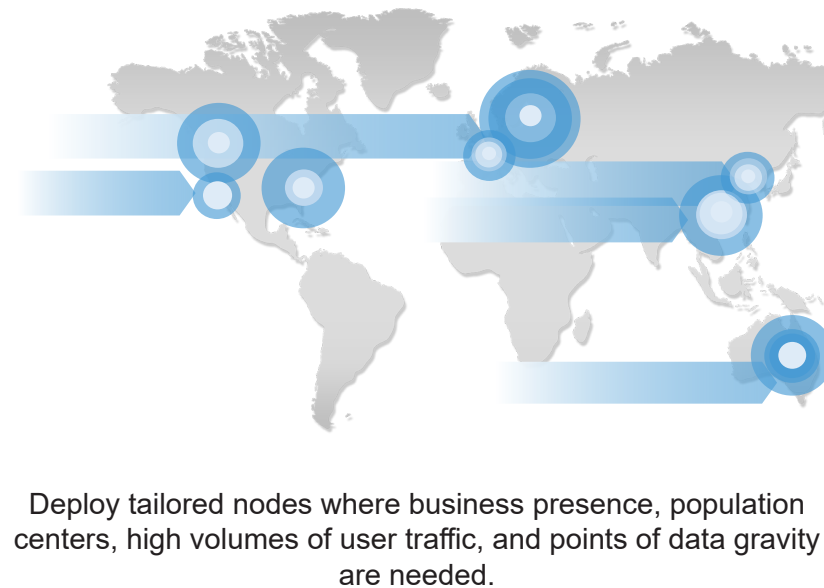
## Edge Node Components



To architect for the digital edge, you need to bring the WAN and LAN together and create a digital edge node. Each digital edge node is tailored for network and traffic types that have to be localized, segmented and optimized — at specific geographic locations. Build the nodes in step function and deploy them in metro-based zones where there is density in users, traffic and data. Directly connect the nodes to reduce topological distance and optimize bandwidth.



## Simplify the Topology



## Capabilities

- Low-latency intersection points of ecosystem density (clouds, partners).
- Shortest path from mobile, broadband and internet to edge.
- Localized and segmented traffic.
- Multicloud connectivity options.
- Resilient mesh of geographic edge nodes with E2E SLA control.
- Optimized and globally load-balanced bandwidth and traffic.
- Ability to scale customers and markets globally by adding interconnected digital edge nodes to the mesh.
- Control bandwidth from anywhere.
- Interconnection at the center of the architecture, with IT in control of digital's biggest differentiator.

## Edge Node Deployment



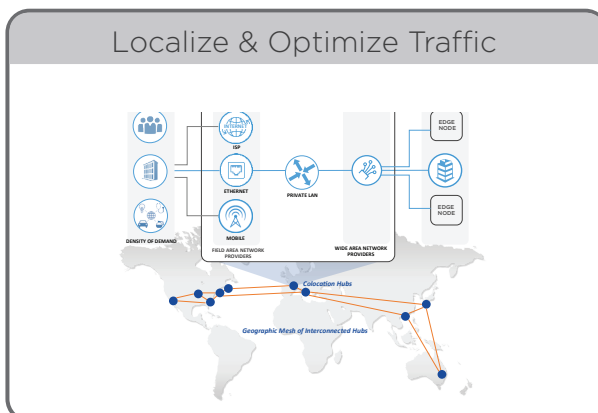
Implementation is a mix of physical devices and virtual (NFV) appliances.

Plan for half a cabinet in your edge node design.

## DESIGN PATTERNS

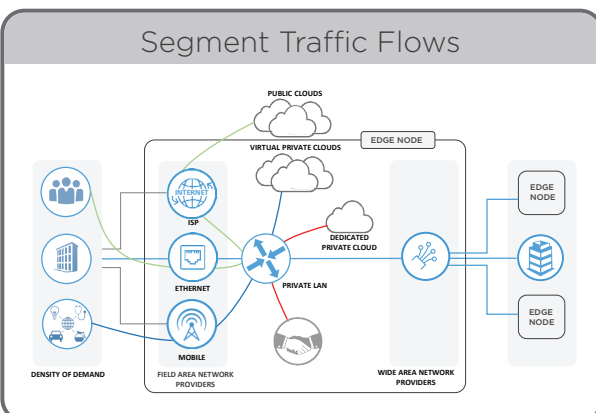
### 1 STEP 1

Redefine the edge by establishing a node in a metro area closer to customers and where business is conducted



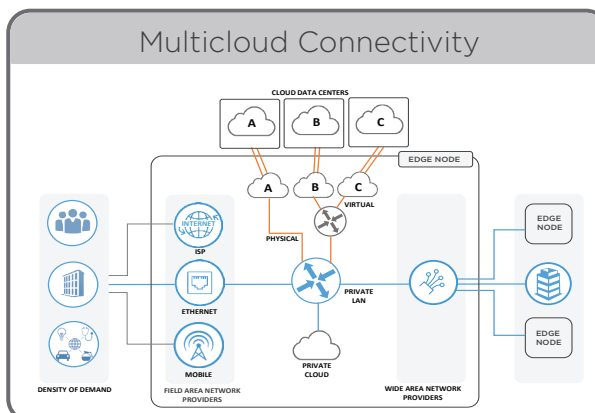
### 2 STEP 2

Prepare for multicloud and partner network integration as well as digital service flow isolation



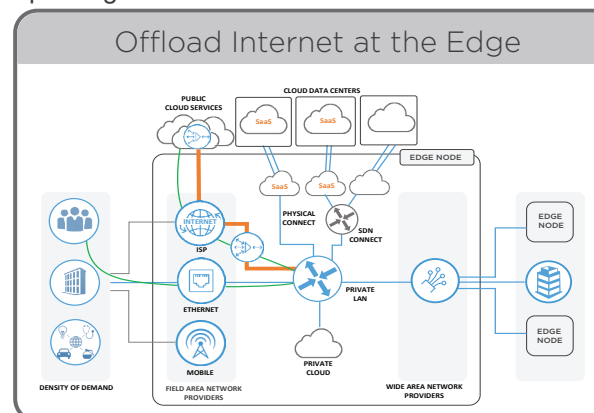
### 3 STEP 3

Service chain clouds, applications and data together across locally cross-connected cloud providers, accessing SaaS services as needed



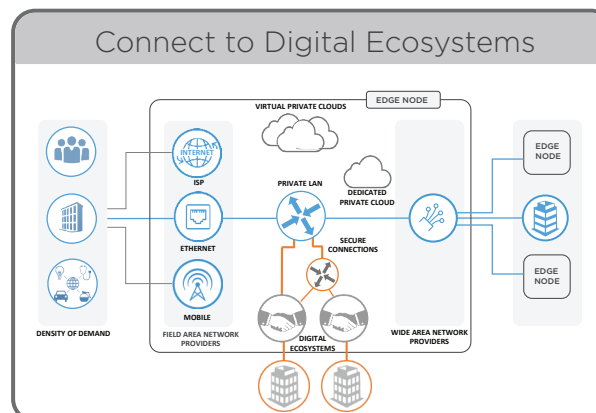
### 4 STEP 4


Bring all traffic into the edge node and benefit from added control and reduced risk, while routing public internet traffic directly with internet peering



### 5 STEP 5

Cross connect to business partners and ecosystems for digital commerce and/or data exchanges





## Problem

The limitations of current networking technology and certain immutable laws of physics just do not allow latency-free transmission over long distances.



## Solution

Models of consolidated access (i.e., network transports) into a colocation hub, interconnection within the hub and building of a geographic mesh (a distributed core) based on competitive connections between hubs (redefines the edge) offer the same architectural and technical capabilities and benefits already enjoyed by digital leaders (e.g., cloud providers, content providers, carriers) that reduce costs and latency while adding functionality and flexibility.



## Constraints

1. MPLS links to geographic locations are expensive and not sustainable with increasing demand for performance and data growth.
2. Digital B2B and cloud connections face similar challenges, and conducting business over the internet increases latency and the potential for cyber risk.
3. Overlay networks (SD-WAN) help by optimizing for the least inefficient path, but underlying limitations have not changed.
4. Distributed application flows across multiple cloud services and multiple businesses (B2B) continue to accelerate, increasing demand for more robust, secure and low-latency interconnectivity.



## Steps

1. Establish a geographic hub based on population and business intersection advantages (as part of a mesh).
2. Solve the last-mile problem via Metro Ethernet instead of MPLS (T1 etc.).
3. Use intra-colocation (cross-connect) networking to enable new connectivity paradigms through business peering and direct cloud connection.
4. Implement competitive inter-colocation access to reduce WAN costs, reduce latency, improve availability and provide new topology offerings.
5. Use the combination of intra- and inter-colocation networking to derive the same benefits in performance and efficiency enjoyed by cloud service providers.
6. Apply traffic management to load balance traffic across your network fabric.



## Forces

- More complex applications require more sources of data for closer-to-real-time processing.
- Mobility is driving more devices (of more types) with high density in metro locations (+IoT devices).
- Diminishing tolerance of high latency or general lack of responsiveness (users vote with their feet).
- Increasing sophistication, threat and impact of cybercrime means risk is proportional to internet use.
- Need for greater market reach without increasing the network surface attack area.
- Companies driven to improve end-to-end controls.



## Results

## Efficiency and Choice Reduce Costs and Risk

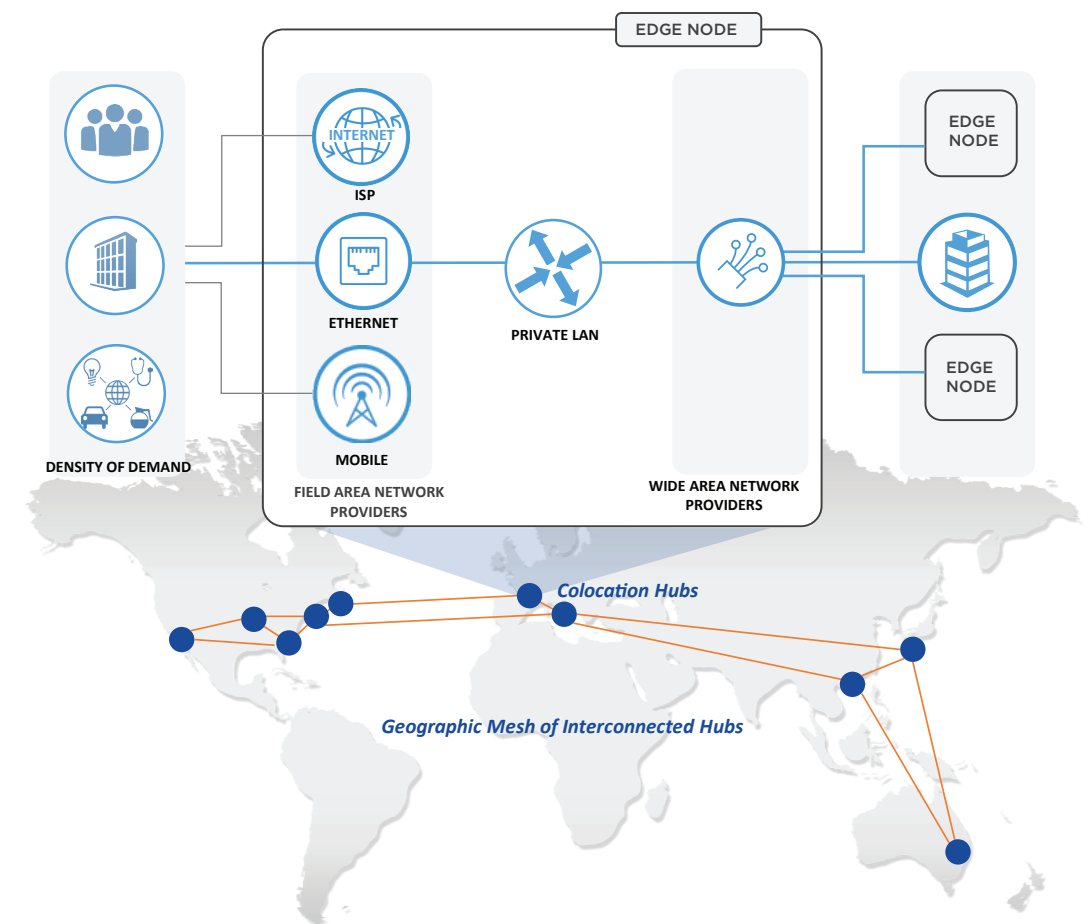
- The availability of more WAN choices: >\$100/Mbps to <\$100/Mbps.
- The availability of more Metro Ethernet: >\$100/Mbps to <\$10/Mbps.
- Interconnection within the hub is <\$1/Mbps.
- Localizing traffic in the hub decreases latency from ~20 ms/~20 hops to <1 ms/hop (or wire speed) with unlimited local bandwidth.
- Network-based attacks are localized and limited to last mile.


## Potential New Challenges

- Increased vendor management (contracts).
- Network management skillsets or services required.



## Reference View





## Problem

There is a need to interconnect multiple clouds and SaaS services, along with more business partners to more end user devices, involving exponentially increasing data, traffic and regulatory scrutiny – outside of the data centers, at the edge.



## Solution

After applying IOA Network Design Step 1 — Localize and Optimize Traffic — you can implement traffic segmentation inside the edge node(s) at the network intersection points. Now traffic segmentation is established locally inside the edge node. WAN links connect the edge nodes into a mesh, and traffic management capabilities exist at each intersection point (on both ends of the WAN). In addition, direct connections to clouds and business partners will also simplify segmentation. Create an interconnection for each segmented flow.



## Constraints

1. Many networks were not designed to interconnect multiple external networks.
2. The majority of business traffic shifting to the edge has also exposed limitations in traditional network design – outside the realm of control.
3. This can lead to a proliferation of WAN networks (both physical and logical) with multiple sets of infrastructure required to police them.
4. Defining acceptable communication paths across these networks is difficult as applications become more portable and services become more interconnected.
5. Ultimately the rate of change (in flows) outpaces the network's ability to adapt and accommodate.



## Steps

1. Plan for how many clouds, partners and field area networks will be involved and select the segmentation approach (or mix) that will efficiently accommodate those demands (mix of appliances and SDN).
2. Begin with a simple workload (and therefore traffic) risk classification (like low, medium, high). Next apply an axis of suitability for internet/public cloud, hybrid or private-only.
3. The intent is to determine coarse-grained policies as traffic continues to move to the edge and application/cloud and multi-tenancy maturity increases. Stamp these policies across all edge nodes. All traffic is collected from field area networks as close to the edge as possible. High-risk traffic is segmented to private-only, whereas the rest can route directly to cloud, and low-risk can also be off-loaded over the internet. Both remain in your control and the decisions are made at the edge.



## Forces

- Business is becoming more interconnected, increasing the need for traffic segmentation.
- As more data is exchanged, sensitivity and regulations associated with the data need to be addressed, which requires separation (e.g., PCI-DSS & HIPAA).
- The growing number of edge points in the network make network policy enforcement increasingly challenging.
- MPLS layer VPNs are cost prohibitive.
- With data and traffic continuing to exponentially grow, even SD-WAN technologies are limited by the underlying data plane capabilities.

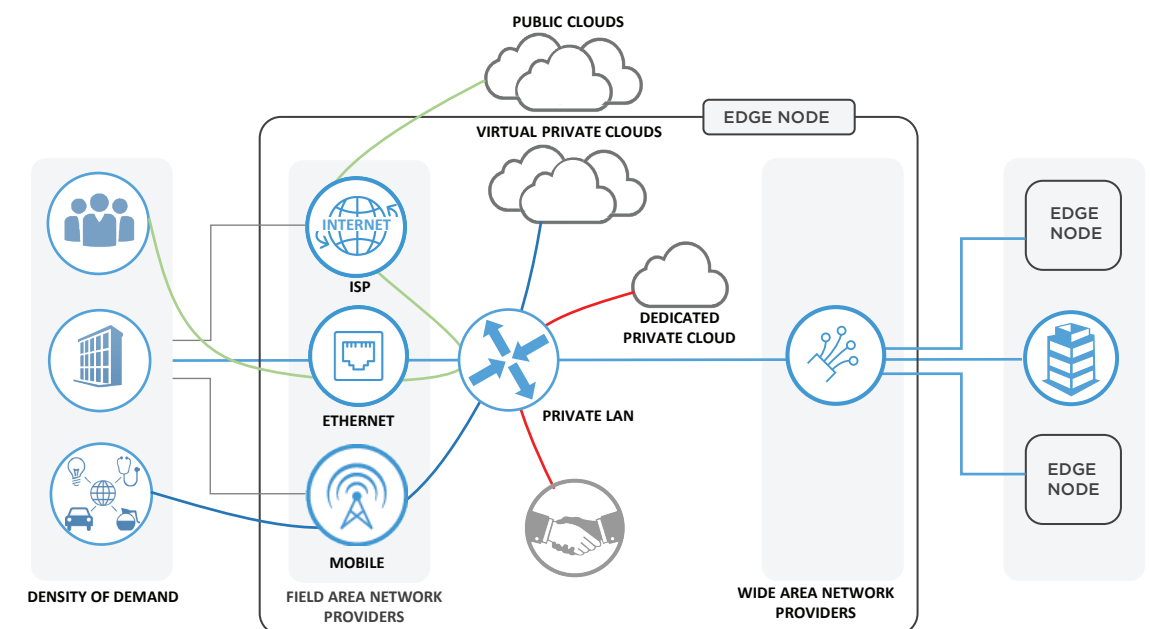


## Results

- Localized traffic segmentation in a LAN-like environment.
- Maximized choice in methods and partners to implement it, with the ability to adopt SDN over time.
- Implemented basic risk policies that are clear for developers to understand.
- Influenced foundational workload placement and resource investment.
- Orchestrated flexibility to change policies and implementation on demand.



## Reference View




Ex. 1 (Green) – Employee accessing a low-risk public SaaS application is internet off-loaded.

Ex. 2 (Blue) – Secure IoT communication to direct connected cloud service.

Ex. 3 (Red) – B2B transaction over dedicated secure cross-connect passing PSS-DCI regulated data.





## Problem

Multicloud connectivity over WAN can be complex, expensive and/or exhibit performance issues. As demand for higher throughput increases, WAN (and SD-WAN) options test the limits of physics. The absence of viable alternatives can put sensitive traffic on the internet which offsets potential cost savings with added risk.



## Solution

After applying IOA Network Design Step 1 — Localize and Optimize Traffic — you can cross connect to multiple clouds locally within the edge node, either physically with a dedicated cable (single tenant) or depending on your edge node's options, by using SDN services to create virtual dedicated cross connects (multi-tenant). Each offers varying bandwidths, and all extremely lower latency than WAN or SD-WAN connections.



## Constraints

1. Extending an existing network to various clouds already affords limited topological options and/or choice.
2. Limited bandwidth makes high throughput cost prohibitive.
3. Distance can make high throughput unachievable regardless of WAN or SD-WAN options.
4. There can be resource costs involved to get links turned up, with contracts and equipment purchases.
5. Internet begins to look like a viable option for sensitive data — if given a better alternative, this would never be considered.
6. In all cases, as traffic and data volumes exponentially increase, performance continues to deteriorate.



## Steps

1. Using your workload classification (see Step 2 — Segment Traffic Flows), determine what connections of what levels of segmentation and bandwidth are needed, to which clouds.
2. Prepare your receiving side in the clouds and/or register for the appropriate cloud services to facilitate the connections.
3. Create or request the actual connections.
4. Right-size bandwidth to usage patterns.
5. Redirect all business traffic for those workloads and SaaS services over secure and dedicated interconnections.
6. Add links as necessary for cloud pilots and tear them down when no longer needed.
7. Leverage the intersection point to link clouds together and conduct inter-cloud transfers over your own low-latency secure backplane.



## Forces

- The pace and rates of change are driving businesses toward ecosystems of advantage, leading to leveraging clouds to deliver IT services and infrastructure.
- Speed and avoidance of IT debt and aging infrastructure are driving new business applications to be delivered on cloud platforms.
- Enterprises don't just use one cloud; they typically use several (including SaaS services) and the numbers continue to increase.
- While enterprises struggle with forming direct connections to so many, the cloud providers will have a harder time handling bespoke connections.

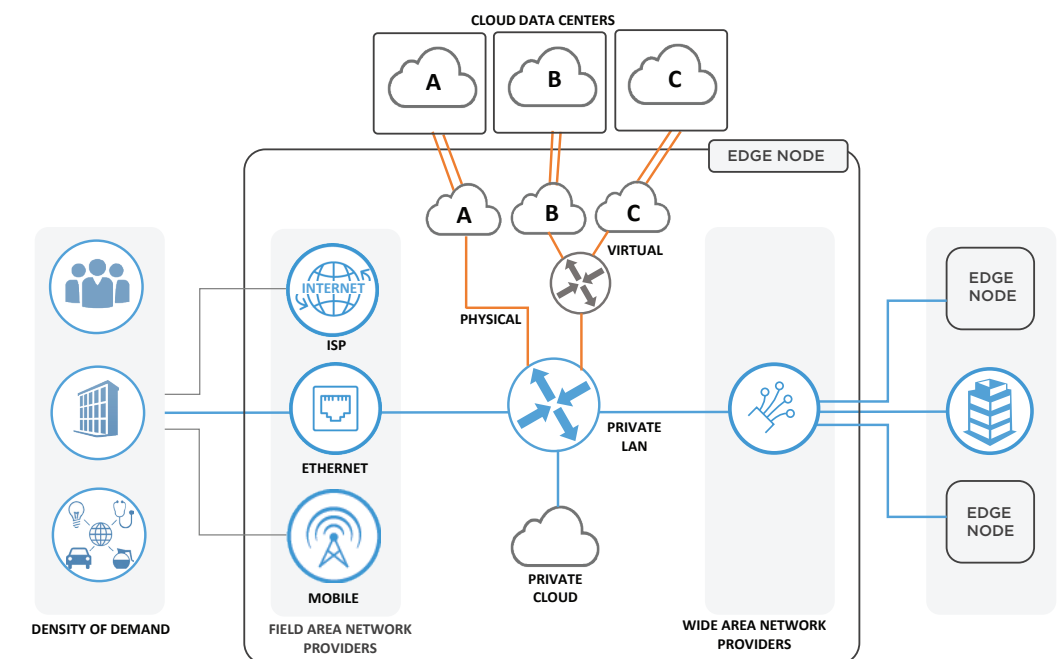



## Results

- Provisioning a new connection is dramatically simplified. Time to new connections is hours—not weeks.
- Connecting to clouds this way enables enterprises to construct cross-cloud service or business flows (service chaining), which requires low latency between services.
- The majority of traffic is localized, with the most optimal path via edge node to users (customers, employees) of cloud-based workloads. Minimize backhaul to corporate data centers.
- You can rightsize the connections if capacity requirements change.
- As use of cloud services changes, you can continue to adapt the topology to keep pace with digital.



## Reference View





## Problem

The exponential growth in traffic and data, as well as the use of cloud services, are placing extreme demands on bandwidth, with no end in sight. SD-WAN buys some time but is still a temporary fix. Moving traffic to the internet can be up to 10x cheaper, but shifts the problem to exponential risk.



## Solution

After applying IOA Network Design Step 1 – Localize and Optimize Traffic – you have already significantly closed the gap in cost between WAN and the internet, and in most cases WAN is now cheaper. Cost, performance and throughput are no longer the primary drivers as internet is no longer better, faster or cheaper. The risk exposure is also gone as traffic stays on dedicated private connectivity (and your controls are applied). However, the growing use of public cloud services and the need for the lowest latency access to public clouds (while still providing some level of protection) remains. The low-risk traffic you classified in IOA Network Design Step 2 – Segment the Traffic — applies here. For that remaining requirement, you can use local internet peering at the edge and decentralize internet breakout points to offload that low-risk traffic to local internet suppliers and local markets. This provides the next cheapest, lowest latency and most secure way to manage that traffic, with the ability to change your mind at any time should the risk profile change.



## Constraints

1. MPLS bandwidth is cost prohibitive and constrained.
2. Backhauling internet traffic over the WAN is consuming bandwidth and impacting security infrastructure user experience.
3. WaaS has become a necessity, but provides only temporary relief — it doesn't change the underlying limitations.
4. Investing instead in a security arms race against cybercrime is likewise not a good risk/reward option. A breach is a matter of when, not if.
5. Tolerance for downtime, data theft, privacy violations and associated regulatory accountabilities prohibit choice to take on risk.
6. Providers who therefore depend on internet business traffic will emphasize safety, while SSL (and VPN) have already been proven vulnerable.



## Steps

1. Within the edge node, establish an intra-colocation connection (cross connect) to the ISP provider.
2. Apply segmentation rules (in Step 2) to identify opportunities and compliant traffic to use this path.
3. Where applicable, redeploy or apply WaaS for internet traffic as (after Step 1) the internet is now the least performant network and optimization is mostly needed there.
4. Review which public internet services you are optimizing for and determine if direct connection (from Step 3) is a more viable option.
5. Redirect the remaining internet traffic flows across the ISP link locally at the edge.
6. Security still applies (covered in the Security Blueprint\*).



## Forces

- Traffic and data are growing exponentially, which will be compounded by IoT and analytics.
- The way people interconnect is changing and shifting to digital engagement. Smart devices, screen resolutions, video over voice, rich content, highly interactive applications and now emerging virtual reality are just the beginning.
- Cyber threats, both foreign and domestic, dark innovation, and growing sophistication of attacks/breaches is adding volatility and increasing cost.
- Business is becoming more digitally dependent and using internet introduces higher risk and impact (both financial and reputational) that can far outweigh the short-term benefits.

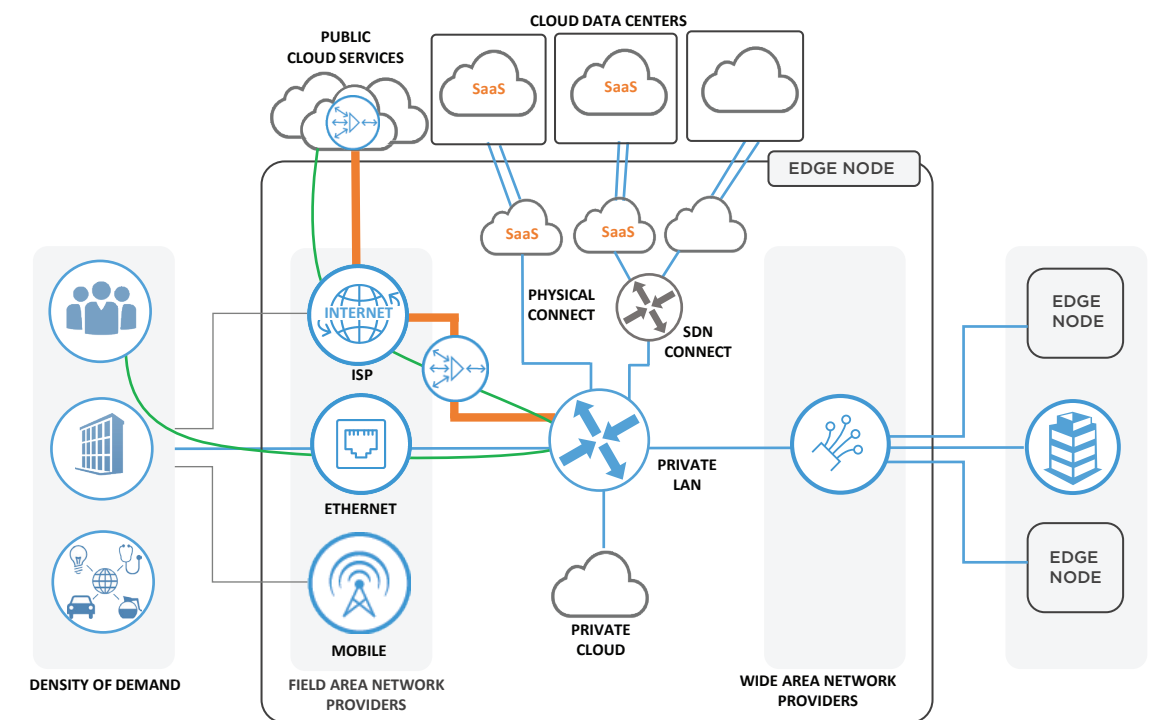


## Results


- Network optimization removed the situation where internet was a lower cost, more performant option.
- The attack surface area remains small and business traffic remains on private, secure connections (and compliant).
- Continue to bring your users/traffic from field area networks into your edge node at the closest point and then accelerate internet traffic through your edge node while retaining control and levels of protection.
- As SaaS services provide options to direct connect, shift those flows over accordingly to reach optimal security and performance in more cost-effective ways.



## Reference View



\* Security Blueprint — IOAKB.com



## Problem

There is an understandable reluctance to exchange sensitive data and services over the public internet. This can expose both companies if either has a breach. In addition, the costs of MPLS connections are the same for B2B, with most businesses being too far apart.



## Solution

Leverage your edge node as a “meet me” place for digital business. Invite key partners into your edge node in order to cross connect (or agree on establishing a mutually beneficial edge location). The benefits of your edge node compound beyond cloud ecosystem access when digital business ecosystem partners are also in the same edge node location. Establish direct, secure connections (intra-colocation) as needed to extend a private LAN to business partners. Your digital exchange runs on secure, dedicated, low-latency bandwidth, removing the barriers to expand the ecosystem and seize business opportunities. This can be done in any and/or all of the global metros and markets where you and your business partners have an edge presence.



## Constraints

1. The same constraints of MPLS apply to B2B connections. The bandwidth is expensive and latency reduces the throughput.
2. Insecure internet connections inhibit ecosystem growth and are a barrier to entry. Digital ecosystems are a big target, and prize, for cyber criminals.
3. Provisioning direct connections can take time, and they are difficult to change as business requirements change.
4. Lack of viable connectivity inhibits an ability to expand into markets safely and profitably. This can be disadvantageous and detrimental to business.
5. Need to easily connect with partners to seize business moments and opportunities before they pass. Current architectures aren't designed to do so.



## Steps

1. You selected your edge node based on business intersection advantages (in Step 1).
2. Similar to cloud interconnection, now you simply implement a cross connect between your two edge nodes within the same location for low-cost, high-speed interconnection.
3. As ecosystems evolve, SDN capabilities can also be applied (by a neutral provider) to facilitate many-to-many connections even more effectively at finer granularity than a dedicated cross connect. These are called exchanges.
4. The ecosystem should encourage more participants in order to amplify the benefits and business capabilities available.
5. This is similar to how financial services exchanges achieve such large digital business scale cost-effectively.



## Forces

- Digital business is driving companies to interweave business processes, resulting in increases in the exchange of digital services and data.
- The digital economy is based on the strength of ecosystems, not individuals.
- As business processes become digitized, business and technology juxtapose. This means the rate of business change is as fast as technology change.
- Companies need to leverage ecosystems in order to maintain competitive advantage or risk falling behind and becoming immaterial in the digital economy.
- This is driving the increasing need to form dynamic, secure, direct B2B connections cost-effectively.



## Results

- Dedicated private connectivity to business partners with nearly unlimited cost-effective bandwidth.
- Acts as a keystone and foundation for higher level security, data and application (B2B) challenges.
- You have implemented proven practices similar to how financial services exchanges achieve such large digital business scale cost-effectively. This model is now being extended to clouds and other business ecosystems.
- Establish simplified compliance transparency around how data and communications are being secured and protected.
- Business is being conducted in private with reduced cyberthreat and SLA exposure.



## Reference View

