Working Requirements for Network Connectivity and DMZ Architectures in Hybrid Cloud Deployments

Enterprise Cloud Customer Council Technical Working Group Position Paper

Requirements and Guidelines for Cloud Service Providers.

May 2017

Preamble

Overview

The Enterprise Cloud Customer Council (E3C) is a group of enterprise companies, primarily financial services, with the vision to become a leading voice for enterprise organizations in the pursuit of cloud interoperability, security, and portability with major cloud industry vendors and standards bodies. The top cloud companies are closely involved and are actively assisting the E3C in achieving its goals. Intel Corporation serves as facilitator for the group.

Mission

The E3C mission is to develop requirements centered on common functional architecture, detailed specifications, and technical solutions that improve enterprise cloud adoption and integration while mitigating regulatory risk.

For the smooth and secure usage of hybrid cloud for applications and data, the council shall work toward

1) Enhancing security of clouds for secure network connectivity and data management that align to enterprise requirements
2) Ease of integration and life cycle management through interoperability across clouds
3) Ensuring enterprises have the same visibility and control of their resources in public cloud as on premises

Process

In order to achieve the E3C mission, the members form technical work groups focused on the top current impediments. The technical work groups work towards understanding common architectures and approaches and then establishing common requirements that form the E3C position papers. After the papers are approved by the E3C council members, consisting of executive level members from each of the member companies, the papers are published on a public website. Papers will be iterated accordingly as needed.

Usage of Position Papers

The position papers are not designed as prescriptive requirements but rather as input for organizations developing roadmaps and/or requirements relating to cloud adoption. These are freely available, however please reference E3C if you utilize these in RFPs, RFQs, publications, etc.

Document Conventions

RFC 2119 conventions apply to this document. For convenience, these definitions are provided here:

1. MUST: This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement of the specification.

2. MUST NOT: This phrase, or the phrase "SHALL NOT", mean that the definition is an absolute prohibition of the specification.

3. SHOULD: This word, or the adjective "RECOMMENDED", mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.

4. SHOULD NOT: This phrase, or the phrase "NOT RECOMMENDED" mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

5. MAY: This word, or the adjective "OPTIONAL", mean that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option MUST be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)

1. Overview

A typical large scale enterprise compute environment provides resources and services to internal users and private B2B services to external users. As the organization shifts various compute resources to different cloud providers, the same controlled access techniques that existed in legacy Data Centers or Private Cloud deployments must be offered for these resources running in the public cloud. This position paper is meant to give high level requirements to the CSPs and illustrate the challenges experienced by a typical enterprise as it designs proper architecture to move workloads into the cloud utilizing IAAS model.

2. Enterprise Challenges at a High Level

A traditional enterprise operates a Data Center environment, in which compute hosts serve this enterprise and are operated by the enterprise or trusted service provider. As some resources start moving from private environments to the CSP, the following are the general challenges.

- **Common Definitions.** When defining every technology design standard, there has to be a set of common definitions defining various elements in the designs. When enterprises design their Hybrid Cloud solutions they are often dealing with a mix of definitions that could vary between various NSPs and CSPs. The definitions are important as they drive the standards and network designs.
- **Cloud Access** As the enterprises create touchpoints between their transport networks and the CSP networks, the enterprises must create various design models that provide these transport network expansions. While this is very similar to designing network connectivity to various B2B and content providers, the exercise is different because users, legacy Data Centers, private and other public must use this infrastructure in the same way internal infrastructure is used.
- **Cloud Network Services.** As the organizations build networks in public clouds, these networks will have various forms of network gateways. These gateways can provide connectivity to other

networks within the same CSP, public Internet or private CSP networks. In addition to connectivity they can provide other network services such as perimeter security. Most CSPs will offer these functions as services implemented natively as part of cloud infrastructure and various vendors can offer them as dedicated NFVs running on virtual compute nodes. While it is a lot more practical and scalable to utilize native CSP NFs, in some situations, primarily when extending network infrastructures outside of the CSP domain, specialized systems that run as compute nodes can be used as NFs. The scaling of these systems is not responsibility of the tenant and not the CSP.

- **Hybrid Cloud Connectivity.** As services are deployed in multiple public and private clouds, there is a need to provide connectivity between virtual compute instances deployed in diverse CSPs, as well as internal Data Centers. These connections can be physical or virtual, but no matter what network infrastructure exits between these environments, the networks must be extended. This expansion involves virtual private networks to insure address space isolation as well as some segmentation and if native services within CSP are used as connectivity gateways, they must be capable creating connectivity to other CSP gateways, as well as inter-operate with third party VFs.
- **Data Protection.** CSP networks are multi-tenant networks. For many enterprises when operating on public network, encryption becomes a requirement for all traffic. As with perimeter security, scaling encryption in concentrated areas that attract can become a challenge both from performance as well as actual security perspective.
- **High Volume Data Transfer between Clouds.** Similar to cloud access, the challenge is to create networks that mimic traditional Data Center Interconnects between the CSPs and customer networks and how to design them to move large amounts of data. Cloud Access is still a problem of interconnectivity and data transfers, but the primary activity is data processing and not transfers. Exactly the opposite happens here, where special capabilities must exist to facilitate the ability to move very large data sets between small numbers of endpoints.
- **Performance SLOs.** A typical cloud tenant expects a predictable performance for compute node communications. This performance is a combination of throughput, latency and jitter. These performance SLOs exist in most private DC/Cloud deployments, even if they are not enforced or verified and become actual SLAs. When migrating to the CSP, the expectation is to have SLA and distinct service levels for both North/South traffic (in and out of the cloud environment or CSP itself) and East/West (within the cloud environment) traffic.
- **Availability SLOs.** As with performance SLOs, private cloud deployments have special availability SLOs. This is required to determine actual availability of applications that may be deployed across multiple clouds or multiple cloud regions.
- **Perimeter Security.** As enterprises connect to third party Application Service Providers, the enterprises maintain traffic filtering at the network perimeter tightly controlling access to and from these ASPs. Connectivity to CSP is often designed in the similar fashion, except that CSP network is effectively an extension of the enterprise Data Center. As the CSP network is a not a network that belongs to the customer, and provides multi-tenancy beyond the customer control, the challenge is to develop solutions that can put the CSP network within the enterprise's secure network perimeter.
- **Zero Trust.** While it is still possible to design CSP connectivity in a way that provides traffic filtering between CSP virtual networks and enterprise internal physical and virtual networks, perimeter security model will not scale from as the capacity and administration requirements will increase. Scaling the environment will require either lots of perimeters, very restricted resource mobility or gradual migration to zero trust model in which all virtual compute

resources are expected to provide localized security services moving towards a micro-security model. Zero Trust can be Limited Trust model, but it is then up to virtual compute nodes to group themselves into micro-secure domains, where they provide endpoint security as well as build virtual private networks when appropriate.

- **Traffic Analysis Services.** Most enterprise compute environments maintain a special infrastructure allowing them to perform packet capture and analysis on traffic going through the Data Center network.
- **B2B Services.** When resources are move into the CSP networks, B2B communication requirements persist. Many B2B customers and providers will continue connecting through legacy touchpoints, but for those who share the same CSP and B2B service, the ability to use special B2B clouds would be a bonus. This is especially important if opportunities exist to improve performance and conserve network resources. B2B service can be accomplished by allowing different tenants to create virtual links between their cloud networks, as well as creating virtual cloud networks that allow virtual compute instances belonging to multiple tenants to have access into the same virtual networks.

These high levels requirements will be discussed in subsequent sections.

3. Common Definitions

While the CSPs build their services different from each other, it is desirable to have common definitions when describing various physical and virtual network functions and their general attributes.

- **Cloud Definition**. CSPs must create a common framework for defining cloud as a network environment. Devices are in the same cloud if they are part of common domain:
    - Connectivity Domain. Ability for virtual compute nodes to connect to each other without any restrictions.
    - Network Security domain. Can be isolated from other connectivity domains
    - Availability Domain. All systems a part of an environment whose availability is the upper bound for all other availability enhancements within the cloud.
    - Mobility Domain. Virtual instances to move within their network environment without any manual network reconfiguration or the need to rebuild the instance from scratch. Ability to scale up and down falls into that category.

- **Cloud Transport Edge**. System owned by the enterprise or cloud provider providing connectivity between the cloud virtual infrastructure and physical transport infrastructure.
- **Cloud Transport Network.** Physical and/or virtual infrastructure between Cloud Transport Edges.
- **Cloud Gateway**. A physical or virtual system providing connectivity between individual cloud connectivity domains.
- **Cloud Transit Network.** Any form of physical or virtual infrastructure providing connectivity between Cloud Gateways and Transport Edges.
- **Cloud Geographical Domain.** A CSP must clearly define what constitutes a network area in which a cloud can be deployed.
- **Physical Cloud Exchange Network.** A physical multipoint network for multiple cloud Transport Edges of multiple tenants to interconnect together using a common infrastructure

- **Virtual Cloud Exchange Network.** A virtual multipoint network for different Cloud Gateways of multiple tenants to interconnect together using a common virtual infrastructure.
- **Virtual Cloud Peering.** A virtual Point-to-Point connection between cloud gateways interconnecting cloud networks. This is a multitenant virtual link (from one tenant to another) as opposed to exchange network, which is a network interconnecting all the tenants without requiring tenants to build links to each other.

4. Cloud Access

Designing transport networks is ultimately tenant's responsibility. Transport network designs would be driven by availability, performance and scale requirements, but in general this connectivity would have two major variations:

- **Private Access Networks.** Tenant building out a transport network and creating a physical touchpoint between its Transport Edge and CSP Transport Edge. This touchpoint provides connectivity between users and data processing services on the tenant's network and various clouds on the CSP network.

- **Public Access Network**. Tenant using public services such as the Internet, special cloud exchanges or general purpose Internet exchanges. It is up to the tenant to decide where to extend the network perimeter, but the expectation is that CSP Transport Edge and Cloud Transit Networks are accessible by other tenants.

- **Cloud Network Functions.** As the organizations build networks in public clouds, these networks will have various forms of network gateways. **CSP overall physical and virtual topology.** The tenant must have a general understanding of what the network looks like between the transport edges and virtual compute services.
- **Service Level Objectives**. Must provide network connectivity SLOs for availability and performance (jitter, latency, packet loss)
- **Provide Performance Data.** On public network must offer the ability to measure performance between the CSP Internet Edges and major public network points of presence.
- **Global and localized access**. For Internet services, may offer both globally accessible service or when connecting to Internet Exchanges, only exchange based services.

The above basic requirements turn the CSP into "network connectivity advisor" something that majority of the Network Service Providers already do. The CSPs are excellent in documenting the guides for creating connectivity, but the CSP must start providing documentation describing its network internals in the same way, a typical ISP does it.

The next layer of virtual infrastructure between CSP Cloud Transport Edge and the actual cloud resources is the Cloud Gateway. In the pure context of network infrastructure, it is the boundary between tenant virtual networks and CSP aggregation infrastructure. The only major requirement for such a system is to act as IP router between tenant virtual networks and either public networks or private transport network. The virtual device is expected to provide both static routing as well as support for BGP protocol.

5. Cloud Network Services

Cloud Network Services can be CSP designed and deployed, or can be installed by the customer as virtual compute nodes. It is desirable to use CSP resources and all components mentioned under Common Definitions must be offered by the CSP. Requirements are the following

- **Inter and Intra Cloud Interconnects.** If the CSP allows building internal networks, it must provide all the components to interconnect them to other physical and virtual networks.
- **Inter-cloud Security Services.** Ability to provide packet filtering at the network level between distinct cloud network domains.
- **Intra-Cloud Security Services.** Ability to provide packet filtering at the network level to restrict connectivity access to individual virtual compute nodes. This becomes important when discussing Zero Trust networking in which hosts or groups of hosts can maintain their own special purpose network perimeter.
- **Security Groups –** Ability to contain a set of virtual compute nodes into a single service with common security policy. It is a combination of creating localized network perimeters as well as creating a system to maintain these network perimeters independent of end system addressing.
- **Service Chaining –**ability to traffic engineer network application flows (all distinct flows making up individual application session) between application end systems and have it traverse intermediate virtual compute nodes along the way that host various services. Intermediate compute nodes can either act as network and application proxies or transparently pass and modify packets
- **Encryption –** Ability to encrypt traffic between CSP network gateways.

6. Hybrid Cloud Connectivity and Data Protection

Hybrid cloud connectivity is usually achieved by building virtual connections between cloud networks hosted in different CSP environments or Private Cloud environments. CSP or tenant administered cloud gateways build these connections. These connections are meant to resemble virtual Point-to-Point links

- **Standard Control/Data Plane Protocols –** Ability to build virtual IP connections between diverse cloud environments using standard protocols. At the very minimum this connection can resemble IP tunnel using static routing to provide reachability between cloud environments. Use of dynamic routing is desired as well, but assuming that a combination of CSP provisioned cloud gateway and virtual link can provide availability levels found in systems with redundant components, it may not be necessary. This type of connectivity should not be designed as transit (CSP 1 <-> CSP 2 <-> CSP 3)
- **Encryption –** Support for common encryption parameters when provisioning these “links

7. High Volume Data Transfers

As the tenant moves the data into the cloud, or between different clouds of different CSPs, the amount of required bandwidth may be by factors of magnitude greater than the amount of bandwidth required for day-to-day operation. Most Data Centers and environments built to connect to them are built for many to many communications optimized for access and data processing. When many become few, it may be desirable to build compute and networking optimized for transfer. The following requirements exist for CSP to address this challenge.

- **Specialized DCI Network Design**. Must have Cloud Gateway and Transport Edge infrastructure optimized for connecting at very high speeds directly to high speed transport.
- **Resource Reservation.** Must have the ability to reserve resources on demand.

The responsibility for providing a very high speed transport is the responsibility of third parties or the enterprise itself. It may be interesting to see direct interconnect relationships between different CSPs, but these should be completely transparent to the enterprise. While such transfers can support encryption, it can only be achieved using specialized physical hardware or at application level.

8. Performance SLOs and SLAs

SLOs can range from best effort service to precise availability and performance parameters for specific application traffic with some conservative guarantees in the middle. They may turn into actual negotiated SLAs. Since multiple applications and service domains go through the same transport network, Cloud Transit Networks and Cloud Transport Edges, then multiple SLA levels would have to be enforced. Cloud and transit networks may have different methods identifying traffic belonging to applications requiring different SLA levels, but the transport network must have a uniform way of identifying those. Some of the requirements for performance

- **CSP to Standard DiffServ Conversion**. Must be able convert from whatever packet or just application analysis mechanism that exists within the CSP to standard IP packet ToS field markings to interoperate with transport DiffServ model.
- **Performance Domains.** CSP must clearly define performance domains. When resources are deployed, the tenant must have the option of determining which domain the resources can reside in and how it impacts application performance.
- **Performance Domains vs. Availability Domains.** CSP must clearly indicate the relationship between performance domains and availability domains. Trying to deploy infrastructure to minimize latency and jitter may impact its availability.
- **Performance Domain SLOs.** Performance SLOs must exist for crossing performance domains. Since geography and transport network SLOs play a huge role when crossing performance domains, the CSP must have the capability to provide performance SLAs not just for individual domains.
- **Throughout.** Throughput SLO must define the percentage of traffic is guaranteed to get between any two virtual compute nodes.
- **Latency.** Latency SLO must be available for a particular performance domain. It should give any tenant deployment choices in terms of how close different services must be to each other.
- **Jitter.** Jitter SLA must be available to allow the tenant to offer near real life performance to some applications.

- **Service Classes**. It is up to the CSP to determine the amount and types of service classes and their individual performance parameters. The tenant must have the option to select any amount or no service classes at all. There must a premium class operating as Strict High Priority, as well as a class that operates as Best Effort.

9. Availability

Availability is a function of MTBF and MTTR and covered in other documents. The general expectation is that the availability of each component is less than or equal to the general availability of the entire cloud service. Two basic requirements exist.

- **Connection availability.** The CSP must provide availability numbers for connectivity between any two virtual compute nodes within the cloud infrastructure including Cloud Gateways.
- **Enhanced availability.** The CSP may offer the option for enhanced availability for certain types of traffic.

10. Perimeter Security

Due to the nature of the cloud service itself, the traditional perimeter security model, in which packet filtering aggregation devices provided separation between many networks and controlled access between them starts breaking down, as the compute becomes more scaled out and distributed and can no longer be in monolithic network constructs administered by the enterprise. This does not mean that the infrastructure is no longer secure, but the fact that all the components of a traditional security aggregation devices are distributed and scaled out as well. The challenge of perimeter security is to move from one all-encompassing monolithic perimeter offering all the protection to all of infrastructure to a set of perimeters around individual infrastructure components organized in various tiers providing protection services. Three major tiers can be identified within the CSP.

- **Tier 3 – Service Tier.** Protection of individual virtual compute endpoints. This is ultimately responsibility of the tenant and defines the most granular protection mechanism, providing and restricting access to individual services. Service domain is a single virtual system or group of virtual systems running a specific service with very service specific access protection mechanisms. The perimeter around service domain is a perimeter around individual virtual compute node.
- **Tier 2 – Connectivity Tier.** Protection of individual virtual networks. This service can be offered either by the tenant in the form of tenant deployed Cloud Gateway or by the CSP in the form of CSP provided Cloud Gateway. Combinations are possible, but protection is aggregate, which is reduced to protecting entire network domains from other network domains.
- **Tier 1 – Infrastructure Tier.** Protection of entire infrastructures, including physical aspects of it. Ultimately the responsibility of whoever provides the infrastructure, so it becomes the CSP in all CSP networks. It involves such global protection mechanisms, such as closing off the infrastructure belonging to the tenant from other major public networks as well as providing various monitoring and reactive services.

It is beyond the scope to analyze all the security tradeoffs, but the ultimate goal of all these tiers is to still replicate the original continuous perimeter security domain and dividing it into tiers and spreading those tiers out is scaling out.

The CSP can develop various product offerings that help with security at any of these tiers, but one of those tiers has to form a virtual secure connectivity domain that that can connect various perimeters together and having the CSP deal with the interoperability of the components of that tier in different CSPs and customer networks may turn out to be very counterproductive. The lower the tier, the better job the CSP can do. Majority of security requirements are actually tenant requirements. The CSP is responsible for infrastructure tier and those requirements are primarily monitoring and reactive, covered bellow and in other sections, as well as optional security features in Connectivity Tier. Tenants that request enhanced security features, as well as various VPN-like features are better off deploying specialized Cloud Gateways. There are some basic security requirements.

- **Cloud Gateway Filters**. CSP must provide the ability to set up network and transport layer filters on Cloud Gateways. These filters may keep track of connection state for state full transport protocols.
- **Packet Redirection.** The tenant must be able to redirect offending traffic to special compute instances that can process offending traffic, clean it up, strip invalid flows and either drop or redirect it to destination application hosts. This service should be available on CSP deployed Cloud Gateways.
- **Denial of Service Traffic Drop or Redirection.** When connecting to the public networks, the tenant must have the ability to signal to the network edge to drop traffic based on network and transport layer information.
- **Tenant Control of all Access into the Cloud.** The CSP may offer other proactive services that restrict access, but they should be enabled or disabled at the tenant's requests.

11. Zero Trust or Limited Trust

Any virtual compute instance is assumed to be deployed in an environment where infrastructure is completely shared. While connectivity can be segregated, breaches are still, in theory possible when virtual compute instances are exposed to connection attempts from unauthorized sources. As with Cloud Gateways, there has to be way for the tenant to deploy security to protect virtual compute hosts from such connections attempts. Doing it by the tenant is beyond the scope of this document, but the CSP must offer various packet filtering services and the ability to support security groups. Intra-Cloud Security becomes critical in supporting limited trust models.

12. Traffic Analysis Services

The basic requirements for Traffic Analysis are the requirements that exist for all private Data Center operators. In its ideal these requirements involve the ability to look at all the traffic, be able to short term analyze it, as well as do more complete long term analysis. The basic requirements for CSPs are:

- **Virtual Traffic monitoring infrastructure.** The CSP must provide capabilities to successfully capture all traffic from any endpoint in the cloud for some period of time and have that capture available for analysis and pass to analytics systems.
- **Non-sampled Flow Accounting.** The tenant must have the ability to see all traffic going to and from all endpoint sin the cloud for enough time to extract flow accounting records, de-duplicate the data and make it available to any security or analytics VNFs. Where the previous requirement is more meant for troubleshooting, this is design more for continuous security monitoring.
- **Publishing Information.** The tenant must be able to publish this (flow) information to various analytics tools. CSP may offer these security analytics tools as a product or use common standard protocols to allow publishing this information to third party tools.
- **Retention of Flow Data.** A database of records of all flows (flow being either a set of packets with distinct transport and network layer header attributes or a set of packets identified as belonging to the same application instance) allowing the operator to have all the record of system communications.
- **Deep Packet Inspection and Metadata** Ability to perform packet inspection to analyze information at other layers above network and transport to be able to sample applications flows. Ability to analyze and put special Metadata on packets that can be used for services like endpoint security and SLA enforcement.
- **Packet Scrubbers.** When Cloud Gateways redirect offending traffic, the CSP may offer DPI services that clean up and redirect the traffic back to targeted destination hosts.

13. B2B Service or Community Clouds

When multiple networks are interested in connecting to each other, they may either utilize the network created by one party or a common network operated either by a commercial third party or a consortium to enable interconnectivity. The first approach is generally designed for situations when there is primary content provider and the second approach is when multiple networks are producers and consumers or resources used by users of these networks. A CSP that has multiple tenants would benefit from developing a similar solution removing the need for customers to send their traffic into the CSP, out from the CSP towards network exchange and back into CSP. This would be a Community Cloud and two types of such clouds should exist.

- **Tenant Controlled Community Clouds.** The CSP must provide the capability for the tenant to set up a typical cloud and then utilize a portal or an API, which would allow guests to become part of that cloud. These clouds would operate without gateways and designed for tenant-to-tenant interconnectivity.
- **CSP Controlled Community Clouds.** The CSP can create either general purpose or special purpose (such as for all qualified FSIs) Community Clouds that different tenants can join and utilize public interconnectivity services.
- **Community Clouds vs. Inter-Cloud "Peering" Virtual Links.** If distinct clouds belonging to different tenants want to interconnect, they can set up virtual links between CSP provisioned Cloud Gateways. This approach requires setting up virtual links for every interconnection, at the

same time insuring a very clean delineation. The same network design tradeoffs apply here as when choosing point to point connections between individual networks vs. networks interconnecting many touchpoints. Both solutions have a place when designing B2B services and the desire for dynamic routing is no longer driven by availability concerns, but by the need to simplify administration.

The general architecture of such a cloud is similar to any form of network exchange, such as Internet Exchange and has the following architectural requirements.

- **Virtual Ethernet Network.** Cloud networks can resemble Ethernet switched networks, full routed domains or a mix of the two, but the Community Cloud network is best positioned to emulate Ethernet switched network, even if it is not a true Ethernet broadcast domain. This is to ensure that all tenant Cloud Gateways can create IP routes using each other as next hops. In fact, the underlying architecture can be anything as long as all cloud tenants can resolve each other using Layer 2 protocols, even if it is a proxy resolution. Virtual IP networks can be created, treated as IP VPNs by the CSP, but then accommodations must be made to support passing various routing data.
- **Point to Multipoint Clouds.** CSP may offer a special service for Community Clouds with the main content provider, where all guests can connect to a single host or set of hosts, but not to each other. It is a debate on whether P2MP clouds are ultimately a use case for Inter-Cloud virtual links.
- **Central Route Servers.** Tenants can exchange routing information about their Cloud Networks using CSP provisioned Route Servers. The CSP may offer such service or allow the tenant who controls the cloud to deploy one.
- **Ability to Restrict Guest Tenants.** Tenants who control their community clouds must have the ability to restrict or remove guest tenants on demand.