ISP Design Guide:

Overview and introduction to separated functions

What are network functions?

Network Functions - The major tasks in the data plane that must be performed by an L2/L3 network device to ensure smooth delivery of the Internet from the border of an ISP down to the subscriber last mile. Examples are border routers, core switches and aggregation routers.

What are operational support functions?

Operational Support Functions - The major tasks in the control and management plane that must be performed by a device or service to facilitate and support the operation of network functions. Examples are DHCP, DNS, Applications/Servers, Billing Systems, Corporate VPNs and connectivity.

Why separate them?

It's tempting for new and even experienced ISPs to pile all of the functions into one router, switch or server - and then add another for "redundancy". This generally creates problems with complexity, failure domains and growth. Separating functions allows for network designs to be modular, repeatable and more scalable. Automation is easier beacuse templating is easier. The end result is better uptime, lower opex, easier growth and lowered risk.

ArchiTechs MANAGED SERVICES
iparchitechs.com
+1(855)645-7684
consulting@iparchitechs.com

Call or e-mail for professional assistance with your network

Core - The job of the network core is to connect all other devices and functions as simply as possible.

Ideally, the core has a very simple L2/L3 config and enough ports to connect the current prod devices and have room for growth.

This is a great place for Layer 3 switches because they are fairly inexpensive these days and come with a variety of port layouts, densities and speeds.

NAT - Network Address translation is increasingly used by service providers as IPv4 has become more scarce.

Typically CG-NAT in a NAT444 configuration to support a dual stack deployment with IPv6 is the most common.

This is a market segment that's grown significantly in the last year due to the bandwidth explosion caused by the pandemic and a move to working remotely. This function can also use NAT64 or 464XLAT for single stack networks that need IPv4 connectivity.

QoE - Quality of Experience or QoE is a term that's become popular within the last 5 years or so. It generally refers to a shaping appliance that has advanced traffic identification capabilities at L7 and data about the health of the network. The appliance will normally sit logically or physically inline between routers in the core and the aggregation layer.

Most QoE appliances use Active Queue Management (AQM) shapers like fq_codel or cake to manage throughput to each subscriber.

Shaping - Shaping of traffic is an important element in end-to-end delivery of bandwidth. It can help to smooth issues with capacity, backhaul quality and even wireless problems inside the subscriber's home or business.

Without going into an enormous amount of detail around shaping, it differs from policing/rate limiting in that it queues a portion of traffic and tries to hit a target rate before it's exceeded.

Whereas policing/rate limiting is a hard limit that drops traffic as soon as the max rate is reached. This function is commonly found on aggregation and last mile routers but can also be a separate appliance which will be expanded on in the QoE description.

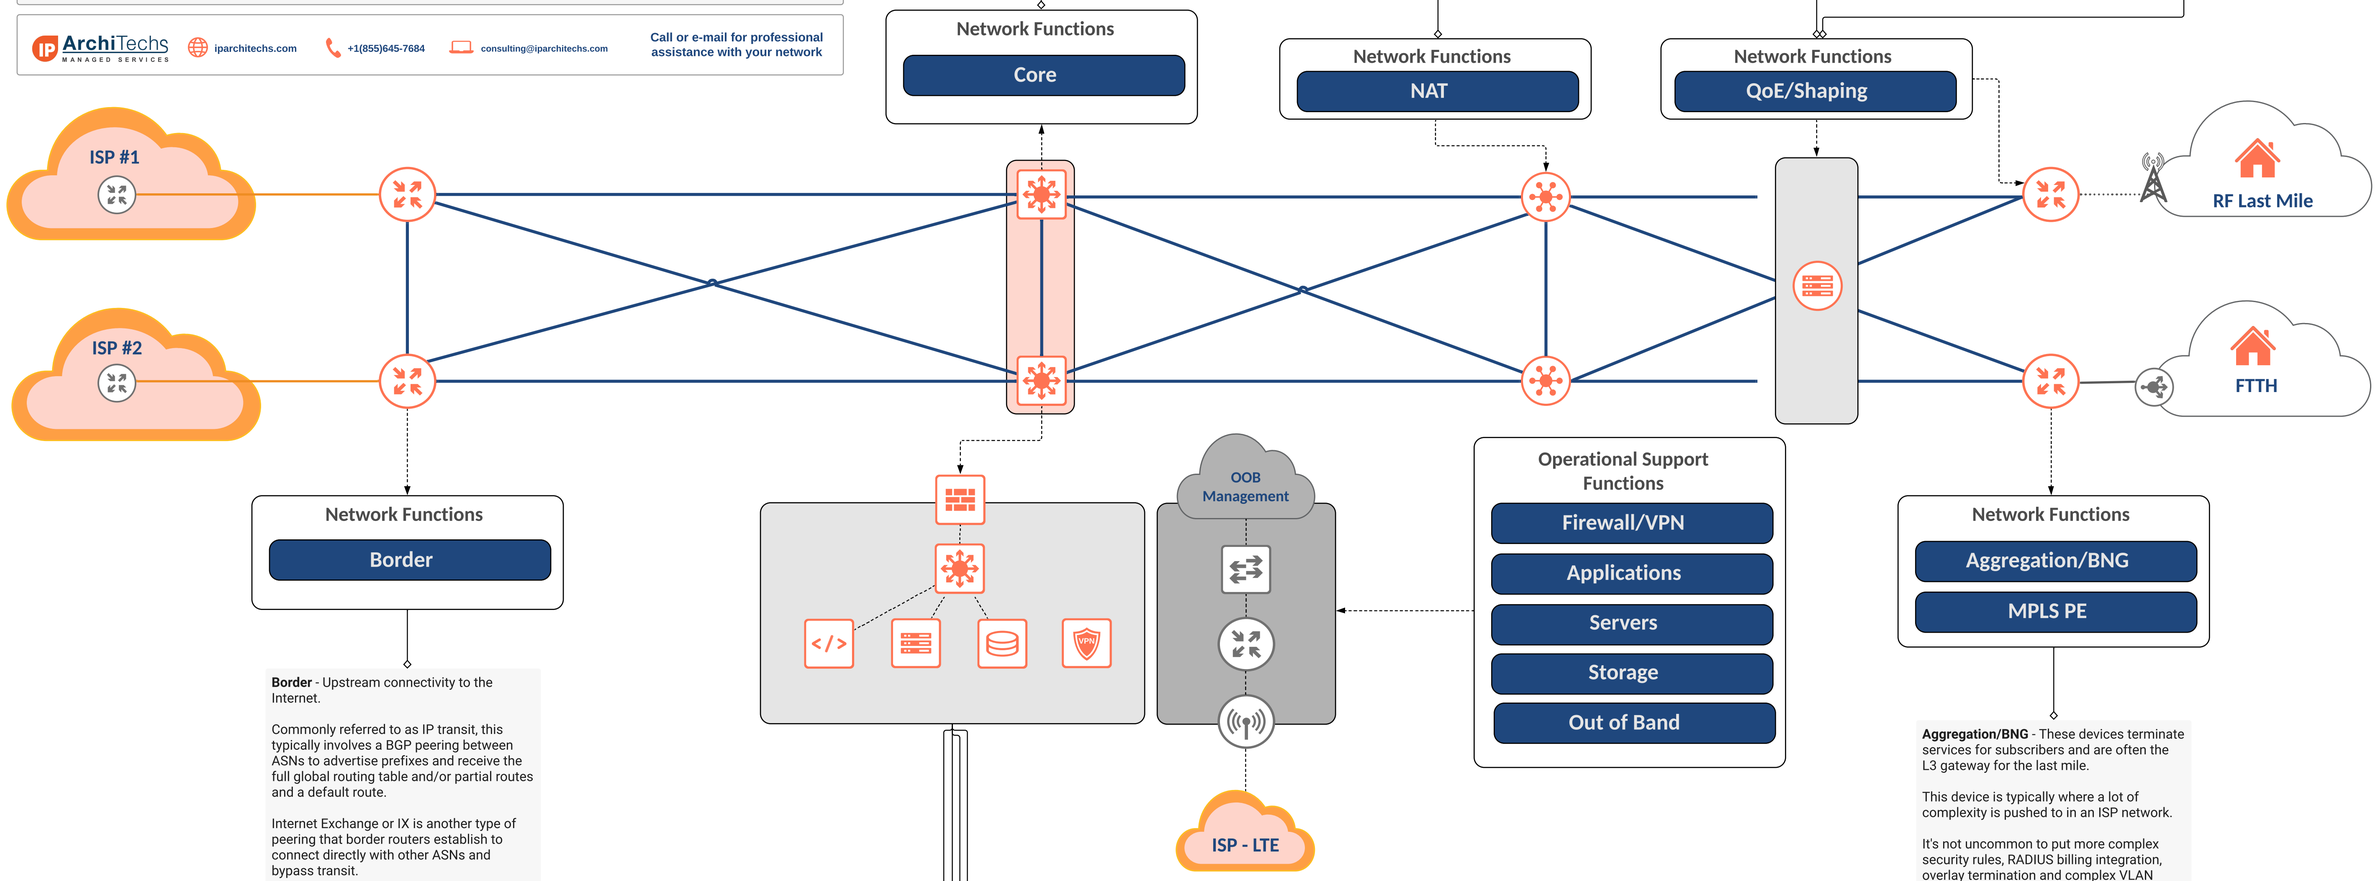

Border - Upstream connectivity to the Internet.

Commonly referred to as IP transit, this typically involves a BGP peering between ASNs to advertise prefixes and receive the full global routing table and/or partial routes and a default route.

Internet Exchange or IX is another type of peering that border routers establish to connect directly with other ASNs and bypass transit.

Peering and failver can become more complex as the ISP grows, so keeping the border layer separate reduces the amount of churn when a peering needs to be turned up or down because the border layer and any failover or path changes are insulated from the customer gateways on the agg/bng routers by the core layer.

This also allows for connection tracking and maintaining "state" to be turned off on the border and core functions so they can defend against low level attacks and put more resources into throughput and not unnecessary "border" firewalls.

Not all startup ISPs will peer with their upstream, so for the purposes of this drawing, Direct Internet Access or DIA is assumed to fall under border as well.

Firewall - One of the more common mistakes that new ISPs make is using the routers that carry subscriber traffic as the security and L3 gateway for applications and servers.

This creates a challenging environment for server/application failover, becuase it's directly tied to subscriber failover and uptime.

It's easy to take down the entire network when working with firewall and NAT rules so creating a separate "data center" zone for operational support functions makes sense.

Using a router as a firewall or an actual UTM firewall helps to separate the applications and services needed to support the ISP from subscriber traffic.

VPN - Separate from the OOB VPN, an in band VPN can be used for both techincal and corporate access to the networks as well as apps and services.

Applications - ISPs often require a variety of applications and systems to operate the network and manage both sales and technical operations for the subscriber.

Examples of ISP applications are monitoring, billing, management controllers, IPAM and DNS/DHCP.

Designing an application infrastructure can be a daunting task for some and if it's not separated from the service provider data plane routers, it can be very hard to grow and change.

With the advent of cloud services, it's sometimes more advantageous to begin with cloud hosted applications or split some of the resources between the cloud and on premises.

By separating applications into their own function within operational support, it's much easier to scale and move applications to and from the cloud without impacting the production data plane.

This also makes security patches much easier as the impact of upgrading is lower when the entire ISP isn't on one device.

Servers/Storage - Most ISPs will run their own servers and storage to support applications and corporate needs.

This can be as simple as a single server with a single disk or as complex as several racks of dedicated compute and storage.

Whether the servers are intended to be used as bare metal, hypervisors or container platforms, building them in the "data center" zone as a separate function keeps the servers from becoming a dependency in the prod ISP routers.

One common and often costly mistake is running virtual routers that are intended to be used in a network function role like border or BNG on top of the same infra that's used to deliver applications.

Network Function Virtualization or NFV is intended to be used on a standalone server/hypervisor that's dedicated to the data plane and will not be running applications or clustering.

The performance tweaks for apps. vs. routers are different and can interfere with each other. Keep these roles separate.

Out of Band (OOB) - This is one of the most overlooked additions to an ISP network. OOB is a critical element to executing low-risk maintenance windows.

It's also incredibly helpful to reduce the number of trips to the data center and can provide access during an unplanned outage.

Out of band networks typically are standalone with a router, switch, console server and possibly a laptop or small server.

IP enabled PDUs are an optional but very desirable item to have in an OOB network to remotely power cycle equipment.

The internet connection used to feed the OOB network should be separate (if possible) from the transit or DIA connections used to connect subscribers.

In this example, LTE securely connects the OOB network with users and might use a mesh VPN like ZeroTier.

Aggregation/BNG - These devices terminate services for subscribers and are often the L3 gateway for the last mile.

This device is typically where a lot of complexity is pushed to in an ISP network.

It's not uncommon to put more complex security rules, RADIUS billing integration, overlay termination and complex VLAN tagging operations - as well as other types of config that wouldn't generally go into a border or core router.

The role can also combine other functions like NAT and Shaping.

This is sometimes called a Broadband Network Gateway or BNG which is a term from the days of PPPoE and usually implies dynamic service provisioning of shaping polices, IP assignment and other attributes via RADIUS integration with DHCP and a billing system.

MPLS PE - Similiar to agg/bng, the role of an MPLS PE is to terminate MPLS services like VPLS and L3VPN and forward the traffic to non-MPLS enabled routers/switches.

Typically the MPLS PE will be using an IGP like OSPF or IS-IS and label distrbution via LDP or SR-MPLS. Some ASNs will also run iBGP to advertise prefixes to the border.

Sometimes the BNG function is mixed with the MPLS PE function and MPLS services are dynamically provisioned, shaped, etc.