High Availability & Resiliency

An overview of **Pelion's** resilient architecture
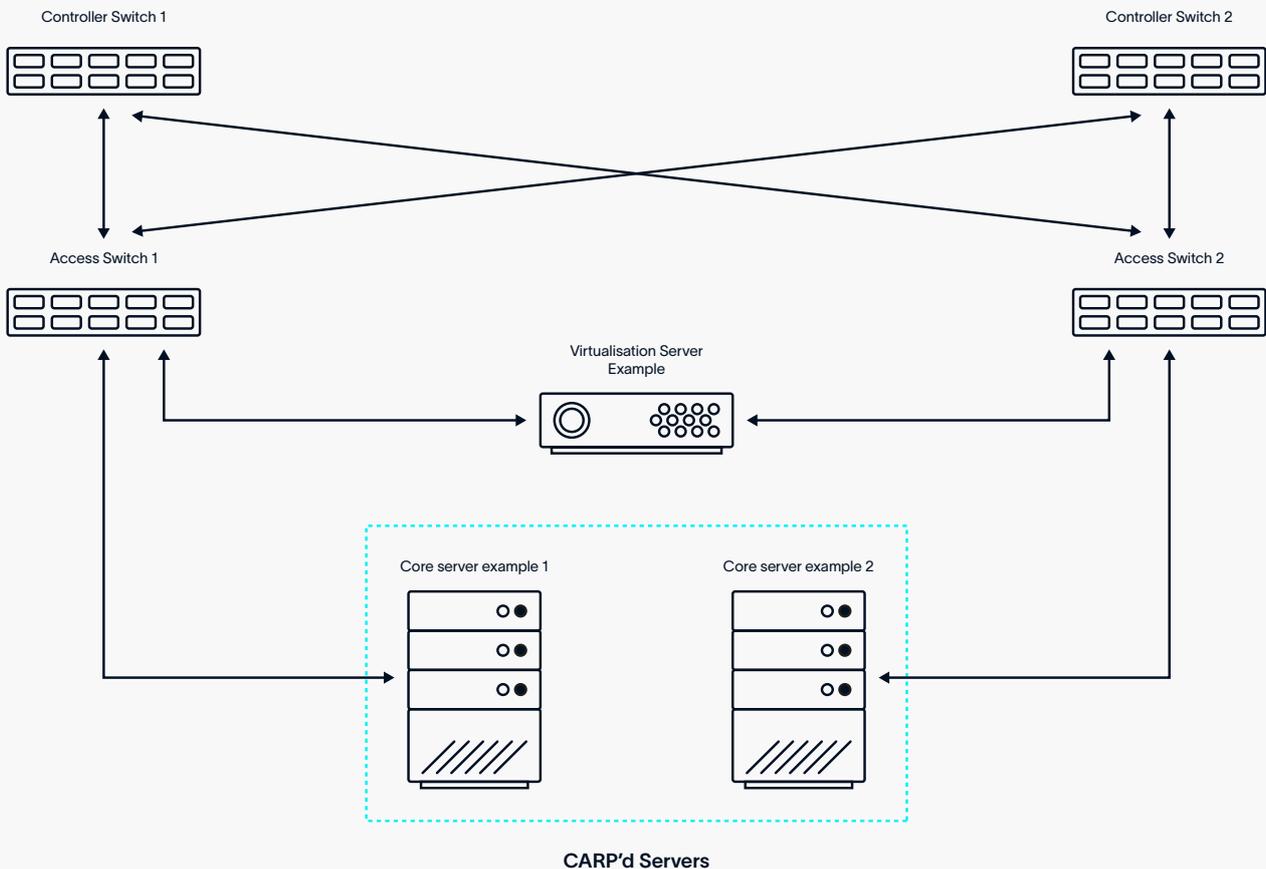
Intro

Pelion uses a mixture of scalable architectures across all types of hardware and datacentres to ensure high availability (HA) and minimal downtime. We work to N+1 in all areas and all services are deployed as N+1. We largely employ active-active techniques, meaning our servers and datacentres are actively sharing workloads, and each place there is a workload it has at least one backup component to take over if needed, and at least two routes there.

Single site resiliency

Equipment

- Everything into switching layer has a primary/secondary route

- Core routers CARP'd (means they are an HA pair sharing the same subnet, have IP failover, and are aware of each other) if one dies, the other takes load (core routers are for firewall and routing rules)

- VPNCs CAPR'd if one dies, other takes load (VPNCs are for customer IPSec endpoints)

- Anything without 2x CARP'd servers i.e. data stores or clusters/virtualisation is set up using RAID or similar to continue operations even with failed components.
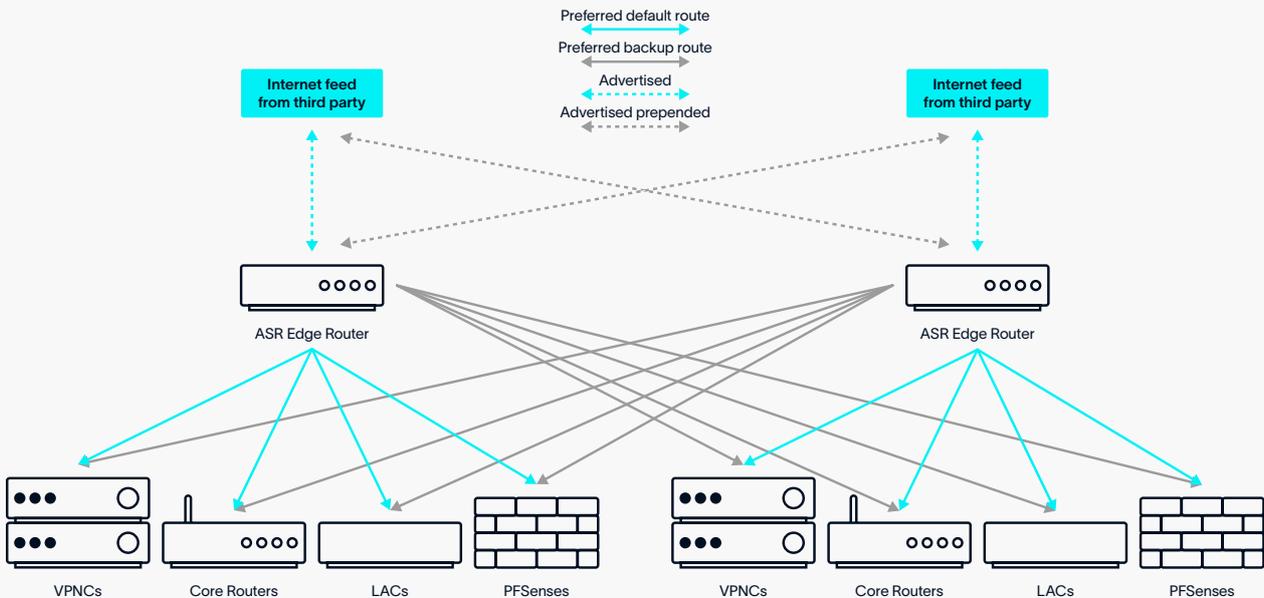
What cabling within site looks like.

Controller Switch 1   Controller Switch 2

Access Switch 1   Access Switch 2

Virtualisation Server
Example

Core server example 1   Core server example 2

CARP'd Servers

Internet redundancy

- All services in Site A will be redirected to the internet in Site B if the primary link goes down

What internet redundancy looks like:

Power

- Each datacentre is allowed to use only 80% of the power into the rack in order to avoid failures, or we will up the power allowance with the datacentre

- Each datacentre has 2x PDUs in a rack

- Every piece of equipment (server, router or switch) has two power cables: one into PDU1 and one into PDU2

- Every piece of equipment has two power packs.

Operator redundancy

- Every operator is integrated twice and automatic failover controlled with BGP

- All RADIUS is dual site capable and load balanced

- 4x types of operator redundancy:

    - Leased line into each site for data and RADIUS

    - IPSecs into site for data and RADIUS

    - Leased line into the 3rd datacentre and forwarded to our core routers in the other sites and either a leased line/IPSec elsewhere for data and RADIUS

    - L2TP integration including data over one of the above and RADIUS from 4x load balanced LACs across the sites.

Load balancing

- Operators are split over the two sites

- RADIUS is load balanced across both sites via the cross site links no matter which site is that operator's primary

- Customer traffic can authenticate and traverse either site across the cross site links no matter which site their operator's primary.
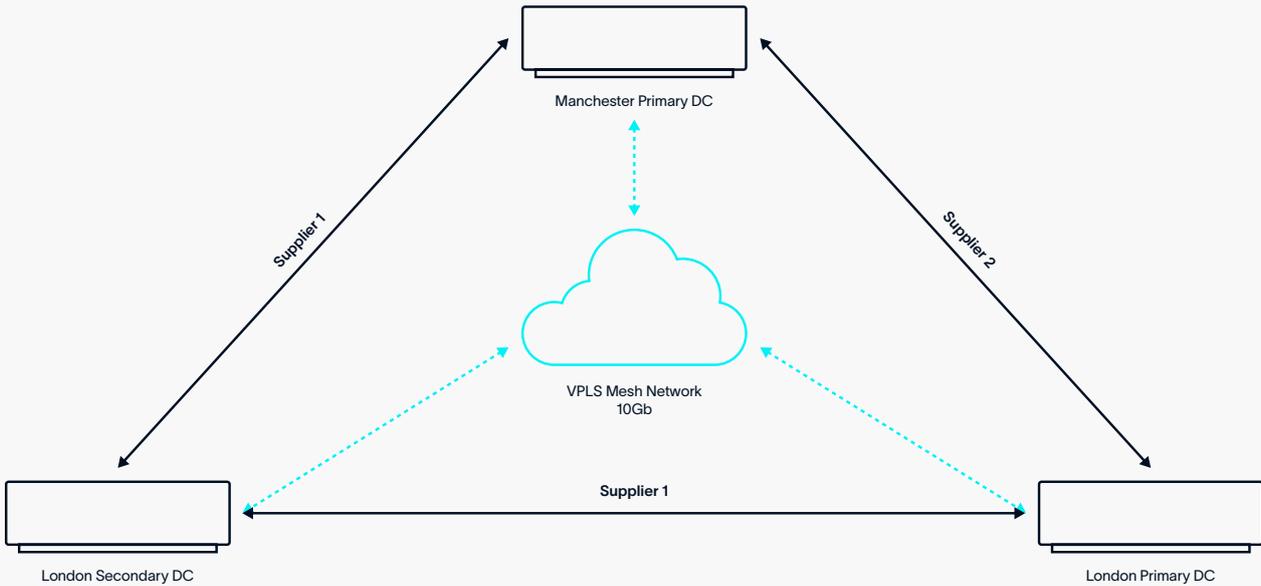
Cross site resiliency

- If one site goes down all operators are able to automatically failover to second site using BGP

- Capacity of one site is always enough to hold all traffic in an emergency (10-25Gb ports at edge, over the cross-site links and on the switching layer)

- All core services are duplicated exactly in each site

- All data collected and securely stored is able to be collected/pushed from the sources in each site across to the other site

- All data is geo-redundant meaning no matter which site the source is from it is replicated else where at least twice

Cross site link resiliency

- Mesh network across all sites means all sites can communicate even in the event of node failure within the mesh

- Physical links through 2 different suppliers, if one goes down then then other takes it on

What a cross site looks like:

Manchester Primary DC

Supplier 1

Supplier 2

VPLS Mesh Network
10Gb

Supplier 1

London Secondary DC

London Primary DC

Datacentre and supplier resiliency

- We have two primary datacentres which work in an active-active manner. We split out which things run across both for capacity and resiliency management

- We have a third spare datacentre used for additional connectivity if and when we see fit

- The primary datacentres are geographically dispersed

- Same with all suppliers – we will always use two separate suppliers for redundancy.

Monitoring

- All of the above is monitored and alerted on so in the event of any type of failure, service is automatically recovered, while internal staff will also immediately notify for triage and action.

Pelion