



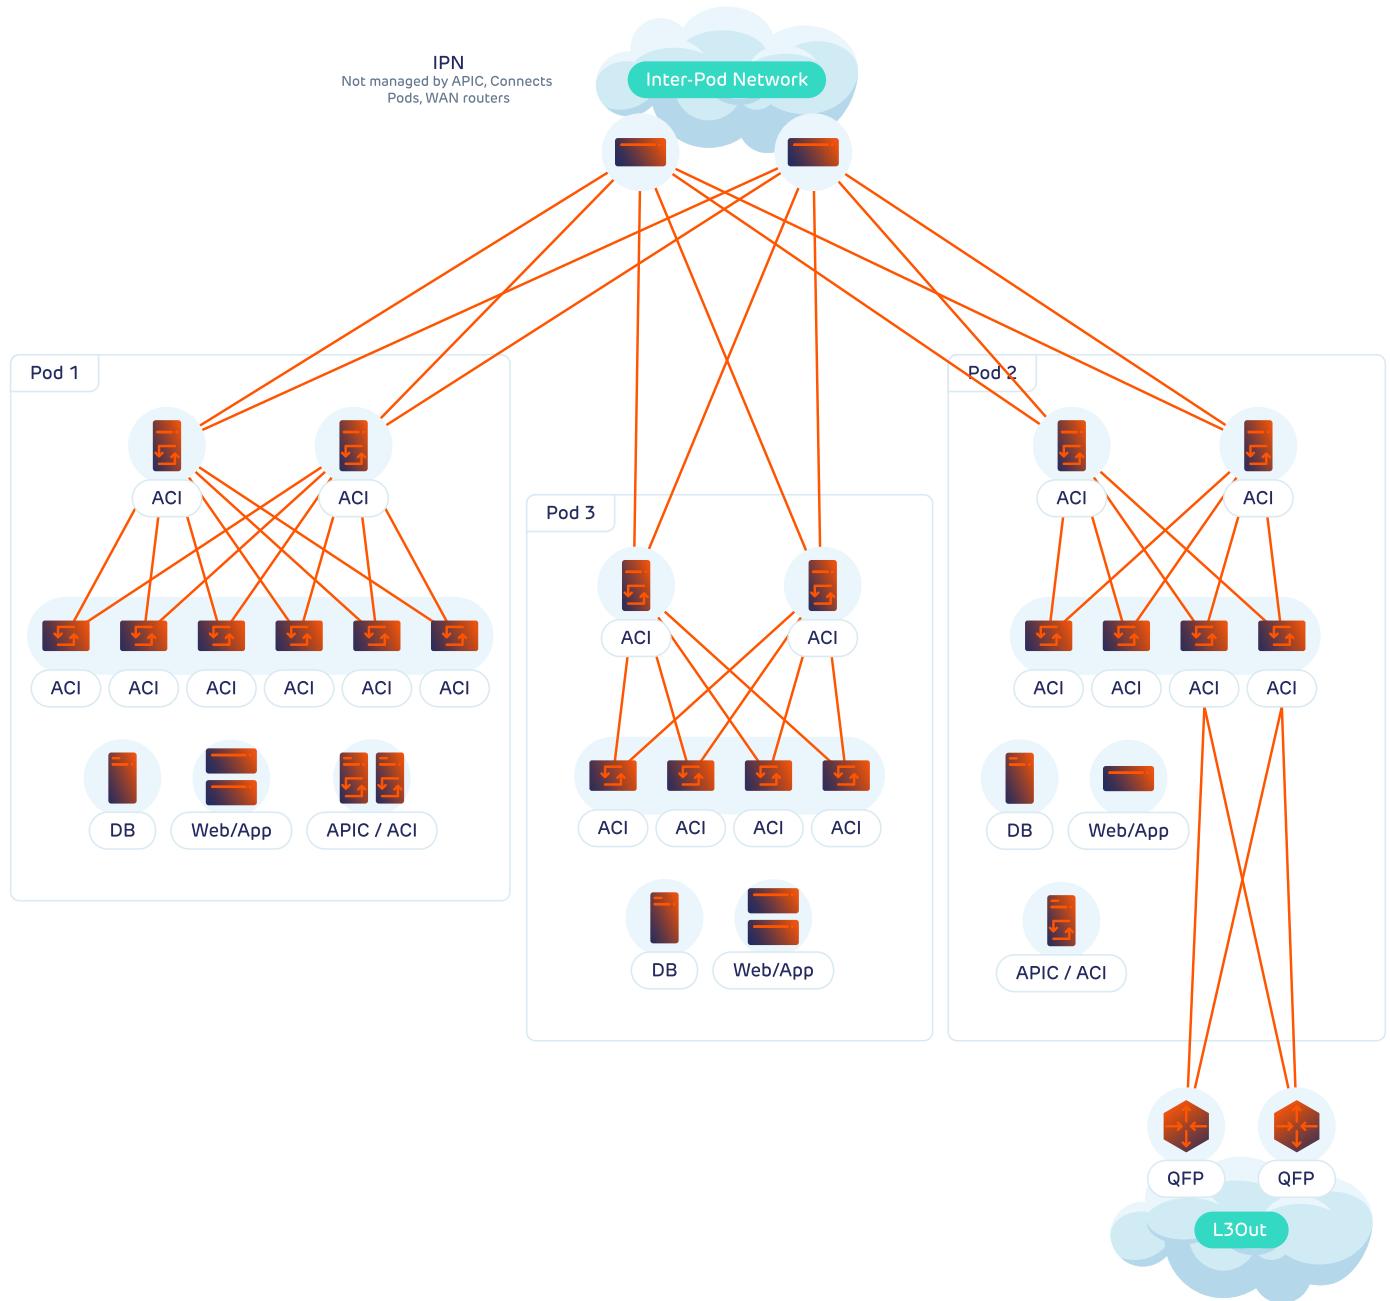
White Paper

Application Centric Infrastructure (SDN for your Data Centre) Demystified

How does Application Centric Infrastructure (ACI) coincide with the context of Local Area Networks (LAN) and Wide Area Networks (WAN)?

As a follow on from the previous software defined white paper, this white paper seeks to attempt to place Application Centric Infrastructure (ACI) terms and concepts into the context of traditional Local Area Network (LAN) and Wide Area Network (WAN) concepts. In this paper we will focus on general terms that crop up in ACI and how a Multi-Pod ACI implementation functions.

ACI Multi-Pod Deployment



The Application Policy Infrastructure Controller (APIC) is the management and policy plane of ACI.

It controls how each switch behaves within the network, which in ACI speak is known as a fabric. The functions of a fabric are based on a set of user created policies, which are usually referred to as the Policy Plane.

The APIC also controls the physical addressing of the switches from a Tunnel Endpoint (TEP) point of view. The APIC runs a Dynamic Host Configuration Protocol (DHCP) server that hands out addresses to the spine and leaf switches, it should be noted that APICs are connected to leaf switches. In Virtual Extensible LAN (VXLAN) a Virtual Tunnel Endpoint (VTEP) is used in the terminology to describe the VXLAN tunnel source and destination.

TEPs form the tunnel source and destination addresses to carry overlays

An Endpoint Group (EPG) is a collection of hosts (that can be micro-segmented) that live in a Bridge Domain, EPGs work in almost the same way as Security Group Tags (SGT) in TrustSec, in fact they do if you integrate Identity Services Engine (ISE) with ACI. Think of EPGs as static / dynamic groups in ISE or the policy part of Virtual Networks (VN) in Software Defined Access (SD-A).

A Bridge Domain (BD) is just a Virtual LAN (VLAN) that share common layer 2 broadcast domains for the EPGs to talk to each other (or not). From a layer 3 perspective a BD can have associated subnet objects (it needs at least one, with the exception of a pure L2 VLAN) and needs to be in a Virtual Routing and Forwarding (VRF) instance. But both the VRF, BD and subnet are separate entities so you need to associate them with each other, it's common to have multiple BDs and subnets in a VRF, just like a traditional LAN.

A BD can be set to operate in flood mode for unknown unicast frames or in an optimised mode which eliminates flooding for these frames. When operating in flood mode, Layer 2 unknown unicast traffic is flooded over the multicast tree of the bridge domain (Group IP Outer (GIPo)). For the bridge domain to operate in optimised mode you should set it to hardware-proxy. In this case, Layer 2 unknown unicast frames are sent to the spine-proxy anycast VTEP address.

Each BD has associated a separate multicast group (GIPo) to ensure granular delivery of multi-destination frames only to the endpoints that are part of a given BD, as within the ACI fabric a BD is not a true broadcast domain like ethernet.

Get the White Paper

Want to keep
on reading?

To get access to this white paper
and the full series, head to:

www.forfusion.com/white-papers