SOLUTION OVERVIEW

A Scalable Solution to IP Address Overlap

Whether on-premises or in the cloud, F5 Distributed Cloud Mesh controls IP address overlap, supporting seamless service discovery and advertisement to ensure full reachability with clean separation between private networks.

KEY BENEFITS

Simple and scalable
Works via a single click outbound and app-centric inbound, eliminating the overhead of one-by-one address re-mapping.

Efficient fleet-wide management
Get seamless bulk configuration without sacrificing localization. Manage all sites as a fleet while still allowing one-off exceptions.

Easy to understand
A full stack end-to-end solution that eliminates confusion from manual renumbering in policy and observability.

Optimized for app-to-app automated connections
Integrated services ensure that apps perform service advertisement and discovery regardless of IP address issues.

Comprehensive and seamless
Multi-site orchestration ensures zero overlooked elements or holes in coverage. Everything works.

Understanding IP Address Overlap

Enterprise organizations are seeking to leverage the benefits of combining previously separate projects. But sometimes those systems can't connect because their networks were built with address ranges that overlap and collide. F5 Distributed Cloud Mesh fixes IP address overlap, supporting seamless service discovery and advertisement to ensure full reachability with clean separation between private networks.

AN OVERVIEW OF IP ADDRESS OVERLAP

As it becomes increasingly easy to connect formerly separate networks and segments, there is a growing risk that the addresses of newly connected segments will conflict with existing networks. In a directly routed network, every segment must be configured with a unique IP subnet address to ensure that the rest of the network knows how to send packets there. If a subnet address overlaps with a different network segment, the network routing systems will not be able to determine which segment is the "real" destination. Services to and from those segments will become unreachable, even if they were working before.

On the public Internet, overlap isn't a widespread problem because all network address ranges are centrally assigned. Within a network, an organization can expand the number of available internal IP addresses by reusing a special set designated for private use. Those private IP addresses can be used by any number of organizations because they're not routable on the Internet, but each address must be treated as unique within the organization. It should only be used once, or problems will occur with traffic both to and from that address.

Historically, IP address overlap has only been common when two previously separate networks are connected, often caused by organizational mergers and acquisitions. If both networks use internal private IP addresses, there's a chance that some of these private addresses get assigned in both networks.

More recently, multi-cloud networking tools provide an agile method to create virtual connections between network segments, such as virtual private clouds (VPCs) or virtual networks (VNets). Traditional solutions for IP overlap are more difficult in these environments, whether due to scale, control, or multi-layer issues.

Therefore, expansion into multi-cloud and edge deployment architectures is making IP address overlap an increasingly frequent problem, even within the same cloud.

KEY FEATURES

Automatic egress Source Network Address Translation (SNAT)
Connections outbound from IP overlap realms automatically use addresses that are globally routable in the private network.

East-west service discovery and advertisement
Apps within the IP overlap realm can be automatically discovered and re-advertised with routable addresses for app-to-app connections.

North-south application delivery
Services are decoupled from their origin IP address, with integrated security and app performance instrumentation for delivery to clients and users with full stack observability.

Flexible consumption
Orchestrated configuration across multiple subnets and/or multiple sites allows delivery either as a reverse proxy (service IP is remote from client) or transparent proxy (service IP is local to client).

Single-site simplicity, multi-site scalability
Centralized SaaS console for easy control of automation in the public and private cloud, with a consistent interface extensible to multi-site orchestration across many different cloud providers.

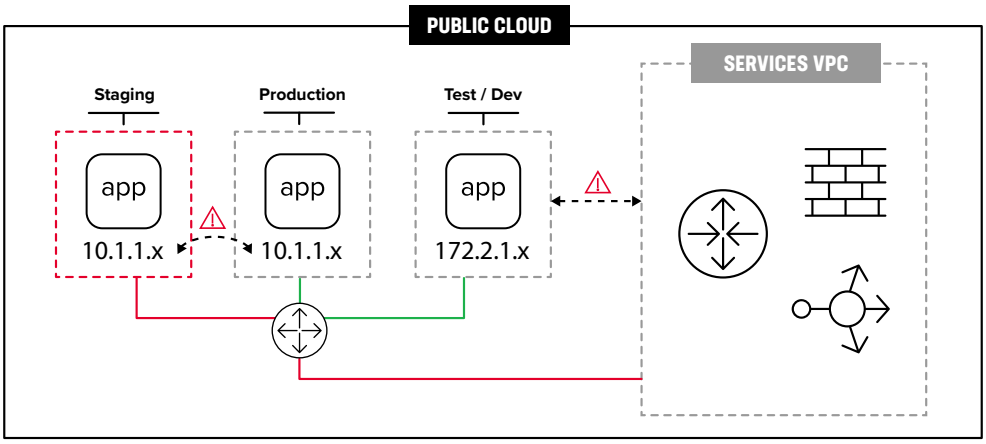

Figure 1: A depiction of IP address overlap and the issues it creates

AN EXAMPLE OF IP ADDRESS OVERLAP

In the figure above, each remote network segment has a unique public IP internet address, each with a private internal IP subnet. The campus and branch locations have similar but different internal subnet addresses with room for expansion, implying central allocation. The public cloud network segments are using internal subnet addresses from private IP address allocation pools. The problem occurs if any of those newly connected segments were configured with the same internal subnet addresses.

- **Traditional networking:** In a traditional networking environment, no problems would occur because the internals of each segment are hidden from the outside. Any services hosted in public or private clouds would be exposed to the outside via load balancer, NAT pinhole, or similar presentation method.
- **Multi-Cloud networking:** Using multi-cloud networking to create direct L3 routing connections between the public and private cloud network segments would cause service failures with the public-cloud virtual private cloud (VPC) and the private cloud. Traditional north-south applications hosted in those segments would potentially be unreachable. For distributed microservices applications, east-west traffic will trigger internal errors rendering these services unreachable, even if the user-facing ingress is not affected.

This problem is not caused by multi-cloud networking. Multi-cloud networking, when implemented properly, enhances agility by allowing rapid creation of virtual direct connections. The problem occurs when agility uncovers an abstraction leak of compromised private addresses. The solution resides in adopting a scalable pattern that fixes and prevents

EXPANSION INTO MULTI-CLOUD AND EDGE DEPLOYMENT ARCHITECTURES IS MAKING IP ADDRESS OVERLAP AN INCREASINGLY FREQUENT PROBLEM, EVEN WITHIN THE SAME CLOUD.

future occurrences.

TRADITIONAL IP OVERLAP APPROACHES

- **Prevention:** The most reliable way of dealing with IP address overlap is prevention through central allocation of network addresses. On the Internet, that's Internet Assigned Numbers Authority (IANA), and for most organizations it's their IT department. Most larger organizations use specialized tools for IP Address Management (IPAM), although some locations still manually track addresses using tools like spreadsheets. This approach works well until something happens beyond the control of IT, such as a merger that requires combining with another network, or a "shadow IT" project that needs to be integrated, or a requirement for direct connection to a previously external network like the public cloud.
- **Renumbering:** The traditional remedy for overlap is to "renumber" or change the IP subnet for any segments that overlapped with another. This means reconfiguring that segment's router, every device on the segment, and every external reference to those devices. It is tedious and error-prone, but a one-time effort—at least, until the next merger. It also works if IT has the power to make changes, which is not always possible in multi-cloud and edge environments.
- **Network Address Translation:** The more recent and more common remediation for IP overlap is to use Network Address Translation (NAT). This makes a segment merely appear to have a different subnet to the rest of the network. NAT effectively turns the segment into its own private network inside the organization's overall network, just like the organization's network is a private network inside the public Internet.

MULTI-CLOUD NETWORKING AND MODERN APPLICATIONS

While the effort for NAT has historically been much easier than renumbering, the trade-off is that NAT creates two separate realms: the inside and the outside. When all traffic is traditional north-south traffic, maintaining the separation has negligible impact. However, east-west traffic adds service advertising and discovery to let app internals intercommunicate. With NAT, the service coordination layer needs to recognize what's an "inside" resource versus what's an "outside" resource.

Large environments, like clouds, especially require a scalable pattern. Otherwise, each VPC or VNet must become its own "inside" and treat all others as "outside" in order for services to

SINCE DISTRIBUTED CLOUD MESH PROVIDES OBSERVABILITY FOR NETWORKING, SECURITY, APPLICATIONS, AND SERVICES, THERE'S NO COMPROMISE BETWEEN LARGE-SCALE MANAGEMENT AND INDIVIDUAL VISIBILITY.

remain reachable.

F5 Distributed Cloud Mesh: A Scalable Solution for IP Address Overlap

F5® Distributed Cloud Mesh is a SaaS-delivered, app-centric network and security service to unify management and simplify interconnection of networks within one or more public and private clouds.

SOLVING FOR IP ADDRESS OVERLAP IN MULTI-CLOUD NETWORKING

F5 Distributed Cloud Mesh fixes IP overlap with a simple, scalable solution that solves and prevents problems from Day One onward. Each virtual network segment is connected to the others by a transparent proxy, rather than a router. Connections that start within a virtual network segment (VNS) are modified with Source Network Address Translation (SNAT) on their way out, changing their address to one that's routable within the overall network.

For connections that should be allowed to enter the virtual network segment, the proxy can discover each app or service that should be available to other networks, then advertise that service to the rest of the network. Traditionally this function would be represented by a firewall policy for source and destination addresses. The difference here is that the proxies for each virtual network segment act together like a distributed firewall, keeping IP addresses of virtual network segments isolated from each other in a way that's transparent to all the apps in each segment. All configuration is orchestrated and automated because a useful tool should make work easier, not harder.

If necessary, the method and policy for each segment can be individually controlled based on the traffic needs. If more nuance is required, ingress and egress can be handled by an API gateway for enhanced visibility and policy enforcement. It's even possible to break the abstraction and create direct routed access, with the caveat that it disables IP overlap protection.

SOLVING FOR IP ADDRESS OVERLAP IN MODERN APPLICATIONS

From a scalability standpoint, Distributed Cloud Mesh offers significant advantages. All sites can be managed as a fleet, which enables the administrator to treat them like a single logical object to configure, with the site-specific details handled by automation. This reduces the complexity of managing large numbers down to the simplicity of a single firewall policy, a single network policy, a single service policy, etc., reflected across potentially hundreds or

thousands of sites. The automation enables apps to discover each other and cross-connect, decoupled from IP addresses so the presence or absence of overlap doesn't matter.

Additionally, Distributed Cloud Mesh provides application-aware observability for networking, security, applications, and services—so there's no ambiguity from IP addresses, and no compromise between large-scale management and individual visibility.

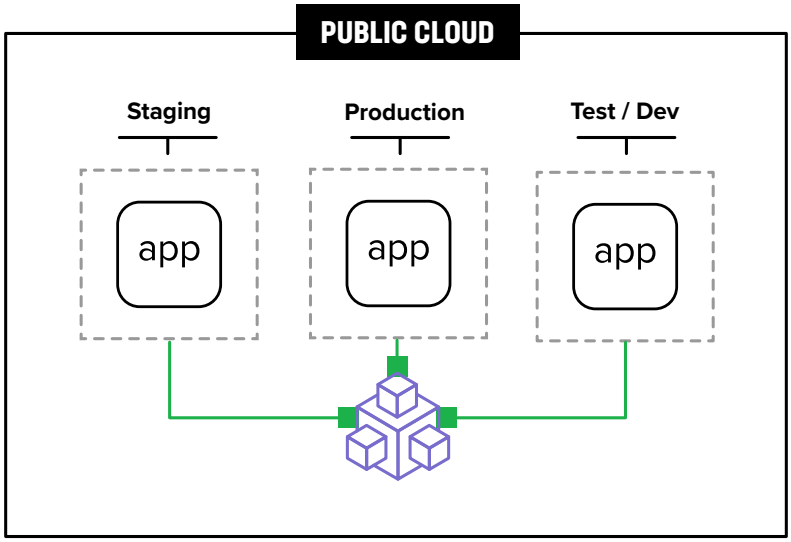

Figure 2: How IP address overlap is solved through app-layer networking

Conclusion

As networks grow, IT will inevitably face problems created by old workarounds that were previously considered best practices. F5 has helped IT and Operations teams deliver applications throughout the lifecycle of the modern Internet—and with Distributed Cloud Mesh it's ready to help with the ongoing process of digital transformation without creating additional problems for the future. F5 can help you with application delivery at scale.

IP address overlap is available as part of F5 Distributed Cloud Mesh. Learn more and test it yourself by signing up for a free trial.

Interested in talking to an F5 Distributed Cloud sales specialist? Contact sales@f5.com today.