Data Sheet

F5 Distributed Cloud Customer Edge (CE) Deployable Software

Key Benefits

Reduce Network Complexity
Simplify network connectivity across clouds, on-premises data centers, and edge locations, without needing to manually provision networking services like VPN, NAT, or load balancers.

Improve Hybrid and Multicloud Security
Provide consistent network and application security by applying security policies that are cloud agnostic and can be applied anywhere. Advanced application layer protection for containers and API endpoints safeguard against a broad spectrum of attacks and threats.

Simplify Service Cluster Connectivity
Seamlessly connect app service clusters across disparate Kubernetes distributions in public cloud and on-premises environments.

F5 Distributed Cloud Services at the Customer Edge

F5© Distributed Cloud Services© provide a full-stack (L3 through L7) SaaS-based platform to connect and secure networks and applications in a cloud-agnostic manner. It strips away the complexity of networking and security enforcement across cloud providers, on-premises data centers, and edge sites. F5 Distributed Cloud Services simplify operations for SecOps, DevOps, and NetOps teams, and streamline network segmentation, security and network policy deployment and enforcement, while enabling visibility across distributed environments.

Customers can consume services from the platform in a variety of ways:

1. **Traditional SaaS delivery.** Services are delivered via the console across the F5 Global Network. The F5 Global Network is comprised of multiple points-of-presence (PoPs) or regional edge (RE) sites that are distributed around the world connected by a private backbone that is managed by F5.
2. **Local delivery.** In instances where applications are subject to more stringent security requirements that limit the ability to advertise them to the public Internet, services can be delivered from a locally deployed Customer Edge (CE) software package. A CE is a self-contained software stack that enables network and app security capabilities, with built-in software lifecycle management capabilities.
3. **Hybrid SaaS delivery.** Enabling a combination of the above delivery mechanisms enable the platform to service the broad spectrum of internal and public-facing apps using a single console to connect and enforce security compliance.

A CE allows customers to deploy a lightweight software package locally, while enabling connectivity to the regional edge sites for service visibility and performance from the console.

Features

Network Auto-discovery
Discover existing networks using BGP and integrate with public cloud APIs to discover subnets and routes on cloud VPCs.

Cloud Orchestration
Automate configuration of routes and gateways to abstract away the complexities of different clouds and provide a consistent experience across any public cloud.

Automated Encrypted Transport
Connect multiple CEs over the public Internet or via the F5 Global Network using IPSec tunnels and native TLS encryption across workloads.

Distributed Application and API Security
Isolate traffic from different VLANs and VPCs into multiple segments, acting as a distributed DMZ and enabling local security for apps.

Centralized Console-based Management
Centralized management from the Distributed Cloud Console, enabling a single portal where users can create, update, scale, and delete CE instances as needed.

Application Hosting
Distributed Cloud App Stack is integrated within each CE to provide a consistent managed Kubernetes platform for hosting applications on the cloud, on-premises, or at the edge.

What is a CE?

A CE is a lightweight software package that can be deployed as a virtual machine (VM) or as a standalone containerized service in any environment. It orchestrates the local control plane and data plane components to route, encrypt, and secure traffic. A CE operates as a highly available edge gateway that can be deployed on any site in your network, and extends the network to that site, without the need to establish physical network connectivity. Additional functionalities include:

- Provides access to the regional edge (RE) sites on the F5 Global Network
- Delivers Distributed Cloud Services locally. Multiple CEs can be deployed to create a service mesh-like fabric which facilitates inter app communications independently of the underlying L3 network connectivity between sites
- Enables consistent enforcement and management of security policies between hybrid and multicloud environments.
- Provides a managed Kubernetes platform which allows hosting of apps as containers or VMs.
- Service mesh-like functionality of load balancing and security features for individual sites.

CE Deployment Requirements

A CE can be deployed as a VM instance on public cloud, as a VM on VMware and KVM, directly on a bare metal server, and in any regular or cloud managed Kubernetes cluster as a Pod.

F5 recommends deploying a minimum of 3 CEs for high availability in production environments. Additional worker nodes can be added to improve L7 performance for any deployment or extra hosting capacity for App Stack deployments.

Minimum resources required per node:

- 4 vCPUs
- 14 GB RAM
- 80GB of disk space

F5 offers CEs in three different sizes, depending on performance requirements:

- Small (4 vCPU 16GB RAM)
- Medium (8 vCPU 32GB RAM)
- Large (16 vCPU 64GB RAM)

Instructions for Deploying CEs

While every environment will have different requirements, there are a few core steps involved for creating Sites in cloud environments, and any on-premises, edge, bare metal, KVM, or VMware Site. A Site is any location where a CE is deployed.

Deploying a CE to public cloud:

1. Navigate to the Site Management section in Network Connect on the Distributed Cloud Console.
2. Create a new cloud site object in the Distributed Cloud Console aligned to the provider of your choice.
3. Enter your account credentials and cloud networking details in the subsequent screens on the Distributed Cloud Console.
4. Initiate deployment through the console. An automated process will deploy the CE and create the related objects needed.

It typically takes 15 minutes to create the Cloud Site, and another 15 minutes to provision connectivity.

Deploying a CE to an on-premises, edge, bare metal, KVM, or VM location:

1. Navigate to the Site Management section in Network Connect on the Distributed Cloud Console.
2. Select Secure Mesh Sites and create a new Secure Mesh Site.
3. Download the latest CE image from the F5 documentation site for your specific environment.
4. Deploy the CE image as a VM (in VMware or KVM environments) or bootstrap the bare metal server to install the CE software.
5. Connect to the CE Console and configure the specific Site and networking details for the newly installed CE
6. Approve the registration request in the Distributed Cloud Console.

This process can take several hours, depending on how long it takes to download each image and install.

F5 has compiled deployment instructions for cloud, on-premises data centers, Kubernetes clusters, and bare metal servers in the Site Management documentation.

Upgrades

While the REs and the Global Controller on the F5 Global Network are upgraded automatically with every new maintenance release from F5, the CEs are not automatically upgraded.

Upgrading CEs remains the responsibility of the customer and can be done in the Distributed Cloud Console. Similar to the upgrade window for the platform, customers have up to 6 months to upgrade a CE to the latest version before it becomes mandatory.

Offline Survivability

CEs can continue to deliver services in locations where there is intermittent connectivity. At edge locations where network connectivity may be intermittent, Offline Survivability allows a CE to continue operating through network disruption in headless mode, by maintaining certificates and allowing management through the local control plane for up to 7 days.

In Offline Survivability mode, there are three components that ensure the CE remains active:

- **Routing:** Routes are exchanged via BGP within a Site, and across Sites in a Mesh Group or a DC Cluster Group. In Offline Survivability mode, the local control plane allows local traffic load balancing for the Site to continue. If two or more Sites in a Mesh Group have Offline Survivability enabled, and the Site Mesh Group is a Full Mesh type group with control plane enabled, load balancing across local and remote endpoints in those Sites continues to function, even when connectivity with the RE is lost. The same is also true for Sites within a DC Cluster Group.
- **Identity Management:** Certificates for services are issued via a Certificate Authority local to the Site when they start or restart without connectivity to the Global Controller. If services restart, they get new certificates and continue functioning.
- **Secret Management:** Secrets that are decrypted from the platform when connectivity to Global Controller was intact are cached locally on the site. This enables services to obtain decrypted secrets even when the connectivity is lost.

When offline In Offline Survivability mode, logs older than 5 minutes, and metrics older than 2 hours, are lost. Learn more about how to enable Offline Survivability.

More Information

Learn more about site management.

©2024 F5, Inc. All rights reserved. F5, and the F5 logo are trademarks of F5, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com.
Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, expressed or implied, claimed by F5, Inc.
DC 04.2024 | DS-GTM-1310417565