Avesha

Avesha

Dec 15, 2022

Application Fabric for Infrastructure control

-Prasad Dorbala

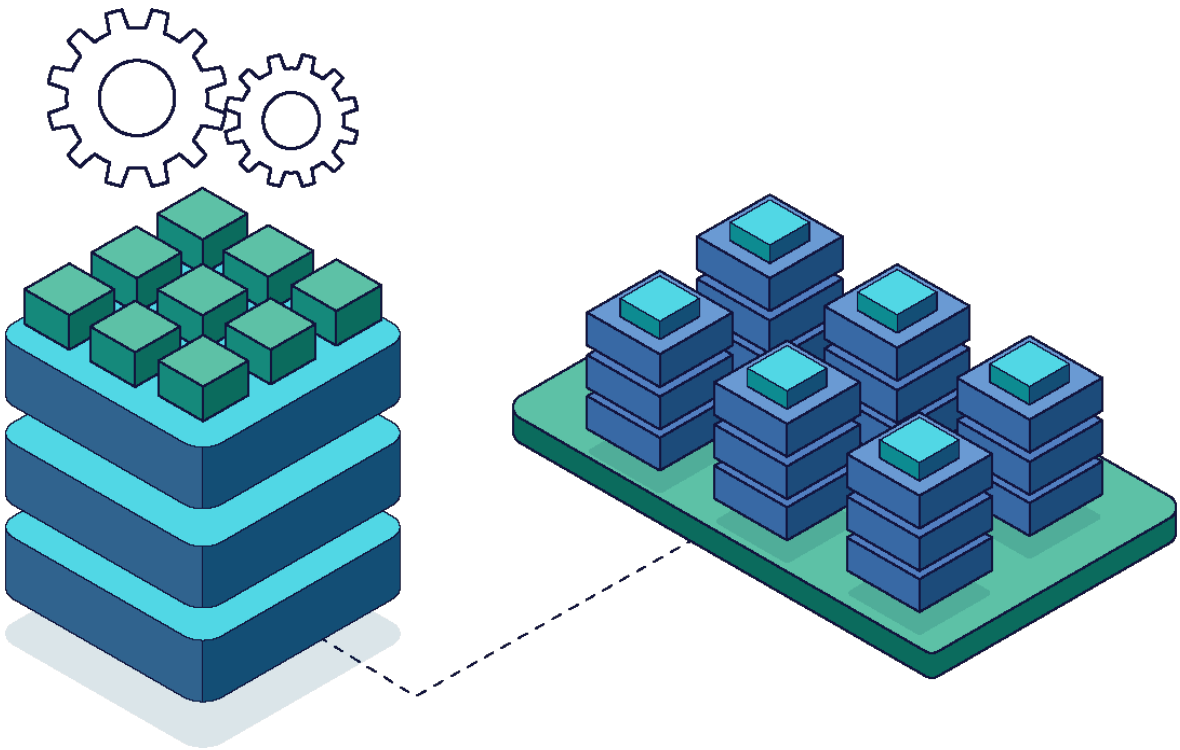

Introduction

A very common problem encountered by application developers is the networking roadblock they encounter with the platform team when they must deploy their applications in multiple clouds, be it for resiliency, regional data compliance, or collaboration. All current solutions require manual intervention ranging from provisioning transit gateways to creating service "exposures" in firewalls to achieve what is normally taken for granted inside a cluster.

A new breed of multi-cloud fabric providers have surfaced aiming to solve the connectivity problem between clouds using a combination of tunnels, cross-connects and meet-me points as needed to offer the best-in-class connectivity. However, a missing piece is the control plane to automate and drive the connectivity to meet an application's data distribution needs including the dynamic topology and the QoS needed for that distribution. Current solutions consider an entire cluster as a unit of inter-networking, which poses enormous complexity of exposing the union of clusters as an attack surface and needs additional scaffolding to secure such application workloads.

In this paper, we explore the role of an application fabric – as a concept that captures the essence of a declarative distributed deployment model that enables a high-velocity infrastructure, which is essential for business applications for modern applications.

Application Fabric Principles

Avesha has set as its goal the single purpose of automating the infrastructure in a cloud-native declarative manner as an extension of existing Kubernetes cluster services of trust domains and service availability but across multiple clusters regardless of the underlying Kubernetes version, kernel, and other variations that are commonly encountered in real world deployments. Securing such trust domains with least privilege and constraining to expose only required application domain inside a cluster or group of clusters is foundational to application fabric principle.

A key enabler for Avesha's Service-Aware architecture is the "application slice" defined by kubeSlice – an open source project (github.com/kubeslice) – that encapsulates all the components of a service including its resources (compute, storage and networking), service availability and trust domain over multiple clusters. The application slice provides a single controller to manage the entire estate with comprehensive visibility, security and automation that makes a distributed deployment a self-serve model via a yaml file or GUI.

The application fabric system has a controller which manages life cycle management of application slices across the fleet of clusters be it in a single region, across regions, or across clouds.

Application Slice

One of the core principles of securing fleet of applications is to reduce attack surface. Attack surface and attack vectors reduction is core principle of application slice. Application slice construct will let users define specific application namespaces in distributed clusters to form a slice. Application slice defines Role Based Access Control (RBAC) rules which helps to implement least access privilege principle for zero-trust policy. Application slice ensures right people have right access to right defined resources in a cluster or group of clusters. Application slice implements resource allocation to reduce denial of service attacks and starvation and chatty neighbor problems. Application slice creates an overlay network to form a dynamic application fabric with end-to-end encryption for secure transport across clusters. This unique behavior is one of a kind for increased velocity in deploying application and grow elastically across multiple cloud or cloud to datacenter deployments.

Distributed Topology

Kubernetes offers the ability to elastically grow and shrink pods in a cluster. An application slice extends this idea to a distributed network of clusters where the application is now distributed across multiple clusters. The elasticity of the infrastructure refers to the extent to which the application needs to be distributed. Examples of such application distribution are seen in gaming, metaverse, autonomous vehicle and retail use cases.

Interface to the Multi Cloud Network (MCN) Fabric

Figure 1 below shows the relationship of the Application Fabric to the MCN Fabric. The Application Fabric describes the application's deployment topology and is controlled by the Slice Controller. The Application Fabric conveys to the MCN Fabric below the desired infrastructure state via the "API" of the MCN Fabric. Such an API is in an early stage and is getting realized at the time of this writing.

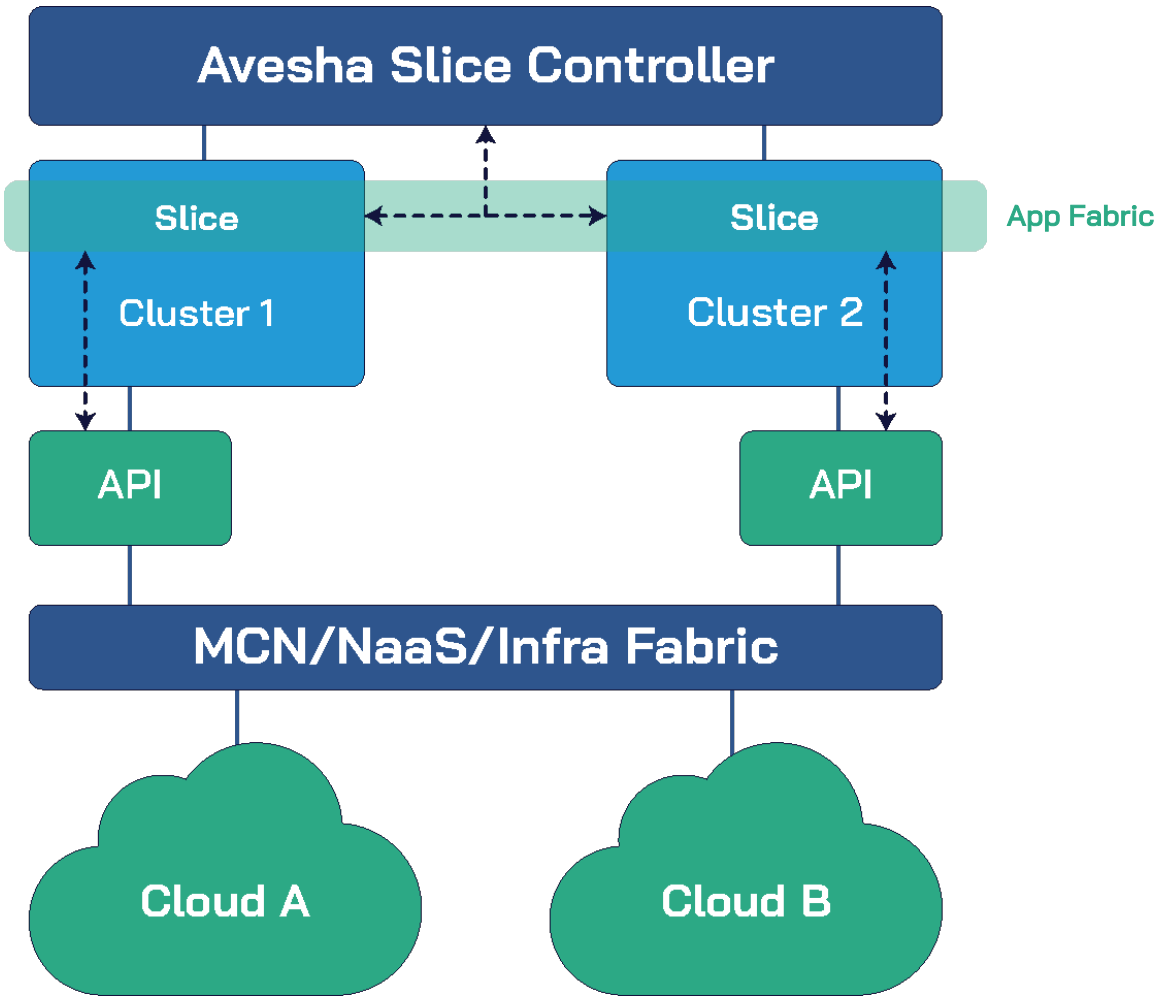

Figure 1: Relationship between Application and MCN Fabrics

An application only needs to specify the QoS type (low latency or min bandwidth) and the locations at which it needs its microservices to be available via declarative metadata in its yaml file. This information is conveyed to the controller that in turn constructs an application slice topology and gives instruction to individual cluster-level application slice components to connect to their neighbor's via the underlying MCN fabric including QoS and bandwidth specifications. In 5G networks, the network slice is an example of an interface to which application slices can map to obtain the desired QoS.

The topology and loading of the application slice can evolve over time based on the needs of the application and its users. It is the role of the application slice to translate these varying needs into concrete instructions to the underlying MCN fabric. This is where AI/RL comes in and helps to solve a complex multi-objective optimization problem.

AI/RL Optimization of the Application Slice

An application slice may experience varying traffic loads at each of its microservices at different locations. Moreover, the pod capacity estimation of these microservices may differ from time to time based on the instance types and any application updates. Both the traffic and the pod capacity are learned and combined together into an RL model that in turn optimizes each microservice at each location in order to provide a global holistic optimization

of the workload distribution across end points of the application slice. The application-level metrics of latency and failure rates are used across the application slice in order to drive the RL engine.

Multi Cloud Control Plane

Application Slice Network Policy

Network segmentation is trusted and proven principles for organization to create barriers that block attacks and reduce attack surface. Network segmentation strategy is used by enterprises to segregate and isolate networks to reduce attack surface. In Kubernetes network policy is a way to define isolation. Application slice uses network policy construct in Kubernetes to implement isolation principle across application slice. This automated framework improves defining consistent policies across elastic application slice, manages for any configuration drifts across application slice footprint with a single source of truth. Policy consistency is one of the key principles for cyber security strategy. Furthermore, it is important to have consistent policy definition for CJIS, HIPAA and other compliance workloads running in Kubernetes environments

Application Slice Security

Security is an aspect of management and needs to be addressed at all layers – starting from physical and lower protocol layers all the way to through the network and ultimately the application layer. It is in the interface between the latter two layers that the application slice creates value by offering a templatized security boundary based on best practices of cloud-native security (including namespaces, resource limits and RBAC policy controls) that is further extended to multiple clusters including their interconnection security using a secure overlay network that ensures end-to-end encryption and trust enablement. This latter point greatly enhances security for multi-cloud networks because only trusted entities can access the overlay network via a common certificate authority. Note that the present alternative for inter-cluster connectivity leaves open many security vulnerabilities from misconfigurations of

firewalls and service exposures via public-facing Internet API gateways that need to be audited for each new inter-cluster service interconnection.
Having a well-architected and templatized approach relieves the security team of getting burdened with a backlog of inter-cluster connectivity requests.

Application Fabric Use Cases

An Application Fabric applies to many important situations ranging from Disaster Recovery (DR), Regional Data Compliance, Low Latency Applications to Interconnectivity Needs from M&A situations. Below, we examine how an Application Fabric can help in such applications.

Disaster Recovery (DR)

Many enterprises have a business continuity aspiration during or after a disaster. However, many have their IT infrastructure primarily situated in a region of a single cloud vendor. When this region is hit by an outage, the enterprise cannot serve its customers unless it has another deployment available that is isolated from the original deployment. Ideally, this other deployment, which we shall call a DR deployment is in a different cloud vendor and in a different geographical area. The Application Fabric can make the task of orchestrating a redundant DR deployment easy by providing a well-defined boundary with the same security envelope and resource limits as well as the continuation of the namespace from the main deployment. Moreover, it is very easy to further extend the resiliency to multiple DR deployments if needed to further enhance the robustness of recovery. Also, the DR deployments may be operated on an active-passive manner to save costs albeit at the expense of the time needed to bring up the deployment at the time of the recovery. Finally, it is possible to have different DR deployments for different applications based on their DR geography.

Regional Data Compliance

Global Enterprises such as Retail and Banking have the challenge of ensuring regional data compliance based on privacy standards such as GDPR and PIPL (China's version). These laws require data to remain resident within the geographical boundaries of the data's provenance. Consequently, it is the workload that needs to move to the data to operate on it. Clearly, an application fabric can easily enable this type of movement of the relevant workload to run within the geography of the data and yet connect back to the remaining workloads via its application slice. In addition, workloads can move to multiple locations to operate on the local data and yet operate within the same application slice with no change to any part of the

application. Access to the workloads running in the remote location is via a geo-centric DNS name and within the policy of the application slice. A cluster can be added or removed from the slice at any time to ensure strict compliance with data compliance audit requirements.

Low Latency

Gaming and Metaverse applications require that the server be located close to the client to overcome any speed of light latency issues. Tournament-style gaming requires consistently low-latency to achieve good game matching outcomes. Likewise, Metaverse applications need low-latency server interactions to keep motion lag from interfering with end user experience. An application fabric can assist in both directing client requests to the closest server and placement of the client workload at the closest and most available server using AI /RL to assist in this task as described earlier under the section on AI/RL optimization.

Mergers and acquisitions

On a Merger or Acquisition, an Enterprise IT department must deal with the need to run a workload in several different environments that it has inherited from the acquired entity, which may have developed its own cloud deployment in an independent manner not necessarily compatible with the acquiring entity. In these situations, an application fabric empowers the Enterprise IT by offering a simple and on-demand cross-deployment manner of running workloads that can talk to each other easily. In addition, all the security, resource, and network policies of this workload can be uniformly enforced at all locations easily.

Application Fabric Use Cases

The key benefits of an Application Fabric can be summarized as follows:

- Single Source of Truth: Despite the location of individual clusters in the fabric, the policies can be controlled from a single controller for the entire application slice. In addition, each application slice can have its own controller. Also, the controllers themselves can have a redundant controller to ensure resiliency. The advantage of a single source of truth is that drift between configurations can be avoided at different locations.

- Application-specific Connectivity: Each application can have its own topology and specific distribution of its workloads based on its requirements. Having a simple declarative way of specifying these requirements enable an application to get its purpose-driven topology whether it is for low-latency, regional data compliance or other reason. The Application Fabric provides a simple and fast way to meet application workload distribution requirements without having to involve IT tickets resulting in 40% gain in development velocity.

- Optimization: Using an Application Fabric can significantly reduce cluster sprawl by using application slices to create multi-tenancy instead of spinning up dedicated clusters for each team or application. The resulting cluster utilization can double from 30% to 60% very easily leading to lower cost and carbon footprint.

Conclusion

In this paper, we describe the role of an Application Fabric to achieve comprehensive control of the deployment of applications in a multi-cloud world. The Application Fabric provides the "control plane" for such application deployments and helps bring benefits of speed, cost and simplicity. There are several applications of Application Fabrics in the areas of Disaster Recovery, Low Latency and Regional Data Compliance in the verticals of Gaming, Retail, Banking among others.