Architecture Requirements for Scalable Zero Trust Microsegmentation

After the installation and configuration tasks of a microsegmentation implementation are complete, there are two major functions executed when running the solution. The first is traffic monitoring, and the other is policy management. These are the two functions which affect the performance of the solution as the number of protected endpoints are increased.

Traffic monitoring involves taking the traffic data sent by the agents on the endpoints and using it to build a map which shows the communications between all the assets in the environment, the users, the endpoints and servers, and even external hosts to which the internal devices are communicating. Traffic can then be categorized as conforming or nonconforming according to the defined polices. In the administrator's graphical user interface, lines representing the connections between the illustrated objects are color-coded to indicate compliance using green, yellow and red. Traffic monitoring is a compute-intensive function. It requires near-real time analysis of all the traffic information sent to the policy engine as it occurs, to normalize flows and build a real-time hi-fidelity application dependency map. This map is used as the basis on which policy is defined. The architecture of the microsegmentation solution must deliver scalable performance in the traffic monitoring function to deliver on the security promise of zero trust microsegmentation.

Policy management is the computation executed in the central microsegmentation engine which processes the inventory of systems, their tags, and the administrator-defined communications policies. It calculates the compliance of the traffic using the tags and defined policies and then sends instructions to the policy enforcement points, the host-based firewalls or microservice proxies for containerized applications. Especially in cloud applications, which are very dynamic, every time a process is instantiated or terminated the policy engine needs to calculate and update rules for that and all surrounding workloads. For a successful microsegmentation solution, the policy management function must scale as the numbers of workloads and endpoints in the landscape increases.

Scalability of a microsegmentation solution becomes an issue as the number of mapped endpoints increases. Larger enterprises with many thousands of endpoints require significant computing power in executing a microsegmentation strategy, because of the two functions described above. How much processing power? For an on-premise, server-based microsegmentation solution, anecdotal evidence tells us that 96 GB of RAM, many high-end Intel processors, and terabytes of hard drive space is a typical server specification. Multiple clusters of such servers are needed to support large numbers of protected endpoints.

The virtues of cloud computing are well known and well accepted now. While in the past some organizations hesitated to adopt cloud computing for critical applications, at this point the deployment of mission-critical enterprise applications using one of the major hyperscaler infrastructures has become the norm. Agility of configuration, elasticity of processing, and overall cost savings in lieu of maintaining an on-premise data center to support applications are undeniable benefits of using cloud infrastructure. Software-as-a-Service brings the benefits of constant up-to-the-minute updates of software without disruption to user's applications, and continuous innovation/continuous delivery of new functionality. This type of agility is especially important in the cyber defense realm where the adversary is smart and very agile; the cyber defense strategy must be agile as well.

Using a cloud SaaS infrastructure to deliver compute-intensive microsegmentation functions is a natural technical fit. Currently, many organizations prefer this approach because of the benefits discussed above. But this raises a question; is your microsegmentation technology vendor primarily investing their development resources in innovation and upgrades for their on-premise solution, or their cloud-native SaaS version? Many vendors offer both types of deployment of their solution, but only their on-premise versions are truly scalable. This is because the development focus of their software product was initially for on-premise deployment, often driven by early-phase customers who several years ago had security concerns about cloud deployment in general. That on-premise architecture, when ported to the cloud, may be less scalable and delivered with fewer features. If organizations choose to leverage Software-as-a-Service deployment of zero trust microsegmentation, for all the advantages discussed above, leaders should ascertain that the fundamental architecture of the cloud solution is scalable to accommodate the increasing numbers of endpoints for future needs. Factors that could influence scalability include:

- Is the SaaS application architected with separate microservices that can be elastically scaled up when demand arises, or is the architecture actually a monolithic application ported to the cloud?
- Does the architecture support direct downloading of agents from the cloud, or is an on-premise management server required?
- Likewise, is an on-premise policy enforcement server required?

The challenge for on-premise deployment of microsegmentation is that costs rise at a multiple of the increasing number of endpoints, because of the required FTE staffing and the server infrastructure required. The need for an on-premise deployment for scalability also means that the innovation and continuous delivery of new features and upgrades is less agile, compared to a truly scalable SaaS solution. These challenges may not arise during an initial, smaller scale on-premise implementation of microsegmentation. However, when planning a zero trust microsegmentation initiative, leaders should consider not just the current number of endpoints to be protected, but the potential for increasing numbers of endpoints in their enterprise through organic growth, reorganization or merger and acquisition.