illumio

RFP TEMPLATE FOR SELECTING A MICRO-SEGMENTATION SOLUTION

This Request for Proposal template is comprised of common requirements and questions from buyers when selecting a solution for micro-segmentation. Use it to identify your solution criteria, clarify unknowns, and prioritize potential trade-offs.

SECTION 1: COMPUTING ENVIRONMENT

1.01 Does the solution support workloads deployed on bare-metal (physical servers)? Are there any limitations?

1.02 Which OS platforms are supported (Windows, Linux, AIX, Solaris, etc.)? Are there any limitations?

1.03 Does the solution support workloads deployed as virtual machines on any virtualization platform or hypervisor? Which hypervisors are supported? Do you have feature parity among hypervisor support? Please explain any feature parity if it exists.

1.04 What public clouds are supported (AWS, EC2, S3, Azure, GCP)? Are there any limitations with each provider that the solution supports? Do you have feature parity among various providers?

1.05 What private cloud stacks are supported (IBM SoftLayer, Nutanix, etc.)? Are there any limitations?

1.06 Does the solution support workloads deployed as containers? Which container platforms are supported? Describe how the solution supports, visualizes, and secures containers.

1.07 Does the solution work on any network infrastructure without dependencies on hardware, such as VMware servers or Cisco switches, or requiring changes to the network?

1.08 Is the solution able to recognize and secure homegrown and legacy applications? If yes, how does it do this?

SECTION 2: VISIBILITY

2.01 Does the solution provide an operations-level "big picture" view of locations with interactive drill-down capabilities for segmentation administration?

2.02 Does the solution provide an application-centric view of the connectivity and relationships of applications and their workload components? Note: This is an application dependency map, not a network map.

2.03 Does the solution successfully show process-level visibility of workloads and their associated traffic across various size environments – from 100s to 100,000+ workloads?

2.04 Does the solution provide live visibility with updates on communication, workload information, and application traffic? How is it delivered?

2.05 Does the visibility provide information gathered directly from the workload (e.g., network interfaces, hostname, etc.)?

2.06 Does the visibility provide port information for each network flow?

2.07 Does the solution use visibility information to provide a way to develop and monitor micro-segmentation policy that is clear and easy for non-security teams to understand at a glance?

2.08 Describe how users could leverage the visibility information to design and approve the micro-perimeter design.

2.09 Describe how users create micro-segmentation policies.

2.10 Describe how users can leverage visibility to allow application owners to validate and agree with security and compliance teams on the workloads/applications that are in scope for the micro-segmentation project.

2.11 Does the solution use visibility to meet regulatory compliance?

SECTION 3: SEGMENTATION RULES AND POLICIES

3.01 How is a security policy defined in the solution? Describe how segmentation policies are created and maintained.

3.02 Is cloning/modifying/deleting an existing security monitoring policy supported?

3.03 Does the solution support import/export of policies?

3.04 Does the solution support the assignment of a firewall policy to and the removal of a firewall policy from a specific workload group?

3.05 Does the solution support removal of all firewall policies when the agent is removed from a workload?

3.06 Does the solution allow users to define multidimensional policy for groups of workloads based on multiple metadata attributes (e.g., application, environment, location, role, etc.)?

3.07 Does the solution allow segmenting based on port?

3.08 Does the solution support dynamic ports?

3.09 Does the solution provide process-aware policy enforcement for Windows workloads? Note: Windows processes use dynamic ports, requiring network-based firewalls to open a large range of ports to ensure legitimate application traffic is not potentially blocked. Process-aware policy enforcement closes this gap by allowing security policy to automatically adjust to required ports as they are dynamically opened and closed.

3.10 Does the solution allow segmenting based on vulnerability information? Can your product do this without disruption to functions or breaking the application?

3.11 Does the solution allow segmenting based on user and group membership information for both VDI desktops and Microsoft Active Directory?

3.12 Can the solution integrate into the software development life cycle, embedding policy to ensure a workload remains secure across the different stages of its life cycle?

3.13 Does the solution provide a uniform policy model across private, public, and hybrid clouds?

3.14 Does the solution provide templates as an option to easily segment common commercial off-the-shelf applications (e.g., Microsoft Active Directory and Exchange)?

3.15 Does the solution dynamically adapt to changes in the environment, including IP address changes, application scale-up/down, VM load balancing, workload migrations across data centers or public cloud, user mobility, and disaster recovery scenarios? Please describe.

3.16 Does the solution auto-recommend security policies for all flows (including intra-app and inter-app flows) based on discovered application communications?

3.17 Does the solution have mechanisms to build, visualize, and evaluate policy impact before enforcing rules to ensure applications do not break?

3.18 Does the solution provide the ability to tune security policy to identify and reduce the risk of vulnerabilities without breaking applications?

3.19 Does the solution alert on potentially blocked traffic while in testing mode and alert on blocked traffic when in enforcement?

3.20 Does the solution provide visual feedback that the environment is operating under the defined policy?

3.21 Is the solution aligned with a whitelist Zero Trust security model?

SECTION 4: VULNERABILITY MANAGEMENT

4.01 Does the solution integrate third-party vulnerability scanners including Qualys and Rapid7?

4.02 Does your solution incorporate vulnerability data into visualizations to help with the focus and prioritization of security policy?

4.03 Does the solution provide insights into vulnerabilities beyond a single workload?

4.04 How does the solution measure the risk of a vulnerability?

4.05 Does the solution provide numerical accounting of active vulnerabilities and counteracting measures to promote quantitative risk mitigation and reporting?

4.06 Does the solution provide a compensating control to reduce risk when patching is not an option?

4.07 Can the solution recommend optimal policies to minimize risk by constraining or blocking vulnerable ports based on visibility mechanisms?

SECTION 5: ENFORCEMENT

5.01 Describe how the solution provides enforcement of policy. Where is the point of enforcement in the solution? If there are multiple points of enforcement, please explain why and how.

5.02 Describe how the solution identifies policy violations.

5.03 Does the solution support automatic firewall policy enforcement on a workload when the agent first starts?

5.04 Are enforcement points stateful or stateless?

5.05 Does the solution provide enforcement using native OS capabilities or it is out of band from the kernel?

5.06 Does the solution modify the kernel in any way?

5.07 Does the solution instrument enforcement of policies within other networking devices (e.g., load balancers and data center switches)?

5.08 How does the solution support operating environments like IBM AIX and Oracle Solaris for enforcement?

5.09 Does the solution enforce both inbound and outbound segmentation policies at either end of a connection path, both the destination and the source?

5.10 Does the solution have a quarantining mechanism for workloads that violate policies?

5.12 Provide scalability limits in terms of workloads to protect.

5.13 Can the solution enforce encryption of data in motion?

5.14 Does the solution support end-to-end encryption?

5.15 Can the solution enforce authentication of machine identity prior to establishing connection?

SECTION 6: DEPLOYMENT AND ADMINISTRATION

6.01 Can the solution be run on-premise with bare-metal or virtual machines, in public cloud, or hosted cloud?

6.02 Do the solution support SaaS deployment model?

6.03 What are the pricing options?

6.04 Is the solution strictly software based (i.e., it does not require any hardware)?

6.05 Does the solution have any dependency on any network infrastructure? Do you require any changes to the network, either physical or virtual, to support the solution?

6.06 If the solution includes an agent, describe installation of the agents across enterprise scale of workloads. Does it have the capability to do mass installation?

6.07 If the solution includes an agent, what is the upgrade and uninstall process?

6.08 Can the agent be installed in "monitor only" mode if desired?

6.09 Does the solution support role-based access control (RBAC) for different administrative roles, allowing application teams to write segmentation rules for their applications that require approval before being provisioned? If so, please explain the granularity levels.

6.10 Does the solution provide a mechanism to separate policy authoring from policy provisioning?

6.11 Does the solution support SAML for authenticating users with an Identity Provider?

6.12 Does the solution provide an audit trail of all configuration changes?

6.13 Does the solution provide a record of all traffic flows between workloads, with export functionality?

6.14 Describe the experience for implementing micro-segmentation in large and diverse production enterprise environments.

6.15 Have you implemented the solution with a customer of similar size and in the same industry? Please provide reference information.

5.16 What types and levels of training do you provide or recommend? If applicable, describe training materials offered.

6.17 Do you offer global worldwide professional services to expedite deployment?

SECTION 7: API AND INTEGRATIONS

7.01 Does the solution work without requiring custom vendor software in the kernel space of the workloads?

7.02 Does the solution have an API? Describe the API, documentation, and how it can be used.

7.03 Can the solution integrate with DevOps tools for IT orchestration and provisioning such as Chef, Puppet, Ansible, etc.?

7.04 Does the solution integrate with SIEMs such as Splunk, QRadar, and ArcSight? Can it export traffic logs and audit events in vendor-specific native formats to SIEM tools and parse the logs for security events?

7.05 Does the solution integrate with IT Ops tools such as ServiceNow or other CMDBs?

SECTION 8: PERFORMANCE, SCALABILITY, AND AVAILABILITY

8.01 Describe how global visibility is enabled. Does the solution scale to 250,000+ workloads?

8.02 Describe what resources are required to enable global visibility.

8.03 Does the solution support a large number of rules on the workloads without significant performance implications?

8.04 Describe how global security policies are enabled. Does the solution dynamically adapt to changes in the environment regardless of scale?

8.05 Describe the high availability and resiliency characteristics of your solution.

8.06 Does the solution have a central portal available to manage multiple data center deployments from one place?

8.07 Describe the required overhead on a network to support your product.

8.08 If the solution requires an agent, is communication between agent and management controller encrypted?

8.09 What happens if communication between the agent and management controller is lost?

8.10 Does the solution respond to connections initiated by non-agents?

8.11 If the solution requires an agent, describe how it handles tampering.

8.12 Is the solution purpose-built for micro-segmentation? If not, how can it be optimized for micro-segmentation use cases?

SECTION 9: COMPLIANCE

9.01 How can a user leverage the platform to determine and validate what is in scope for compliance measures (e.g., PCI compliance)?

9.02 How does the solution support compliance audits?

9.03 Does the solution support export of compliance records?

9.04 Does the solution allow querying of compliance records based on a single workload or a group of workloads?

ABOUT ILLUMIO

Illumio, the leader in micro-segmentation, prevents the spread of breaches inside data center and cloud environments. Enterprises such as Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite use Illumio to reduce cyber risk and achieve regulatory compliance. The Illumio Adaptive Security Platform® uniquely protects critical information with real-time application dependency and vulnerability mapping coupled with micro-segmentation that works across any data center, public cloud, or hybrid cloud deployment on bare-metal, virtual machines, and containers. For more information about Illumio, visit www.illumio.com/what-we-do or follow @Illumio.

- Engage with Illumio on Twitter
- Like Illumio on Facebook
- Follow Illumio on LinkedIn
- Subscribe to the Illumio YouTube Channel

CONTACT US

For more information about Illumio ASP and how it can be used to achieve visibility behind the firewall, email us at illuminate@illumio.com or call 855-426-3983 to speak to an Illumio representative.

Illumio, Inc. 920 De Guigne Drive, Sunnyvale, CA 94085 Tel (669) 800-5000 www.illumio.com