algosec
BUSINESS-DRIVEN
SECURITY MANAGEMENT

Micro-segmentation From Strategy to Execution

EBOOK

Table of Contents

What is Micro-segmentation

Micro-segmentation is a technique to create secure zones in networks. It lets companies isolate workloads from one another and introduce tight controls over internal access to sensitive data. This makes network security more granular.

Micro-segmentation is an "upgrade" to network segmentation.

Companies have long relied on firewalls, VLANs, and access control lists (ACL) to segment their network. Network segmentation is a key defense-in-depth strategy, segregating and protecting company data and limiting attackers' lateral movements.

Consider an intruder who enters a gated community. This does not mean the intruder should have free reign into all of the houses in the community because, in addition to the outside gate, each house has locks on its door. Micro-segmentation takes this a step further – even if the intruder breaks into a house, they cannot access all the rooms.

Why Micro-segment?

Organizations frequently implement micro-segmentation to block lateral movement. Two common types of lateral movements are insider threats and ransomware.

Insider threats are employees or contractors gaining access to data that they are not authorized to access.

Ransomware is a type of malware attack in which the attacker locks and encrypts the victim's data and then demands a payment to unlock and decrypt the data. If an attacker can take over one desktop or one server in your estate and deploy malware, you want to make sure that the malware can't spread throughout the entire data center in order to reduce the "blast radius."

The cost of downtime is
23x greater
than the average ransom requested in 2019.

Downtime costs due to ransomware are up by
200%
year-over-year.

(Source: Datto's Global State of the Channel Ransomware Report)

The SDN Solution

With software-defined networks, such as Cisco ACI and VMware NSX, micro-segmentation can be achieved without deploying additional controls such as firewalls. Because the data center is software-driven, the fabric has built-in filtering capabilities. This means that you can introduce policy rules without adding new hardware.

SDN solutions can filter flows both inside the data center (east-west traffic) and flows entering or exiting the data center (north-south traffic).

The SDN technology supporting your data center eliminates many of the earlier barriers to micro-segmentation.

Yet, while a software-defined fabric makes segmentation possible, there are still many challenges to make it a reality.

What is a Good Filtering Policy

A good filtering policy has three requirements:

01

Allows all business traffic

The last thing you want is to write a micro-segmented policy and have it break necessary business communication, causing applications to stop functioning.

02

Allows nothing else

By default, all other traffic should be denied.

03

Future-proof

"More of the same" changes in the network environment shouldn't break rules. If you write your policies too narrowly, when something in the network changes, such as a new server or application, something will stop working. Write with scalability in mind.

How do organizations achieve these requirements? They need to know what the traffic flows are – what should be allowed and denied.

This is difficult because most traffic is undocumented. There is no clear record of the applications in the data center and what network flows they depend on. To get accurate information, you need to perform a "discovery" process.

A Blueprint for Creating a Micro-segmentation Policy

To create your micro-segmentation policy, you need to:

Implement the discovery process

Define logical segments

Create the filtering policy

Avoid connectivity disruptions during implementation

Enforce micro segmentation: "D-Day"

Continually maintain your network

Manage change requests

Keep reading to find out how.

Micro-segmentation Blueprint
Discovery

You need to find out what traffic needs to be allowed and then you can decide what not to allow. Two common ways to implement a discovery process are traffic-based and content-based.

Traffic-Based Discovery

Traffic-based discovery is the process of understanding traffic flows: Observe the traffic that is traversing the data center, analyze it, and identify the intent of the flows by mapping them to the applications they support.

You can collect the raw traffic with a traffic sniffer/network tap or use a NetFlow feed.

Content-based or Data-Based Approach

In the content-based approach, you organize the data center systems into segments based on the sensitivity of the data they process. For example, an eCommerce application may process credit card information which is regulated by the PCI DSS standard. Therefore, you need to identify the servers supporting the eCommerce application and separate them in your filtering policy.
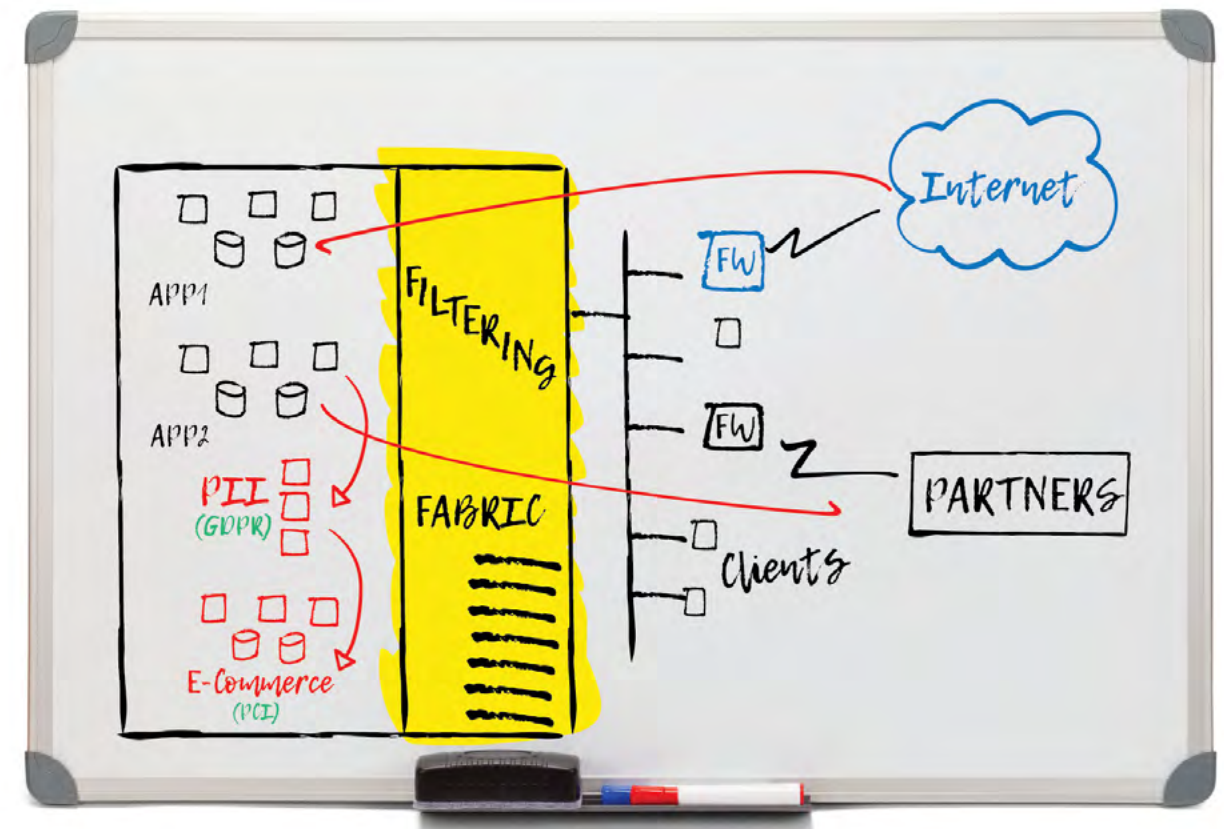
Discovering traffic flows within a data center

01
What is Micro-
segmentation

02
Why Micro-segment

03
The SDN
Solution

04
What is a Good
Filtering Policy

05
A Blueprint
for Creating a
Micro-segmentation
Policy

06
Enabling
Micro-segmentation
with AlgoSec

07
About
AlgoSec

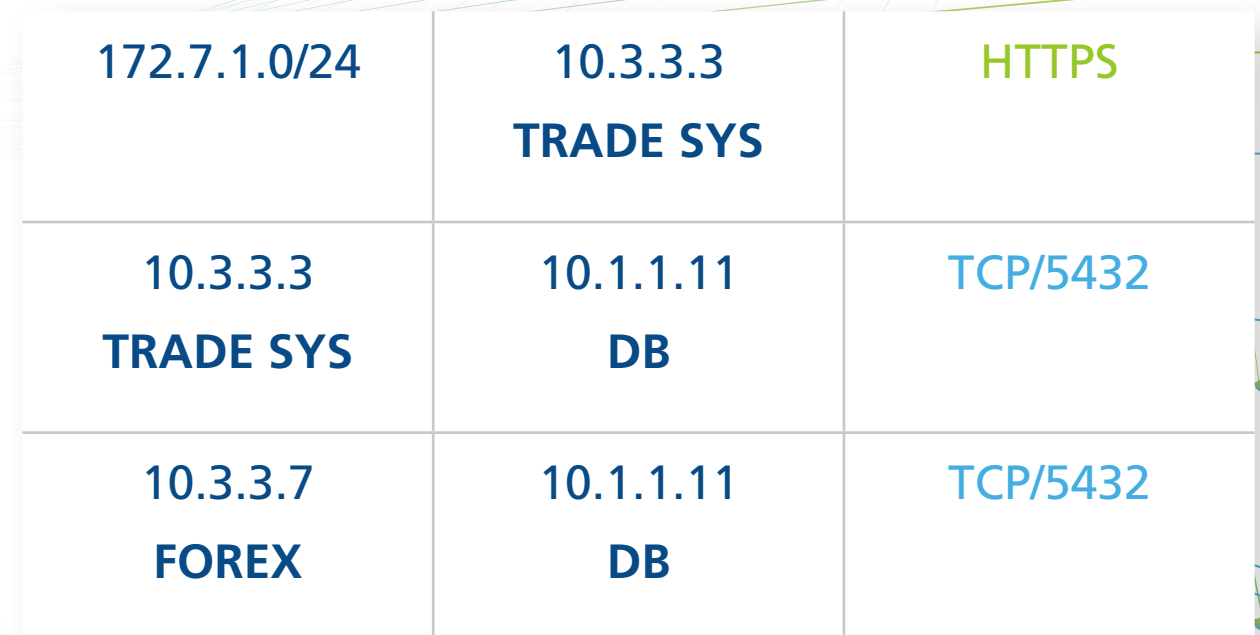# Micro-segmentation Blueprint
Using Netflow for Traffic Mapping

The easiest traffic source to base application discovery on is NetFlow. Most routers and switches can be configured to emit a NetFlow feed without needing to deploy agents throughout the data center.

The flows in the NetFlow feed are clustered into business applications based on recurring IP addresses and correlations in time (e.g., if an HTTPS connection from a client at 172.7.1.11 to 10.3.3.3 is observed at 10 AM, and a PostgreSQL connection from the same 10.3.3.3 to 10.1.1.1 is observed 0.5 seconds later, it's clear that all three systems support a single application, which can be labeled with a name such as "Trading System").

NetFlow often produces thousands of "thin flow" records (one IP to another IP), even for a single application. In the example above, there may be a NetFlow record for every client desktop. It is important to aggregate them into "fat flows" (e.g., that allow all the clients in the 172.7.1.0/24 range). Besides avoiding an explosion in the number of flows, aggregation also allows a higher-level understanding – and future-proofs the policies against fluctuations in IP address allocation.

Using the discovery platform in the AlgoSec Security Management Suite to identify the flows in combination with information from your firewalls can help you decide where to put the boundaries of your segments and which policies to put in these filters.

172.7.1.0/24	10.3.3.3 **TRADE SYS**	HTTPS
10.3.3.3 **TRADE SYS**	10.1.1.11 **DB**	TCP/5432
10.3.3.7 **FOREX**	10.1.1.11 **DB**	TCP/5432

Identifying traffic flows in common based on shared IP addresses

Micro-segmentation Blueprint
# Defining Logical Segments

Once you have discovered the business applications whose traffic is traversing the data center (using traffic-based discovery) and have also identified the data sensitivity (using a content-based approach) you are well positioned to define your segments.

Bear in mind that all the traffic that is confined to a segment is allowed. Traffic crossing between segments is blocked by default – and needs to be explicitly allowed by a policy rule.

There are two potential starting points:

1. Segregate the systems processing sensitive data into their own segments: you may have to do this anyway for regulatory reasons.

2. Segregate networks connecting to client systems (desktops, laptops, wireless networks) into "human-zone" segments: Client systems are often the entry points of malware, and are always the source of malicious insider's attacks.

Then, place the remaining servers supporting each application into a separate segment. Doing so will save you the need to write explicit policy rules to allow traffic that is internal to only one business application.
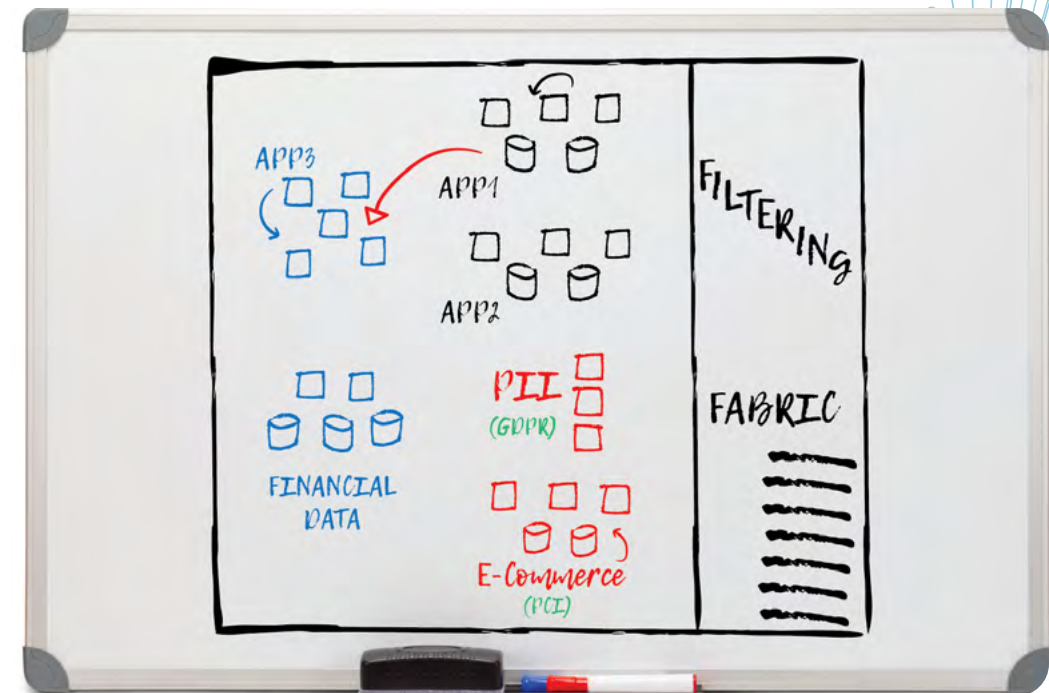
Example segment within a data center
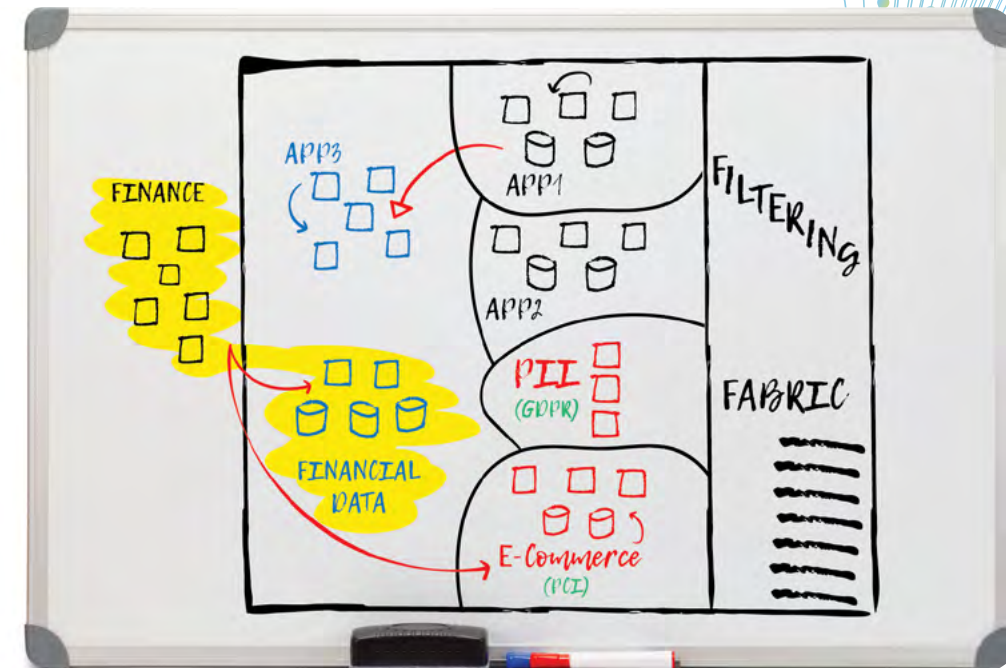
Micro-segmentation Blueprint
Creating the Filtering Policy

Once the segments are defined, we need to write the policy. Traffic confined to a segment is automatically allowed so we don't need to worry about it any more. We just need to write policy for traffic crossing micro-segment boundaries.

Eventually, the last rule on the policy must be a default-deny: "from anywhere to anywhere, with any service – DENY." However, enforcing such a rule in the early days of the micro-segmentation project, before all the rest of the policy is written, risks breaking many applications' communications. So start with a (totally insecure) default-allow rule until your policy is ready, and then switch to a default-deny on "D-Day" ("deny-day"): we'll discuss D-Day shortly.

What types of rules are we going to be writing?

- Cross segment flows - Allowing traffic between segments: e.g., Allow the eCommerce servers to access the credit-card data.

- Flows to/from outside the data center – e.g., allow employees in the finance department to connect to financial data within the data center from their machines in the human-zone, or allow access from the Internet to the front-end eCommerce web servers.

User outside the data center need to access data within the data center

Micro-segmentation Blueprint
Default Allow – with Logging

To avoid major connectivity disruptions, start your micro-segmentation project gently. Instead of writing a "DENY" rule at the end of the policy, write an "ALLOW" rule – which is clearly insecure – but turn on logging for this ALLOW rule. This creates a log for any connection that is matched by the default-allow rule. Initially you will receive many logs from the default-allow rule: your goal in the project is to eliminate them.

To do this, you go over the applications you discovered earlier, write the policy rules that support each application's cross-segment flows, and place them above the default-allow rule. This means that the traffic of each application you handle will no longer match the default-allow (it will match the new rules you wrote) – and the amount of default-allow logs will decrease.

Keep adding rules, application by application, until the final allow rule is not generating any more logs. At that point, you reach the final milestone in the project: D-Day.

Micro-segmentation Blueprint
Preparing for "D-Day"

Once logging from the default-allow rule ceases to indicate new flows that need to be added to your filtering policy, you can start preparing for "D-Day." This is the day that you flip the switch and change the final rule from "default ALLOW" to "default DENY." Once you do that, all the undiscovered traffic is going to be denied by the filtering fabric, and you will finally have a secured, micro-segmented, data center. This is a big deal!

However, you should realize that D-Day is going to cause a big organizational change. From this day forward, every application developer whose application requires new traffic to cross the data center will need to ask for permission to allow this traffic: they will need follow a process, open a change request, and wait for it to be implemented. The free-wheeling days are over.

You need to prepare for D-Day. Consider steps like:

- Get management buy-in

- Communicate the change across the organization

- Set a change control window

- Have "all hands on deck" on D-Day to quickly correct anything that may have been missed and causes applications to break

Micro-segmentation Blueprint
Change Requests & Compliance

Notice that after D-Day, any change in application connectivity requires filing a "change request". When the information security team is evaluating a change request – they need to check whether the request is in line with the "acceptable traffic" policy.

A common way to organize the high-level policy is to use a table, where each row represents a segment, and every column represents a segment. The content of each cell in the table lists all the services that are allowed from the "row" segment to the "column" segment.

Keeping this table in a machine readable format, such an excel spreadsheet, allows software systems to run a what-if risk-check, which compares each change-request with the acceptable policy, and flag any discrepancies before the new rules are deployed.

Such a what-if risk-check is also important for regulatory compliance: regulations such as PCI and ISO27001 require organizations to define such a policy, and to compare themselves to it: demonstrating the policy is often part of the certification or audit.

AlgoSec Enables
Micro-segmentation

The AlgoSec Security Management Suite (ASMS) makes it easy to define and enforce your micro-segmentation strategy inside the data center, ensure that it does not block critical business services and meet compliance requirements.

AlgoSec's powerful AutoDiscovery capabilities help you understand the network flows in your organization. You can automatically connect the recognized traffic flows to the business applications that use them. Once the segments are established, AlgoSec seamlessly manages the network security policy across your entire hybrid network estate. AlgoSec proactively checks every proposed firewall rule change request against the segmentation strategy to ensure that the change doesn't break the segmentation strategy, introduce risk, or violate compliance requirements.

AlgoSec enforces micro-segmentation by:

- Generating a custom report on compliance with the micro-segmentation policy

- Identifying unprotected network flows that do not cross any firewall and are not filtered for an application

- Automatically identifying changes that will violate the micro-segmentation strategy

- Automatically implementing network security changes

- Automatically validating changes
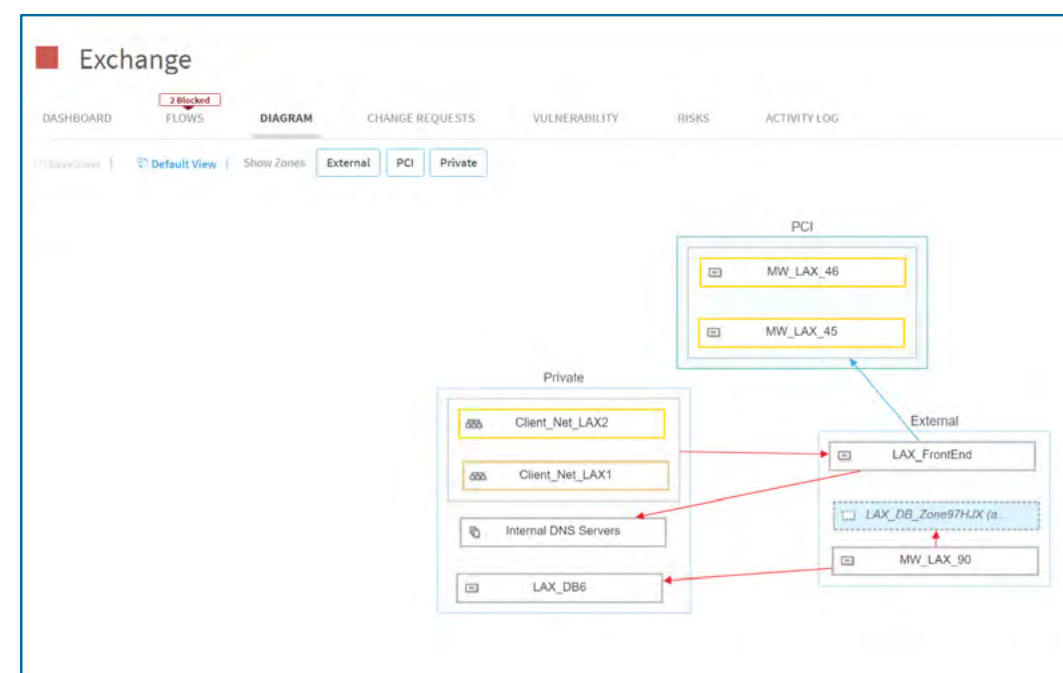
Security zones in an AlgoSec AppViz

Want to learn more?

Get a personal demo

About AlgoSec

AlgoSec enables the world's largest organizations to align business and security strategies, and manage their network security based on what matters most - the applications that power their businesses.

Through a single pane of glass, the AlgoSec Security Management Solution provides holistic, business-level visibility across the entire network security infrastructure, including business applications and their connectivity flows - in the cloud and across SDN and on-premise networks. With AlgoSec users can auto-discover and migrate application connectivity, proactively analyze risk from the business perspective, tie cyber-attacks to business processes and intelligently automate time-consuming security changes - all with zero-touch, and seamlessly orchestrated across any heterogeneous environment.

Over 1,800 leading organizations, including 20 of the Fortune 50, have relied on AlgoSec to drive business agility, security and compliance. AlgoSec has provided the industry's only money-back guarantee since 2005.

Want to learn more about how AlgoSec can help enable micro-segmentation?

Schedule a demo.

● ● ● ● ● ●

For more information, visit www.AlgoSec.com.