VCF TechCon

Powered by VMUG

Start *vandaag* nog met microsegmentatie!

Robert Cranendonk

VMUGNL
VMware User Group

Start with micro segmentation *today*!

Robert Cranendonk

VCF TechCon 2025

VCF TechCon
Powered by VMUG

Robert Cranendonk MSc.

- 🐣 Since 1990
- 🏢 From 2015 working with NSX
- 🎓 2024 MSc. Cyber Security Engineering
- 😎 IT Consultant @ itq
- 🎯 Broadcom Knight, vExpert
- 🌐 significant-bit.com

Terminology

- Zero Trust
  - *"All devices and network flows are not trusted by default"*
- Micro Segmentation
  - *"Logically divide the data center into distinct security segments down to the individual workload level"*
- North-South traffic
  - *"Traffic going in and out of the NSX fabric"*
  - *"Physical to virtual"*
- East-West traffic
  - *"Traffic within the NSX fabric"*

Misconceptions

- Microseg = Microseg?
  - Application Centric Security
- Need to know all flows
- All or nothing
- Big bang, big problems!
- Physical workloads are a problem

Application Ring Fencing

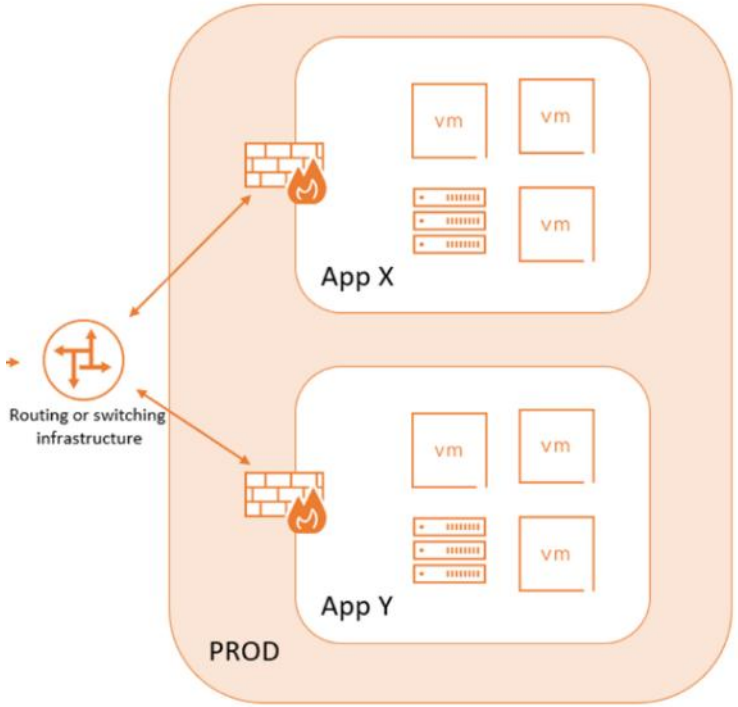

Filtering only on Application group as a whole

Rules apply to ALL VMs in app

Intra-app traffic default allowed

Micro Segmentation

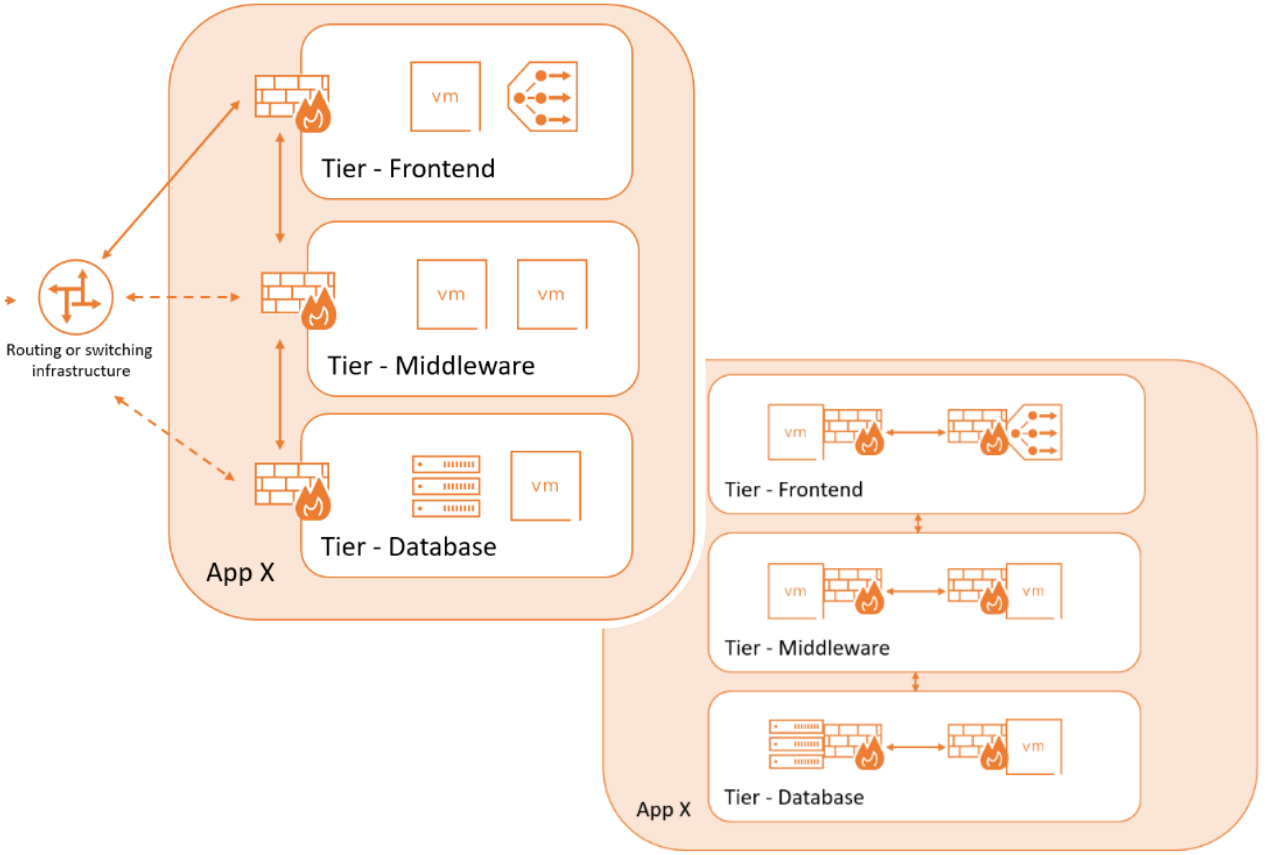

Filtering on Tier or individual VM

Rules only apply to Tier or VM

Intra-app traffic not default allowed

Rule Evaluation

ETHERNET (1) | EMERGENCY (0) | INFRASTRUCTURE (6) | ENVIRONMENT (9) | APPLICATION (3)

Rule evaluation stops at the first hit

Image source: Broadcom

Rule Evaluation

Jump-to-application skips the remainder and goes to the top of the 'Application' section and continues from there

Image source: Broadcom

Filter Direction

Incoming Traffic

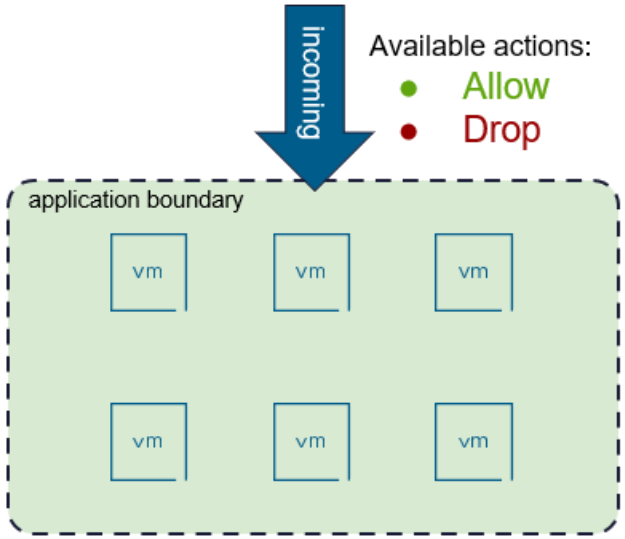

Outgoing Traffic

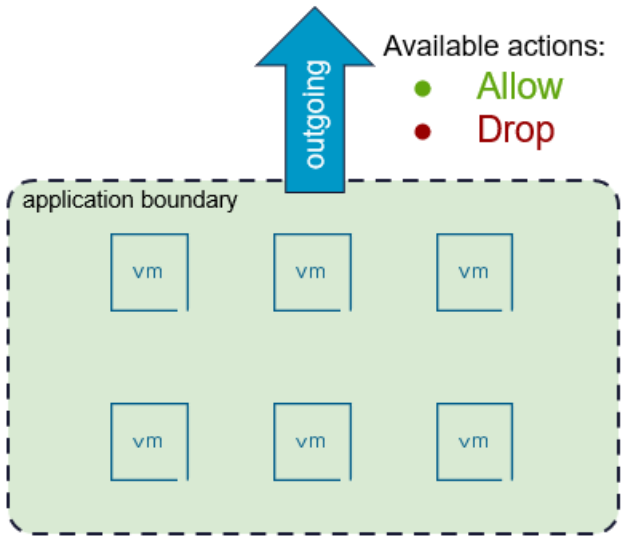

Intra-application Traffic

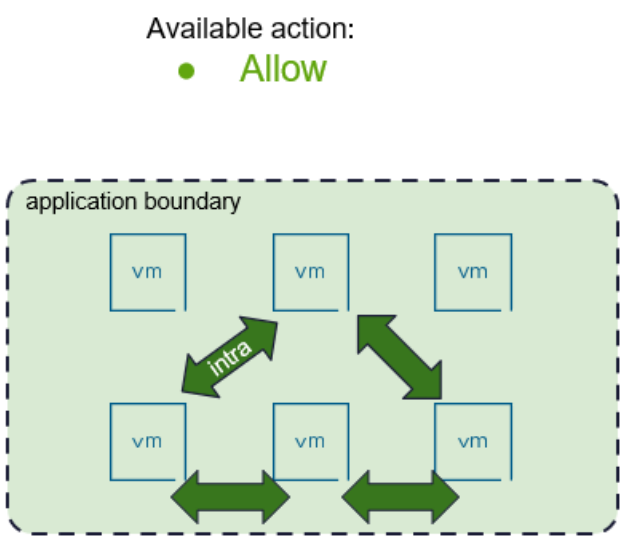

Image source: Broadcom

Overall concept

- VM identification agnostic to underlying network
  - Tags & Security Groups
- Security “onion” (next slide)
  - Layers
  - Datacenter hierarchy
- Allow some stuff – deny the rest
  - Easier troubleshooting
  - No need for be-all, end-all any-any-deny (but recommended!)
- Build outside-in
  - Shared services -> environments -> applications -> (tiers?)

Security Onion

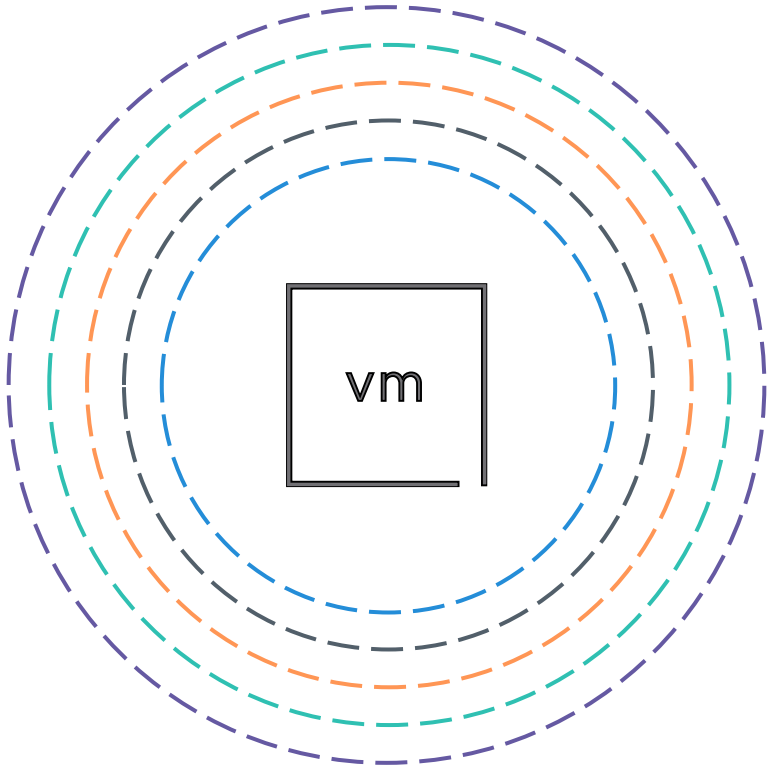

- Region
- Zone
- Environment
- Application
- Tier

Security Onion

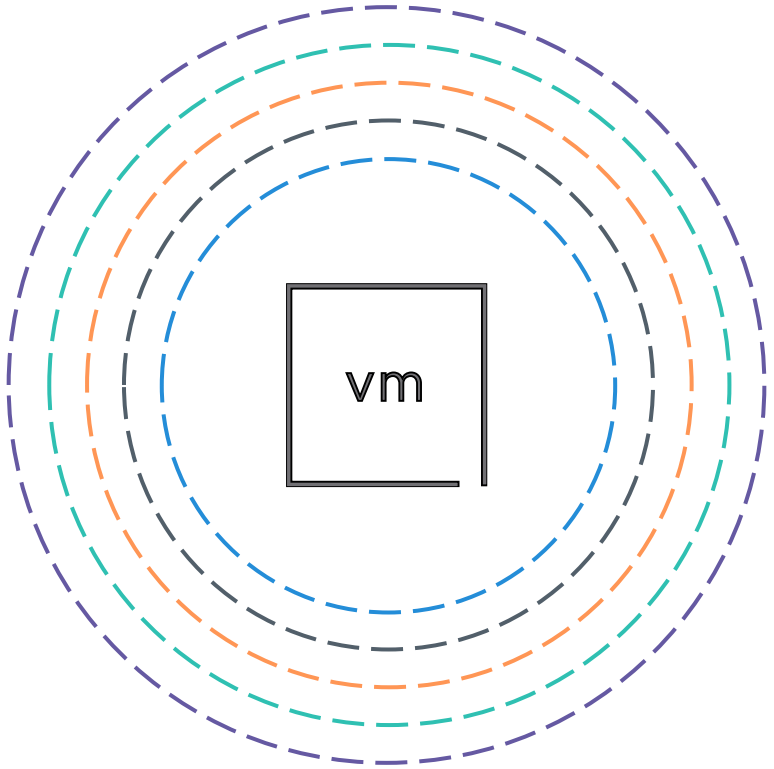

EU

Bunnik

Prod

Wordpress

MySQL

All VMs located in Bunnik, in the EU

Security Onion

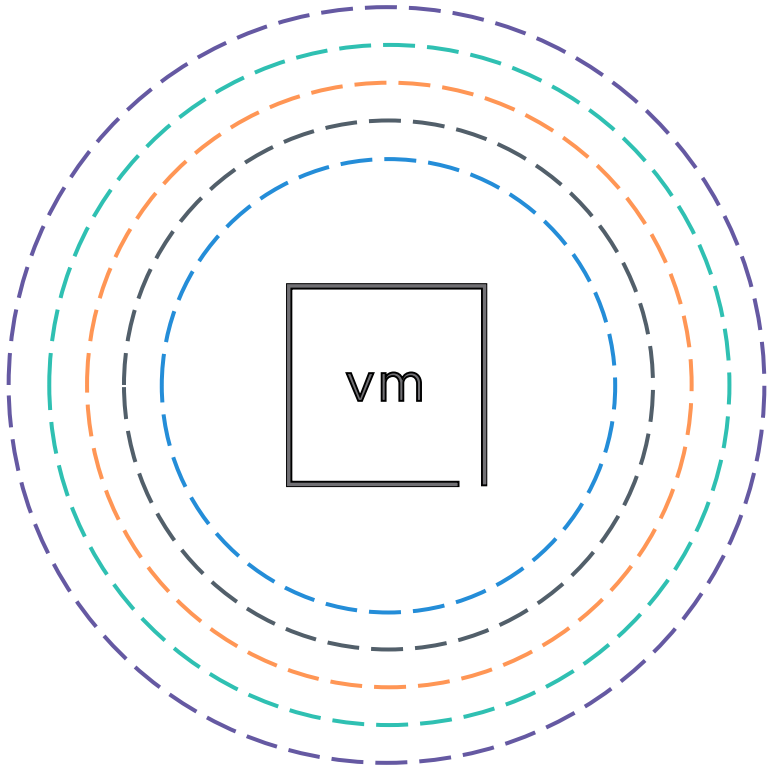

EU

Bunnik

Prod

Wordpress

MySQL

All Wordpress VMs in production

Security Onion

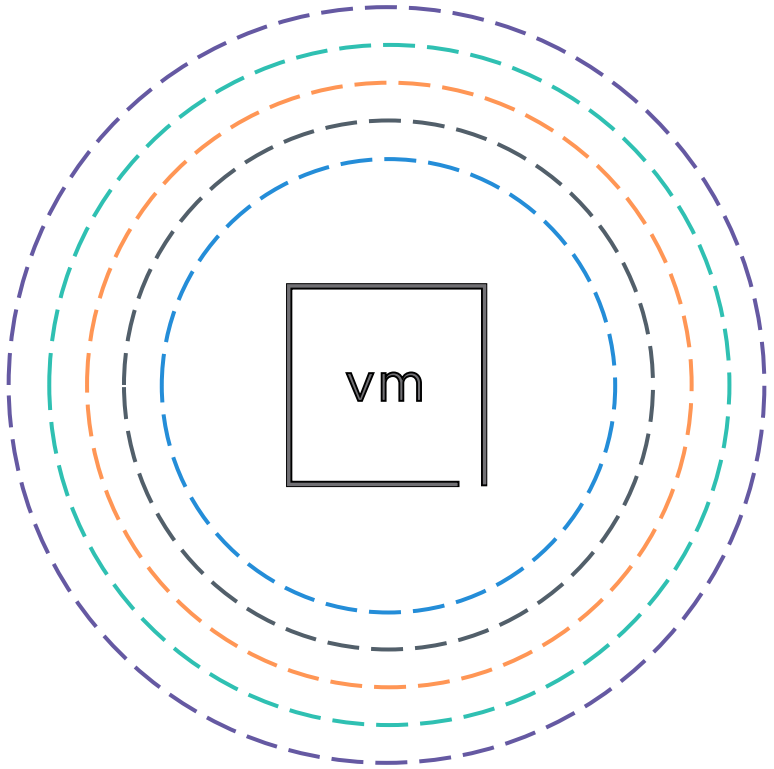

EU

Bunnik

Prod

Wordpress

MySQL

VMs running MySQL in the Wordpress app in production, located in Bunnik in the EU region

Security Onion

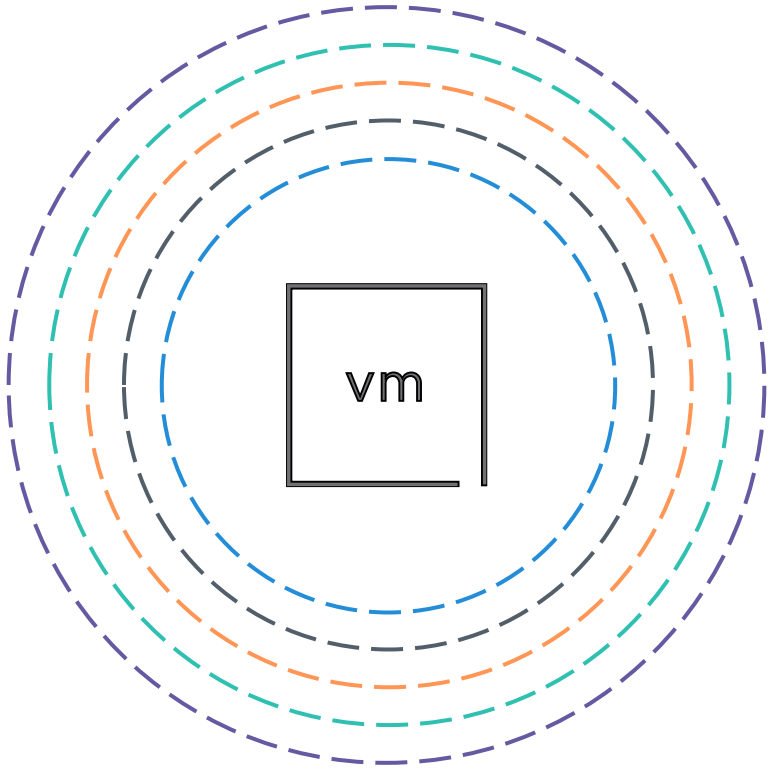

EU

Bunnik

Prod

Wordpress

MySQL

All VMs running MySQL for Wordpress

Datacenter Hierarchy

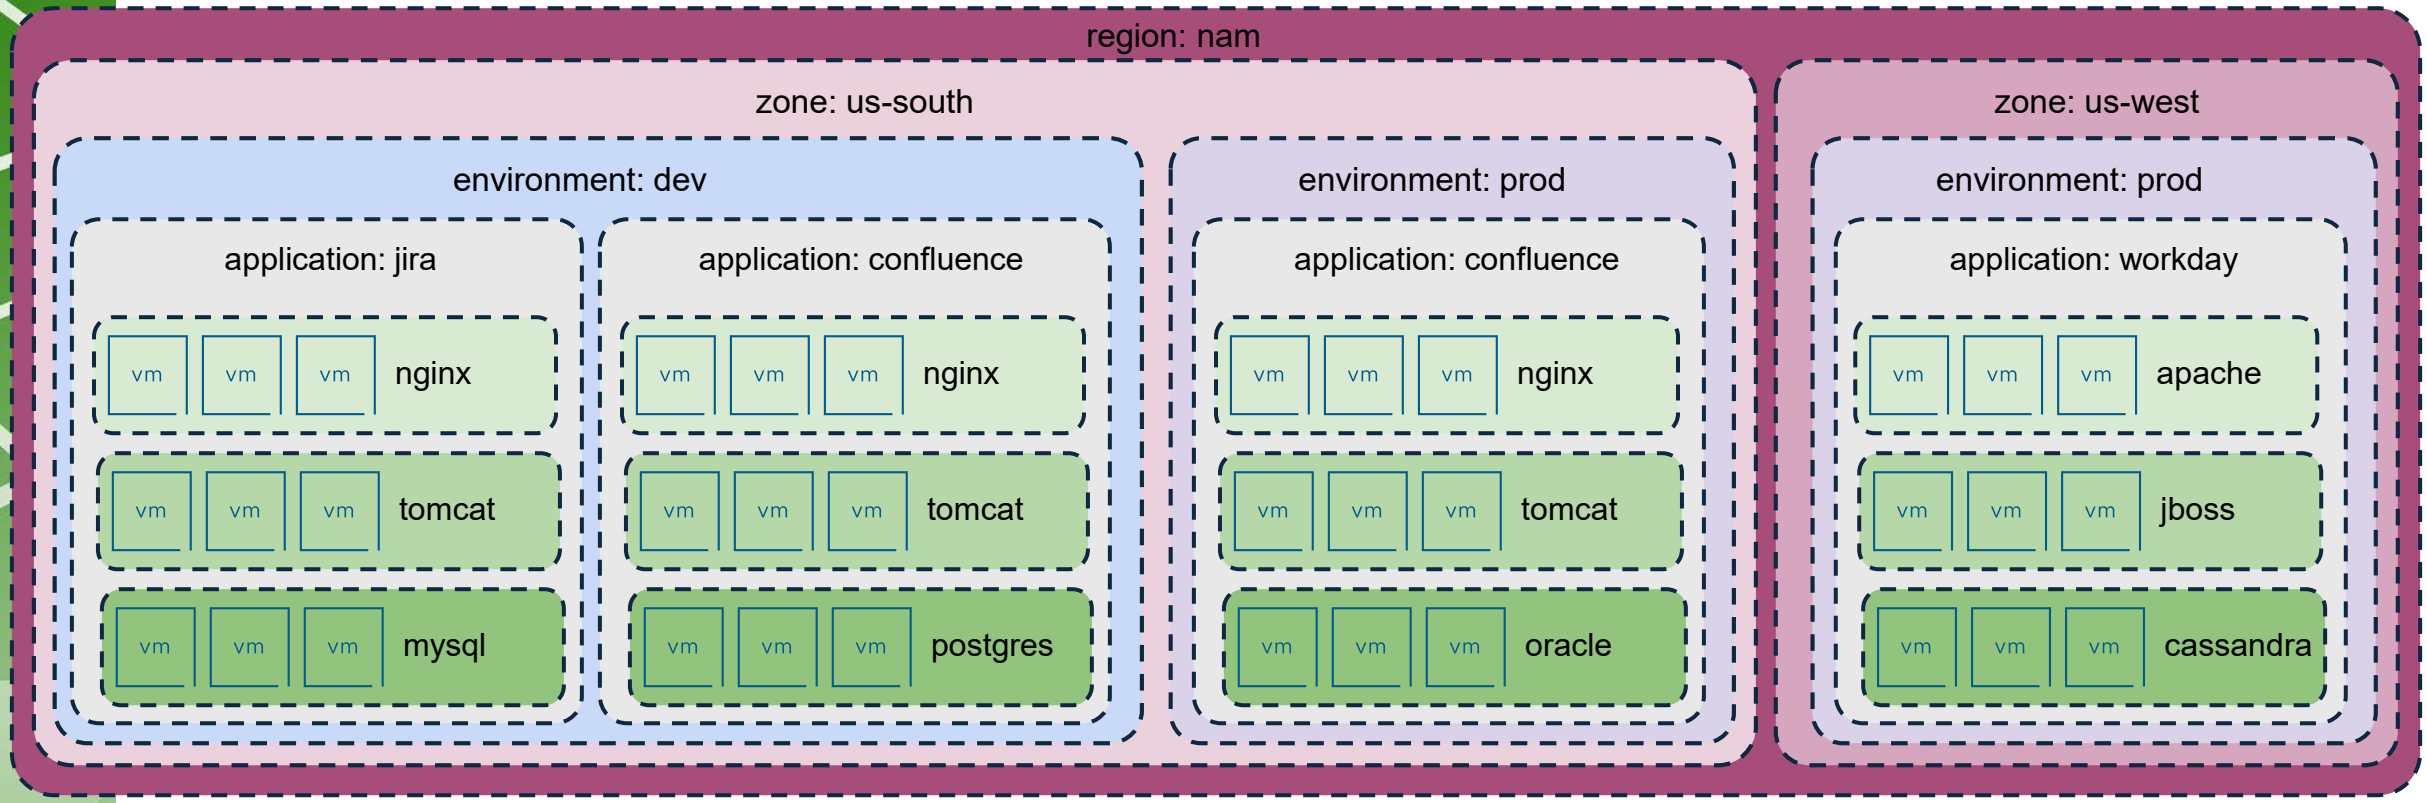

Image source: Broadcom

Step 1: Tag, Tag, Tag

This is the most annoying part of the project

- Know thyself! Or... your environment. (know thyronment? Thyvironment?)
- All VMs need to be tagged
  - One tag per category
- It depends:
  - Manual
    - Existing tags, CMDB
  - Algorithmically
    - Naming convention, folders
  - Magically
    - K-nearest Neighbours, flow analysis

Step 2: shared services

- Examples?
- DNS, NTP, AD, Monitoring, XDR, Repo, etc.
- Broad rules, generic access
- Initial set of rules for each new VM
  - Even without tags!

From this:

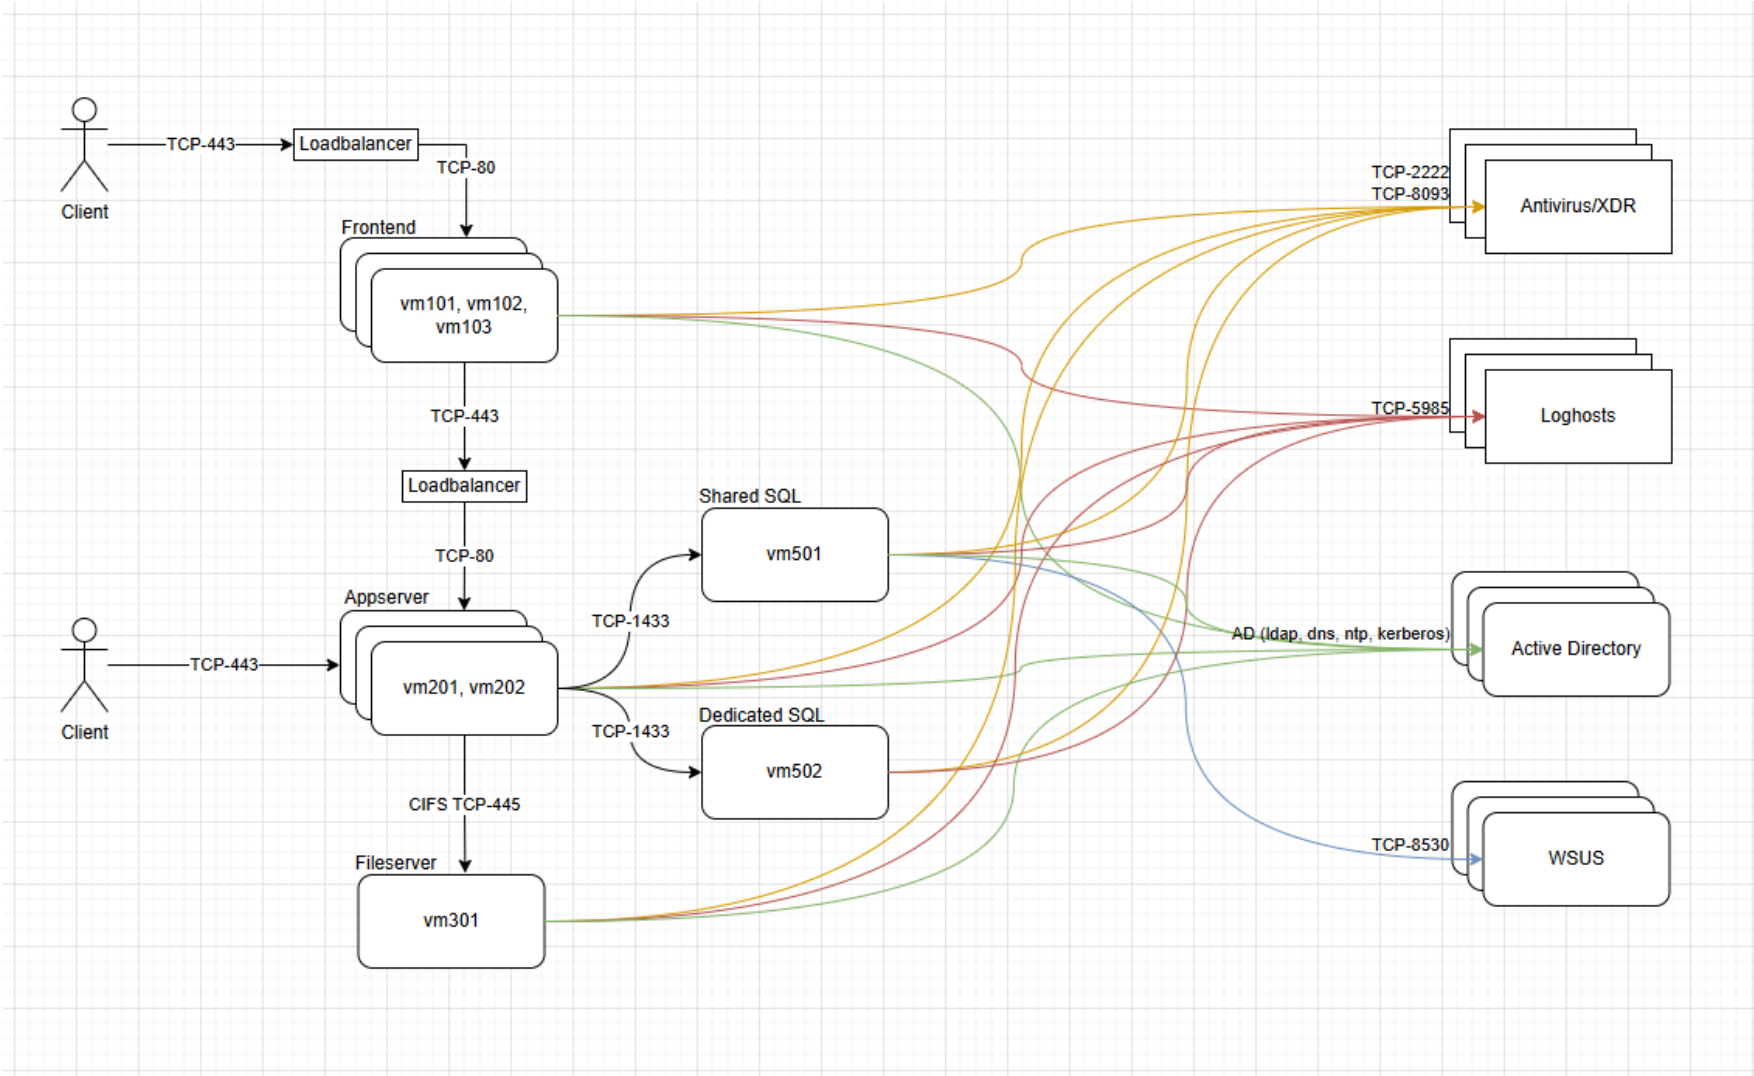
Client
TCP-443
Loadbalancer
TCP-80
Frontend
vm101, vm102, vm103
TCP-443
Loadbalancer
TCP-80
Appserver
vm201, vm202
Client
TCP-443
TCP-1433
Shared SQL
vm501
TCP-1433
Dedicated SQL
vm502
CIFS TCP-445
Fileserver
vm301
TCP-2222
TCP-8093
Antivirus/XDR
TCP-5985
Loghosts
AD (ldap, dns, ntp, kerberos)
Active Directory
TCP-8530
WSUS

To this:

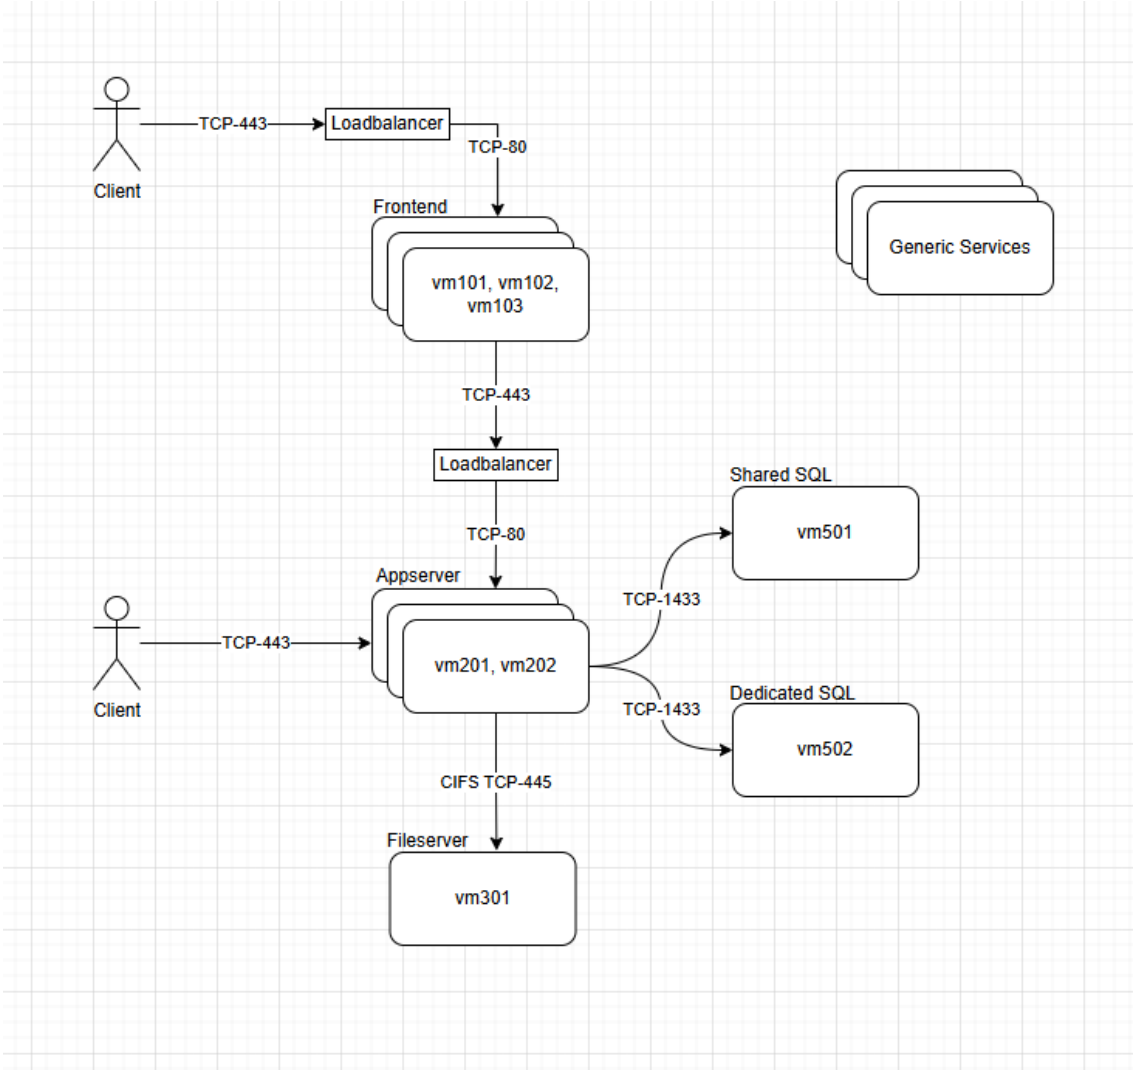
Client
TCP-443
Loadbalancer
TCP-80
Frontend
vm101, vm102, vm103
Generic Services
TCP-443
Loadbalancer
TCP-80
Shared SQL
vm501
Appserver
Client
TCP-443
vm201, vm202
TCP-1433
TCP-1433
Dedicated SQL
vm502
CIFS TCP-445
Fileserver
vm301

Rules

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

From	To	Service	Action
Any	SG_Infra_NTP	NTP	Allow
Any	~~SG_Infra_NTP~~	NTP	Drop
Any	SG_Infra_DNS	DNS	Allow
Any	~~SG_Infra_DNS~~	DNS	Drop
SG_Windows	SG_Infra_AD	SS_LDAPS	Allow
SG_Windows	~~SG_Infra_AD~~	SS_LDAPS	Drop

Security Group containing all Windows VMs

Security Set containing services related to AD

Step 3: Environments

- Create a security matrix like so

From\To	**Dev**	**Prod**	**HighSec**
Dev	Allow	Allow	Allow
Prod	Allow	Allow	Allow
HighSec	Allow	Allow	Allow

INFRASTRUCTURE | ENVIRONMENT | APPLICATION

Environment Pair	From	To	Service	Action
Dev ↔ Dev				
	Dev	Dev	Any	Allow
Prod ↔ Prod				
	Prod	Prod	Any	Allow
HighSec ↔ HighSec				
	HighSec	HighSec	Any	Allow

- Lock down environment transitions

From\To	**Dev**	**Prod**	**HighSec**
Dev	Allow	Drop	Drop
Prod	Drop	Allow	Drop
HighSec	Drop	Drop	Allow

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Environment Pair	From	To	Service	Action
Dev ↔ Dev				
	Dev	Dev	Any	Allow
Dev ↔ Prod				
	Dev	Prod	Any	Drop
Dev ↔ HighSec				
	Dev	HighSec	Any	Drop

- Allow for reality

From\To	**Dev**	**Prod**	**HighSec**
Dev	Allow	Drop with Exceptions	Drop
Prod	Drop with Exceptions	Allow	Drop with Exceptions
HighSec	Drop	Drop with Exceptions	Allow

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Environment Pair	From	To	Service	Action
Dev ↔ Dev				
	Dev	Dev	Any	Allow
Dev ↔ Prod				
	Dev-AppA	Prod-DatabaseX	Any	Allow
	Dev	Prod	Any	Drop
Dev ↔ HighSec				
	Dev	HighSec	Any	Drop

- Prepare for step 3: Application security

From\To	**Dev**	**Prod**	**HighSec**
Dev	Jump to Application	Drop with Exceptions	Drop
Prod	Drop with Exceptions	Jump to Application	Drop with Exceptions
HighSec	Drop	Drop with Exceptions	Jump to Application

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Environment Pair	From	To	Service	Action
Dev ↔ Dev				
	Dev	Dev	Any	Jump to Application
Dev ↔ Prod				
	Dev-AppA	Prod-DatabaseX	Any	Jump to Application
	Dev	Prod	Any	Drop
Dev ↔ HighSec				
	Dev	HighSec	Any	Drop

Step 4: Applications

- Reminder:
  Jump to Application = skip current category, start atop Application
- Reminder:
  Strategy – Application ring fencing or micro segmentation?

- Key point: monitor before lockdown

Application ring fencing

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Application	From	To	Service	Action
EU_Bunnik_Dev_Wordpress	Applied to: EU_Bunnik_Dev_Wordpress			
	EU_Bunnik_Dev_AppX	Any	HTTPS	Allow
	EU_Bunnik_Dev_AppY	Any	MySQL	Allow
	EU_Bunnik_Dev_Wordpress	Any	Any	Allow
	Any	Any	Any	Drop

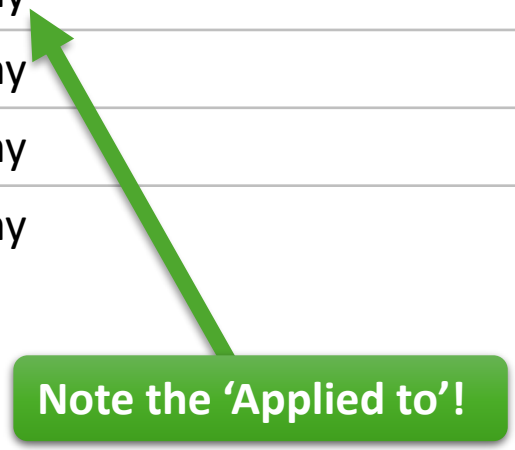

Micro segmentation

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Application	From	To	Service	Action
EU_Bunnik_Dev_Wordpress	Applied to: EU_Bunnik_Dev_Wordpress			
	Any	Nginx	HTTPS	Allow
	Nginx	Apache	HTTPS	Allow
	Apache	MySQL	MySQL	Allow
	Any	Any	Any	Drop

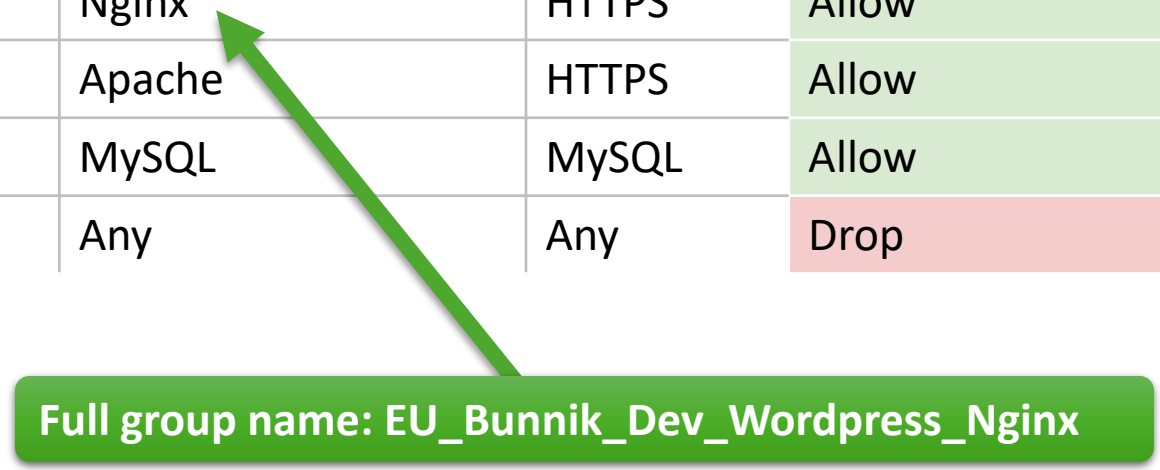

Logging & Learning

INFRASTRUCTURE → ENVIRONMENT → APPLICATION

Application	From	To	Service	Action
EU_Bunnik_Dev_Wordpress	Applied to: EU_Bunnik_Dev_Wordpress			
	Any	Nginx	HTTPS	Allow
	Nginx	Apache	HTTPS	Allow
	Apache	MySQL	MySQL	Allow
	Any	Any	Any	Allow & Log

After learning: Drop & Log!

More about logging & learning

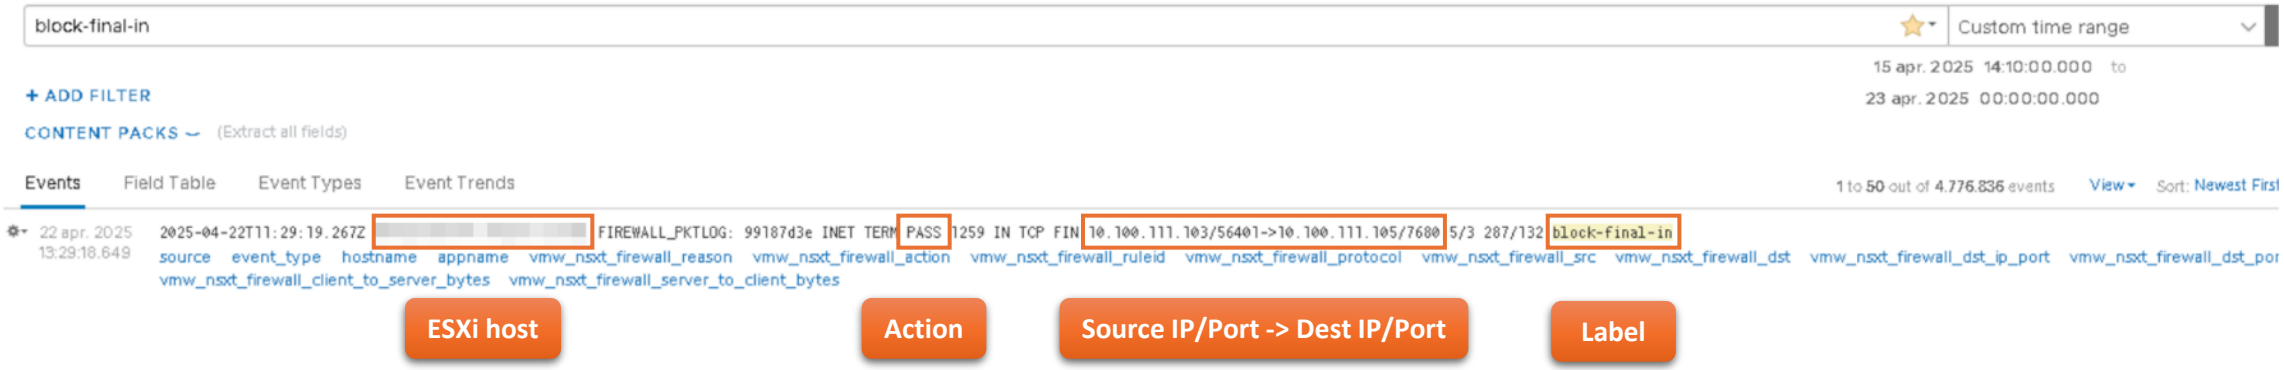

- Exposes **all** flows
- Trigger alerts on hit:
  - It means bona fide traffic that must be allowed
  - It means malicious traffic that must be investigated

Either way: action is needed!

Security Journey with Security Services Platform

Customer Challenges

Where should I start?

What is my current security posture?

How can I secure so many applications?

How can I do it safely and quickly?

Security Journey

New

vDefend Console

Stage 1: Security Assessment

Deploy Security Intelligence

Understand the current security posture

SSP 5.0

Stage 2: Infrastructure Services Protection

Automated infrastructure security

Provide immediate security benefits

SSP 5.1

Stage 3: Environment-level Protection

Isolation without need for deep application knowledge

Define global rules and manage exceptions

SSP 5.1

Stage 4: Application-level Isolation

Automated on a per-application basis

Based on learning from multiple customers

SSP 5.1

Built into vDefend – Enables Rapid Self-deployment

Slide courtesy of Broadcom

Security Journey

Show Journey Status on Homepage

vDefend Security Journey Starts Here!

The Security Journey prescribes a multi-stage segmentation workflow to speed up zero trust implementation with VMware vDefend Distributed Firewall. Continuous monitoring at each stage highlights any deviations and allows recalibration of security posture.

Instruction: Track progress by clicking the stage indicator ◯ to mark it "Reviewed".

Last calculated on Oct 18, 2025, 3:13:20 PM

1 out of 4 stages reviewed

STAGE 1 Assess Segmentation Posture

STAGE 2 Secure Infrastructure Shared Services

This stage enables protection to workload traffic that communicates with shared services such as DNS, NTP, Syslog and LDAP Servers. Import service information from CSV files or validate auto-detected services. Review and publish Distributed Firewall.

ⓘ 8 New Infrastructure Servers are detected.

↳ Update System with Infrastructure Services (Optional)

Refer to **Sample CSV** for format and examples. Then, import your CSV file to Segmentation Planning. The workflow will recommend groups and firewall rules to be implemented.

SEGMENTATION PLANNING

↳ Monitor and Generate Infrastructure Rules

The system discovers infrastructure services by analyzing traffic and **recommends groups and firewall rules** to be implemented.

MONITOR INFRASTRUCTURE SERVICES

CSV

View Sample CSV

Import one CSV, or import in stages for infrastructure services, environments, and applications.

Tip

To jump-start Infrastructure Services security efforts, review the "Project: Secure Infrastructure Services" section in your Security Segmentation Report. This section provides a summary of progress towards securing Infrastructure Services.

Guide for Data Center Segmentation & How to Prepare Your Data

Additional Resources

Slide courtesy of Broadcom

vmw vDefend | Security Services Platform

Administrator

Home | Security Journey | Monitor & Plan | Inventory | System

- Overview

Security Intelligence

- Security Explorer
- Segmentation Planning
- Segmentation Monitoring
- Policy Recommendations

Threat Monitoring

- Network Detection & Res...
- Malware Prevention
- Suspicious Traffic

Reports

- Security Segmentation R...
- Flow Export

Segmentation Planning

Organize your data center assets—such as infrastructure services, environments, applications— into a defined hierarchy and generate security policies. After publishing the data center assets, you can monitor traffic and refine policies to strengthen your security posture.

SEE DATA CENTER SEGMENTATION GUIDELINES

Data Center Segmentation Steps

STEP 1	STEP 2	STEP 3	STEP 4
Import File	Map File Columns To Assets	Preview File Data	Review & Publish Assets & Policies

File Import

Upload Checklist

1. **Download the Sample File**, fill out the asset type, identifiers, hierarchy scopes, save, and import CSV.
2. **File Format** View Details
3. **Recommended:** Match column headers exactly as in the sample file for automatic mapping.

Note: Data center asset imports add new assets. Changes take effect only after publishing. Existing assets stay unless manually deleted in NSX. Missing hierarchy paths in new files might cause sync issues. Review before importing.

Get Started With The Sample File

DOWNLOAD

SELECT A FILE OR DRAG AND DROP IT HERE

Supported file format: .csv

Maximum file size: 500 MB

Slide courtesy of Broadcom

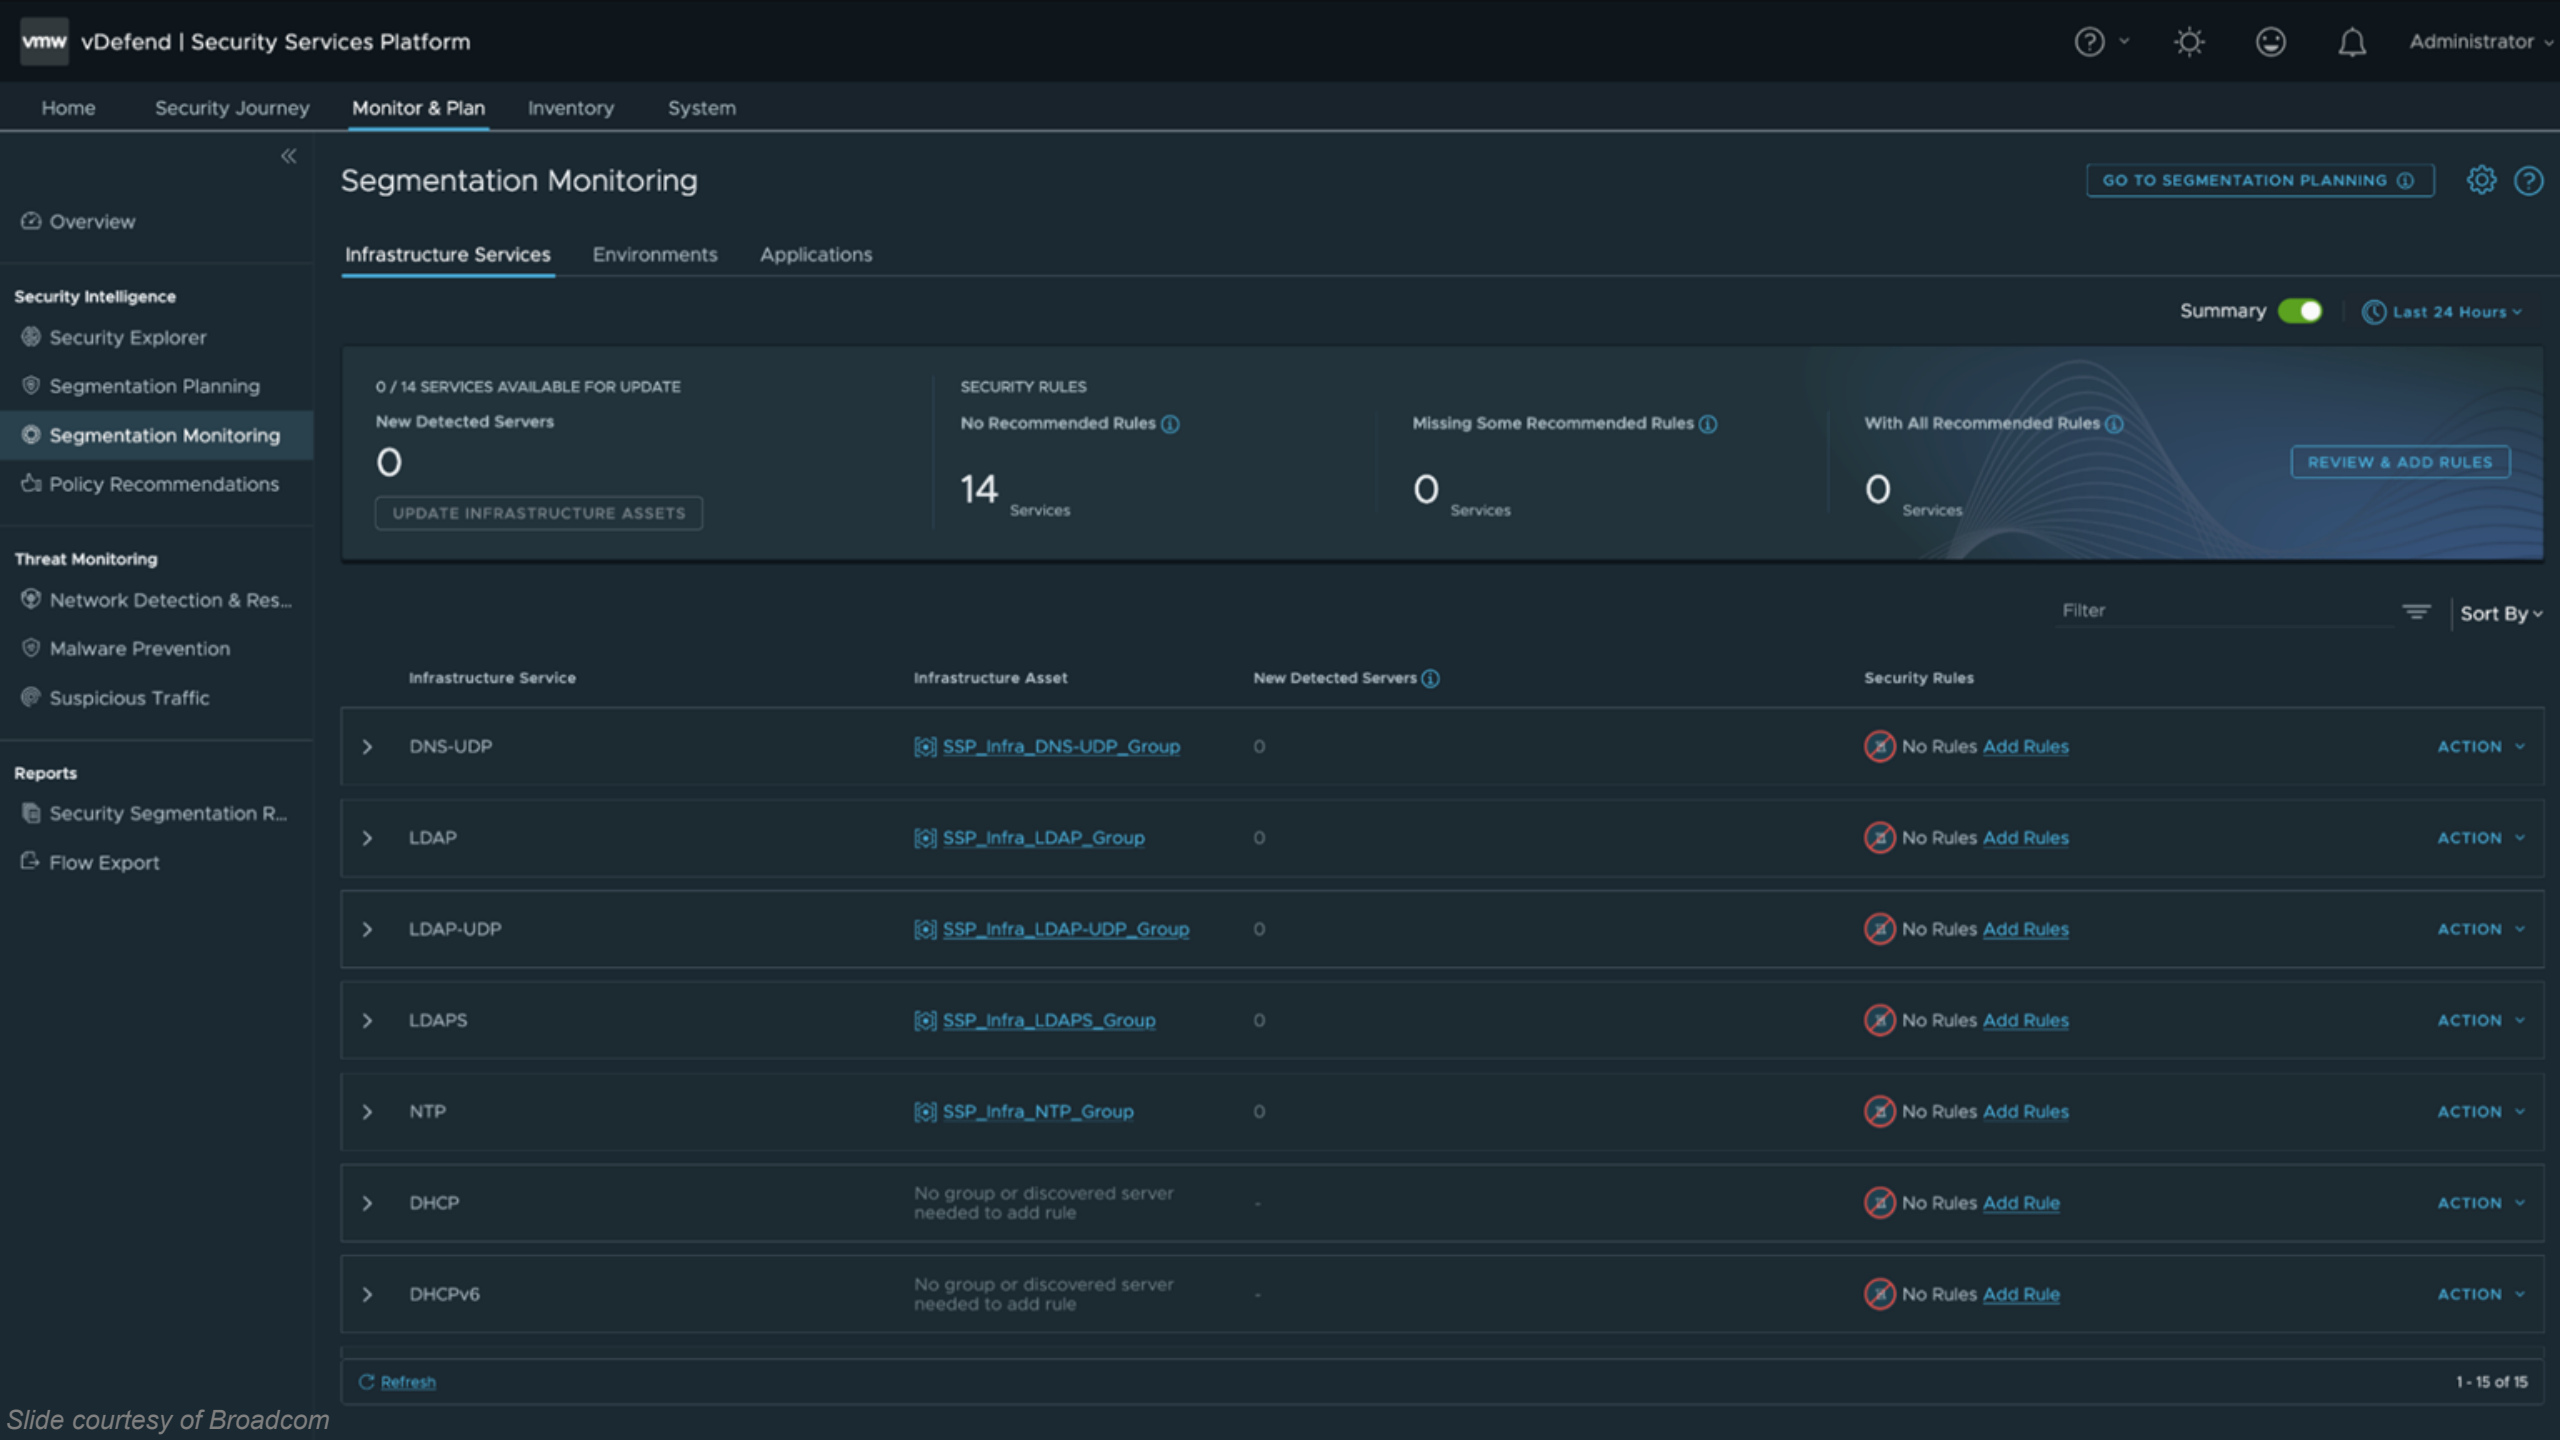

Slide courtesy of Broadcom

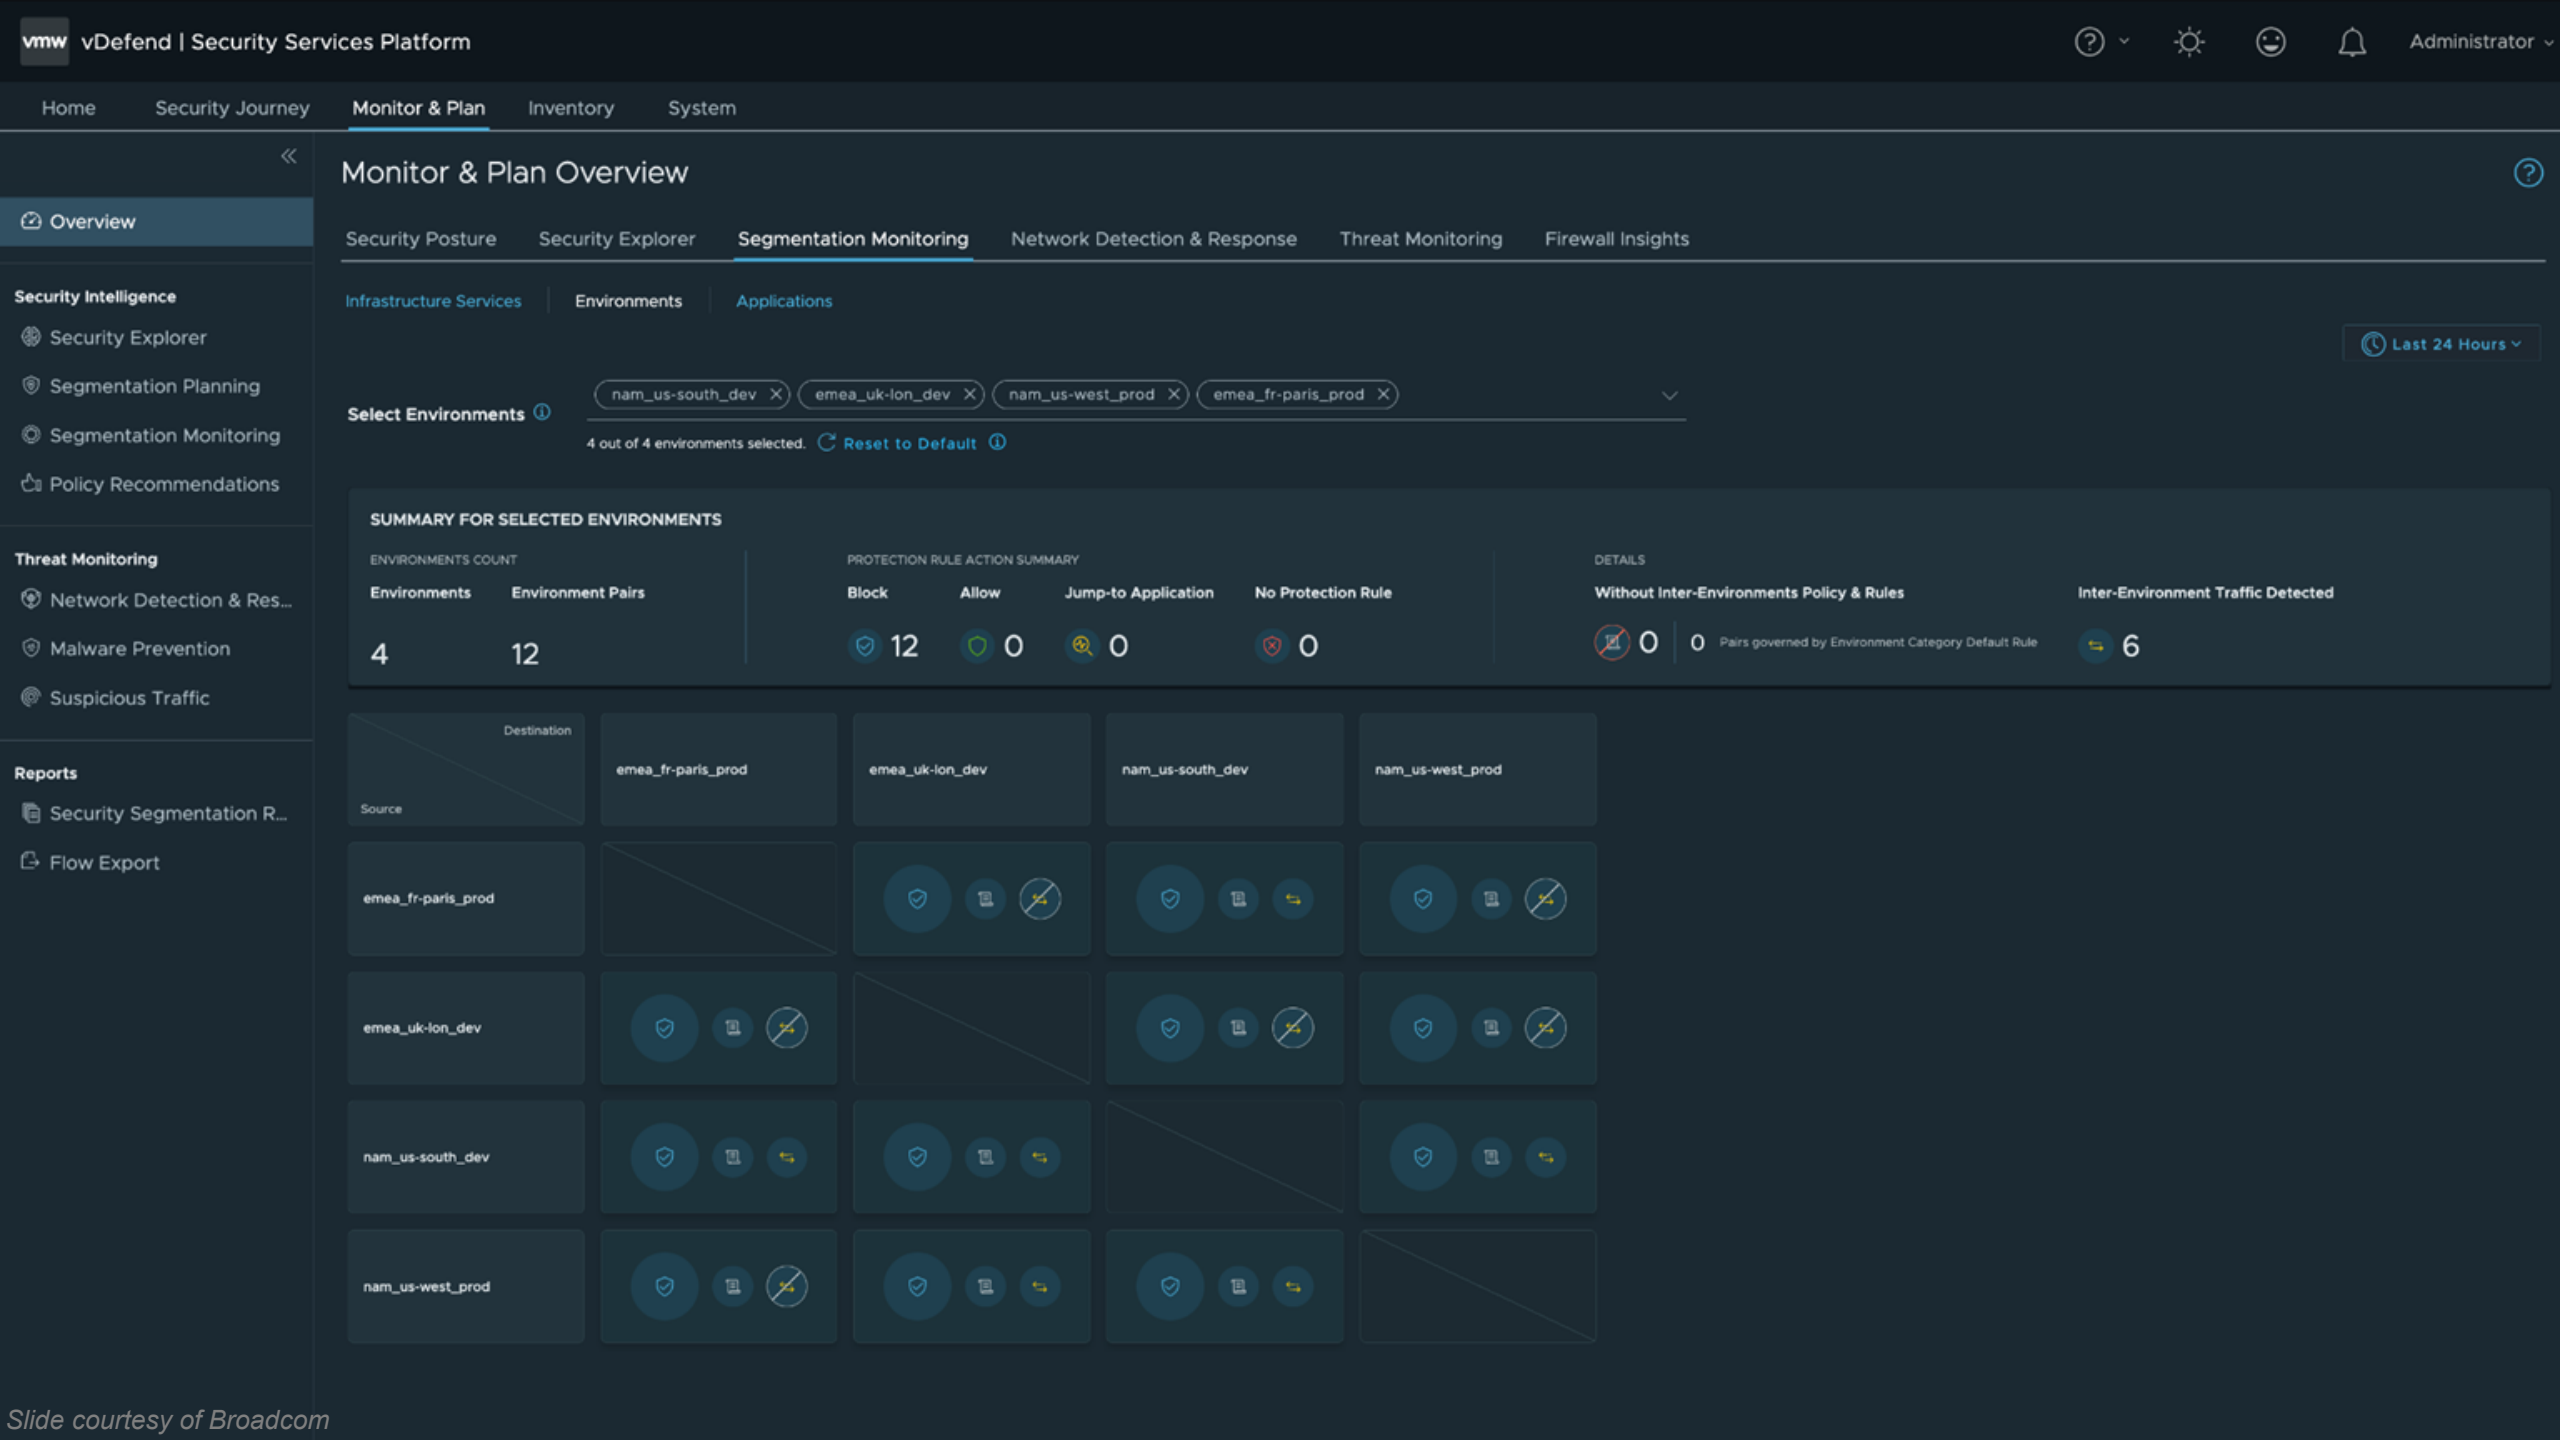

Slide courtesy of Broadcom

Robert's tips for success

- Define your end goal:
  - Microseg, App Centric, something else?
  - Any-any deny at the end or no?
- Go step-by-step
  - You *cannot* do this quickly, it takes time
- Focus! Get a project owner
- Shared services first, high risk next
- All VMs need tags - one tag per category max!
  - VM that serves multiple? Give it the most-secure tag
  - Example: VM serves dev & prod? Then tag as prod

So, get started today... How?

- Tag your environment!
  - Create an export
  - Give each VM a home
  - Work together with application owners
- Define your shared services
  - What applies to for all VMs?
  - Which services are needed regardless of application?
- Minimum of tools needed:
  - vDefend
  - Syslog
  - Excel (probably)

daa t
osegm tatie!

Robert Cranendonk

VMUGNL
VMware User Group

VCF TechCon

Powered by VMUG

Start *vandaag* nog met microsegmentatie!

Robert Cranendonk

VMUGNL
VMware User Group

Extra

vDefend Security Policies Across Multi-Tenancy Scope

Distributed Firewall Rules Precedence

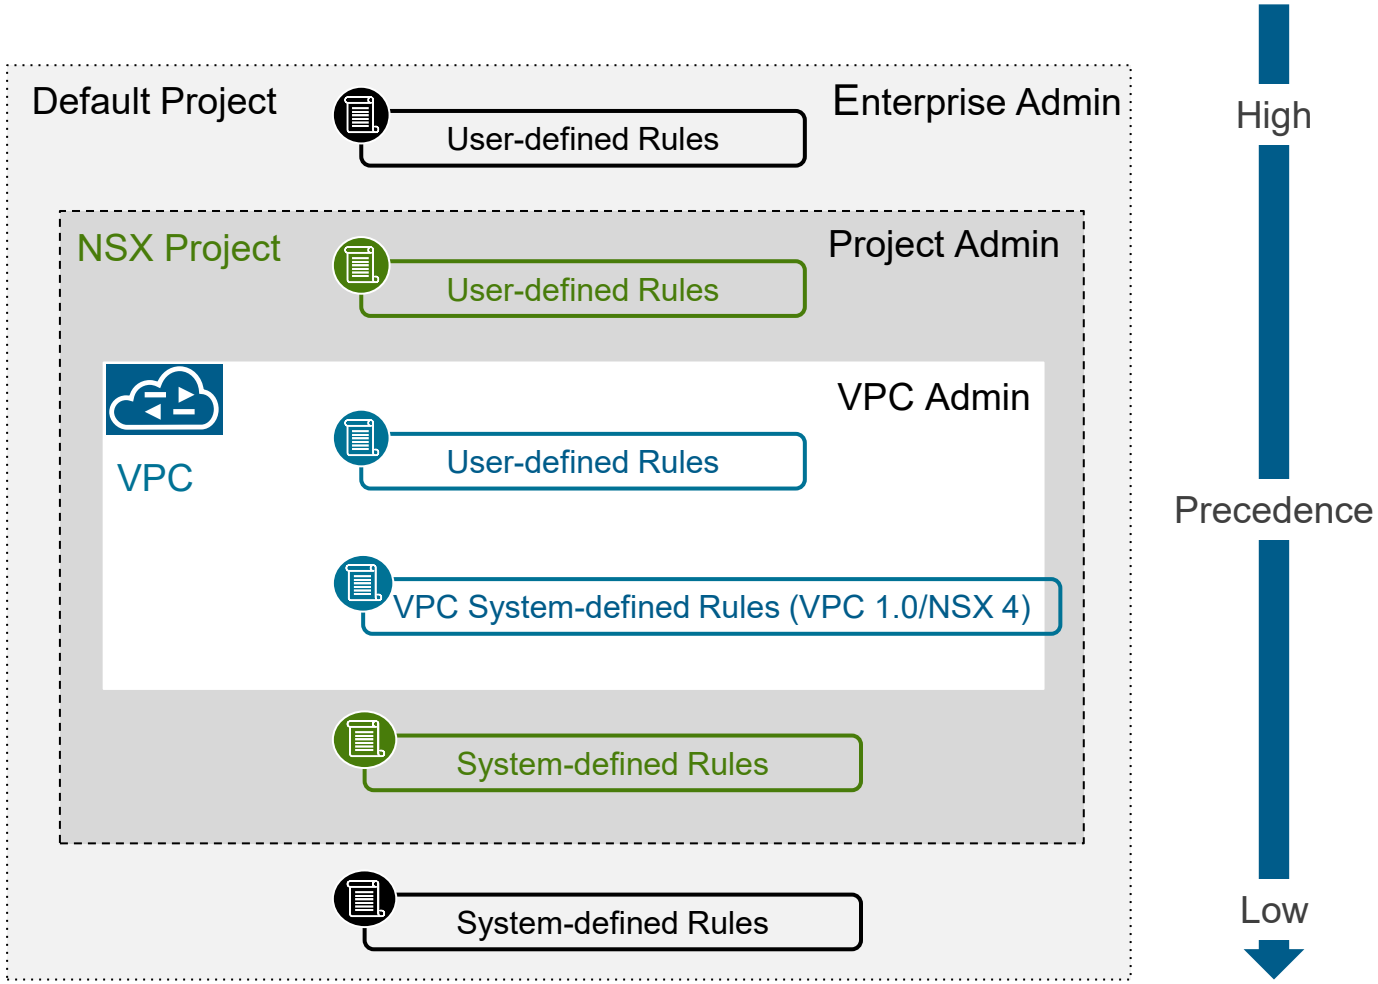

DFW Security policies are defined at the multi-tenancy space (Default Project, NSX Project, and VPC) by respective user personas

Default Project and NSX Project user-defined security policies get enforced before the VPC security policies

Default Project and NSX Project system-defined (default) security policies get enforced after the VPC security policies

VPC policies by default are hidden in the Project view, but can be shown if desired

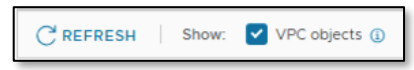

Slide courtesy of Broadcom

vDefend Self-Service Security

Persona-specific capabilities with CTGW connectivity

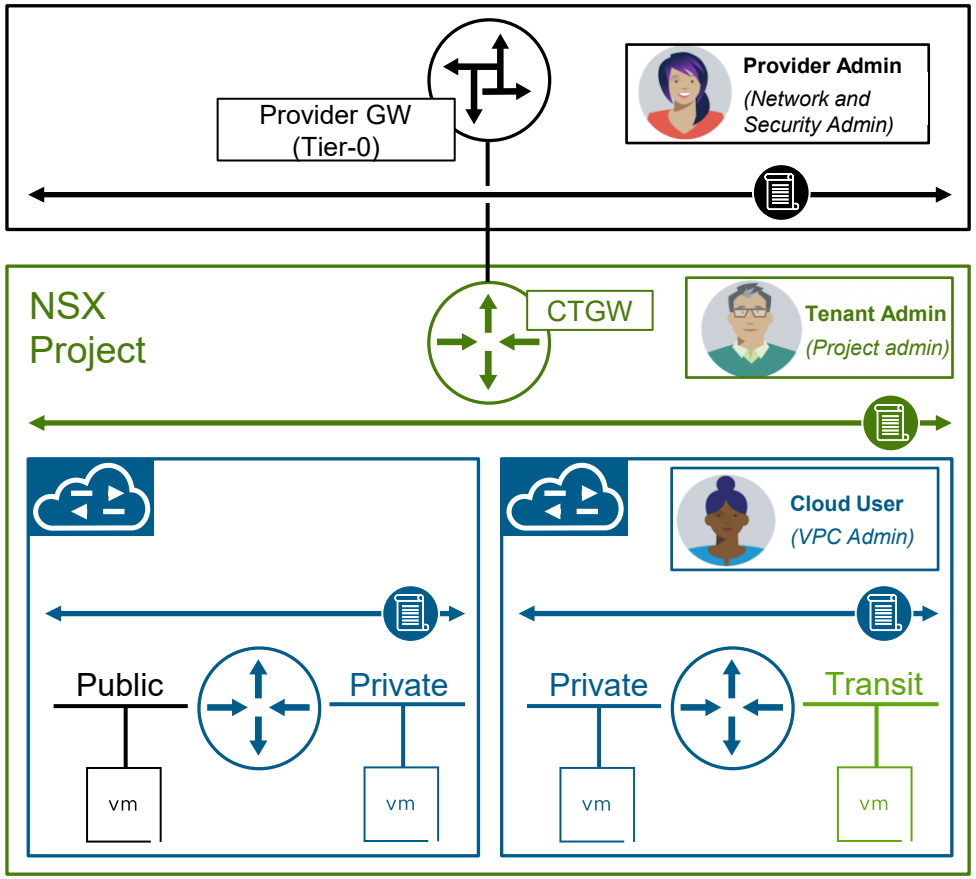

Provider Admin:

- Gateway Firewall on A/S Tier-0
- NSX Project system-defined DFW Policy (on/off)
- Default Project' DFW Policy
- Resource share with Projects

Tenant Admin

- Project DFW Policy
- Project Distributed IDS/IPS
- Project-level Groups, Tags, Services, and Profiles definition
- Resources share with VPCs

Cloud User

- VPC Groups definition
- VPC E-W Firewall
- VPC N-S Firewall

Slide courtesy of Broadcom

vDefend Firewall Multi Tenant Logging

Logs segregation for Projects and VPCs

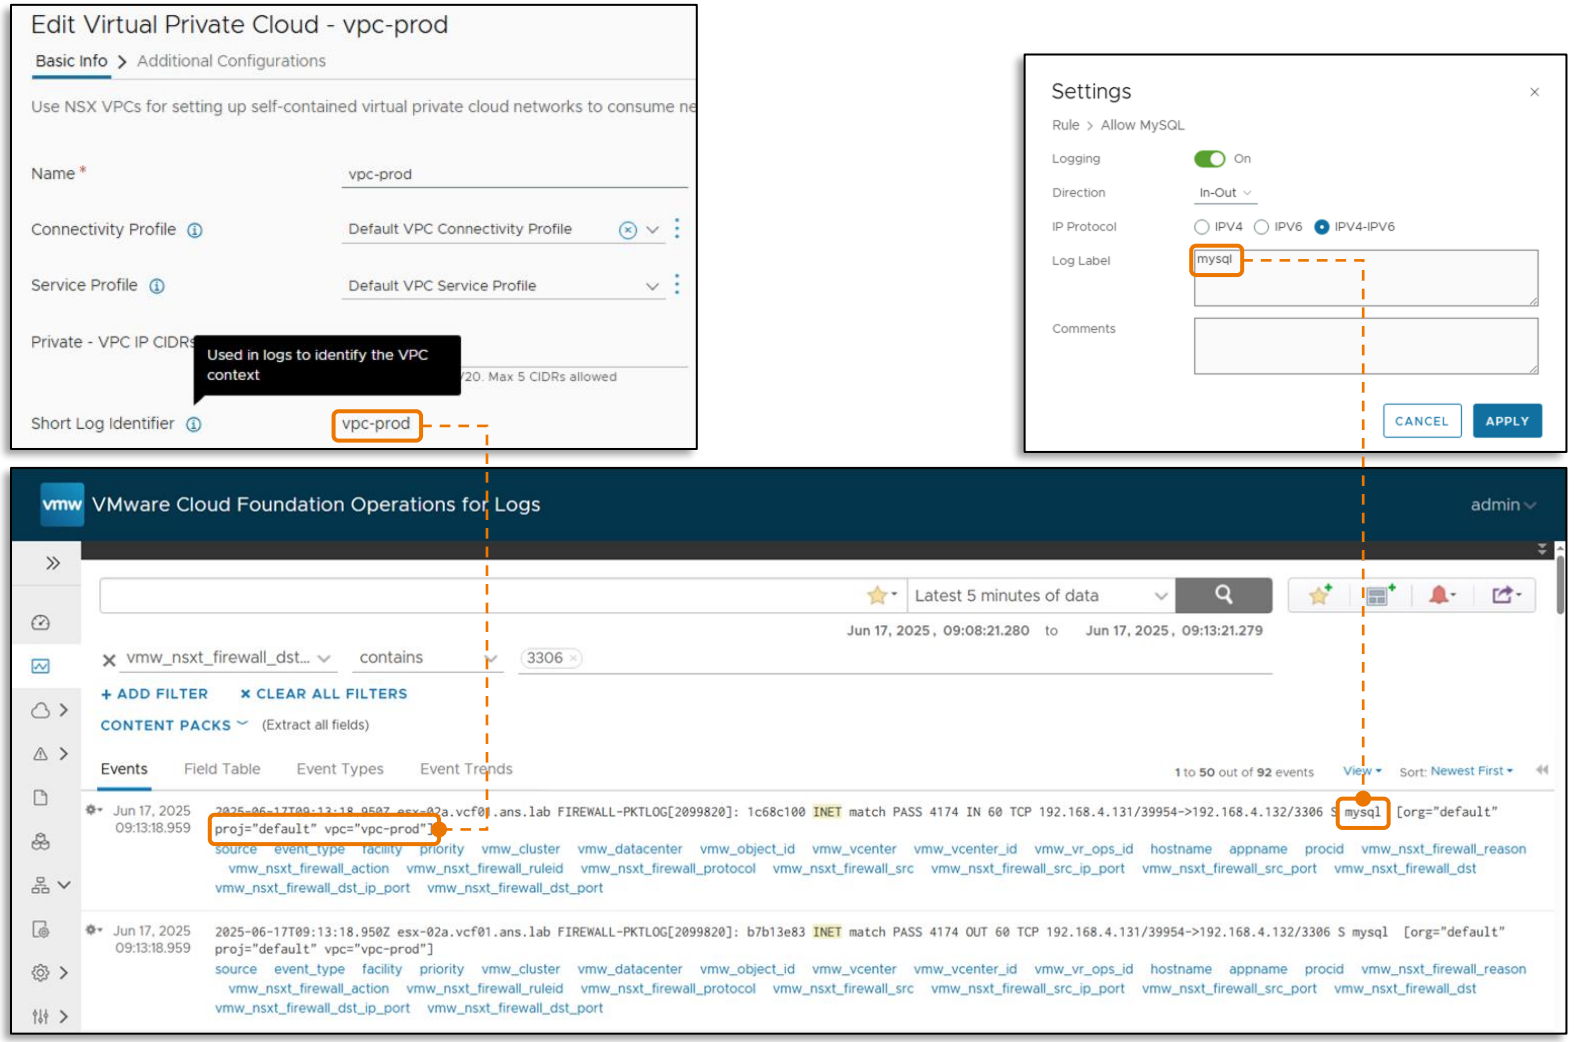

- NSX Project and VPC unique Short Log Identifiers
- Appended to the DFW and GFW rules with logging enabled
- Easy logs filtering for a specific Project and VPC workload
- Set at the creation time and can't be changed later
- If not specified, a system one is generated
- Rule-level log labels are also appended to the firewall logs

Slide courtesy of Broadcom