Slide 1
Hello, welcome everyone to this hands-on course, "Hands-On Network Threat Detection and Response Using ACAS."

My name is Edgardo Montesdeoca from Montimage.

I will be guiding you through this curriculum. This course is part of the CyberSuite project, which is co-funded by the European Union and focuses on implementing innovative Security-as-a-Service solutions to protect modern digital infrastructures.

Slide 2
Our journey is organized into several key phases.

We will begin with a general course introduction and then move into Module 1, where we cover the fundamentals of network threats and anomaly detection.

In Module 2, you will get your first look at the ACAS platform.

Module 3 dives into the technical details of Deep Packet Inspection and feature extraction, while Module 4 compares rule-based and deep learning detection methods.

We conclude the instructional portion with Module 5, focusing on Explainable AI (XAI), before moving to a final capstone assessment and course wrap-up.

Slide 3
Let's start with the course introduction section.

Slide 4
To understand why we need advanced tools like ACAS, we must look at the current state of network security.

As of 2022, approximately 95% of network traffic is encrypted, and alarmingly, over 85% of attacks now hide within that encrypted traffic.

Traditional signature-based and rule-driven systems often struggle because they cannot easily inspect encrypted payloads.

While decrypting traffic for inspection is possible, it is extremely resource-intensive, often impractical, and raises significant privacy concerns.

This creates an urgent need for AI-driven solutions that can analyse metadata and behavioural signals in real-time without needing to decrypt the content.

Slide 5

By the end of this course, you will have a solid understanding of both rule-based and machine learning approaches to anomaly detection.

You will be proficient in navigating the ACAS platform, performing Deep Packet Inspection (DPI), and analysing complex protocol hierarchies.

Furthermore, you will learn how to extract flow-based features used by ML models and, perhaps most importantly, how to apply Explainable AI methods like SHAP and LIME to interpret why an AI model flagged specific traffic as a threat. Later, we will present these techniques in more detail.

Slide 6
This course is specifically designed for IT and cybersecurity students, SOC analysts, and network security engineers who want practical, hands-on experience with AI-driven detection.

While you should have a basic understanding of networking—such as IP addresses, ports, and protocols—you do not need any prior experience with machine learning or AI.

We have designed the curriculum so that every advanced concept is introduced clearly before we put it into practice.

Slide 7
The curriculum is built for flexibility, totalling less than 15 hours of effort that you can spread over several days.

Each of the five modules consists of a short lesson followed by a quiz to test your knowledge. For instance, you will spend about two hours getting started with ACAS, four hours on detection approaches, and another three hours on the final Capstone assessment, where you will apply everything you've learned in a realistic scenario.

Slide 8
Recommended procedure to follow is: View each video in order (or go through the slide set and read the text).

For each video, pause it if some term or concept is not clear for you, and search the web or ask a chat bot for more details. Active participation is important for acquiring higher knowledge.

For Modules 2 to 5 and Capstone, follow the instructions to carry out the exercises or answer the quiz. At the end, provide your feedback. This is important for us to be able to improve the course.

Slide 9
Practical application is the heart of this course.

You will use a web-based portal to explore key functionalities including DPI, model evaluation, and XAI-driven incident analysis.

Our goal is for you to assess detection accuracy and performance under realistic conditions.

You can find all the necessary lab materials and the platform link by scanning the QR code on this slide, which will take you to our dedicated site.

Slide 10
To earn your digital certificate of completion from the CyberSuite Academy, you must pass each module’s quiz with a score of at least 7 over 10 and complete the final capstone assessment.

Before we begin, a brief note on ethics: when working with network traffic, you may see sensitive information.

You must treat all data as confidential and follow your organization's data handling policies. The PCAP files we provide for the labs have been sanitized for educational safety.

Slide 11
We will now begin Module 1, Network Threats and Anomaly Detection Fundamentals, which provides the theoretical foundation for everything that follows.

This module focuses on understanding the core concepts of network threats and how we can use various detection methods to identify them.

Slide 12
In this lesson, we will define what network threats are and why they are so dangerous for organizations today.

We will describe common attack types like DDoS and data exfiltration, explain how Deep Packet Inspection functions as a "microscope" for network traffic, and compare the different approaches to anomaly detection, ranging from simple rules to complex machine learning.

Slide 13
A network threat is any malicious activity or event that can compromise the confidentiality, integrity, or availability of networked systems and data.

Unlike vulnerabilities, which are weaknesses in a system, threats are active attempts to exploit those weaknesses.

Organizations are increasingly challenged by the sheer volume (millions of packets per second), sophistication, and speed of modern cyber threats.

Furthermore, with more than 80% of network traffic now encrypted, human analysts can no longer manually inspect every packet.

As a result, automated and intelligent systems have become essential for effective modern network defence.

Slide 14
It is important to recognize the unique "fingerprints" of different types of attacks. The table on this slide summarizes common network attacks and their typical characteristics.

For example, a DDoS attack is characterized by a high volume of traffic or packet rates designed to overwhelm system resources.

An intrusion attempt often manifests as repeated failed login attempts or unusual port access patterns, while malware activity may be identified through Command and Control (C2) beaconing. Data exfiltration is frequently indicated by large outbound data transfers using unusual protocols, and reconnaissance activities are commonly characterized by sequential port scanning.

Slide 15
Deep Packet Inspection or DPI goes beyond traditional firewalls by examining both the header and the payload of a packet.

This allows us to identify protocols regardless of the port being used—such as finding HTTP traffic on an unusual port like 8080.

DPI lets us extract content, reconstruct entire communication sessions, and find subtle anomalies that deviate from normal behavioural patterns.

In the ACAS platform, DPI provides the foundation for our traffic statistics and protocol visualizations.

Looking at the packet structure, you can see how different levels of filtering work. While simple filters only look at Layer 2 or Layer 3 headers, DPI inspects all layers, including the payload.

This deeper view is what allows ACAS to generate detailed traffic statistics and perform bidirectional flow visualizations, giving analysts the full context of what is happening on the wire.

Slide 16
There are two main ways we detect anomalies.

Rule-based detection uses "signatures" to match traffic against known attack patterns. It is excellent for catching known threats with very few false positives.

Machine learning, on the other hand, learns complex patterns from data.

While rules are easy to understand and audit, machine learning is far superior at identifying unknown or "zero-day" attacks that have no existing signature.

Slide 17
To look closer at the trade-offs: rule-based systems are fast and computationally efficient, but they require constant manual updates and can be easily evaded if an attacker slightly modifies their behaviour.

Machine learning models, such as CNNs and Autoencoders, can handle high-dimensional data and detect very subtle attack patterns.

However, they are often seen as "black boxes" that are harder to interpret and require high-quality training data to be effective.

Slide 18
This leads us to the category of Network Threat Detection and Response, or NDR.

Unlike traditional systems that might only alert you to a problem, an NDR platform provides a full suite of tools for investigation and response.

ACAS is a modern NDR platform that combines the best of both worlds: it uses rule-based detection for reliability, deep learning for advanced threats, and Explainable AI to make sure the analyst understands the "why" behind every alert.

Slide 19
To summarize Module 1: network threats range from DDoS to sophisticated exfiltration.

DPI is our fundamental tool for understanding network behaviour at the packet level.

While both rule-based and machine learning approaches have their pros and cons, modern platforms like ACAS integrate these methods to provide a comprehensive, explainable defence against both known and unknown threats.

Slide 20
Now that we have the fundamentals down, Module 2, Getting Started with ACAS, will introduce you to the ACAS platform itself.

We will look at its architecture and show you how to start your first analysis.

Slide 21
In this lesson, you will learn exactly what problems ACAS solves for security teams.

We will walk through the platform's key features, show you how to create an account and navigate the interface, and explain how to properly upload PCAP files for analysis.

Slide 22
Security analysts today face a massive "data overload" problem; traditional tools give alerts but no context, leading to hours spent chasing false positives.

ACAS (the Advanced Cybersecurity Analytics Service), developed by Montimage, solves this by automating traffic analysis.

It extracts meaningful features for ML, detects anomalies using multiple methods, and uses SHAP and LIME to provide natural language explanations and actionable recommendations through an LLM-powered assistant.

Slide 23
The ACAS architecture is a sophisticated pipeline.

It starts with a React-based user interface and a backend API.

When a PCAP is uploaded, the DPI Engine processes it, and the Feature Extraction module pulls out over 60 different metrics.

This data then flows into the Rule-based and ML models for detection.

Finally, the XAI Module interprets the results, and the LLM Assistant provides guidance to the user.

Slide 24
Every feature in ACAS is designed with a specific benefit in mind.

DPI helps you understand detailed behavior, while automated feature extraction saves you hours of manual data preparation.

The prebuilt ML models allow you to detect anomalies without having to train your own neural networks, and the XAI analysis ensures you can explain exactly which features—like packet length or timing—driven a specific security prediction.

Slide 25
You can access the platform at ndr.montimage.eu.

Setting up an account is simple: navigate to the URL, sign up with your email or use Gmail/GitHub for a quick login, and verify your credentials.

While some features are available for public exploration, we highly recommend registering to access the full range of detection and analysis tools.

Slide 26
For those interested in integration, ACAS is built on over 70 Swagger-based APIs.

This allows developers and advanced users to programmatically manipulate the monitoring tools, access reports, upload trace files, and even manage the model building and prediction processes directly.

Slide 27
ACAS uses the industry-standard PCAP formats.

For our web-based platform, there is a 10 MB file size limit to ensure fast processing. We provide several sample files, including normal traffic and various attack scenarios like DDoS and intrusion attempts.

As a best practice, start with these smaller samples to familiarize yourself with the interface before moving to your own captures.

Slide 28
Your first hands-on task is Platform Exploration.

Log in to the portal and navigate through the different sections. Upload one of our sample PCAP files and examine the DPI results.

Pay close attention to the traffic statistics and protocol distribution to see how ACAS visualizes the raw data you just uploaded.

Slide 29
In summary, ACAS is a comprehensive NDR platform that integrates advanced analysis and explainability. Its architecture is built around a powerful DPI engine and multiple detection modules.

You now know how to access the platform at ndr.montimage.eu and how to manage the PCAP files that serve as the primary input for our analysis.

Slide 30
In Module 3, Deep Packet Inspection & Feature Extraction, we move from theory to practice regarding how we actually look inside network traffic.

We will focus on the data that fuels our machine learning models.

Slide 31
This lesson will teach you how to interpret DPI results and analyse traffic statistics like packet counts and timing.

You will learn to visualize bidirectional flows to see endpoint conversations and understand how we extract over 60 different features to describe network behaviour numerically.

Slide 32
When you upload a file, the ACAS DPI engine processes every single packet.

It doesn't just look at the whole file; it breaks it down into summaries, hierarchies, and flows.

This granular processing is what allows us to see not just *how much* traffic is moving, but *exactly what* that traffic is doing and how it changes over time.

Slide 33
Traffic statistics provide a high-level overview of captured network data. Several key metrics are listed in the table.

When analysing these statistics, look for unusual volume spikes that may indicate a DDoS attack, duration anomalies where network flows are either suspiciously short or unusually long, and irregular IP distribution patterns.

For example, if many different source IP addresses are communicating with a single destination, it may be a classic indicator of a scanning or flooding attack.

Slide 34
The protocol hierarchy shows you the relative proportions of different traffic types.

In a healthy network, you might expect a mix of TCP (HTTP/HTTPS) and UDP (DNS/QUIC).

If you suddenly see a massive percentage of an unexpected protocol—like a huge surge in ICMP or an unusual tunnelling protocol—that is a red flag that warrants an immediate investigation.

Slide 35
A flow is a complete conversation between two endpoints, identified by a 5-tuple of source/destination IPs, ports, and the protocol.

In these visualizations, you want to look for asymmetric flows where one side is sending much more than it receives, which is a common indicator of data exfiltration.

You should also watch for unusual ports or persistent, long-duration connections that might be Command and Control channels.

Slide 36
Machine learning models cannot "read" a raw PCAP file; they need numerical data. ACAS automatically transforms raw packets into over 60 flow-based features.

These features capture essential characteristics like payload lengths and time between packets.

Good feature engineering is what allows the model to distinguish between a normal user browsing the web and a malicious script scanning for vulnerabilities.

Slide 37
For this lab, you will upload a sample file and perform a deep dive into the DPI and Feature Extraction results.

We want you to identify the largest flows, review the dominant protocols, and examine at least 10 different extracted features to see how they describe the traffic mathematically.

Slide 38
Module 3 has shown us that DPI produces the vital statistics and flow visualizations needed to spot threats. We've learned that asymmetric or long-duration flows are key indicators of trouble and that ACAS automates the complex task of extracting the features required for advanced AI analysis.

Slide 39
We are now entering Module 4, Anomaly Detection with Rule-Based and Deep Learning Approaches, where the actual detection happens.

We will look at how we combine traditional logic with modern deep learning.

Slide 40
In this lesson, you will learn to configure detection rules, understand how CNNs and Autoencoders work for cybersecurity, and learn how to interpret performance metrics like precision and recall.

We will also analyse the specific strengths and trade-offs of each approach.

Slide 41
Rule-based detection uses predefined patterns or signatures to identify known threats.

Rules work by continuously monitoring network traffic as a stream of events and checking whether specific patterns of those events match predefined conditions.

Each rule defines constraints on event types, attributes (like IPs, ports, or protocols), and sometimes timing or ordering between events.

When the observed traffic satisfies all the conditions of a rule, such as a sequence of actions happening in a specific order or multiple suspicious events occurring within a time window, the rule engine triggers an alert with a defined severity and message.

Slide 42
On this slide, you can see a rule in XML format designed to detect a Nikto web scan by looking for its specific User-Agent string in the HTTP header.

While rules like this are predictable and auditable, their main limitation is that they cannot detect unknown attacks. They only catch what you have explicitly told them to look for.

Slide 43
Convolutional Neural Networks (CNNs), which were originally made for images, are excellent at detecting spatial patterns in network features.

They use supervised learning, meaning they are trained on labelled datasets to classify traffic as either "normal" or a specific "attack type" based on patterns they have learned.

Stacked Autoencoders work differently. They are trained to "reconstruct" normal traffic. When they encounter anomalous traffic, they can't reconstruct it accurately, leading to a high reconstruction error.

This makes them fantastic at detecting novel attacks that the system has never seen before, as anything that isn't "normal" will trigger an alert.

Slide 44
As shown in the architecture on the right, ACAS begins with feature extraction using the open-source MMT tool, which parses raw network traffic in PCAP format and extracts more than 60 flow-based features.

These features include protocol field values and statistical metrics related to bytes, timing, and flow behaviour, all of which can be computed without decrypting the traffic.

The extracted features are then processed by a deep learning model that combines Stacked Autoencoders (SAEs) and Convolutional Neural Networks (CNNs).

Two SAEs are trained separately—one on normal traffic and the other on malicious traffic. Their outputs are concatenated and passed to a one-dimensional CNN, which performs the final classification.

ACAS features a flexible architecture and has demonstrated strong performance across multiple datasets.

For example, it has achieved up to 99% accuracy in botnet detection when evaluated on both public datasets and private datasets containing normal traffic as well as malicious samples collected through honeypots and attack-generation tools.

Slide 45
To know if our models are actually working, we use a confusion matrix.

Precision tells us how many of our alerts were actually real attacks, while Recall tells us how many of the total attacks we actually caught.

The F1-Score is our best metric for balanced evaluation, as it combines both precision and recall into a single number.

Slide 46
Choosing the right approach depends on your goals.

For known patterns, rules are best because they are fast and accurate.

For zero-day attacks, you need an Autoencoder. For high-value targets, we recommend a "Defence-in-Depth" strategy that combines all methods.

If you are struggling with alert fatigue, a combination of rule-based logic and Explainable AI is often the most effective solution.

Slide 47
Now it’s your turn. Select a PCAP file and run both the rule-based and ML-based detection engines.

Compare the metrics and observe which abnormal flows were detected by each method.

This will give you a firsthand look at how these different approaches complement each other.

Slide 48
To wrap up Module 4: rule-based detection is fast and precise for known threats, while CNNs and Autoencoders provide the power to find novel and complex patterns.

By using standard performance metrics, you can find the right balance for your specific security needs.

Slide 49
Our final module, Module 5, Advanced Explainable AI Analysis, addresses the biggest hurdle in AI security: the "Black Box" problem.

We will learn how to make our AI's decisions transparent.

Slide 50
In this lesson, we will explore why Explainable AI is essential for building trust in cybersecurity.

You will learn to apply SHAP for global and local explanations and LIME for individual flow analysis.

Finally, you will learn to interpret these visualizations to see exactly which features are driving the model's decisions.

Slide 51
Machine learning models, particularly deep learning models, are often referred to as “black boxes” because they can make highly accurate predictions without providing clear explanations for their decisions.

When a model simply labels a network flow as an "attack" without explaining its reasoning, the result is difficult to act upon. Security analysts need to understand why a flow was flagged in order to assess the threat and respond appropriately.

Explainable Artificial Intelligence (XAI) addresses this challenge by making model decisions more transparent and interpretable.

XAI provides the accountability and transparency required for auditing, regulatory compliance, and operational trust.

It enables analysts to justify detection decisions to stakeholders and helps prevent false positives from undermining confidence in the system.

To address both trust and accountability concerns, ACAS integrates XAI techniques that allow analysts to understand not only what was detected, but also why the model

reached its conclusion.

Slide 52
SHAP (which means **SHapley Additive exPlanations**) is based on concepts from game theory. It treats features as “players” that contribute to a “prize,” which in this case is the model's prediction.

SHAP calculates the fair contribution of each feature to the final outcome, providing a quantitative measure of feature importance.

SHAP is particularly useful for global explanations, as it helps identify which features have the greatest influence across the entire dataset.

Looking at the SHAP output example on the right, we can see that flow_packets_per_second has the highest impact on the model's predictions.

This indicates that it is the most important feature from a global perspective, contributing more to the model's decisions than any other feature across all analysed samples.

Slide 53
LIME (that stands for **Local Interpretable Model-agnostic Explanations)** focuses on explaining individual predictions rather than the overall behaviour of a model.

It works by making small perturbations to the input features of a specific network flow and observing how the model's prediction changes.

Based on these observations, LIME generates a simple and interpretable explanation showing which features contributed positively or negatively to the predicted outcome.

As shown in the LIME output example, the high packet rate and SYN packet count strongly support the "Attack" classification.

In contrast, the presence of backward packets and larger packet sizes contribute slightly against the attack prediction.

This local explanation helps analysts understand the specific factors that influenced the model's decision for a particular network flow.

Slide 54
While both are useful, they have different strengths.

SHAP is mathematically consistent and best for a deep, global understanding of the model.

LIME is generally faster and excellent for a quick explanation of a single alert. In high-stakes investigations, we often recommend using both to get the most complete picture possible.

Slide 55
For your final lab, you will select a detected anomaly and run both SHAP and LIME analyses on it.

Based on the insights you gain—such as seeing that a high packet rate was the primary driver for the alert—you will formulate a response plan for how you would handle that incident.

Slide 56
In summary, XAI transforms "black box" models into transparent tools.

By following a workflow of Detect → Explain → Validate → Act, you can use SHAP and LIME to turn AI predictions into actionable security intelligence.

Slide 57
It’s time for the Capstone, the final Hands-On Assessment. This is your opportunity to demonstrate that you can perform a complete, end-to-end network threat analysis using everything we have covered in this course.

Slide 58
You will step into the role of a SOC analyst at a mid-sized company.

Your system has captured traffic from a suspected incident.

Your task is to analyse that PCAP file and produce a report for management that identifies the threats, explains how you found them with evidence, and provides clear remediation recommendations.

Slide 59
To succeed, start with DPI, move to feature extraction, run both rule and ML-based detection, and then use LIME and SHAP to explain your findings.

Your final report should be a comprehensive summary of all these observations. For this follow the step-by-step process:

In Step 0, we will capture a small PCAP file containing attacks, such as brute-force attempts or DDoS traffic. You can use a freely available PCAP file by searching in the web or capture your own using wireshark.

In Step 1, an initial analysis and DPI is performed.

In Step 2, feature extraction analysis is performed.

Step 3, follows by carrying out the detection analysis using both rule-based and ML-based approaches.

Step 4, shows how to perform explainability analysis using both LIME and SHAP methods.
And, finally, Step 5 provides a report on all findings and observations.

Slide 60
By completing this assessment, you are proving that you have mastered the skills of a modern SOC analyst.

You have demonstrated the ability to identify malicious activity, validate it with multiple methods, and produce an evidence-based security report.

Slide 61
Now let's wrap-up and evaluate the course.

Slide 62
Congratulations! You have successfully completed the course.

You now have the skills to recognize network threats, perform deep packet inspection, and use Explainable AI to justify machine learning predictions in a real-world security environment.

Slide 63
Before you go, please take five minutes to complete our anonymous survey.

Your feedback on module clarity and the effectiveness of the labs helps us improve the ACAS tool for everyone.

Remember to apply your new skills responsibly and only on systems you are authorized to test.

Slide 64
Thank you for your attention and hard work throughout this course.

If you have any further questions or would like to learn more about our work, please reach out to us at Montimage using the contact information provided here.

Good luck with your future in cybersecurity!