SIMOS

Implementing Cisco Secure Mobility Solutions

Volume 1 **Student Guide**
Version 1.0

Part Number: 97-3334-01

https://t.me/learningnets



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS" AND AS SUCH MAY INCLUDE TYPOGRAPHICAL, GRAPHICS, OR FORMATTING ERRORS. CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

© 2014 Cisco Systems, I

https://t.me/learningnets

Welcome Students

Note Students, this letter describes important course evaluation access information.

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise that you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions. Therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. Please complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL, directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web.

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,
Cisco Systems Learning

https://t.me/learningnets

https://t.me/learningnets

Table of Contents

Course Introduction I-1

Overview I-1
Course Goal and Objectives I-2
Course Flow I-3
Additional References I-4
Your Training Curriculum I-5

Fundamentals of VPN Technologies and Cryptography 1-1

The Role of VPNs in Network Security 1-3
VPN Definition 1-4
Key Threats to WANs and Remote Access 1-6
Cisco Modular Network Architecture and VPNs 1-10
VPN Types 1-13
VPN Components 1-16
Summary 1-18
VPNs and Cryptography 1-19
Secure Communication and Cryptographic Services 1-20
Cryptographic Algorithms 1-22
Confidentiality Using Cryptographic Controls 1-25
Integrity Using Cryptographic Controls 1-28
Authentication Using Cryptographic Controls 1-32
Nonrepudiation Using Cryptographic Controls 1-34
Keys in Cryptography 1-35
Public Key Infrastructure 1-41
Next-Generation Encryption 1-48
Dependencies in Cryptographic Services 1-49
Cryptographic Controls Guidelines 1-50
Summary 1-51
Module Summary 1-53
References 1-53
Module Self-Check 1-55

Deploying Secure Site-to-Site Connectivity Solutions 2-1

Introducing Cisco Secure Site-to-Site Connectivity Solutions 2-3
Site-to-Site VPN Topologies 2-4
Site-to-Site VPN Technologies 2-6
IPsec VPN Overview 2-7
Internet Key Exchange v1 and v2 2-9
Encapsulating Security Payload 2-12
IPsec Virtual Tunnel Interface 2-13
Dynamic Multipoint VPN 2-14
Cisco IOS FlexVPN 2-15
Summary 2-17

https://t.me/learningnets

Deploying Point-to-Point IPsec VPNs on the Cisco ASA 2-19
Overview of Point-to-Point IPsec VPNs on the Cisco ASA 2-20
Configuration Tasks for Basic Point-to-Point Tunnels on the Cisco ASA 2-22
Enable IKE on an Interface 2-24
Configure IKE Policy 2-26
Configure PSKs 2-28
Choose Transform Set and VPN Peer 2-30
Choose Traffic for VPN 2-33
Configuring Site-to-Site VPN with Connection Profiles Menu 2-34
Verify and Troubleshoot Basic Point-to-Point Tunnels on the Cisco ASA 2-37
Summary 2-41
Deploying Cisco IOS VTI-Based Point-to-Point IPsec VPNs 2-43
Overview of Cisco IOS VTIs 2-44
Configuring Basic IKE Peering 2-48
Verify IKE Peering 2-52
Configure Static VTI Point-to-Point Tunnels 2-55
Verify Static VTI Point-to-Point Tunnels 2-59
Configure Dynamic VTI Point-to-Point Tunnels 2-63
Verify Dynamic VTI Point-to-Point Tunnels 2-69
Summary 2-72
Deploying Cisco IOS DMVPNs 2-73
Overview of Cisco IOS DMVPN 2-75
DMVPN Solution Components 2-76
GRE 2-78
NHRP 2-82
DMVPN Operations 2-84
Types of Authentication 2-90
Configure DMVPN on Hub 2-91
Configure DMVPN on Spoke 2-95
Configure Routing in DMVPN 2-99
Verify DMVPN 2-103
Summary 2-109
Module Summary 2-111
Module Self-Check 2-113
Deploying Cisco IOS Site-to-Site FlexVPN Solutions *3-1*
Introducing Cisco FlexVPN Solution 3-3
FlexVPN Overview 3-4
Public Key Infrastructure 3-5
FlexVPN Architecture 3-7
FlexVPN Configuration Overview 3-8
FlexVPN Capabilities 3-10
IKEv2 vs. IKEv1 Overview 3-11
IKEv2 Message Exchange 3-12
IKEv2 DoS Prevention 3-13

https://t.me/learningnets

IKEv1 and IKEv2 Comparison 3-14
FlexVPN Use Cases 3-15
Summary 3-16
Deploying Point-to-Point IPsec VPNs Using Cisco IOS FlexVPN 3-17
Point-to-Point FlexVPN 3-18
FlexVPN Configuration Blocks 3-19
IKEv2 Smart Defaults 3-20
Basic FlexVPN Scenario 1 3-22
Configure Basic FlexVPN: Scenario 1 3-24
Basic FlexVPN Scenario 2 3-25
Configure Basic FlexVPN: Scenario 2 3-26
Advanced FlexVPN Scenario 3-27
Configure Advanced FlexVPN 3-31
Verify FlexVPN Configurations 3-36
Summary 3-38
Deploying Hub-and-Spoke IPsec VPNs Using Cisco IOS FlexVPN 3-39
Cisco IOS Hub-and-Spoke FlexVPN 3-40
IKEv2 Configuration Payload 3-41
Locally Managed Hub-and-Spoke Scenario 3-42
Configure a Spoke in a Hub-and-Spoke Scenario 3-43
Configure a Hub in a Hub-and-Spoke Scenario 3-44
Configuration Exchange 3-45
Verify and Troubleshoot Hub-and-Spoke FlexVPN 3-47
Summary 3-49
Deploying Spoke-to-Spoke IPsec VPNs Using Cisco IOS FlexVPN 3-51
FlexVPN Spoke-to-Spoke Shortcut Deployments 3-52
NHRP in FlexVPN 3-53
Configure a Spoke in a Spoke-to-Spoke Shortcut Scenario 3-55
Configure a Hub in a Spoke-to-Spoke Shortcut Scenario 3-56
Verify and Troubleshoot Spoke-to-Spoke Shortcut Switching 3-57
Summary 3-59
Module Summary 3-61
References 3-61
Module Self-Check 3-63
Deploying Clientless SSL VPNs 4-1
Clientless SSL VPN Overview 4-3
Cisco Clientless SSL VPN 4-4
Cisco Clientless SSL VPN Use Cases 4-6
Cisco Clientless SSL VPN Application Access Methods 4-7
Secure Sockets Layer and Transport Layer Security 4-8
SSL Session Setup and Key Management 4-9
SSL Server Authentication 4-11
SSL Client Authentication 4-13
SSL Transmission Protection 4-14

https://t.me/learningnets

Cisco ASA Remote Access Configuration Concept 4-15
Cisco ASA Connection Profiles 4-17
Cisco ASA Group Policies 4-20
Summary 4-22
Deploying Basic Cisco Clientless SSL VPN on Cisco ASA 4-23
Basic Cisco Clientless SSL VPN 4-24
Cisco ASA SSL Server Authentication 4-26
SSL VPN Clients Authentication 4-27
Clientless SSL VPN URL Entry and Bookmarks 4-29
Configuration Scenario 4-31
Configuration Tasks 4-32
Configure Basic Cisco Clientless SSL VPN 4-34
Verify Basic Cisco Clientless SSL VPN 4-38
Summary 4-40
Deploying Application Access in Cisco ASA Clientless SSL VPN 4-41
Cisco Clientless SSL VPN Application Access Methods 4-43
Cisco Clientless SSL VPN Application Access Solution Components 4-44
Application Plug-Ins 4-45
Application Plug-Ins Available on Cisco ASA Security Appliances 4-46
Plug-In Configuration Scenario 4-47
Configure Clientless SSL VPN Plug-Ins 4-48
Verify Clientless SSL VPN Plug-Ins 4-50
Smart Tunnels 4-52
Smart Tunnel Configuration Options and Scenario 4-53
Configure Smart Tunnels 4-54
Verify Smart Tunnels 4-56
Troubleshoot Clientless SSL VPN 4-59
Summary 4-61
Deploying Advanced Authentication and Authorization in Clientless SSL VPN 4-63
User Authentication and Access Privilege Management 4-64
Client Authentication Scenario Using Local CA 4-65
Configure Local CA Client Authentication 4-66
Verify Client Authentication Using Local CA 4-67
Client Authentication Using External CA 4-69
Client Authentication and Authorization Using AAA Server 4-70
Configure Client Authentication and Authorization Using AAA Server 4-71
Verify Client Authentication and Authorization Using AAA Server 4-77
Multiple Client Authentication Scenario 4-78
Configure Multiple Client Authentication 4-79
Verify Multiple Client Authentication 4-81
Summary 4-82
Module Summary 4-83
Module Self-Check 4-85

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

https://t.me/learningnets

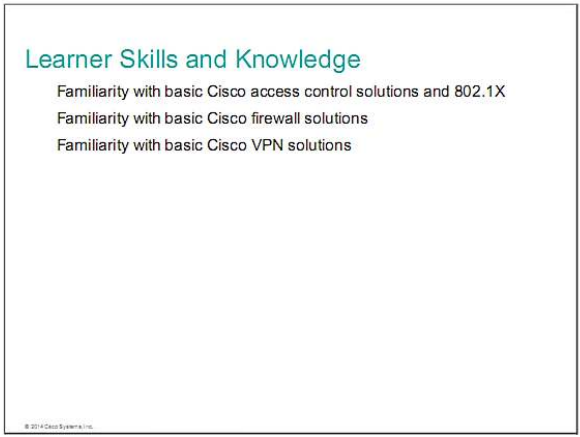
Course Introduction

Overview

This subtopic provides an overview of the course.

Learner Skills and Knowledge

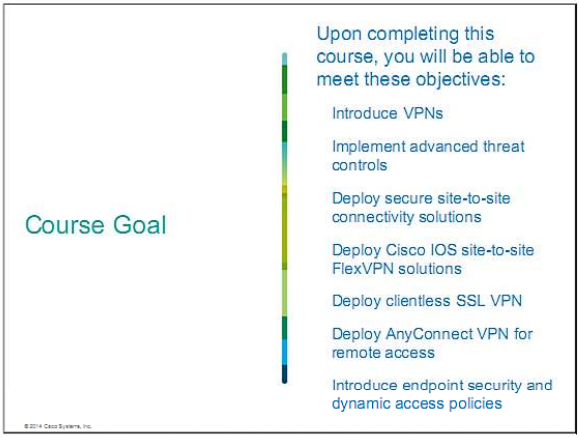
This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.



https://t.me/learningnets

Course Goal and Objectives

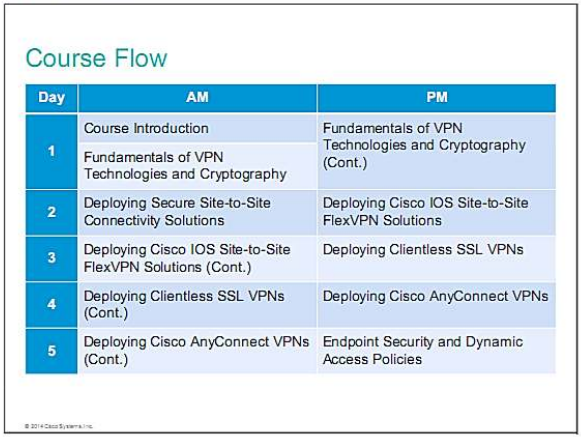
This topic describes the course goal and objectives.



https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.

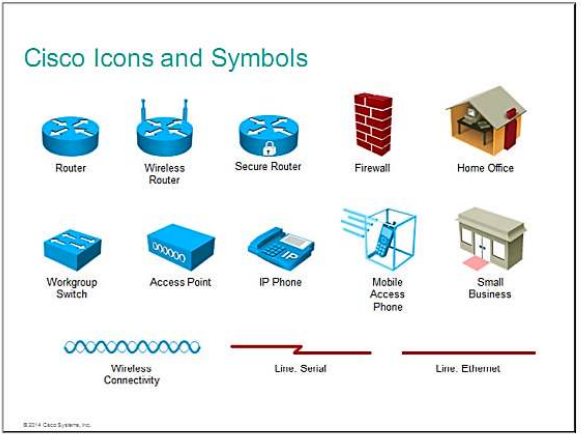


The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

https://t.me/learningnets

Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.



Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the Cisco Internetworking Terms and Acronyms glossary of terms at http://docwiki.cisco.com/wiki/Category:Internetworking_Terms_and_Acronyms_(ITA).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Your Training Curriculum

This topic presents the training curriculum for this course.



Note Cisco provides three levels of general certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused certifications for designated areas such as cable communications and security. There are many paths to Cisco certification, but only one requirement—passing one or more exams demonstrating knowledge and skill. For details, go to www.cisco.com/go/certifications.

You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, CCNP® Security and CCNP® Voice, or CCSP™). It provides a gathering place for Cisco certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit http://www.cisco.com/go/certifications.

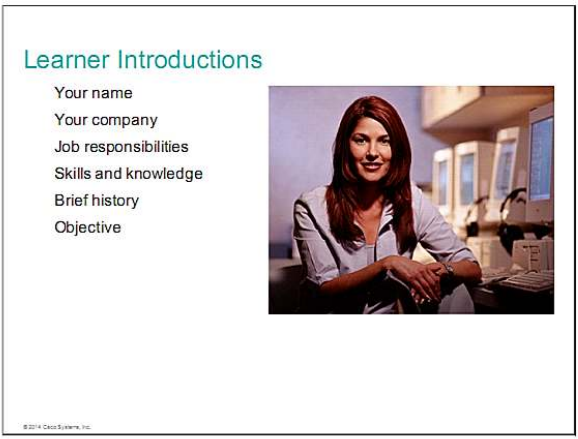
Training Curriculum

This subtopic presents the training curriculum for this course.

https://t.me/learningnets



For more information on certifications, visit http://www.cisco.com/go/certifications.



https://t.me/learningnets

Module 1

Fundamentals of VPN Technologies and Cryptography

Introduction

An IPsec VPN is an integral part of the security architecture of most organizations. It is used to connect branch offices, remote employees, and business partners to the resources of the organization. Providing confidentiality, integrity, and endpoint authentication, VPNs are ubiquitous and provide data loss prevention mechanisms for data in transit at multiple levels. From SSL VPNs to IPsec VPNs, site-to-site VPNs or remote access options, this security control is now embedded in networks and applications and should be made available in a transparent and manageable fashion. This module introduces the cryptographic elements used by VPNs, including symmetric and asymmetric algorithms, and describes the components, deployment options, and operational framework of VPN technologies.

Upon completing this module, you will be able to meet these objectives:

- Describe the role of VPNs in network security
- Describe cryptography solutions, algorithms, and protocols

https://t.me/learningnets

https://t.me/learningnets

Lesson 1

The Role of VPNs in Network Security

This lesson introduces the concepts of VPNs and why they are used. We examine issues with sending traffic across public networks and what VPNs can do to protect this traffic. The lesson introduces connection methods for VPNs, types of VPNs, things to consider when using VPNs, VPN components, VPN designs and issues in context to Cisco modular network architecture.

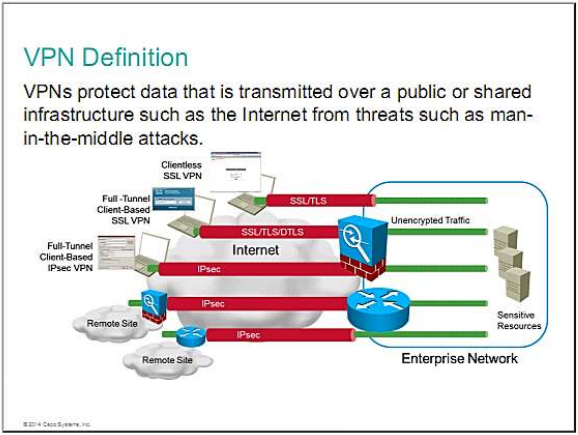
Upon completing this lesson, you will be able to meet these objectives:

- Provide an overview of VPNs.
- Describe types of attacks that expose data and company assets to various risks.
- Describe Cisco Modular Network Architecture.
- Describe different types of commercially deployed VPNs.
- Describe important components that are typically part of a VPN implementation.

https://t.me/learningnets

VPN Definition

This topic introduces the concepts of VPNs.



By definition, as per RFC 2828—a VPN is "a restricted-use, logical (i.e., artificial or simulated) computer network that is constructed from the system resources of a relatively public, physical (i.e., real) network (such as the Internet), often by using encryption (located at hosts or gateways), and often by tunneling links of the virtual network across the real network."

Stated more simplistically, a VPN can be defined as:

Virtual - Logical networks, independent of physical architecture.

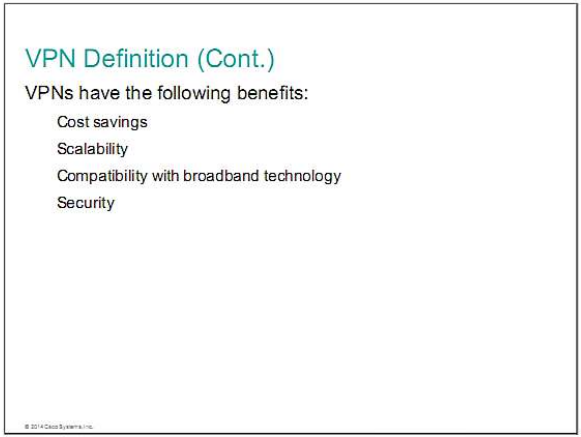
Private - Independent of IP addressing and routing schemes (noncryptographic approaches). Secure confidentiality, message integrity, authentication, privacy (cryptographic approaches).

Network - Interconnected computers, devices, and resources grouped to share information.

A VPN carries private traffic over a public or shared infrastructure (such as the Internet). The most common and effective VPN technology is applied at the network layer of the OSI model to encrypt traffic flow among specific users, applications, or IP subnet pairs. VPN at the network layer is transparent to intermediate network devices and independent of network topology.

A VPN is an integral part of the Cisco Modular Network Architecture, which aims to deliver defense-in-depth by strategically positioning Cisco products and capabilities throughout the network and by using the collaborative capabilities between the platforms.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



VPNs have many benefits:

Cost savings: VPNs enable organizations to use cost-effective third-party Internet transport to connect remote offices and remote users to the main corporate site, thus eliminating expensive dedicated WAN links and modem banks. Furthermore, with the advent of cost-effective high-bandwidth technologies, organizations can use VPNs to reduce their connectivity costs while simultaneously increasing remote connection bandwidth.

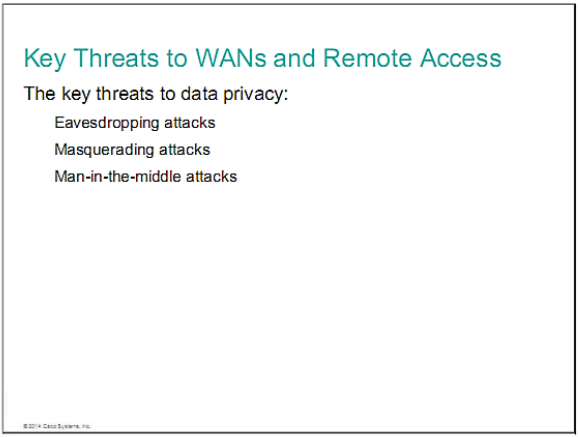
Scalability: VPNs enable corporations to use the Internet infrastructure within ISPs and devices, which makes it easy to add new users. Therefore, corporations are able to add large amounts of capacity without adding significant infrastructure.

Compatibility with broadband technology: VPNs allow mobile workers, telecommuters, and people who want to extend their workday to take advantage of high-speed, broadband connectivity, such as DSL and cable, to gain access to their corporate networks, providing workers significant flexibility and efficiency. Furthermore, high-speed broadband connections provide a cost-effective solution for connecting remote offices.

Security: Optionally, VPNs provide the highest level of security by using advanced encryption and authentication protocols that protect data from unauthorized access.

Key Threats to WANs and Remote Access

This topic describes key threats to WANs and Remote Access networks.

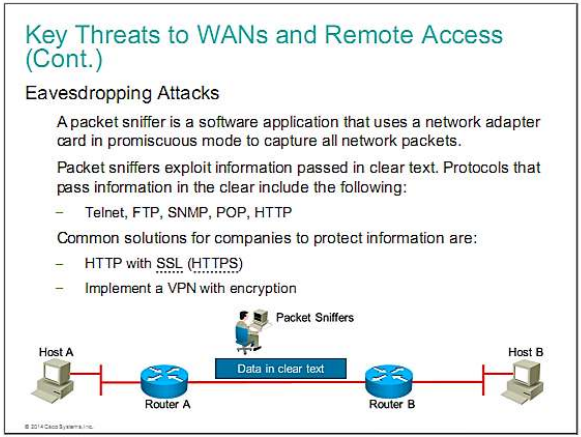


VPNs were developed initially to deal with security issues of transmitting clear text data across a network. Clear text data is information that can examined and understood by any person, including the source, destination, and anyone in between. Examples of applications that send traffic in a clear text format are Telnet, file transfers via FTP or TFTP, e-mail using the POP or SMTP, and many others. Unethical individuals, such as hackers, can take advantage of applications that send clear text data to execute the following types of attacks:

- Eavesdropping
- Masquerading
- Man-in-the-middle

Each type of attack exposes your data and company assets to various risks.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



The most common type of attack with clear text data is eavesdropping. In an eavesdropping attack, a person examines the contents of packets as they are transmitted between two devices. Some types of applications and protocols are susceptible to eavesdropping attacks, including Telnet, POP, HTTP, TFTP, FTP, SNMP, and others.

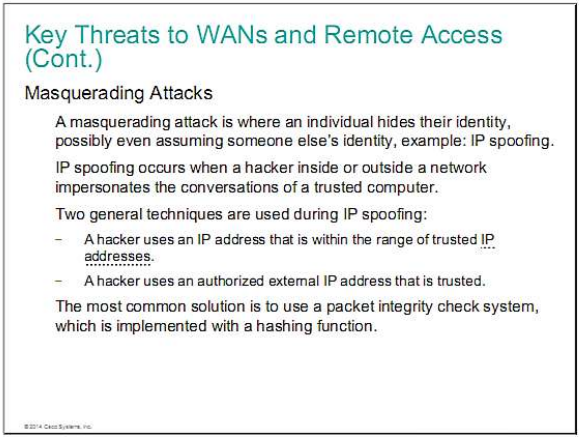
With all of these applications and protocols, authentication information, like usernames and passwords, is passed, in clear text, between two devices. A hacker can use this information to execute access and other types of attacks.

One of the most common solutions for companies to protect information is to use HTTP with SSL (HTTPS) to encrypt user-sensitive information. Another common solution for corporate and partner access is to implement a VPN with encryption. Encryption scrambles the clear text information into what appears as a random string of characters; only the destination will be able to decipher the information. Encryption can be implemented with one of the two following methods:

Link Encryption - The entire frame (such as a PPP or HDLC frame) is encrypted between two devices; this is used on point-to-point connections of directly connected devices.

Packet Payload Encryption - Only the packet payload is encrypted, which allows this form of encryption to be routed across a Layer-3 network, such as the Internet.

Encryption is commonly used for external connections that traverse public networks; however, for certain types of sensitive data, you might want to encrypt it as it traverses your intranet. Of the two solutions, packet payload encryption is the most common one you'll see used in a VPN solution. The reason packet payload encryption is most common is that in many situations, the data must transfer multiple hops, and thus packet payload encryption is the most scalable: only two devices need to handle the encryption/decryption process and the intermediate devices just route the encrypted packets.

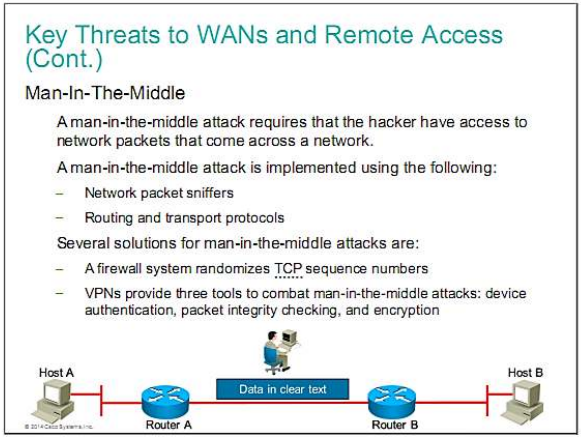


A masquerading attack is where an individual hides their identity, possibly even assuming someone else's identity. In a networking environment, this is accomplished by changing the source addressing information in packets. In the TCP/IP world, this is commonly referred to as spoofing. An attacker using spoofing typically combines this with either a DoS or unauthorized access attack.

Using a robust firewall system to restrict the kinds of packets into your network is a necessity. However, a firewall system will allow traffic from authorized external systems, even if it happens to be VPN traffic. Therefore, some type of packet authentication check is required. For example, you need to determine if packets are coming from a valid source, rather than from an attacker executing a masquerading attack.

The most common solution is to use a packet integrity check system, which is implemented with a hashing function. Hashing functions allow you to verify the source of transmitted packets. Because hashing functions use a one-way hash with a shared key, only the devices that have the key will be able to create and verify the hash values. With VPNs, the most common hashing functions used are MD5 and SHA.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



A man-in-the-middle attack can take on many forms, including the two most common:

Session replay attacks - With a session replay attack, the attacker, sitting between two devices, captures the packets from the session. The attacker will then try to use the captured packets at a later time by replaying (resending) them. The attacker's goal is to gain access to the remote system with the same packets. In some instances, the attacker will change the contents of the packets to assist in this process.

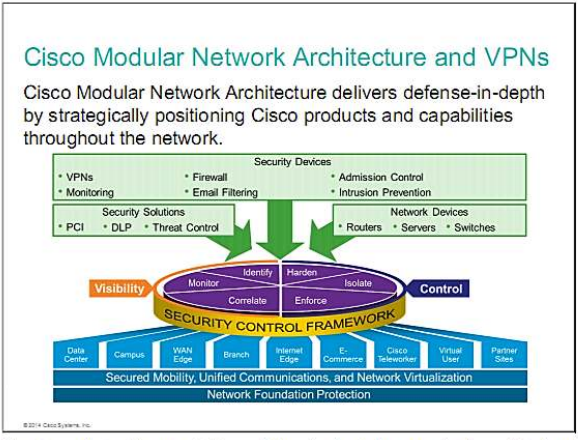
Session hijacking attacks - In a session hijacking attack, the attacker will attempt to insert himself into an existing connection and then take over the connection between two devices. To execute this attack, the attacker will have to perform masquerading, where the attacker is pretending to be the source and destination devices. Plus, the attacker must have access to the packets flowing between the source and destination devices.

There are several solutions for man-in-the-middle attacks. For instance, to prevent the hijacking of TCP sessions, you could have a firewall system randomize TCP sequence numbers, ensuring that it becomes almost impossible for the attacker to predict future sequence numbers for the session. The Cisco Adaptive Security Appliance and other available devices can perform this function. However, an attacker can use other methods to take over a session.

The best solution to this kind of problem is to use a VPN. VPNs provide three tools to combat man-in-the-middle attacks: device authentication, packet integrity checking, and encryption.

Cisco Modular Network Architecture and VPNs

This topic describes VPNs in context to Cisco Modular Network Architecture.



Cisco Modular Network Architecture delivers defense-in-depth by strategically positioning Cisco products and capabilities throughout the network and by using the collaborative capabilities between the platforms. A wide range of security technologies are deployed in multiple layers, under a common strategy and administration. Products and capabilities are positioned where they deliver the most value, while facilitating collaboration and operations.

Cisco Modular Network Architecture is delivered in the form of design blueprints and security solutions:

Design blueprints—Cisco validated designs and security best practice guides. Prescriptive design guidance is provided in CVDs which cover the various PINs present in an enterprise network, such as campus, WAN edge, branches, and data center. Design guidance is also provided for technologies such as unified communications, network virtualization, and network foundation protection, which are present in multiple places in the network. The selection of platforms and capabilities within those designs, is driven by the application of the Cisco Security Control Framework.

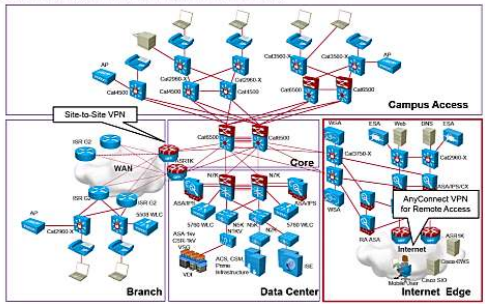
Security solutions—The Cisco Security Control Framework and the design blueprints provide the foundation for industry security solutions which address the requirements of specific industries such as retail, financial, healthcare, and manufacturing.

The Cisco security services are embedded as an intrinsic part of the architecture. The Cisco security services support the solution lifecycle and the security products included in the designs.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Modular Network Architecture and VPNs (Cont.)

Cisco Modular Network Architecture



The Cisco Modular Network Architecture follows a design where the overall network infrastructure is divided into functional modules, each one representing a distinctive PIN. Functional modules are then subdivided into more manageable and granular functional layers and blocks, each serving a specific role in the network.

Each module is carefully designed to provide service availability and resiliency, to facilitate regulatory compliance, to provide flexibility in accommodating new services and adapt with the time, and to facilitate administration.

The following is a brief description of the design modules where VPNs are placed.

Enterprise Internet Edge—The Internet edge is the network infrastructure that provides connectivity to the Internet, and that acts as the gateway for the enterprise to the rest of the cyberspace. The Internet edge services include public services DMZ, corporate Internet access and remote access VPN. The Cisco Modular Network Architecture incorporates an Internet edge design that allows users at the campuses to safely access E-mail, instant messaging, web-browsing, and other common services over the Internet. The Cisco Modular Network Architecture Internet edge design also accommodates Internet access from the branches over a centralized Internet connection at the headquarters, in case the organization's policies mandates it.

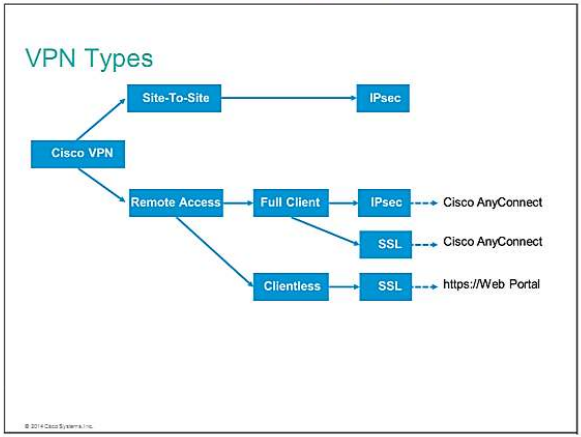
Enterprise WAN Edge—The WAN edge is the portion of the network infrastructure that aggregates the WAN links that connect geographically distant branch offices to a central site or regional hub site. The WAN can be either owned by the same enterprise or provided by a service provider, the later being the most common option. The objective of the WAN is to provide users at the branches the same network services as campus users at the central site. The Cisco Modular Network Architecture includes a WAN edge design that allows branches and remote offices to securely communicate over a private WAN. The design accommodates the implementation of multiple WAN clouds for redundancy or load balancing purposes. In addition, an Internet connection may also be used as a secondary backup option.

Enterprise Branch—Branches provide connectivity to users and devices at the remote location. They typically implement one or more LANs, and connect to the central sites via a private WAN or an Internet connection. Branches may also host local data, voice, and video services. The Cisco Modular Network Architecture includes several branch designs that allow users and devices to securely access the branch resources. The Cisco Modular Network Architecture branch designs accommodate one or two WAN clouds, as well as a backup Internet connection. Depending on the enterprise access policies, direct Internet access may be allowed while in other cases Internet access may be only permitted through a central Internet connection at the headquarters or regional office. In the later case, the Internet link at the branch would likely be used solely for WAN backup purposes.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

VPN Types

This topic describes different types of commercially deployed VPNs.



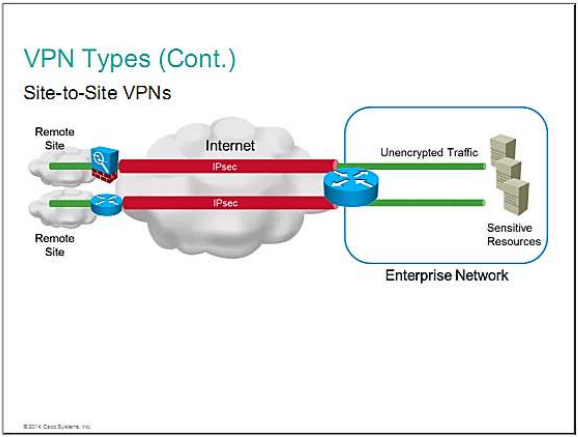
VPN are classified according to the following criteria:

Based on deployment mode - Site-to-site VPN and remote-access VPN

Based on Open Systems Interconnection layer - Layer 2 VPN (legacy protocols such as Frame Relay or ATM, and Layer 2 MPLS VPN), Layer 3 VPN (IPsec and MPLS Layer 3 VPN), and Layer 7 VPN (SSL VPN)

Based on underlying technology - IPsec VPN, SSL VPN, MPLS VPN, other Layer 2 technologies such as Frame Relay or ATM, and hybrid VPNs combining multiple technologies

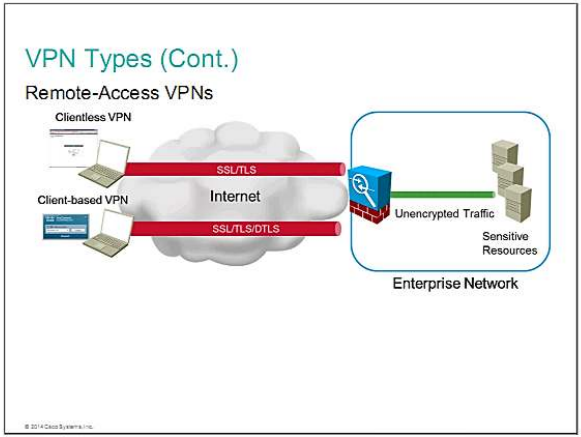
The two basic VPN deployment models typically use either IPsec or SSL technologies to keep the communications secured.



A site-to-site VPN, is an extension of a classic WAN network. Site-to-site VPNs connect entire networks to each other; for example, they can connect a branch office network to a company headquarters network. In the past, a leased line or Frame Relay connection was required to connect sites, but because most corporations now have Internet access, these connections can be replaced with site-to-site VPNs.

In a site-to-site VPN, hosts do not have Cisco VPN Client software; they send and receive normal TCP/IP traffic through a VPN "gateway," which could be a router, firewall, or Cisco Adaptive Security Appliance 5500 Series Adaptive Security Appliance. The VPN gateway is responsible for encapsulating and encrypting outbound traffic for all the traffic from a particular site and sending it through a VPN tunnel over the Internet to a peer VPN gateway at the target site. Upon receipt, the peer VPN gateway strips the headers, decrypts the content, and relays the packet toward the target host inside its private network.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

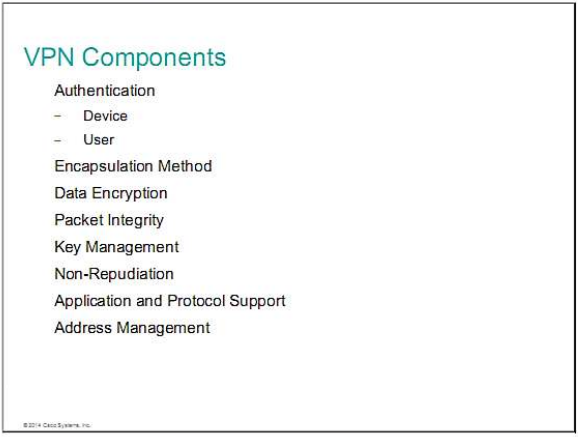


In the past, corporations supported remote users by using dial-in networks and ISDN. With the advent of VPNs, a mobile user simply needs access to the Internet to communicate with the central office. In the case of telecommuters, their Internet connectivity is typically a broadband connection such as DSL or cable. Remote-access VPNs, can support the needs of telecommuters, mobile users, and extranet consumer-to-business traffic. Remote-access VPNs connect individual hosts who must access their company network securely over the Internet.

Hosts in a remote-access VPN commonly have Cisco AnyConnect or are using Clientless SSL VPN. The Cisco AnyConnect uses a virtual network adapter in the host computer to send traffic back to the head office. The Clientless SSL VPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. Clientless SSL VPN provides easy access to a broad range of web resources and both web-enabled and legacy applications from almost any computer that can reach HTTPS Internet sites.

VPN Components

This topic describes VPN components.



There are important components that are typically part of a VPN implementation. Not every VPN implementation will include any or all of these components. Based on the requirements listed in your security policy, you might not need all of these components. You need to examine security policy to determine which VPN implementation has the necessary components to meet security policy's requirements.

VPN implementation will cover the following:

Authentication - One concern you might have is to somehow verify a device's or user's identity before allowing it to establish a VPN connection to your network. There are two general categories of authentication:

- **Device -** Device authentication allows you to restrict VPN access to your network based on authentication information that a remote VPN device provides. Typically this is one of the following two types of authentication: Pre-shared key or keys, and Digital signature or certificate.
- **User -** Many VPN implementations will add an additional layer of authentication, called user authentication, to verify whether or not a VPN connection is allowed by a user using a specific device. Normally this is employed in remote access VPNs. In some instances, depending on the VPN implementation type the VPN might perform both device and user authentication. With user authentication, the user must supply a username and a password. This password might be a static password or a one-time password (through the use of token cards).

Encapsulation Method - Another component a VPN must define is an encapsulation method: how user information, like data, is to be encapsulated and transported across a network. Encapsulation also defines what application(s) or protocol(s) can be placed in the payload of a VPN packet. Some VPN implementations will encapsulate only application layer information, whereas others can encapsulate an entire Layer-3 packet or Layer-2 frame.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Data Encryption - Data encryption is used to solve eavesdropping issues. Data encryption basically takes user data and a key value and runs it through an encryption algorithm, producing what looks like a random string of characters. Only a device with the same key value can decrypt the information. Many encryption algorithms exist, such as DES, 3DES, AES, Blowfish, RSA, IDEA, SEAL, and RC4, to name a few; however, not every VPN implementation supports all encryption algorithms. Typically, two or three algorithms are supported.

Packet Integrity - Some VPN implementations give you the option of performing packet integrity checking, or what some people commonly refer to as packet authentication. With packet authentication, a signature is attached to the packet. The signature is created by taking contents from the packet and a shared key and running this information through a hashing function, producing a fixed output, called a digital signature. This signature is then added to the original packet and the new altered packet is sent to the destination. The destination verifies the signature; and if the signature is valid, the destination will decrypt the packet contents. Verifying a hashed signature requires far fewer CPU cycles than does the decryption process. Two of the more common hashing functions used for packet integrity checking are SHA and MD5.

Key Management - Management of keys becomes important with VPN connections. For instance, how are keys derived? Are they statically configured or randomly generated? How often are keys regenerated to increase security?

Non-Repudiation - Non-repudiation can be a component of a VPN implementation. In the VPN world, non-repudiation involves two components: authentication and accounting. Authentication verifies a device's or user's identity before allowing it to establish a VPN connection to your network. Accounting is the recording of the VPN session. This could include the identities of the two devices establishing the connection, how long the connection was used, how much information was transmitted across it, what types of information traversed the connection, and so on. This can then be used later to detect access attacks and for management purposes, such as creating baselines and looking for bandwidth issues.

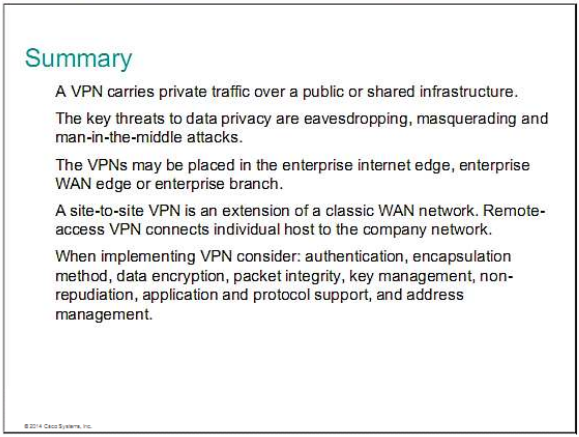
Application and Protocol Support - When choosing a VPN implementation, you'll need to first determine what kinds of traffic need to be protected. For example, if you only have IP traffic in your network, most VPN implementations will be available to you.

Address Management - Address management is an issue only with remote access connections. Keeping track of which internal address is assigned to which remote access client can be problematic. Therefore, you need to weigh this factor when choosing a VPN implementation. Some VPN implementations might lack address management, or have very poor address management, affecting your choice in a VPN implementation.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets

Lesson 2

VPNs and Cryptography

Today, the Internet provides the most efficient and commonly used information highway for communication and information exchange. With millions of people communicating on this highway, privacy has become an extremely important issue. Secure communication is becoming pivotal in every network design. For this reason, cryptography is one of the essential elements of today's information systems, providing secure access to information with greater reliability, authenticity, accuracy, and confidentiality.

This lesson provides an overview of cryptography solutions and builds foundation knowledge of cryptographic algorithms and protocols.

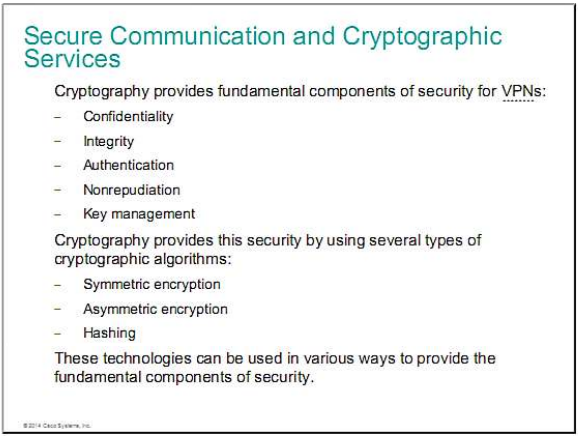
Upon completing this lesson, you will be able to meet these objectives:

- Provide an overview of cryptography.
- Describe cryptographic algorithms.
- Describes how confidentiality is achieved by using cryptographic controls.
- Describe how integrity is achieved by using cryptographic controls.
- Describe how authentication is performed by using cryptographic controls.
- Describe how nonrepudiation is performed by using cryptographic controls.
- Describes keys in context to cryptography.
- Describe Public Key Infrastructure (PKI).
- Describes next-generation encryption.
- Describe dependencies in cryptographic services.
- Describe cryptographic controls guidelines.

https://t.me/learningnets

Secure Communication and Cryptographic Services

This topic provides an overview of cryptography.



Cryptography is the study of information hiding and verification. It includes the protocols, algorithms and strategies to securely and consistently prevent or delay unauthorized access to sensitive information and enable verifiability of every component in a communication.

With a cryptographic system, the confidentiality and integrity of information can be achieved by using various methods that employ cryptography, such as encryption and decryption techniques, hash functions, digital signatures, key management techniques, and various other systems.

Cryptographic services are the foundation for many security implementations. The key services provided by cryptography are as follows:

Confidentiality: The assurance that no one can read a particular piece of data except the receivers explicitly intended.

Integrity or data authentication: The assurance that data has not been altered in transit, intentionally or unintentionally.

Peer authentication: The assurance that the other entity is who he, she, or it claims to be.

Nonrepudiation: A proof of the integrity and origin of data. The sender can't repudiate that he or she is the person who sent the data.

Key management: The generation, exchange, storage, safeguarding, use, vetting, and replacement of keys.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

VPN services use a combination of cryptographic technologies and algorithms to accomplish their goals. In the router-to-router VPN, also known as site-to-site VPN, IP packets use symmetric encryption algorithms to encrypt the payload, with keys negotiated by key management protocols. They also use asymmetric encryption algorithms to create digital signatures and authenticate the VPN peers, and use hashing functions to provide checksum-type integrity checks.

The three types of cryptographic algorithms are:

Symmetric algorithms uses a single key to encrypt and decrypt traffic.

Asymmetric algorithms also known as public-key cryptography uses two keys: one for encryption and one for decryption.

Hash functions are a form of cryptographic security that differ from encryption. Hashing condenses a message into an irreversible fixed-length value, or hash.

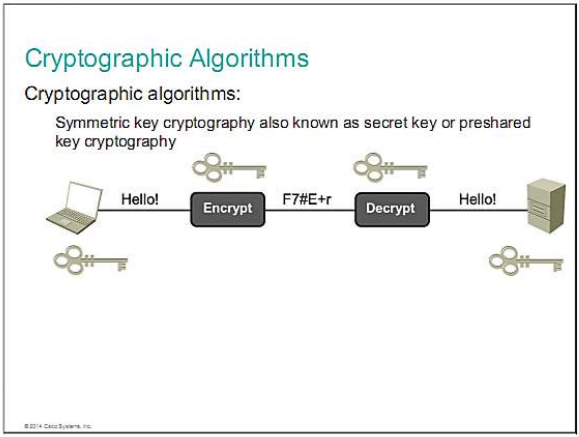
Note Encryption is the use of an algorithmic process that uses a secret key to transform plain data into a secret code, to prevent anyone except the intended recipient from accessing the information. Encryption obscures information to make it unreadable to unauthorized recipients. Encryption provides a means of secure communication over an insecure communications medium.

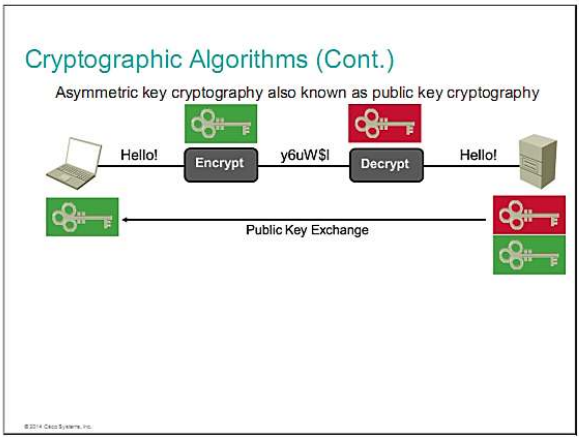
This combination of algorithms and cryptographic methods, when used as a unit to negotiate the security settings to accomplish confidentiality, integrity, and authentication, is known as a cipher suite.

https://t.me/learningnets

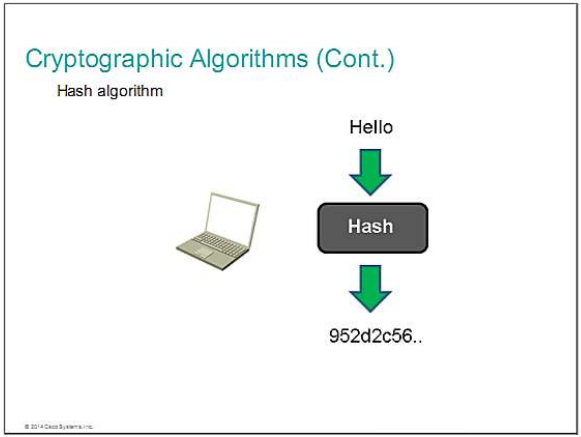
Cryptographic Algorithms

This topic describes cryptographic algorithms.





https://t.me/learningnets
© 2014 Cisco Systems, Inc.



All three types of cryptography schemes have unique function mapping to specific applications. For example, the symmetric key cryptography approach is typically used for the encryption of data providing confidentiality, whereas asymmetric key cryptography is mainly used in key exchange and nonrepudiation, thereby providing confidentiality and authentication. The hash algorithm (noncryptic), on the other hand, does not provide confidentiality but provides message integrity, and cryptographic hash algorithms provide message integrity and identity of peers during transport over insecure channels.

Cryptographic algorithms, in general, are divided into the following:

Symmetric key algorithms (also known as secret key or preshared key cryptography): Uses a single key for both the encryption and decryption process. Examples include 3DES and AES. 3DES, which consists of three sequential DES encryption-decryptions, is a legacy algorithm. This designation means 3DES provides a marginal but acceptable security level, but its keys should be renewed relatively often. Because of its small key size, DES is no longer secure and should be avoided. RC4 should be avoided as well. AES with 128-bit keys provides adequate protection for sensitive information. AES with 256-bit keys is required to protect classified information of higher importance.

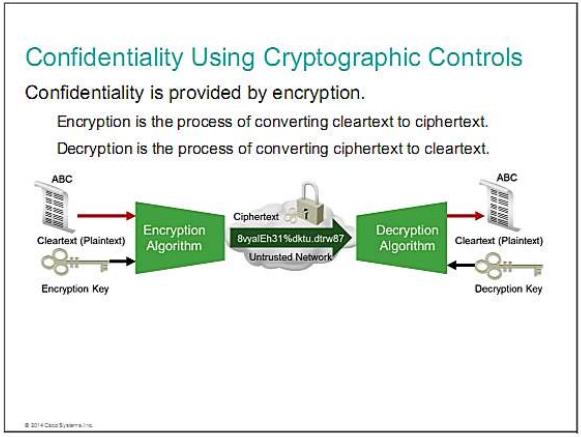
Asymmetric/Public-Key algorithms (also known as public key cryptography): Uses a two-key pair, one key for the encryption and another for the decryption process. Examples include DSA and the RSA algorithm. These keys are usually called the private key, which is secret, and the public key, which is publicly available. The private and public keys are cryptographically related. The private key cannot be derived from the public key. The private key can be used only by its owner, and the public key can be used by third parties to perform operations with the key owner. The RSA algorithms for encryption and digital signatures are less efficient at higher security levels, as is the integer-based DH algorithm. There are subexponential attacks that can be used against these algorithms. To compensate, their key sizes must be substantially increased. In practice, this means that RSA and DH are becoming less efficient every year. DH, DSA, and RSA can be used with a 2048-bit modulus to protect sensitive information. Smaller DH, DSA, and RSA key sizes, such as 768 or 1024, should be avoided.

Hash algorithms provide a constant-sized output for any input, and their most important property is irreversibility. Hash algorithms are also called digital fingerprinting algorithms. Irreversibility and collision resistance are necessary attributes for successful hash functions. Examples of hash functions are SHA-1 and SHA-256. MD5 is a hash function that is insecure and should be avoided. SHA-1 is a legacy algorithm and thus is adequately secure. SHA-256 provides adequate protection for sensitive information. On the other hand, SHA-384 is required to protect classified information of higher importance. HMAC is a construction that uses a secret key and a hash function to provide a MAC for a message. HMAC is used for integrity verification. HMAC-MD5, which uses MD5 as its hash function, is a legacy algorithm (note that MD5 as a hash function itself is not secure). It provides adequate security today, but its keys should be renewed relatively often. Alternatively, the NIST-recommended HMAC function is HMAC-SHA-1.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Confidentiality Using Cryptographic Controls

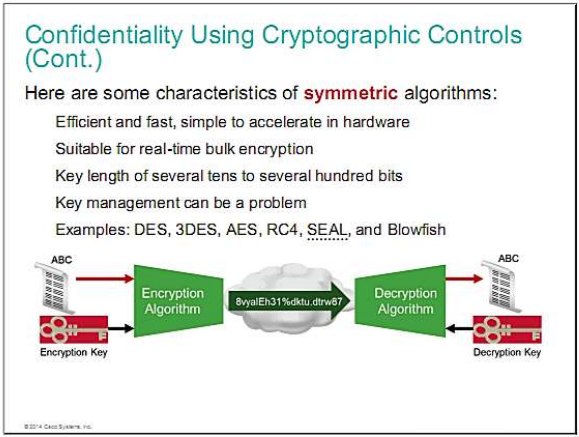
This topic describes confidentiality using cryptographic controls.



Encryption is the process of disguising a message in such a way as to hide its original contents. With encryption, plaintext (readable data) is converted to ciphertext (encrypted data). Decryption reverses the process. The purpose of encryption is to guarantee confidentiality, so that only authorized entities can read the original cleartext message.

The IPsec protocols can provide this functionality for all packets that are routed over an untrusted network. The encrypting IPsec peer takes a packet with the cleartext payload, encrypts the payload into ciphertext, and forwards the packet to the untrusted network. Its IPsec partner receives the ciphertext payload packet and decrypts the payload into the original cleartext. The two IPsec peers share the same encryption and decryption algorithm and proper keys.

The SSL and TLS protocols provide an encrypted channel on top of an existing TCP session. For example, HTTPS (HTTP over SSL and TLS) provides, among other services, confidentiality of the session between a web browser and a web server, using symmetric cryptography.



Encryption algorithms can be divided into symmetric and asymmetric algorithms, based on the manner in which keys are applied in the cryptographic process.

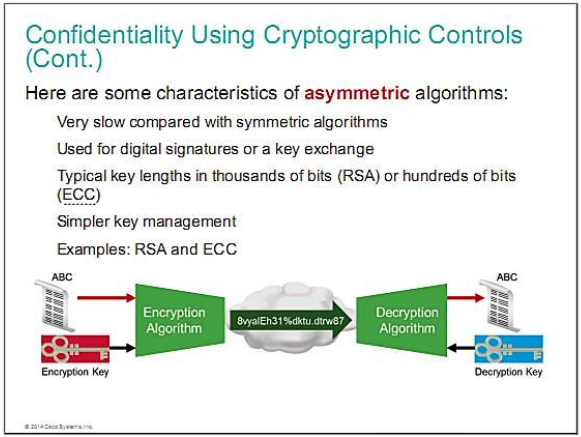
Symmetric encryption algorithms (sometimes also called secret key algorithms) are algorithms for which the encryption and decryption keys are the same. The sender and the receiver must therefore share the same secret encryption and decryption keys before communicating securely. The security of a symmetric algorithm therefore rests in the secrecy of the symmetric key; by obtaining the key anyone can encrypt and decrypt messages.

All symmetric algorithms have some important common features:

- They are, in their essence, based on simple mathematical operations.
- They are efficient in software and can be easily hardware-accelerated.
- Due to their speed, they are most often used to encrypt large volumes of data in real time.
- The key length in symmetric encryption algorithms ranges generally from several tens to several hundred bits; 128, 168, 192, and 256 bits are common examples.
- The key must be known to authorized parties only and it must be exchanged between these parties through a secure channel before any encryption can occur.

The best-known symmetric encryption algorithms are DES, 3DES, AES, the Rivest Ciphers (for example, RC4), SEAL, and Blowfish.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



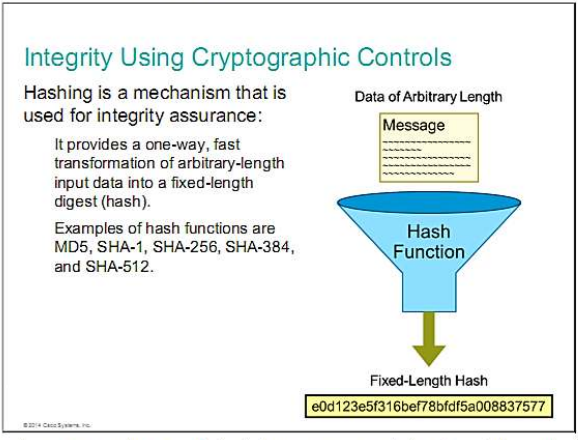
Asymmetric encryption algorithms, sometimes also called public key algorithms, are designed in such a way that the key that is used for encryption is different from the key that is used for decryption. The decryption key cannot be calculated from the encryption key in any reasonable amount of time and vice versa. It is therefore possible to advertise the encryption key in public and to simplify the key exchange.

Asymmetric encryption algorithms have some distinctive common features:

- Their design is based on difficult computational problems, such as factoring extremely large numbers or computing discrete logarithms of extremely large numbers.
- Asymmetric algorithms are relatively slow, up to 1000 times slower than symmetric algorithms.
- Due to their low performance, asymmetric encryption methods are usually used in low-volume cryptographic services such as digital signatures, authentication, or key-exchange protocols.
- The key length for asymmetric algorithms (for example, RSA) ranges generally from several hundred to several thousand bits; 768, 1024, 1536, and 2048 bits are common examples. ECC requires smaller key sizes to achieve a comparable level of security when compared to RSA, and the key size is from 160 to 521 bits.
- With asymmetric algorithms, key management tends to be simpler compared with symmetric algorithms, because, to provide confidentiality, one of the two keys (the encryption key) can usually be made public.

Integrity Using Cryptographic Controls

This topic describes integrity using cryptographic controls.

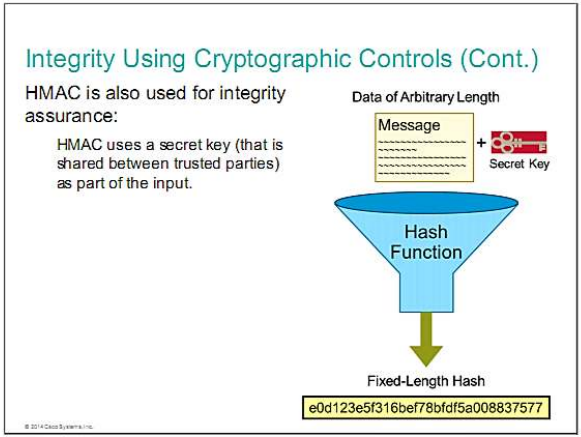


Integrity services protect against unauthorized changes to protected data. One of the mechanisms used for integrity assurance is hashing. Hashing is based on a special class of one-way mathematical functions, which should be very easy to compute, but its inverse function should be computationally infeasible.

The hashing process uses such a hash function to accept an arbitrary amount of data and produce a fixed-length (usually 128, 160, 256, or 512 bits) digest (also called a hash, fingerprint, or thumbprint) of input data. This digest is cryptographically strong; that is, it is impossible to recover input data given only its digest. Essentially, the digest resulting from hashing some data uniquely identifies that data. These functions are usually very efficient and fast.

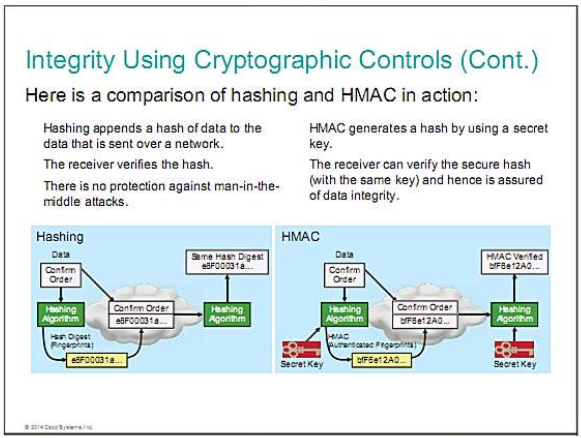
Some well-known hash functions are MD5, SHA-1, SHA-256, SHA-384, and SHA-512.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



HMAC algorithms enhance hash functions with an additional secret key as the input to the hash function. The secret key is symmetric—only the sender and the receiver of the HMAC-protected message know the secret key. Therefore, only parties who have access to that secret key can compute the digest of a keyed HMAC function.

The same hash functions are also used as HMAC algorithms when they are coupled with a secret key.

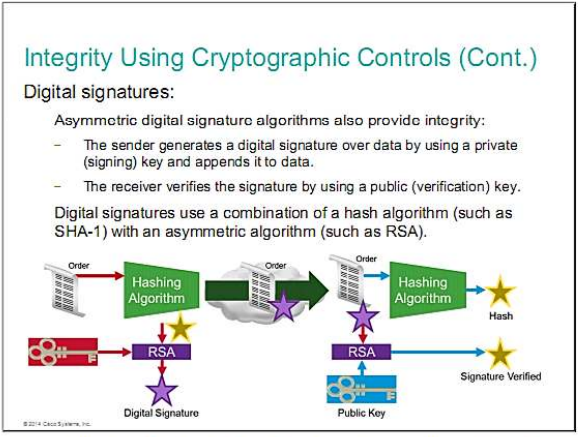


If the sender of the message wants to ensure that the messages will not be altered on their way to the receiver, there are two options: to use hash or HMAC.

When using hash, the sender uses the message as the input to a hashing algorithm and computes its fixed-length digest, or fingerprint. The sender attaches this fingerprint to the message—the message and hash are in cleartext—and sends it to the receiver. The receiver removes the fingerprint from the message and uses the message as input to the same hashing algorithm. If the hash that is computed by the receiver is equal to the one that is attached to the message, the message has not been altered during transit.

With this approach, when the message traverses the network, a potential attacker could perform a man-in-the-middle attack by intercepting the message, changing it, recalculating the hash, and appending the new hash to the message.

With HMAC, the sender takes the data and the secret key, uses a hashing algorithm, and calculates the fixed-length HMAC digest of the message. This authenticated fingerprint is attached to the message and sent to the receiver. The receiver removes the fingerprint from the message and uses the message with the secret key as the input to the same hashing function. If the fingerprint that is calculated by the receiver is equal to the received fingerprint, the message has not been altered during transit, because no one except the sender or the receiver has the secret key that is required to calculate a proper HMAC fingerprint of message data.



Another mechanism that is used for data integrity assurance is the digital signature. Digital signatures are usually based on an asymmetric cryptographic algorithm and use different keys to create a signature (to digitally sign data) and to verify a signature.

When a sender wants to sign some data, the sender uses a signature algorithm with the private (signing) key. The signer makes a hash of the document using, for example, SHA-1, and the resulting hash value uniquely identifies the document and all its contents. The signer transforms only the hash value with the asymmetric algorithm (RSA) and a private key. The transformed hash value is appended to the document and the document is sent to the receiver.

The verifier (receiver) obtains the public key of the signer. The verifier decrypts the signature with the public (verification) key of the signer. This process of decryption with the public key unveils the assumed hash value of the signer. The verifier makes a hash value of the received document, without its signature, and compares this hash to the decrypted signature hash value. If the hash values match, the document has not been changed after signing.

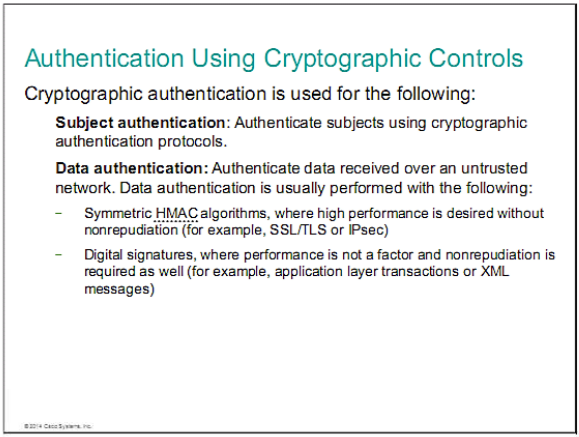
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Because digital signatures are based on asymmetric algorithms, they are relatively slow compared with symmetric encryption, hashing, and HMAC.

https://t.me/learningnets

Authentication Using Cryptographic Controls

This topic describes authentication using cryptographic controls.



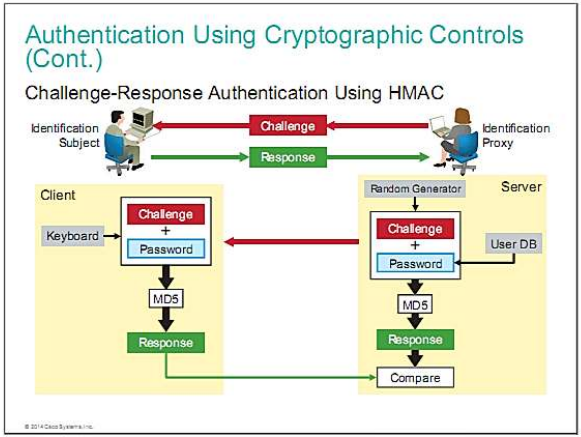
In the context of cryptography, authentication can mean the following:

Subject authentication, in which cryptographic authentication protocols provide assurance about the subject identity (for a user, host, and so forth) through a cryptographic exchange between the authenticating parties

Data authentication, in which cryptographic protection protocols provide authenticity assurance to data that is transferred over an untrusted network and in which the receiver must be able to reliably identify the sender of protected data

For this purpose, you can use HMAC and digital signature algorithms, because these algorithms provide authentication as well as integrity assurance by default. Security systems use HMAC algorithms when there is a need for high-performance authentication, such as authenticating every packet of an IPsec tunnel, and when nonrepudiation is not required. Digital signatures usually process smaller quantities of data (such as the client browser authenticating the secure web server when using HTTPS), especially when the system also needs to provide nonrepudiation services.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

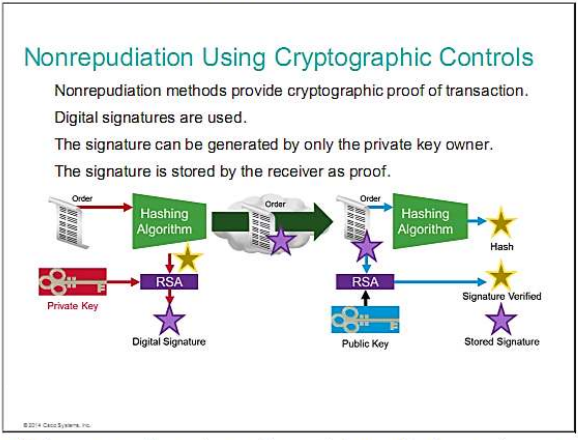


In a challenge-response authentication protocol, the user and server share a symmetric secret, which is usually the password of the user that is authenticating. The server sends a challenge, which is just a random string of characters; the subject (user) responds to the challenge, hashing the secret (the password) with the challenge by using a hash algorithm (for example, MD5), and the user sends the transformed challenge back to the server.

The server, in the meantime, has performed the same operation by using the password of the user, and the server compares the received response with the locally calculated response. If the two match, the user has proven the possession of the secret (password) and is allowed to log in.

Nonrepudiation Using Cryptographic Controls

This topic describes nonrepudiation using cryptographic controls.



The nonrepudiation cryptographic service provides proof of transaction in a security system. In essence, every message that is generated can have a proof of transaction attached to it, which a receiver can store to provide to a third party if needed.

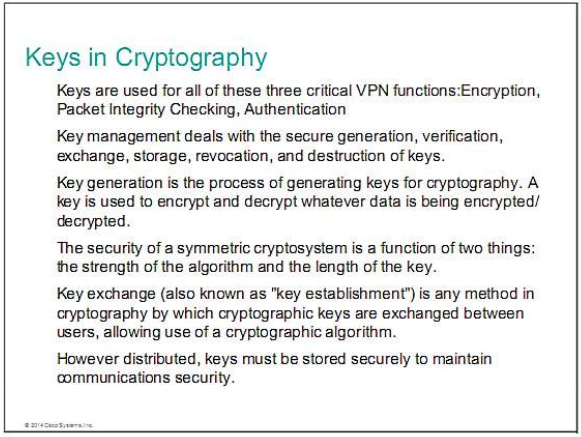
Nonrepudiation is easier to understand by looking at HMAC functions. HMAC also provides authenticity and integrity guarantees but fails to fulfill the requirements of nonrepudiation. When using HMAC, the secret key is shared between the sender and receiver. This means that the receiver is also able to generate any messages that the sender could. Therefore, taking received data and its HMAC fingerprint to a third party does not provide any proof that a specific party sent this data—the HMAC fingerprint could be generated by any of the two. When using digital signatures, both parties will have a unique private signature key that is not shared with anyone else. Because no one else is able to generate messages with this signature, nonrepudiation becomes possible.

In some countries, digital signatures are already considered to be equivalent to handwritten signatures, if certain provisions are met, such as proper protection of the certificate authority (the trusted signer of all other public keys) and the proper protection of private keys of users.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Keys in Cryptography

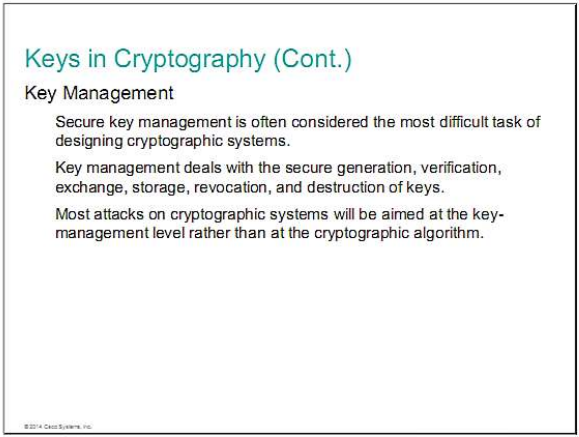
This topic describes keys in context to cryptography.



In cryptography, a key is a piece of information (a parameter) that determines the functional output of a cryptographic algorithm or cipher. Without a key, the algorithm would produce no useful result. In encryption, a key specifies the particular transformation of plaintext into ciphertext, or vice versa during decryption. Keys are also used in other cryptographic algorithms, such as digital signature schemes and message authentication codes.

In network security, keys serve a multi-functional process. For example, keys are used for all of these three critical VPN functions:

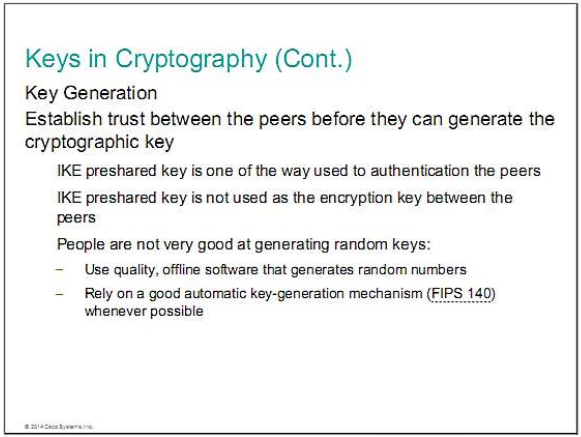
- Encryption
- Packet integrity checking
- Authentication



Key management is often considered to be the most difficult and complex part of designing a cryptographic system, and many cryptographic systems have failed due to mistakes in their key management. Modern cryptographic systems require the services of key-management procedures and rely on their strength and reliability for key generation (that is, the creation of cryptographic keys, either manually or using random-number generators), key verification (as some key values should never be used), key storage (as keys must be protected from disclosure or alteration), key exchange (when keys are exchanged over untrusted channels), key revocation (to notify users that a particular key should no longer be trusted), and key destruction.

Most attacks against cryptographic systems are usually not aimed at breaking the algorithm or by brute-force key cracking. Instead, attackers will try to uncover the complexities of the key-management system and practices, in which the human factor generally plays a significant role.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



Peers need to establish trust between them before they can generate the cryptographic key. For example, IPsec VPNs that use the IKE protocol for key management may require a PSK in all VPN devices for initial IKE peer authentication. In this scenario, the network administrator must generate a key and configure it inside a cryptographic device. IKE preshared key is one of the way used to authentication the peers and it is not used as the encryption key between the peers.

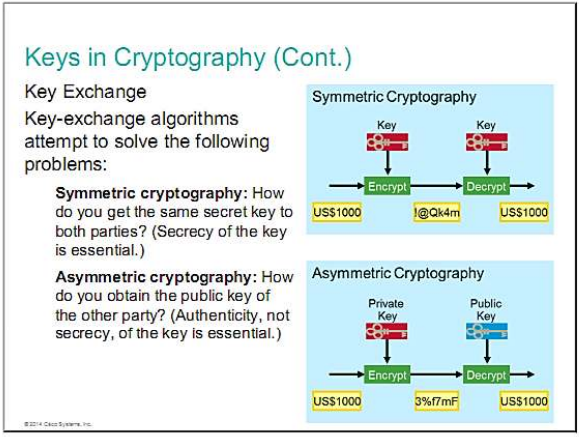
Humans are not very good at generating random numbers. When a system requires manual key configuration, it is recommended to use some quality random-number generator software and use its results as the required cryptographic keys. However, it is usually best if all key generation is left to an automatic random-number generator, if it is embedded in the cryptographic system, especially if the random-number generator has been shown to be of high quality.

These are key length recommendations:

The security of a symmetric cryptosystem is a function of two things: the strength of the algorithm and the length of the key. The former is more important, but the latter is easier to demonstrate.

In cryptography, key size or key length is the size measured in bits of the key used in a cryptographic algorithm (such as a cipher). An algorithm's key length is distinct from its cryptographic security, which is a logarithmic measure of the fastest known computational attack on the algorithm, also measured in bits. The security of an algorithm cannot exceed its key length (since any algorithm can be cracked by brute force), but it can be smaller. For example, Triple DES has a key size of 168 bits but provides at most 112 bits of security, since an attack of complexity 2112 is known. This property of Triple DES is not a weakness provided 112 bits of security is sufficient for an application. Most symmetric-key algorithms in common use are designed to have security equal to their key length. No asymmetric-key algorithms with this property are known; elliptic curve cryptography comes the closest with an effective security of roughly half its key length.

Note Encryption systems are often grouped into families. Common families include symmetric systems (e.g. AES) and asymmetric systems (e.g. RSA); they may alternatively be grouped according to the central algorithm used (e.g. elliptic curve cryptography). As each of these is of a different level of cryptographic complexity, it is usual to have different key sizes for the same level of security, depending upon the algorithm used. For example, the security available with a 1024-bit key using asymmetric RSA is considered approximately equal in security to an 80-bit key in a symmetric algorithm.



Various situations require a key exchange between parties, usually over an untrusted channel. For example, the symmetric nature of HMAC or symmetric encryption algorithms (AES, 3DES, and so forth) requires that both parties share the same secret key and requires a secure key-exchange method that preserves key confidentiality when exchanged over the untrusted channel. In public key cryptography, the public keys of subjects need to be made available to other subjects, and the origin (authenticity) and integrity of the public key must be guaranteed.

There are various key-exchange methods and protocols available, with the most common listed below:

Manual key exchange: This method of exchanging secret (symmetric) and nonsecret (public) keys is simplest, and it does not require any technical infrastructure for the exchange to occur. However, it does not scale, and it often relies on the human operator to perform the procedure securely.

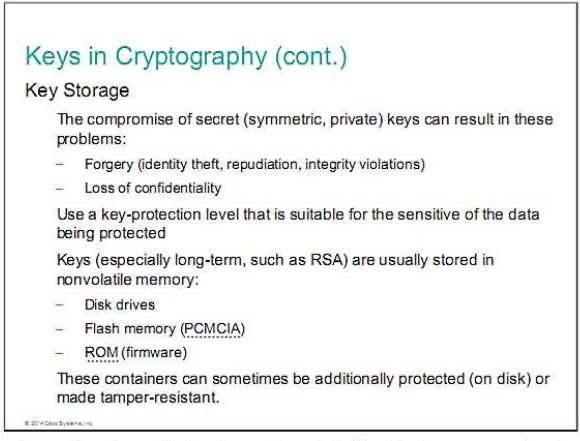
Diffie-Hellman algorithm: This method enables an automatic secure key exchange over insecure channels and is very frequently used in modern key management to provide keying material for symmetric algorithms, such as 3DES or keyed SHA-1 HMAC. It is by default vulnerable to a man-in-the-middle attack and must be authenticated to counter this threat. The algorithm is recommended for an automatic key exchange.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Asymmetric encryption: This method is usually used for a symmetric key exchange. In this case, one of the parties generates the secret key, encrypts it using the public key of its peer, and sends the peer this encrypted message. The peer can decrypt it using its private key and reveal the embedded symmetric key, which it can use to secure a channel that is protected by some symmetric algorithm. For example, by using HTTPS to provide confidentiality of the session between a web browser and a web server, the web browser generates a random secret key and encrypts it by using the public key of the web server and then sends it to the web server, which uses its private key to decrypt the random shared key. This shared key can then be used to encrypt actual HTTP traffic between the web browser and the web server.

Public key infrastructure: The PKI is a system for the creation, storage, and distribution of digital certificates that can verify that a particular public key belongs to a certain entity. The PKI creates digital certificates that map the public keys to entities, securely stores these digital certificates in a central repository, and revokes them if needed. The PKI uses a CA. The primary role of the CA is to digitally sign and publish the public key that is bound to a given entity. This action is completed with the private key of the CA, so trust in the public key of the entity relies on the trust in the validity of the CA key.

Therefore, all entities in the system must possess the public key of the CA in order to verify the CA signature inside the individual entity digital certificates. This provides a linearly scalable mechanism for public key distribution, which is recommended for use in large networks involving asymmetric cryptography.



The secure storage of cryptographic keys is very important. If an attacker can access the storage media the system is using, authorized keys can be observed or changed.

Always use a method of key protection that is suitable for the sensitivity of data these keys are protecting. For example, because disclosing the private key of the CA is a significant breach of security, the key is generally well protected through the use of encrypted files or even tamper-resistant storage devices.

Long-term keys (such as RSA) also can be stored offline to survive power cycling. Correct methods and guidelines for secure key storage in nonvolatile memory can eliminate the problem of a possible key compromise.

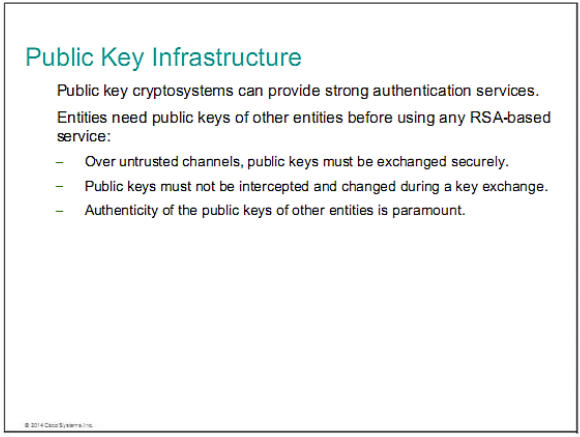
Nonvolatile storage examples are the following:

- Hard drives (for example, storing private RSA keys on a personal computer)
- Flash memory (sometimes in the form of a PCMCIA card)
- ROM (for example, encryption keys, which are hard coded in hardware)

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Public Key Infrastructure

This topic describes PKI.

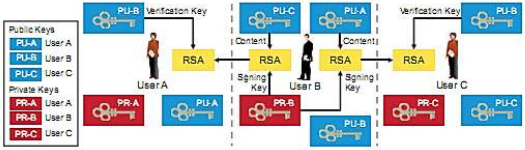


Public key cryptography, which most commonly uses the RSA algorithm, uses key pairs of public and private keys for its encryption and digital signature processes. Both procedures require an exchange of public keys between communicating entities. When verifying a digital signature, the verifier must securely obtain the public key of the signing party. If this exchange is performed over untrusted channels, the public key must not be intercepted and substituted by an attacker (a process that is known as a man-in-the-middle attack). Although the key is public information, its authenticity and integrity must be guaranteed to ensure secure transport.

Public Key Infrastructure (Cont.)

User B can act as a trusted introducer because user B is trusted by both user A and user C.

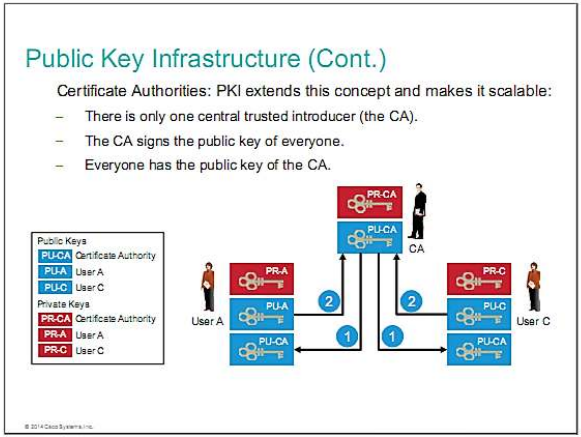
- User B signs the public key of user A with the private key of user B and sends it to user C.
- User B signs the public key of user C with the private key of user B and sends it to user A.
- Users A and C can verify the signature because they already have the public key of user B.



In the example, users A, B, and C have each created a public and a private key pair. This configuration can allow the channel between users A and B and the channel between users B and C to be used to connect users A and C directly into their own point-to-point trust relationship. Because user B has a secure channel with users A and C, user B can act as a trusted introducer. When users A and C need to securely exchange their public keys, they can do so with the help of user B, if user B can be trusted to correctly perform this procedure.

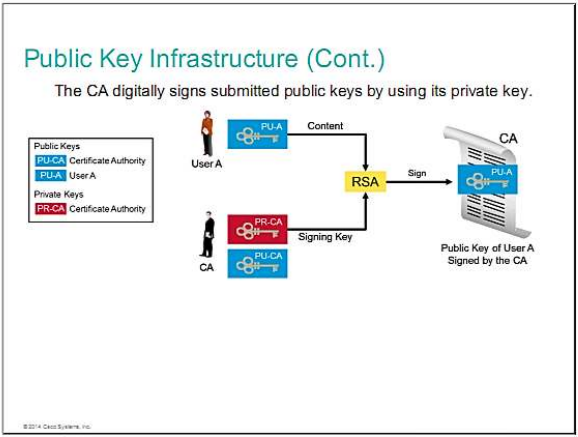
1. User B digitally signs the public key of user A (which user B has available locally) and sends it to user C.
2. User C can verify the signature of user B (because user C has the public key of user B) and will consider the public key of user A to be authentic, if user C trusts user B.
3. User B also digitally signs the public key of user C (which user B has available locally) and sends it to user A.
4. User A can verify the signature of user B (because user A has the public key of user B) and will consider the public key of user C to be authentic, if user A trusts user B.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



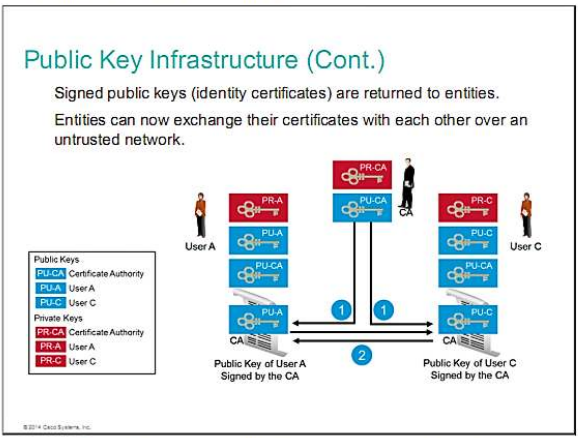
The use of a trusted third-party protocol with public key cryptography is also based on the digital signing of public keys and extends the concept of the trusted introducer. In this case, however, there is only one central trusted introducer, which is the CA. The CA signs all the public keys in its population: for example, servers, users, routers, and so on. Everybody trusts the CA; that is, everyone has the public key of the CA, which is used to verify messages from the CA. The CA is a trusted third party—someone whom all participants trust and have a secure channel with, based on digital signatures.

Every user in the system trusts the CA. This trust is accomplished by digital signing—what the CA signs is considered to be trusted. To verify future signatures of the CA, each user of this system must first obtain the public key of the CA, which is freely distributed among users. Users must obtain this public key in a secure manner and be sure that the key is authentic. This step is usually called authenticating the CA. To become a part of the system, all end users then enroll with the CA; that is, they submit their public keys and their names to the CA. These initial steps must be manually authenticated or performed over a trusted transport network.



After users submit their public keys and names, the CA verifies the identity and public key of each enrolling user. If they are correct (authentic), then the CA uses its private key to digitally sign the submitted information (name and public key).

This process creates an identity certificate for the submitter: a piece of information that binds the name of a PKI member to its public key in a standard format.



The identity certificates, which contain the public keys and name of each end user, are bound together by the signature of the CA. The certificates are returned to the entity, which installs them and uses them until they expire or are revoked.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Because every entity now has its own identity certificate—containing its name and public key, signed by the CA—and the public key of the CA, each entity can verify any data that is signed by the CA.

Now the entities can establish point-to-point relationships, independent of the CA, by exchanging information about themselves in the form of verifiable identity certificates. This ability means that end users can, after enrolling with the CA and having it sign their public keys, mutually exchange certificates over an untrusted network and use the digital signature of the CA as the protection mechanism for the public key exchange. The signature of the CA is trusted because it can be verified (the entities have the public key of the CA) and the CA and its operations are (presumably) secure.

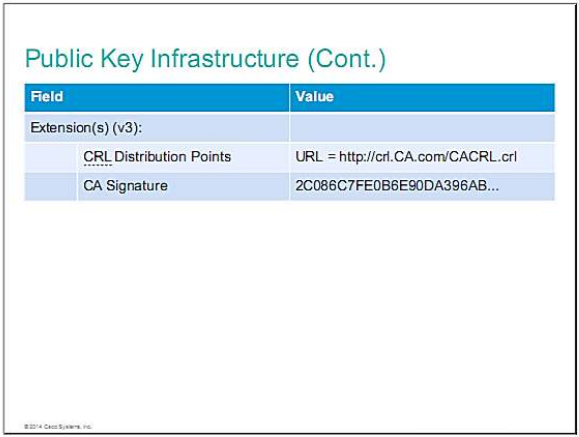
The public key of the CA is also distributed in the form of a certificate, which the CA issues. This certificate is called a self-signed CA certificate because the signer and the holder are the same entity.

Public Key Infrastructure (Cont.)

The table contains examples of the information that is in an identity certificate.

Field	Value
Certificate Format Version	Version 3
Certificate Serial Number	12457801
Signature Algorithm Identifier for CA	RSA with SHA-1
Issuer X.500 Name	C = US O = Cisco CN = CA
Validity Period	Start = 04/01/10 Expire = 04/01/15
Subject X.500 Name	C = US O = Cisco CN = CCMCluster001
Subject Public Key Information	756ECE0C9ADC7140...

© 2014 Cisco Systems, Inc.



An identity certificate is a document that binds the name of an entity to its public key. The CA signs the identity certificate so that all end entities can verify it. Certificates are not secret information and do not need to be encrypted—their only protection is a digital signature, which proves their authenticity and integrity.

The names of the certificate holder, the public key of the holder, and the digital signature of the CA are typical data that is found in an X.509 version 3 identity certificate. Other fields include the certificate serial number, certificate expiration data, algorithms that are used to generate the signature, and extensions, such as CRL distribution points.

X.509 is a well-known standard that defines basic PKI data formats, such as certificate and CRL formats, to enable basic interoperability. The infrastructure of the Internet already extensively uses this format:

- In secure web servers, for website authentication in the SSL protocol
- In web browsers, for services that implement client certificates in the SSL or TLS protocols
- In user mail agents that support mail protection using the S/MIME protocol
- In IPsec VPNs, where certificates can be used as a public key distribution mechanism for IKE RSA-based authentication

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Public Key Infrastructure (Cont.)

CRLs	OCSP	AAA Server Certificate Authorization
A list of revoked certificate serial numbers distributed as a time-stamped, CA-signed file	Revocation information is immediately pushed to an online database	Proprietary Cisco technology that is an alternative to OCSP
PKI entities regularly poll the CRL repository to receive the current CRL	Entities can query the OCSP server at any time to check for validity of the received certificate	Entities can query the AAA server at any time to check for validity of the received certificate
A window of opportunity for the attacker while the new CRL is not yet propagated	Not widely deployed	Not integrated with the PKI—requires a separate authorization database

© 2014 Cisco Systems, Inc.

A potential problem with PKI is when a key is compromised. If a certain RSA private key has been compromised, all other entities have to be signaled not to trust that key (and its corresponding public key) anymore.

The following solutions are available to solve the problem of a key compromise:

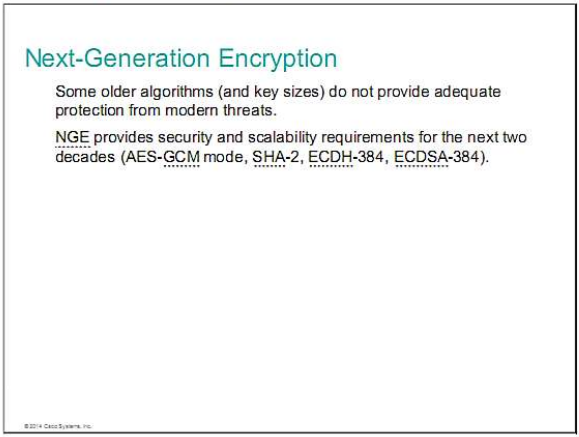
CRLs: A PKI can offer a simple solution by using a CRL, which contains a list of all certificates that are no longer valid (that is, revoked). The CRL is signed by the CA and is time-stamped with a defined lifetime. It is stored on an HTTP server, an LDAP-accessible directory, an SCEP-enabled web server, or some other available location. It is the responsibility of the PKI end user to obtain a fresh CRL after the old one has expired and to compare any certificates it wants to use against the current CRL.

OCSP: The main weakness of CRLs is the fact that they may contain stale information, as they are issued periodically, usually every couple of hours. If a key is compromised in the middle of that period, and a new CRL is issued soon afterward, there is a window of vulnerability when this compromised key can be used by the attacker, while the endpoints still use the old CRL (until it expires), which does not contain the compromised key. OCSP is a protocol for real-time verification of certificates against a database of revoked certificates. After receiving the certificate of another user, the end user or device queries the OCSP server to verify whether the received certificate has not been revoked. OCSP is not widely used yet in the network infrastructure. OCSP eliminates the weakness of CRLs, but it requires a highly available OCSP server infrastructure.

AAA-based certificate authorization: Cisco supports a proprietary method that uses the AAA infrastructure, which you can use as an alternative to OCSP. Using the TACACS+ or RADIUS protocol, some Cisco devices can query a capable AAA server to determine revocation information in real time. However, this method requires the maintenance of a separate revocation database on the AAA server, because it is not integrated with common PKI revocation mechanisms.

Next-Generation Encryption

This topic describes next-generation encryption.



Many cryptographic algorithms have been designed and used for many protocols and functions. Advances in computing systems and in the science of cryptanalysis have made it necessary to adopt newer, stronger algorithms and larger key sizes. Older algorithms are supported in current software and hardware in order to ensure backward compatibility. However, some of these older algorithms and key sizes no longer adequately protect from modern threats and should be replaced with newer ones.

In 2005, the United States National Security Agency specified a set of cryptographic algorithms that devices must support in order to meet federal standards for cryptographic strength. Collectively, these algorithms are defined as the NSA Suite B. Refer to RFC 4869 for more information about NSA Suite B cryptographic suites for IPsec.

The NSA Suite B cryptography specification is rapidly becoming a standard for cryptographic algorithm support.

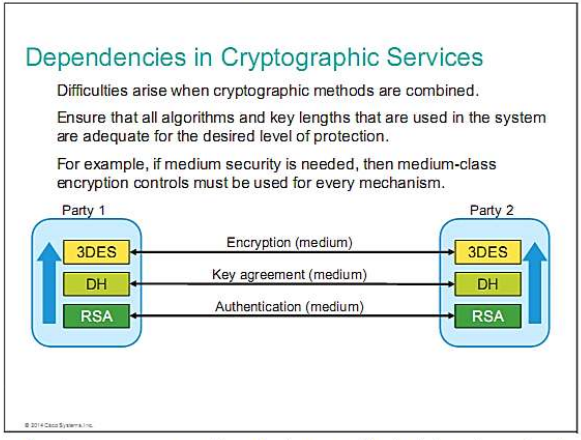
Cisco calls the same set of algorithms NGE. These types of algorithms satisfy the current and future security requirements while using algorithms that scale better.

For example, a Cisco Adaptive Security Appliance with software version 9.0 or Cisco Integrated Services Router with Cisco IOS Software Release 15.3 both provide support for all the NSA Suite B algorithms.

https://t.me/learningnets

Dependencies in Cryptographic Services

This topic describes dependencies in cryptographic services.

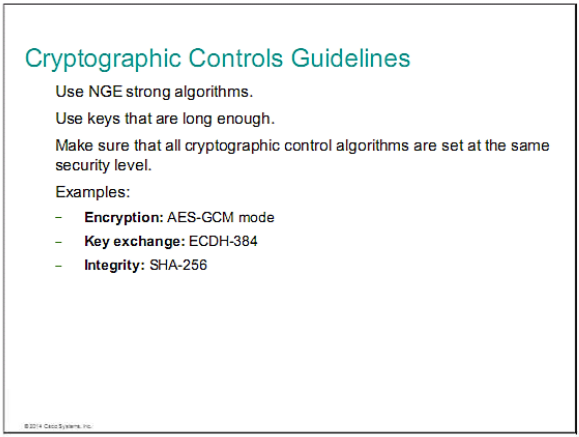


Difficulties arise when many cryptographic methods are combined and depend on each other: for example, when an asymmetric method (RSA) is used to securely transport a symmetric key (for example, RC4). When cryptographic primitives are mixed in this manner, a possibly weak symmetric or asymmetric method that uses a weak algorithm or insufficient key length may be used accidentally to support a strong symmetric or asymmetric method, weakening the overall protection of the system to the least secure algorithm and key component.

In such chains of dependency, all algorithms and their keys must be carefully matched so that no link is weaker than the desired strength of protection. A simple guideline is to use only trusted algorithms and minimum recommended key sizes with all involved methods.

Cryptographic Controls Guidelines

This topic describes cryptographic controls guidelines.



Some guidelines for modern cryptography systems are to use NGE and strong algorithms and keys that are long enough. All cryptographic control algorithms should be set at the same security level (the weakest-link principle).

NGE uses the best algorithms for current and future threats, and it is setting the industry trend. These algorithms are the best standards that can be implemented to meet the security and scalability requirements for the next decade or to interoperate with the cryptography that will be deployed in the future. The NGE algorithms are the result of more than three decades of global achievements and evolution in cryptography.

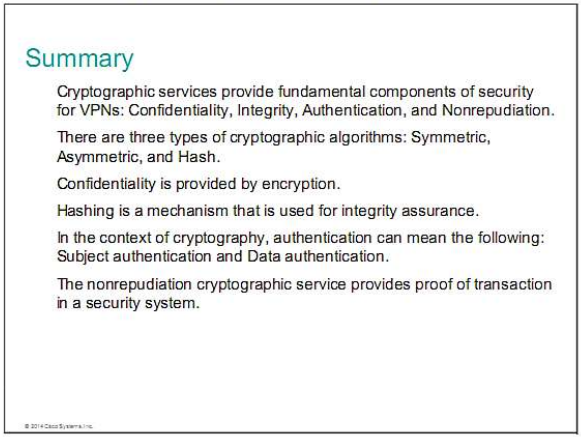
When cryptographic primitives are mixed together, a possibly weak symmetric or asymmetric method that uses a weak algorithm or insufficient key length may be used accidentally to support a strong symmetric or asymmetric method, weakening the overall protection of the system to the least secure algorithm and key component. In such chains of dependency, all algorithms and their keys must be carefully matched so that no link is weaker than the desired strength of protection.

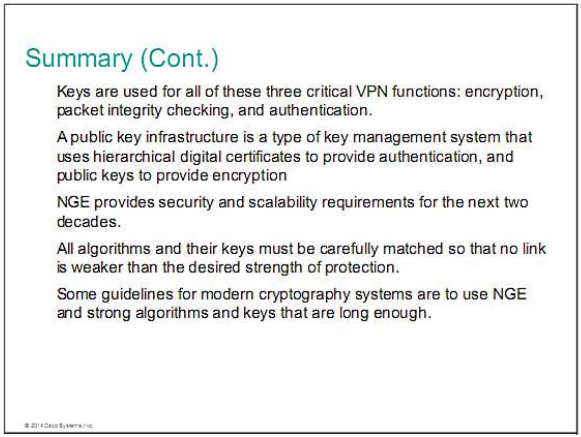
With modern algorithms, the strength of protection depends solely on the length of the key if the algorithm itself is trusted—that is, it is believed not to contain a mathematical shortcut. Choose the key length so that it protects data confidentiality or integrity for an adequate period of time.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

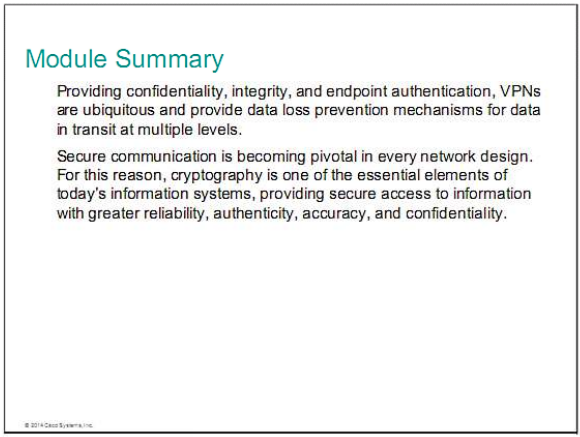




https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Summary

This topic summarizes the key points that were discussed in this module.



References

For additional information, refer to these resources:

Design Zone for Security at http://www.cisco.com/en/US/partner/netsol/ns744/networking_solutions_program_home.html

An Overview of Cryptography (Gary C. Kessler) at http://www.garykessler.net/library/crypto.html

Next Generation Encryption at http://www.cisco.com/web/about/security/intelligence/nextgen_crypto.html

https://t.me/learningnets

https://t.me/learningnets

Module Self-Check

Questions

Use the questions here to review what you learned in this module.

1. Which of the following provide protection against man-in-the-middle attacks? (Choose two.)

   A. TCP initial sequence number randomization
   B. TCP sliding-window checking
   C. Network Address Translation
   D. IPsec VPNs
   E. Secure Sockets Layer

2. Which of the following are valid characterizations of key encryption protocols? (Choose two.)

   A. asymmetric
   B. bidirectional
   C. symmetric
   D. one-way

3. Which encapsulation mode, when deployed in the tunnel mode, provides confidentiality, authenticity, integrity, and antireplay by encapsulating and protecting the entire original IP packet?

   A. Authentication Headers (AH)
   B. Internet Security Association and Key Management Protocol (ISAKMP)
   C. Diffie-Hellman key exchange with Perfect Forward Secrecy (PFS)
   D. Encapsulating Security Payload (ESP)

4. Drew recently received an email message from James. What cryptographic goal needs to be met to convince Drew that James was actually the sender of the message?

   A. integrity
   B. confidentiality
   C. availability
   D. nonrepudiation

https://t.me/learningnets

5. What is the length of the cryptographic key used in the Data Encryption Standard (DES) cryptosystem?

A. 56 bits
B. 128 bits
C. 192 bits
D. 256 bits

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Answer Key

1. D, E
2. A, C
3. D
4. D
5. A

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module 2

Deploying Secure Site-to-Site Connectivity Solutions

Introduction

Enterprises use site-to-site VPNs as a replacement for a classic routed WAN to either connect geographically dispersed sites of the same enterprise or to connect to their partners over a public network. This type of connection lowers costs while providing scalable performance. Site-to-site VPNs authenticate VPN peers and network devices that provide VPN functionality for an entire site and provide transmission security between sites over an untrusted network, such as the Internet or a MPLS WAN. This module describes secure site-to-site connectivity solutions.

Upon completing this module, you will be able to meet these objectives:

- Describe Cisco secure site-to-site connectivity solutions
- Deploy point-to-point IPsec VPNs on the Cisco ASA
- Deploy Cisco IOS VTI-based point-to-point IPsec VPNs
- Deploy Cisco IOS DMVPNs

https://t.me/learningnets

https://t.me/learningnets

Lesson 1

Introducing Cisco Secure Site-to-Site Connectivity Solutions

An IPsec VPN is a VPN that you deploy on a shared infrastructure using IPsec transmission protection technology. You can use a site-to-site IPsec VPN as an alternative to a WAN infrastructure to replace or augment existing private networks that interconnect enterprise sites. IPsec VPNs do not inherently change WAN requirements such as support for multiple WAN Layer 2 protocols, high reliability, and extensive scalability, but instead meet these requirements more cost-effectively and with greater flexibility. IPsec VPNs can use many topologies and technologies, for which this lesson provides an overview.

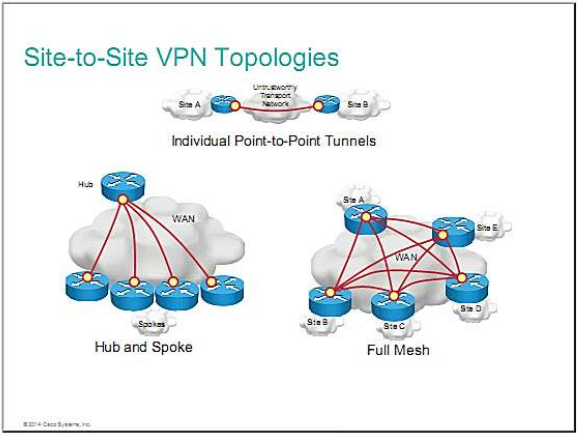
Upon completing this lesson, you will be able to meet these objectives:

- Describe site-to-site VPN topologies.
- Describe site-to-site VPN technologies
- Describe IPsec VPN
- Describe IKE version 1 and version 2
- Describe Encapsulating Security Payload
- Describe IPsec Virtual Tunnel Interface
- Describe dynamic multipoint VPN
- Describe Cisco IOS FlexVPN

https://t.me/learningnets

Site-to-Site VPN Topologies

This topic describes site-to-site VPN topologies.



A network topology is the interconnection of network nodes (typically routers) into a network. With most VPN technologies, this interconnection is largely a logical one because the physical interconnection of network devices is of no consequence to how the VPN protocols create connectivity between network users.

The three typical logical VPN topologies that are used in site-to-site VPNs are as follows:

Individual point-to-point VPN connection: Two sites interconnect using a secure VPN path. The network may include a few such individual point-to-point VPN connections that connect sites that require mutual connectivity.

Hub-and-spoke network: One central site is considered a hub and all other sites (spokes) peer exclusively with the central site devices. Typically, most of the user traffic flows between the spoke network and the hub network, although the hub may be able to act as a relay and facilitate spoke-to-spoke communication over the hub.

Fully meshed network: Every network device is connected to every other network device. This topology enables any-to-any communication; provides the most optimal, direct paths in the network; and provides the greatest flexibility to network users.

In addition to the three main VPN topologies, other more complex topologies can be created as combinations of these topologies. They include:

Partial mesh—A network in which some devices are organized in a full mesh topology, and other devices form either a hub-and-spoke or a point-to-point connection to some of the fully meshed devices. A partial mesh does not provide the level of redundancy of a full mesh topology, but it is less expensive to implement. Partial mesh topologies are generally used in peripheral networks that connect to a fully meshed backbone.

Tiered hub-and-spoke—A network of hub-and-spoke topologies in which a device can behave as a hub in one or more topologies and a spoke in other topologies. Traffic is permitted from spoke groups to their most immediate hub.

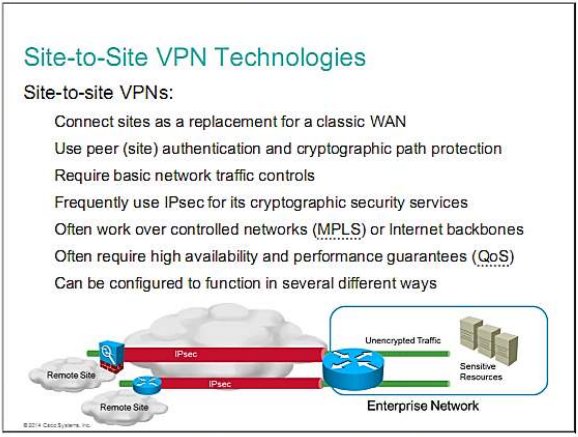
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Joined hub-and-spoke—A combination of two topologies (hub-and-spoke, point-to-point, or full mesh) that connect to form a point-to-point tunnel. For example, a joined hub-and-spoke topology could comprise two hub-and-spoke topologies, with the hubs acting as peer devices in a point-to-point topology.

https://t.me/learningnets

Site-to-Site VPN Technologies

This topic describes site-to-site VPN technologies.



Enterprises use site-to-site VPNs as a replacement for a classic routed WANs to either connect geographically dispersed sites of the same enterprise or to connect to their partners over a public network. This type of connection lowers costs while providing scalable performance. Site-to-site VPNs authenticate VPN peers and network devices that provide VPN functionality for an entire site and provide transmission security between sites over an untrusted network such as the Internet or an MPLS WAN. To control traffic that flows over site-to-site VPNs, VPN devices use basic firewall-like controls to limit connectivity and prevent traffic spoofing. These networks often work over more controlled transport networks and usually do not encounter many problems with traffic filtering in transport networks between VPN endpoints. However, because these networks provide core connectivity in an enterprise network, they often have to provide high-availability and high-performance functions to critical enterprise applications.

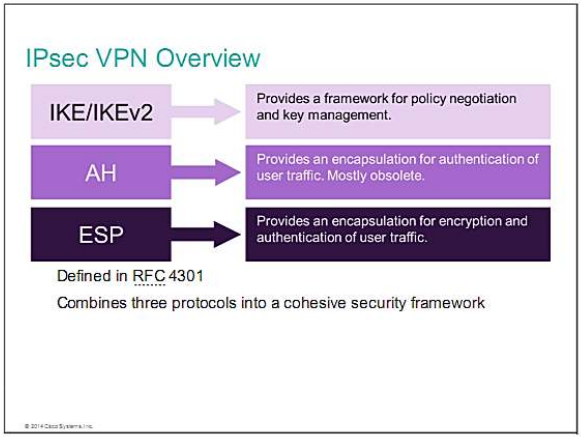
There are several site-to-site VPN solutions, each of which enables the site-to-site VPN to operate in a different way. For example, the DMVPN solution enables site-to-site VPNs without a permanent VPN connection between sites and can dynamically create IPsec tunnels. Another solution, FlexVPN, utilizes the capabilities of IKEv2. Both of these solutions are further explained later in this lesson.

The Cisco router and Cisco ASA security appliances support site-to-site full-tunnel IPsec VPNs.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IPsec VPN Overview

This topic describes IPsec VPN.



IPsec is designed to provide interoperable, high-quality, and cryptographically based transmission security to IP traffic. Defined in RFC 4301, *Security Architecture for the Internet Protocol*, IPsec offers access control, connectionless integrity, data origin authentication, protection against replays, confidentiality, and limited traffic flow confidentiality. These services are provided at the IP layer and offer protection for IP and upper-layer protocols.

IPsec provides security services at the IP layer by enabling a system that chooses required security protocols, determines the algorithm (or algorithms) to use for the service (or services), and puts in place any cryptographic keys that are required to provide the requested services. IPsec can protect one or more paths between a pair of hosts, between a pair of security gateways (usually routers or firewalls), or between a security gateway and a host.

The IPsec protocol provides IP network layer encryption and defines a new set of headers to be added to IP datagrams. These new headers are placed after the IP header and before the Layer 4 protocol (typically TCP or UDP). The new headers provide information for securing the payload of the IP packet.

IPsec combines the following security protocols:

IKE provides key management to IPsec.

AH defines a user traffic encapsulation that provides data integrity, data origin authentication, and protection against replay to user traffic.

ESP defines a user traffic encapsulation that provides data integrity, data origin authentication, protection against replays, and confidentiality to user traffic.

You can use AH and ESP independently or together, although for most applications just one of them is sufficient (ESP is preferred and AH is largely considered obsolete).

Note AH is not supported on the Cisco ASA security appliance.

https://t.me/learningnets

IPsec Security Associations

The concept of a security association is fundamental to IPsec. Both AH and ESP use security associations and a major function of IKE is to establish and maintain security associations.

A security association is a simple description of current traffic protection parameters (algorithms, keys, traffic specification, and so on) that you apply to specific user traffic flows. AH or ESP provide security services to n security association. If AH or ESP protection is applied to a traffic stream, two (or more) security associations are created to provide protection to the traffic stream. To secure typical, bidirectional communication between two hosts or between two security gateways, two security associations (one in each direction) are required.

IKE is a hybrid protocol that is defined by RFC 2408. It uses parts of several other protocols (ISAKMP, Oakley, and Skeme) to automatically establish a shared security policy and authenticated keys for services that require keys, such as IPsec. IKE creates an authenticated, secure connection (defined by a separate IKE security association that is distinct from IPsec security associations) between two entities and then negotiates the security associations on behalf of the IPsec stack. This process requires that the two entities authenticate themselves to each other and establish shared session keys that IPsec encapsulations and algorithms will use to transform cleartext user traffic into ciphertext.

Note Cisco IOS Software uses both ISAKMP and IKE to refer to the same thing. Although these two items are somewhat different, you can consider them to be equivalent.

There are several reasons to implement IKE in your IPsec configuration:

- Scalability
- Manageable manual configuration
- Security association characteristics negotiation
- Automatic key generation
- Automatic key refresh

IPsec Modes

IPsec supports two encryption modes:

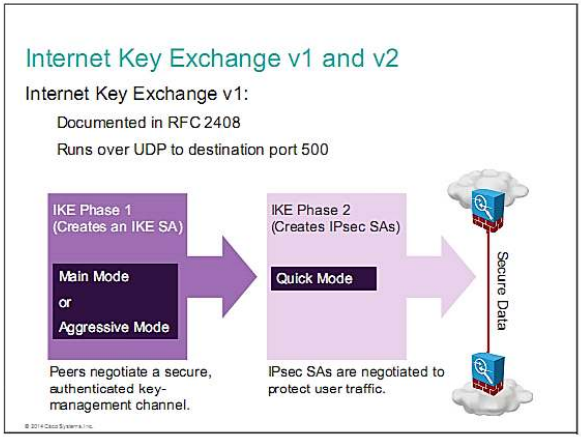
Transport mode encrypts only the data portion (payload) of each packet and leaves the packet header untouched. Transport mode is applicable to either gateway or host implementations, and provides protection for upper layer protocols as well as selected IP header fields.

Tunnel mode is more secure than transport mode because it encrypts both the payload and the header. IPsec in tunnel mode is normally used when the ultimate destination of a packet is different than the security termination point. This mode is also used in cases when the security is provided by a device that did not originate packets, as in the case of VPNs. Tunnel mode is often used in networks with unregistered IP addresses. The unregistered address can be tunneled from one gateway encryption device to another by hiding the unregistered addresses in the tunneled packet.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Internet Key Exchange v1 and v2

This topic describes IKE version 1 and version 2.



IKEv1 operates in two distinct phases:

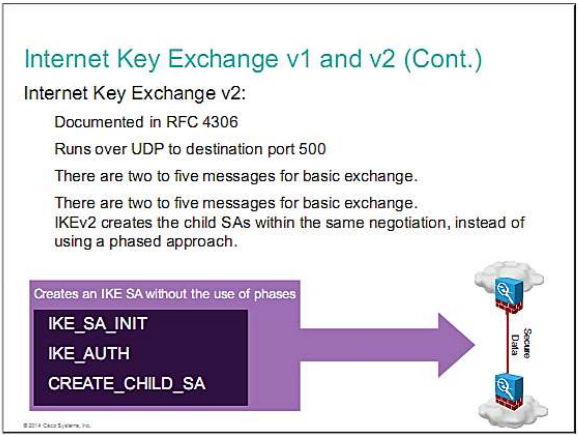
Phase 1: The two IKE peers establish a secure, authenticated channel with which to communicate, and establish shared keying material by using a DH key exchange. This channel is known as the IKE (or ISAKMP) SA. Phase 1 can operate in either main mode or aggressive mode.

- **Main mode** enables a more flexible negotiation of the IKE protection policy and always protects peer identity (an eavesdropper cannot learn about the identity of the entities that communicate by using IKE). A downside of main mode is that it does not support dynamically addressed peers when performing PSK authentication, but it supports them with PKI-facilitated authentication. The only notable exception is when you use wildcard PSKs, which is strongly discouraged.
- **Aggressive mode** does not allow much IKE protection policy negotiation and it does not protect peer identities. The names of communicating peers are sent over the untrusted network in cleartext. The major benefit of aggressive mode is that it supports PSK authentication for dynamically addressed peers, using names (and not IP addresses) to associate credentials to a particular peer.

Phase 2: In Phase 2, additional SAs are negotiated on behalf of services such as IPsec services that need key material or parameter negotiation, or both. By default, the IPsec session keys are derived from the initial keying material that is obtained in the Phase 1 DH key exchange. Optionally, they can be created by using new, independent DH key exchanges if you enable the PFS feature. This Phase 2 exchange is called IKE quick mode, which is one of two modes of IKE Phase 2. The other mode is the GDOI mode that is used by the GET VPN. Cisco IOS routers support GET VPN but the Cisco ASA security appliance does not.

PFS ensures that a given IPsec SA key was not derived from another other secret, such as another key. In the event that an attacker breaks a key, PFS ensures that the attacker is not able to derive any other key. If PFS is not enabled, an attacker can potentially break an IKE SA secret key, copy the IPsec protected data, and then use knowledge of the IKE SA secret in order to compromise the IPsec SAs set up by the IKE SA.

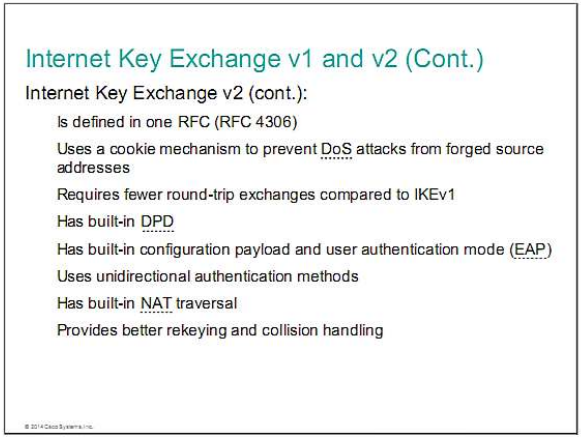
A user can reduce the risk of hackers deciphering a message through the use of larger and larger keys. But, the larger the key, the slower encryption is accomplished, and network performance also decreases. Use of fairly large keys and frequent changes of them is a good compromise. However, the challenge is coming up with ways to generate these new keys. A method to generate a new key that does not depend on the current key is needed. Then, if a hacker knows the current key, he or she will know only a small amount of information. The hacker would have to find out an entirely unrelated key to get to the next part. This concept is called PFS. The way that perfect forward secrecy is done through IKE is called "Diffie-Hellman."



Customers require continued support for remote access and site-to-site VPNs via IPsec. There is a business need to migrate existing Cisco VPN Client customers to the Cisco AnyConnect Client to eliminate the need to support multiple clients in the future but retain the client footprint in the market.

Additionally, some customers require the next-generation IKE protocol, IKEv2, to meet their future security requirements.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

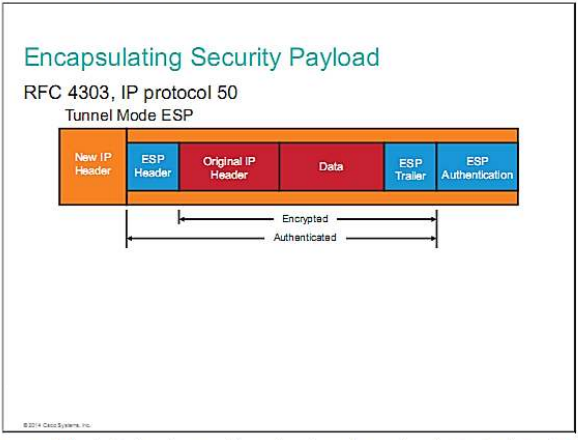


There are notable differences between IKEv1 and IKEv2. IKEv2 provides the following improvements to the IKEv1 protocol:

- Uses a Photuris-style cookie mechanism, which prevents DoS attacks from forged source addresses
- Requires fewer round trips (down to two from five for very basic exchange)
- Has built-in DPD
- Has a built-in configuration payload and user authentication mode (EAP)
- Uses unidirectional authentication methods
- Has built-in NAT traversal
- Provides better rekeying and collision handling

Encapsulating Security Payload

This topic describes Encapsulating Security Payload.



The ESP encapsulation is designed to provide a mix of security services in IPv4 and IPv6. ESP provides confidentiality, authenticity, integrity, and antireplay. To perform these functions, ESP encrypts, sequences, and authenticates the data that you want to protect. ESP then places the encrypted data in the data portion of the IP ESP payload and places the sequencing and authentication tags at the end of the ESP-encapsulated packet.

ESP is defined in RFC 4303, *IP ESP*.

Note The use of ESP increases the IP protocol processing costs in participating systems and increases the communication latency. The increased latency is primarily because of the encryption and decryption that are typically performed for each IP datagram that contains ESP.

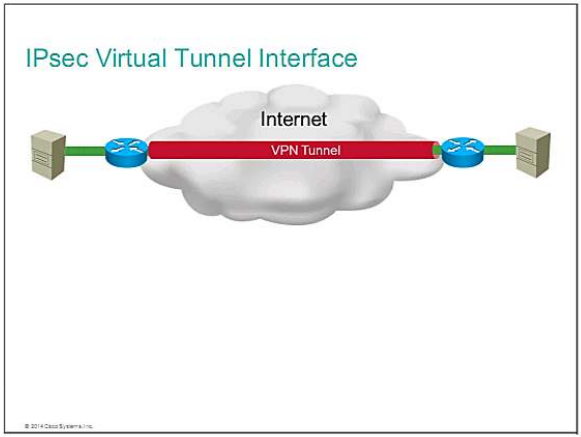
In tunnel mode, the ESP header is inserted after the tunnel IP header and before an encapsulated (Original) IP header. The IANA has assigned IP protocol 50 to ESP. The header immediately preceding an ESP header always contains the value 50 in its Next Header (IPv6) or Protocol (IPv4) field. ESP consists of an unencrypted ESP header that is followed by encrypted data.

Tunnel mode encapsulates and protects an entire IP packet. Because tunnel mode encapsulates or hides the original IP header of the packet, a new IP header must be added for the packet to be successfully forwarded. The encrypting routers own the IP addresses that are used in these new headers. Using tunnel mode leads to additional packet expansion of approximately 20 bytes that is associated with the new IP header.

https://t.me/learningnets

IPsec Virtual Tunnel Interface

This topic describes IPsec VTI.

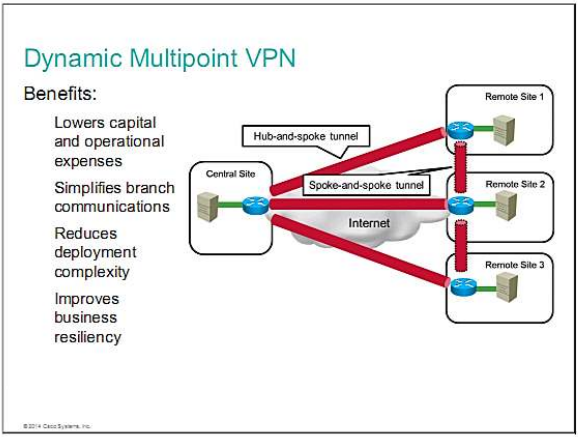


Cisco IPSec VTI is a tool that customers can use to configure IPSec-based VPNs between site-to-site devices. IPSec VTI tunnels provide a designated pathway across a shared WAN and encapsulate traffic with new packet headers, which helps to ensure delivery to specific destinations. The network is private because traffic can enter a tunnel only at an endpoint. In addition, IPSec provides true confidentiality (as does encryption) and can carry encrypted traffic.

With IPSec VTIs, users can provide highly secure connectivity for site-to-site VPNs and can be combined with Cisco AVVID to deliver converged voice, video, and data over IP networks.

Dynamic Multipoint VPN

This topic describes DMVPN.



DMVPN is a Cisco IOS Software solution for building scalable IPsec VPNs. Cisco Dynamic Multipoint Virtual Private Network uses a centralized architecture to provide easier implementation and management for deployments that require granular access controls for diverse user communities, including mobile workers, telecommuters, and extranet users.

Cisco DMVPN allows branch locations to communicate directly with each other over the public WAN or Internet, such as when using VOIP between two branch offices, but doesn't require a permanent VPN connection between sites. It enables zero-touch deployment of IPsec VPNs and improves network performance by reducing latency and jitter, while optimizing head office bandwidth utilization.

Following are benefits of Cisco DMVPN:

Lowers capital and operational expenses - Reduces costs in integrating voice, video with VPN security.

Simplifies branch communications - Enables direct branch-to-branch connectivity for business applications like voice.

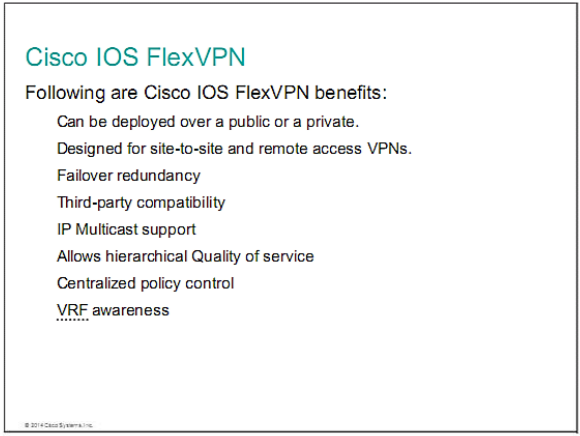
Reduces deployment complexity - Offers a zero-touch configuration, dramatically reducing the deployment complexity in VPNs.

Improves business resiliency - Prevents disruption of business-critical applications and services by incorporating routing with standards-based IPsec technology.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco IOS FlexVPN

This topic describes Cisco IOS FlexVPN.



Large customers deploying IPSec VPN over IP networks are faced with high complexity and high cost of deploying multiple types of VPN to meet different types of connectivity requirements. Customers often have to learn different types of VPNs to manage and operate different types of network. And once a technology is selected for a deployment, migrating or adding functionality to enhance the VPN is often avoided. FlexVPN was created to simplify the deployment of VPNs, to address the complexity of multiple solutions, and as a unified ecosystem to cover all types of VPN: remote access, teleworker, site to site, mobility, managed security services, and others.

As customer networks increase spans over private, public, and cloud systems, unifying the VPN technology becomes essential, and it became more important to address the need for simplification of design and configuration. Customers can dramatically increase the reach of their network without significantly expanding the complexity of the infrastructure by using Cisco IOS FlexVPN. FlexVPN is a robust, standards-based encryption technology that helps enable large organizations to securely connect branch offices and remote users and provides significant cost savings compared to supporting multiple separate types of VPN solutions such as GRE, Crypto, and VTI-based solutions. FlexVPN relies on open-standards-based IKEv2 as a security technology and provides on top of it many Cisco specific enhancements to provide high levels of security, added value, and competitive differentiations.

Cisco IOS FlexVPN is a unified VPN solution and provides the following benefits:

Transport network: FlexVPN can be deployed either over a public internet or a private MPLS VPN network.

Deployment style: Designed for the concentration of both site-to-site and remote access VPNs, one single FlexVPN deployment can accept both types of connection requests at the same time.

Failover redundancy: Three different kinds of redundancy model can be implemented with FlexVPN:

- Dynamic routing protocols over FlexVPN tunnels. Path/head-end selection is based on dynamic routing metrics.

- IKEv2-based dynamic route distribution and server clustering.
- IPsec/IKEv2 active/standby stateful failover between two chassis.

Third-party compatibility: As the IT world transitions to cloud- and mobile-based computing, more and more VPN routers and VPN endpoints from different vendors are required. The Cisco IOS FlexVPN solution provides compatibility with any IKEv2-based third-party VPN vendors, including native VPN clients from Apple iOS and Android devices.

IP Multicast support: FlexVPN natively supports IP Multicast in two ways:

- FlexVPN hub router replicates IP Multicast packets for each spoke.
- If the transport network supports native IP Multicast, the FlexVPN hub router can choose to have the transport network do multicast packet replication after IPsec encryption.

Superior QoS: The architecture of Cisco IOS FlexVPN easily allows hierarchical QoS to be integrated at the per tunnel or per SA basis:

- Per tunnel QoS for each spoke at the FlexVPN hub router.
- Per tunnel QoS dynamically applied to direct traffic between spokes.

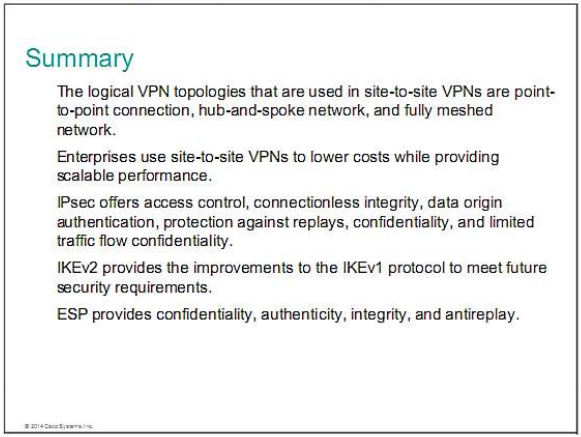
Centralized policy control: VPN dynamic policies such as split-tunnel policy, encryption network policy, VRF selection, DNS server (for remote access), and so on can be fully integrated with the AAA/RADIUS server and applied at a per peer basis.

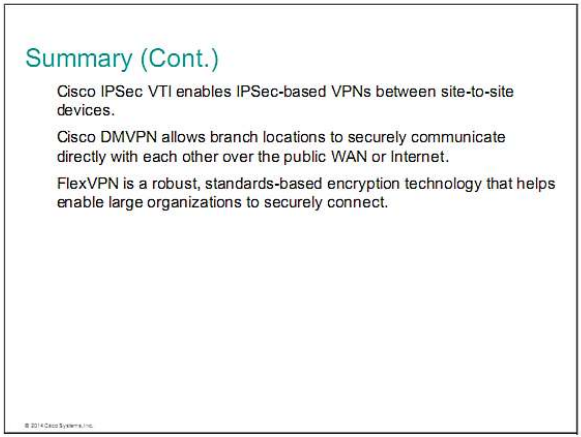
VRF awareness: The Cisco IOS FlexVPN solution can be fully integrated with MPLS VPN networks for service provider type of deployment. Both Inside VRF and front-door VRF are supported. Inside VRF assignment policy can be managed by the centralized AAA server.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.





https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 2

Deploying Point-to-Point IPsec VPNs on the Cisco ASA

The Cisco ASA adaptive security appliance supports site-to-site IPsec VPN deployments, which can be used to protect traffic between remote and central sites. Basic deployments of site-to-site IPsec VPNs use PSKs for authentication, which is recommended for smaller deployments. This lesson describes how to configure and verify basic site-to-site IPsec VPNs on the Cisco ASA appliance.

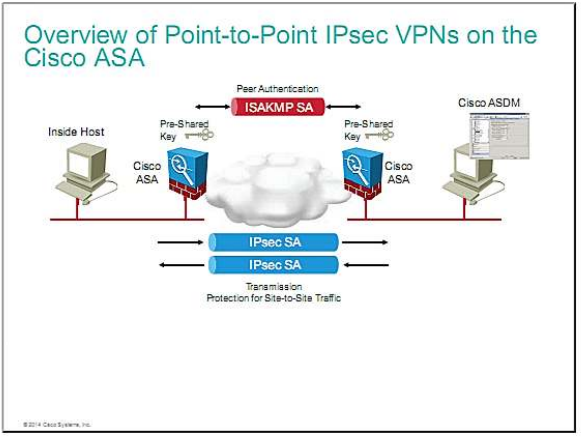
Upon completing this lesson, you will be able to meet these objectives:

- Describe point-to-point IPsec VPNs on Cisco ASA
- Describe how to configure basic point-to-point tunnels on the Cisco ASA appliance
- Describe how to enable IKE on an interface
- Describe how to add IKE policy
- Describe how to configure PSKs for IKE
- Describe how to choose transform set and VPN peer
- Describe how to choose traffic for VPN
- Describe how to enable IKE on an interface.
- Describe how to verify basic point-to-point tunnels on the Cisco ASA appliance

https://t.me/learningnets

Overview of Point-to-Point IPsec VPNs on the Cisco ASA

This topic describes point-to-point IPsec VPNs on the Cisco Adaptive Security Appliance.

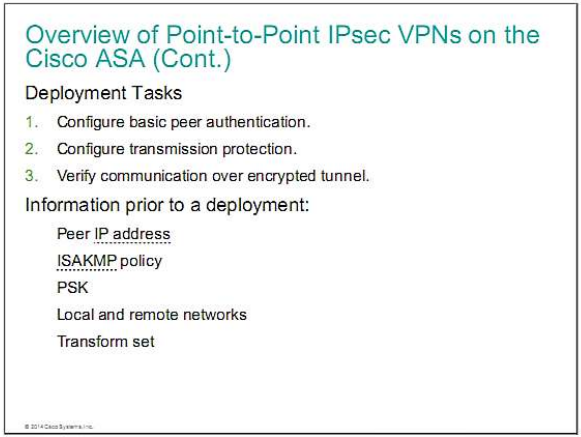


In a basic site-to-site IPsec VPN with a point-to-point topology, two local networks communicate via an encrypted connection between two VPN gateways. One security appliance is installed at each site and all traffic for remote destinations is routed to the security appliance. The security appliances encrypt and encapsulate the traffic. They also perform all IPsec functionality and route all interoffice VPN traffic through the Internet. This setup does not require the installation of any additional software.

When conducting business over a site-to-site VPN tunnel, you must "know" or have verification of who is at the other end of the tunnel. The VPN gateway on the other end of the VPN tunnel must be authenticated before the communications path is considered secure. The last exchange of IKE Phase 1 authenticates the remote VPN gateway peer.

In small networks (basic site-to-site VPNs), PSKs can be used to authenticate remote VPN peers. This authentication method is simple, but you must distribute the PSK to remote sites in advance. Using PSKs to authenticate does not scale well in large networks. The preferred method is the exchange of digital certificates to authenticate remote peers.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



To configure a basic site-to-site IPsec VPN, complete the following overall configuration tasks:

1. Configure basic peer authentication.
2. Configure transmission protection.
3. Verify communication over encrypted tunnel.

Additionally, prior to a deployment you should ensure that you have the following information:

Peer IP address: The peer IP address is the IP address of the opposite Cisco ASA security appliance or Cisco router, usually the outside interface.

ISAKMP policy: This policy provides protection for Phase 1 Internet ISAKMP and IKE SA, including encryption, hash, authentication, and DH group.

PSK: The PSK authenticates the opposite tunnel endpoint.

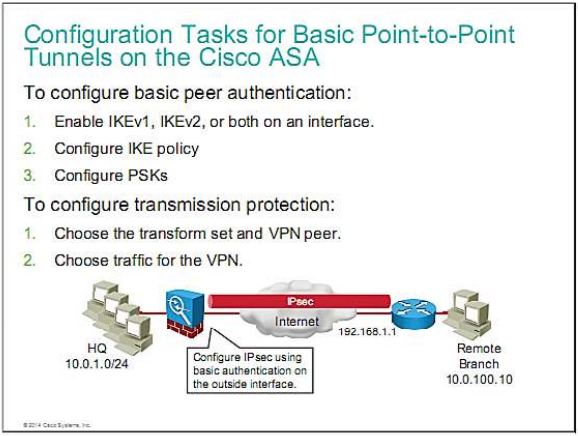
Local and remote networks: This information defines traffic that needs to be encrypted and sent through the tunnel.

Transform set: The transform set provides protection for Phase 2 IPsec SA, including encryption and hashing algorithm (and optional parameters).

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Configuration Tasks for Basic Point-to-Point Tunnels on the Cisco ASA

This topic describes how to configure basic point-to-point tunnels on the Cisco Adaptive Security Appliance.



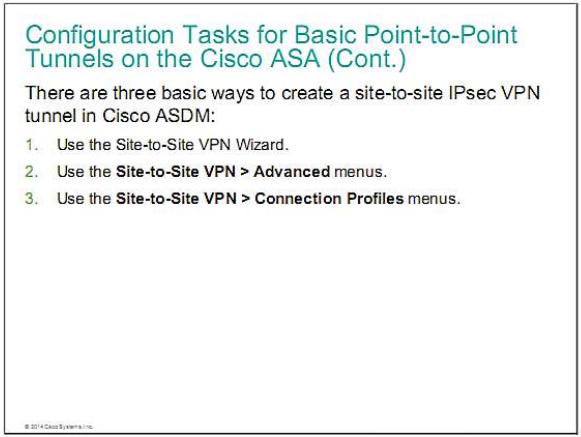
To configure basic peer authentication, complete the following tasks:

1. You must enable IKE for each interface that you want to use for VPN connections.
2. Configure the IKE policy.
3. If you are using PSKs, the next task is to configure the PSKs for the connection. You configure the PSK on the tunnel group. You can add or edit an IPsec site-to-site tunnel group to specify attributes for the IPsec site-to-site connection. In addition, you can choose the IKE peer and user authentication parameters, configure IKE keepalive monitoring, and select the default group policy.

To configure transmission protection, complete the following tasks:

1. Choose the transform set and VPN peer.
2. Choose traffic for the VPN.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



There are three basic ways to add a site-to-site IPsec VPN tunnel by using Cisco Adaptive Security Device Manager:

1. Use the Site-to-Site VPN Wizard. (not covered in this course)
2. Use the **Site-to-Site VPN > Advanced** menus.
3. Use the **Site-to-Site VPN > Connection Profiles** menus. (This option is often the preferred method.)

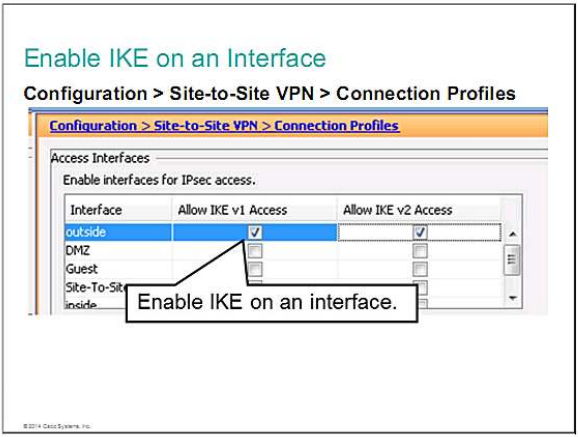
Note	Options two and three both require you to enable IKE on an interface in the **Site-to-Site VPN > Connection Profiles** menu.

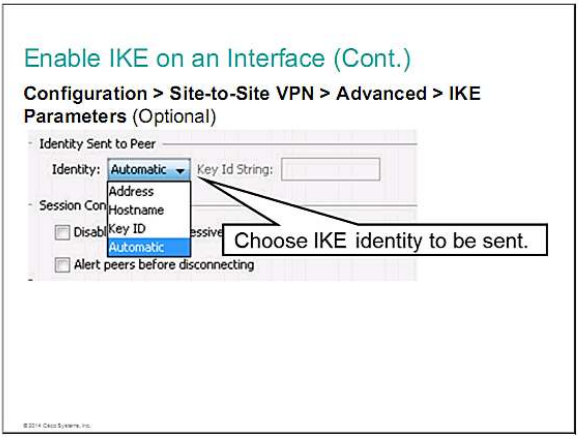
Using the Cisco ASDM **Site-to-Site VPN > Advanced** menus to configure a site-to-site VPN involves the creation of a tunnel group, whereas using the **Site-to-Site VPN > Connection Profiles** menus involves the creation of a connection profile. The terms tunnel group and connection profile are often used interchangeably. Both tunnel groups and connection profiles consist of a set of records that determines tunnel connection policies. These records identify the servers to which the tunnel user is authenticated, as well as the accounting servers, if any, to which connection information is sent. They also identify a default group policy for the connection, and they contain protocol-specific connection parameters.

This lesson first explains how to configure a site-to-site VPN by using the ASDM **Site-to-Site VPN > Advanced** menus. It then explains how to configure a site-to-site VPN by using the ASDM **Site-to-Site VPN > Connection Profiles** menu.

Enable IKE on an Interface

This topic describes how to enable IKE on an interface.





You must enable IKE for each interface that you want to use for VPN connections. Follow these steps to enable IKE on an interface using Cisco Adaptive Security Device Manager:

1. Inside Cisco ASDM, choose **Configuration > Site-to-Site VPN > Connection Profiles**. The Connection Profiles window appears.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

2. Check the **Allow IKEv1 Access** check box, the **Allow IKEv2 Access** check box, or both on an interface to enable the desired IKE versions.
3. Optionally, you can change how IPsec peers identify themselves to each other. Choose **Configuration > Site-to-Site VPN > Advanced > IKE Parameters**. Choose one of the following identification methods:
   - **Address:** Uses the IP addresses of the hosts exchanging ISAKMP identity information.
   - **Hostname:** Uses the FQDN of the hosts that are exchanging ISAKMP identity information. This name comprises the hostname and the domain name. This option is the default.
   - **Key ID:** Uses the string that the remote peer uses to look up the PSK.
   - **Automatic:** Determines IKE negotiation by connection type:

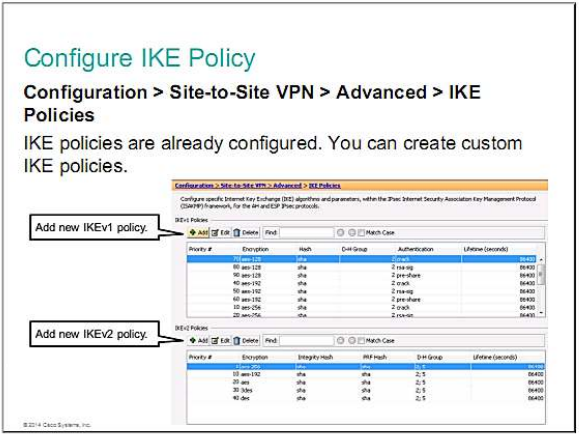
     IP address for PSK

     Certificate DN for certificate authentication
4. Click **Apply** to apply the configuration.

https://t.me/learningnets

Configure IKE Policy

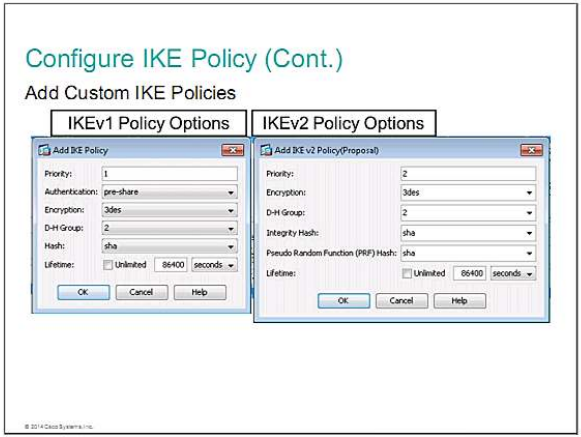
This topic explains how to add an IKE Policy using Cisco Adaptive Security Device Manager **Site-to-Site VPN > Advanced** menus.



Follow these steps to add an IKE policy using Cisco ASDM:

1. Choose **Configuration > Site-to-Site VPN > Advanced > IKE Policies**. The IKE Policies window lists all of the active IKEv1 and IKEv2 policies and their priorities.
2. Click **Add** for the type of IKE policy you want to add (IKEv1 or IKEv2). The Add IKE Policy or the Add IKEv2 Policy (Proposal) window appears.

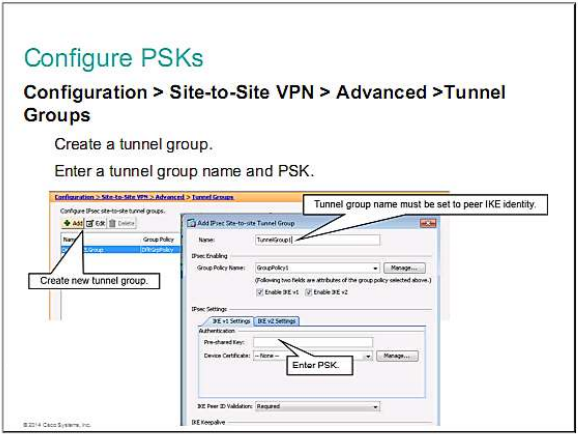
https://t.me/learningnets
© 2014 Cisco Systems, Inc.



2. Configure the desired values for the IKE policy, based on the policy type.
3. Click **OK**.
4. Click **Apply** to apply the configuration.

Configure PSKs

This topic describes how to configure PSKs by using Cisco Adaptive Security Device Manager **Site-to-Site VPN > Advanced** menus.



The security appliance provides two default tunnel groups, one for LAN-to-LAN connections, and one for remote access connections. You can modify these default tunnel groups, but you cannot delete them. You can also create one or more tunnel groups specific to your environment. Tunnel groups are local to the security appliance and are not configurable on external servers.

Tunnel groups specify the following attributes:

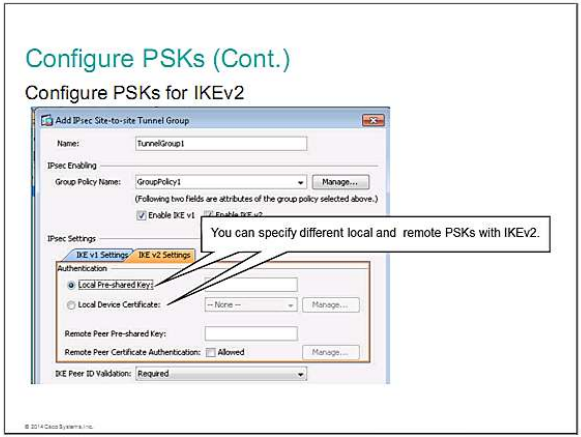
- General parameters
- IPSec connection parameters

You configure the PSK on the tunnel group. You can add or edit an IPsec site-to-site tunnel group to specify attributes for the IPsec site-to-site connection. In addition, you can choose the IKE peer and user authentication parameters, configure IKE keepalive monitoring, and select the default group policy.

Follow these steps to add an IPsec site-to-site tunnel group using Cisco ASDM:

1. Choose **Configuration > Site-to-Site VPN > Advanced > Tunnel Groups**.
2. Click **Add** to add a new tunnel group. The Add IPsec Site-to-Site Tunnel Group window appears.
3. Enter the tunnel group name in the Name field.
4. For IKEv1, click the **IKEv1 Settings** tab and enter the PSK into the Pre-shared Key field.
5. Click **OK**.
6. Click **Apply** to apply the configuration.

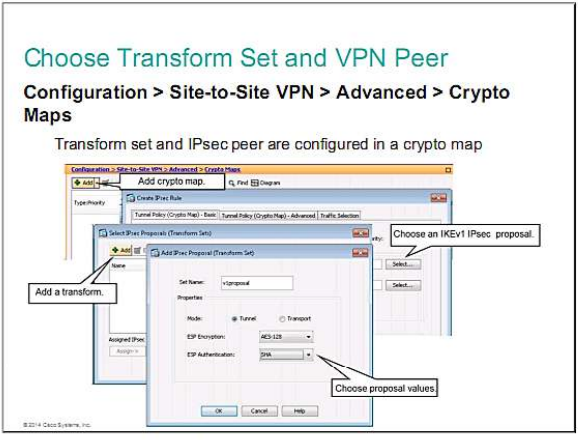
https://t.me/learningnets



For IKEv2, enter the PSK for the local as well as the remote into the Local and Remote Pre-shared Key fields of the IKEv2 Settings tab.

Choose Transform Set and VPN Peer

This topic describes how to choose a transform set and aVPN peer by using Cisco Adaptive Security Device Manager **Site-to-Site VPN > Advanced** menus.



For IKEv1, the purpose of IKE Phase 2 is to negotiate the IPsec security parameters that are used to secure the IPsec tunnel. IKE Phase 2 performs the following functions:

Negotiates IPsec security parameters and IPsec transform sets

Establishes IPsec SAs

Periodically renegotiates IPsec SAs to ensure security

Performs an additional DH exchange (Optional)

IKE Phase 2 has one mode, called quick mode. Quick mode occurs after IKE has established the secure tunnel in Phase 1. It negotiates a shared IPsec transform, derives shared secret keying material that is used for the IPsec security algorithms, and establishes IPsec SAs. Quick mode exchanges nonces that are used to generate new shared secret key material and prevent replay attacks from generating invalid SAs.

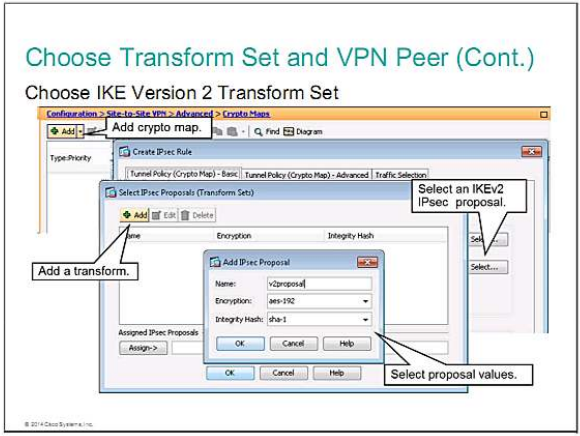
Quick mode also renegotiates a new IPsec SA when the IPsec SA lifetime expires. Quick mode refreshes the keying material that is used to create the shared secret key. The shared secret key is based on the keying material that is derived from the DH exchange in Phase 1.

Follow these steps to create a crypto map for use with IKEv1 using Cisco ASDM:

1. Choose **Configuration > Site-to-Site VPN > Advanced > Crypto Maps**. The Crypto Maps window appears.
2. Click **Add** to configure basic, advanced, and traffic selection parameters for a rule. The Create IPsec Rule window appears.
3. To add an IKEv1 IPsec proposal, click the **Select** button next to the IKEv1 IPsec Proposal field. The Select IPsec Proposals (Transform Sets) window appears.
4. To add a new IPsec proposal, click **Add**. The Add IPsec Proposal (Transform Sets) window appears.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

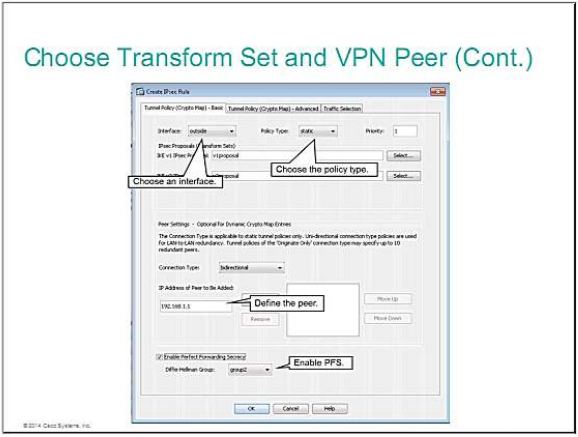
5. In the Add IPsec Proposal (Transform Set) window, configure the following fields: Set Name, Mode, ESP Encryption, and ESP Authentication.
6. Click **OK** to close the Add IPsec Proposal (Transform Set) window.
7. Click **Assign** in the Select IPsec Proposals (Transform Sets) window.
8. Click **OK** to close the Select IPsec Proposals (Transform Sets) window.



For IKEv2, the second portion of the negotiation is where child SAs are created. The existing IKEv2 parent SAs protect this phase. The second phase of IKEv2 exchanges DH values for Perfect Forward Secrecy (PFS), and traffic selectors that define what traffic to encrypt to rekey an existing parent-IKEv2 SA or a child-IPsec SA.

The process of creating a crypto map for use with IKEv2 is similar to that of IKEv1.

In the Add IPsec Proposal window, configure the following fields: Name, Encryption, and Integrity Hash.



Whether you are using IKEv1 or IKEv2, there are common crypto map parameters that you must also define. From the Tunnel Policy (Crypto Map) – Basic tab of the Create IPsec Rule window, you can view and configure the following additional fields and options:

Interface: Choose the interface name to which this policy applies.

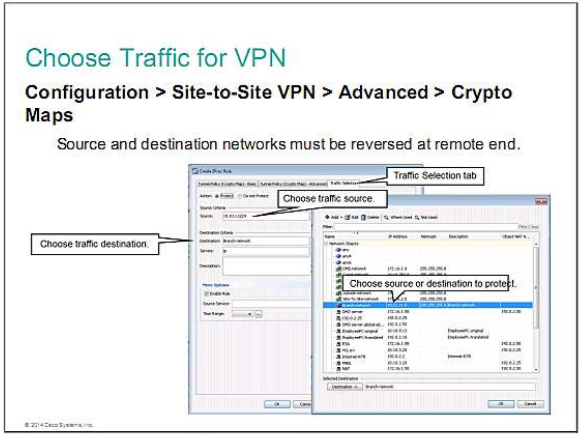
Policy Type: Choose the type, static or dynamic, of this tunnel policy.

Priority: Enter the priority of the policy.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Choose Traffic for VPN

This topic describes how to choose traffic for a VPN by using Cisco Adaptive Security Device Manager **Site-to-Site VPN > Advanced** menus.



The next task is to identify the traffic that the IPsec VPN tunnel should protect. Follow these steps to define this traffic:

From the Create IPsec Rule window or the Edit IPsec Rule window, click the **Traffic Selection** tab.

In the Action row, choose the action that you want this rule to take. There are two possible choices: Protect or Do not Protect.

In the Source and Destination fields, choose the appropriate network addresses.

Click **OK** to close the Create IPsec Rule window or the Edit IPsec Rule window.

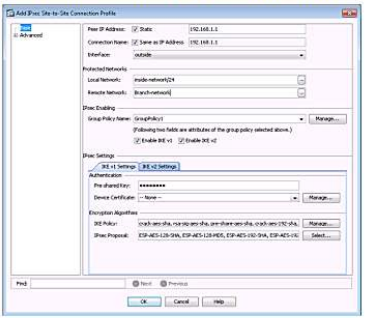
Click **Apply** to apply the configuration.

Configuring Site-to-Site VPN with Connection Profiles Menu

This topic explains how to use the Cisco Adaptive Security Device Manager **Site-to-Site VPN > Connection Profiles** menus to configure a site-to-site VPN.

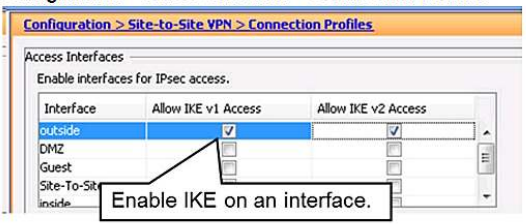
Configuring Site-to-Site VPN with Connection Profiles Menu

Configuration > Site-to-Site VPN > Connection Profiles



Configuring Site-to-Site VPN with Connection Profiles Menu (Cont.)

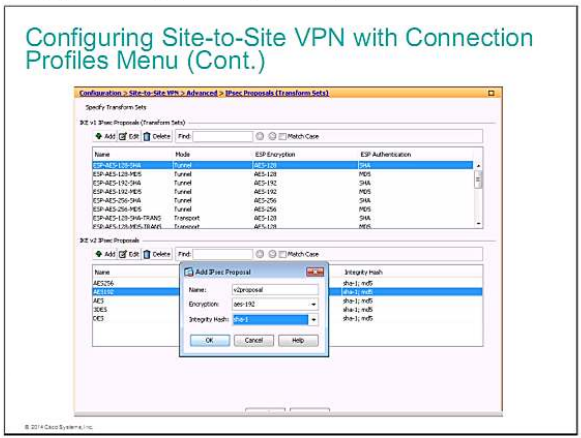
Configuration > Site-to-Site VPN > Connection Profiles



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Follow these steps to configure a site-to-site tunnel using a connection profile in Cisco ASDM:

1. Choose **Configuration > Site-to-Site VPN > Connection Profiles.** The Connection Profiles window is displayed. (not shown in figure)
2. In the Access Interfaces area, check the **Allow IKEv1 Access** or the **Allow IKEv2 Access** check box for the interface that will terminate the VPN tunnel. (not shown in figure)
3. Click **Add** in the Connection Profiles area. The Add IPsec Site-to-Site Connection Profile window opens.
4. Check the **Static** check box to the right of Peer IP Address, and enter the IP address of the peer device in the field provided.
5. From the **Interface** drop-down list, choose the interface on which you enabled IKE.
6. In the Protected Networks area, enter the IP address of the local protected network in the **Local Network** field (or click the ellipsis button to choose from a list), and enter the IP address of the remote protected network in the **Remote Network** field (or click the ellipsis button to choose from a list).
7. Select a group policy from the **Group Policy Name** drop-down list or accept the default setting.
8. Check the **Enable IKEv1** or the **Enable IKEv2** check box.
9. If you are using IKEv1, specify the PSK in the Pre-shared Key field of the IKEv1 Settings tab or choose a certificate from the Device Certificate drop-down list. If you are using IKEv2, enter the PSK for the local as well as the remote in the Local Pre-shared Key and Remote Peer Pre-shared Key fields of the IKEv2 Settings tab, or select a certificate from the Local Device Certificate drop-down list. Check the Allowed check box to allow certificate authentication for IKEv2 connections for the connection profile.
10. Click **OK**. Your new connection profile should be displayed in the Connection Profiles area of the Connection Profiles pane.
11. Click the **Apply** button to apply the changes to the HQ-ASA.

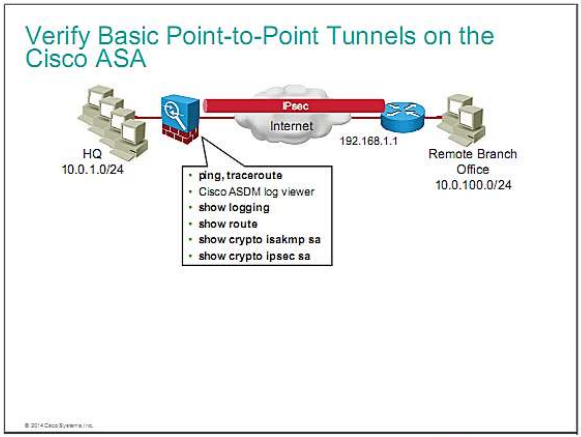


Connection profiles reference an IKE proposal and IPsec proposal (transform set). To tune either of these profiles, choose **Configuration > Site-to-Site VPN > Advanced > IKE Policies** or **Configuration > Site-to-Site VPN > Advanced > IPsec Proposals (Transform Sets).** The figure show the IPsec Proposals (Transform Sets) window, where you can configure IKEv1 IPsec Proposals (transform sets) or the IKEv2 IPsec Proposals.

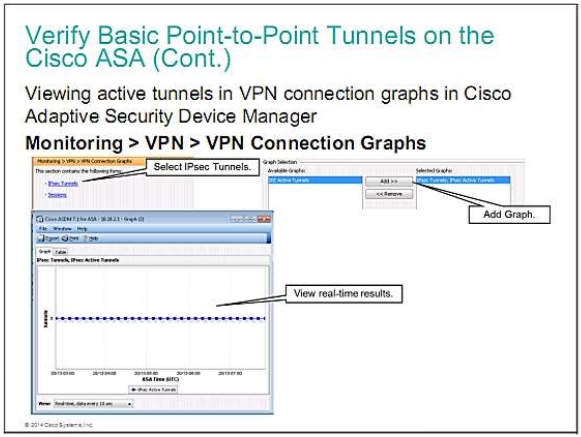
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify and Troubleshoot Basic Point-to-Point Tunnels on the Cisco ASA

This topic describes how to verify basic point-to-point tunnels on the Cisco Adaptive Security Appliance.



This figure presents useful tools for troubleshooting site-to-site VPNs on the Cisco ASA security appliance. When you troubleshoot IPsec site-to-site VPN tunnels, you should troubleshoot both ends of the tunnel.



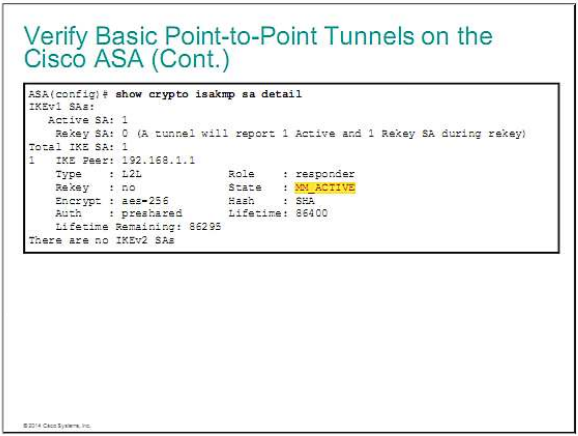
Follow these steps to verify IPsec tunnels using Cisco ASDM:

Choose **Configuration > VPN > VPN Connection Graphs**.

Click **IPsec Tunnels**. The graph selection window will open.

Add IPsec Tunnels into selected area.

Click **Show Graphs** button. The selected graph will open.



This figure shows how to verify established ISAKMP tunnels because of Phase 1 negotiation. The **show crypto isakmp sa detail** command shows the IP address of the IKE peer, type of connection, encryption and hashing algorithm and authentication method, role of the device, and state of the connection. The desired state is `MM_ACTIVE`. In the example, one IKE tunnel is active.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Basic Point-to-Point Tunnels on the Cisco ASA (Cont.)

```
ASA(config)# show crypto ipsec sa detail
interface: outside
    Crypto map tag: outside_map, seq num: 1, local addr: 192.168.1.2
      access-list outside_cryptomap_1 extended permit ip 10.0.1.0
255.255.255.0 10.0.100.0 255.255.255.0
      local ident (addr/mask/prot/port): (10.0.1.0/255.255.255.0/0/0)
      remote ident (addr/mask/prot/port): (10.0.100.0/255.255.255.0/0/0)
      current_peer: 192.168.1.1
      #pkts encaps: 10, #pkts encrypt: 10, #pkts digest: 10
      #pkts decaps: 19, #pkts decrypt: 19, #pkts verify: 19
<portion of output omitted>
      local crypto endpt.: 192.168.1.2/0, remote crypto endpt.: 192.168.1.1/0
      path mtu 1500, ipsec overhead 74, media mtu 1500
      current outbound spi: 49EF9B5C
      current inbound spi : 4A997987
    inbound esp sas:
      spi: 0x4A997987 (1251872103)
         transform: esp-aes esp-sha-hmac no compression
    outbound esp sas:
      spi: 0x49EF9B5C (1240439644)
         transform: esp-aes esp-sha-hmac no compression
```

© 2014 Cisco Systems, Inc.

This figure shows how to verify the IPsec tunnels that are established because of Phase 2 negotiation using the **show crypto ipsec sa detail** command. When you use this command, search for the number of encrypted and decrypted packets, and the number of encapsulated and de-encapsulated packets for individual connections.

Verify Basic Point-to-Point Tunnels on the Cisco ASA (Cont.)

Logging Messages: Unsuccessful IKE

```
ASA(config)# logging console 7
ASA(config)# logging enable
%ASA-5-713257: Phase 1 failure:  Mismatched attribute types for class Group
Description:  Rcv'd: Group 2  Cfg'd: Group 1
%ASA-5-713257: Phase 1 failure:  Mismatched attribute types for class Group
Description:  Rcv'd: Group 2  Cfg'd: Group 1
%ASA-5-713257: Phase 1 failure:  Mismatched attribute types for class Group
Description:  Rcv'd: Group 2  Cfg'd: Group 1
%ASA-5-713257: Phase 1 failure:  Mismatched attribute types for class Group
Description:  Rcv'd: Group 2  Cfg'd: Group 1
%ASA-7-713236: IP = 192.168.1.1, IKE_DECODE SENDING Message (msgid=0) with
payloads : HDR + NOTIFY (11) + NONE (0) total length : 132
%ASA-7-713906: IP = 192.168.1.1, All SA proposals found unacceptable
```

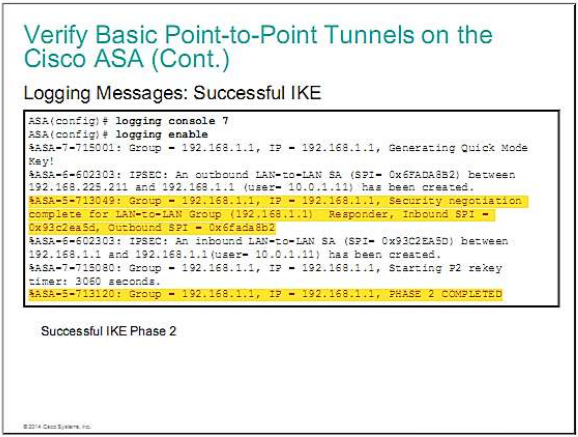
IKE Phase 1 failure due to mismatched IKE proposals

```
%ASA-4-713903: Group = 192.168.1.1, IP = 192.168.1.1, ERROR, had problems
decrypting packet, probably due to mismatched pre-shared key.  Aborting
```

IKE Phase 1 failure due to mismatched PSK

© 2014 Cisco Systems, Inc.

The example shows logging messages when IKE Phase 1 is not successful because of mismatched IKE proposals (different DH groups). The example also shows a logging message that is the result of misconfigured PSKs.

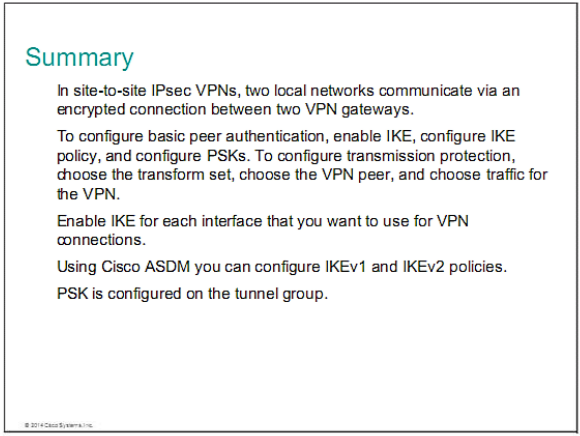


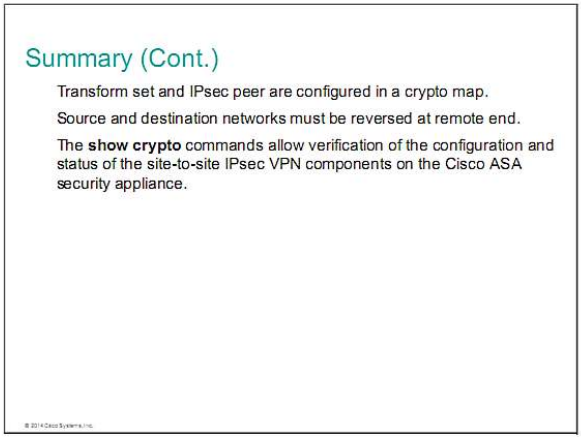
This figure shows logging messages for the successful creation of IPsec SAs and a completed IKE Phase 2.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.





https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 3

Deploying Cisco IOS VTI-Based Point-to-Point IPsec VPNs

The use of IPsec VTIs greatly simplifies the configuration process when you must provide protection for site-to-site VPN tunnels. A major benefit of IPsec VTIs is that the configuration does not require a static mapping of IPsec sessions to a physical interface. IPsec VTIs function like any other real interface so that you can apply many Cisco IOS Software features on them. This lesson will enable you to deploy static and dynamic point-to-point VTI tunnels using Cisco IOS Software.

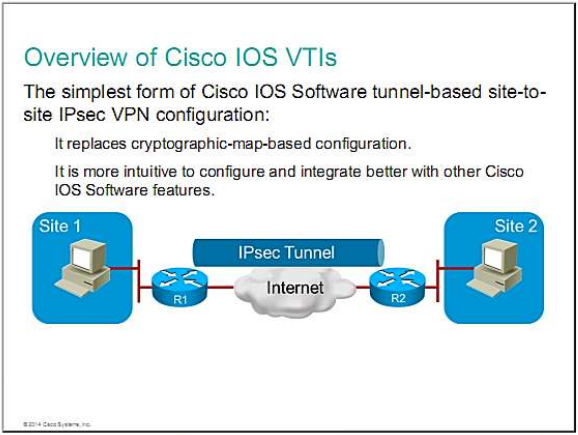
Upon completing this lesson, you will be able to meet these objectives:

- Describe Cisco IOS VTIs
- Configure basic IKE peering
- Verify basic IKE Peering.
- Configure static VTI point-to-point tunnels
- Verify static VTI point-to-point tunnels
- Configure dynamic VTI point-to-point tunnels
- Verify dynamic VTI point-to-point tunnels

https://t.me/learningnets

Overview of Cisco IOS VTIs

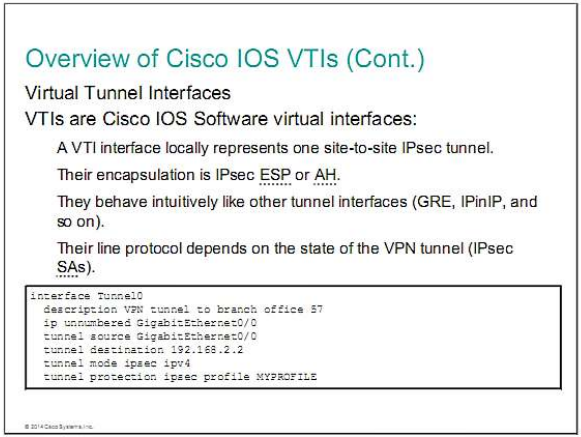
This topic describes a Cisco IOS Software VTI-based site-to-site IPsec VPN.



The use of IPsec VTIs greatly simplifies the configuration process when you must provide protection for site-to-site VPN tunnels and offers a simpler alternative to the use of GRE or L2TP tunnels for encapsulation and crypto maps with IPsec. A major benefit of IPsec VTIs is that the configuration does not require a static mapping of IPsec sessions to a physical interface. The IPsec tunnel endpoint is associated with a virtual interface. Because there is a routable interface at the tunnel endpoint, you can apply many common interface capabilities to the IPsec tunnel.

When you use the VTI approach, you configure a virtual tunnel interface and apply an IPsec profile on it via tunnel protection. It requires fewer configuration lines because crypto maps are automatically generated for each tunnel. Features for plaintext packets are configured on the VTI; features for encrypted packets are applied on the physical, outside interface.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



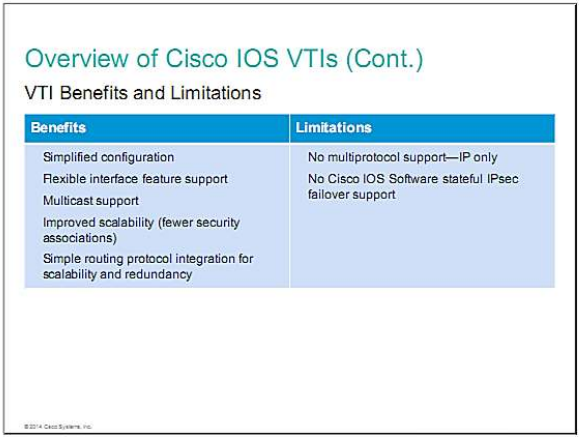
IPsec VTI is another mechanism that is used to support VPNs. IPsec VTIs provide a routable interface type for terminating IPsec tunnels and an easy way to define protection between sites to form an overlay network. A VTI supports native IPsec tunneling and allows interface commands to be applied directly to the IPsec tunnels. The IPsec tunnel endpoint is associated with a virtual interface. Because there is a routable interface at the tunnel endpoint, you can apply many common interface capabilities to the IPsec tunnel. VTIs support interoperability with standards-based IPsec installations of other vendors.

VTIs features include the following:

- They behave as regular tunnels, one for each remote site of the VPN.
- Their encapsulation must be either IPsec ESP or AH.
- Their line protocol depends on the state of the VPN tunnel (IPsec SAs).

The configuration output above shows an IPv4 IPsec tunnel, locally attached to the GigabitEthernet0/0 interface and remotely at the tunneling endpoint of 192.168.2.2. Tunnel traffic is encrypted using algorithms that are specified in the IPsec profile named MYPROFILE.

https://t.me/learningnets



An IPsec VTI has several benefits:

It simplifies configuration: Customers can use the virtual tunnel constructs to configure an IPsec peering, thus simplifying the complexity of the VPN configuration as compared to crypto maps or GRE IPsec tunnels.

Flexible interface feature support: An IPsec VTI is an encapsulation that uses its own Cisco IOS Software interface. This characteristic offers the flexibility of defining features to run on either the physical interface (that operates on ciphertext traffic) or on the IPsec VTI (that operates on cleartext traffic).

Multicast support: Customers can use the IPsec VTIs to securely transfer multicast traffic such as voice and video applications from one site to another.

Improved scalability: IPsec VTIs need fewer established SAs to cover different types of traffic, both unicast and multicast, thus enabling improved scaling.

Provides a routable interface: Like GRE IPsec, IPsec VTIs can natively support all types of IP routing protocols, which provides scalability and redundancy.

The IPsec VTI has the following limitations:

The IPsec VTI is limited to only IP unicast and multicast traffic, as opposed to GRE tunnels, which have a wider, multiprotocol application.

Cisco IOS Software IPsec stateful failover is not supported with IPsec VTIs. However, you can use alternative failover methods such as dynamic routing protocols to achieve similar functionality.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Overview of Cisco IOS VTIs (Cont.)

Deployment Choices

Deployment Choice	Criteria
Use static or dynamic VTI tunnels?	Use dynamic VTI tunnels for the hub in large hub-and-spoke networks. Otherwise, use static VTI tunnels.
Use static or dynamic routing protocol over VTI tunnels?	Use a dynamic routing protocol in large networks and to provide path or peer redundancy with multiple VTI tunnels. Otherwise, use static routing over VTI tunnels.

© 2014 Cisco Systems, Inc.

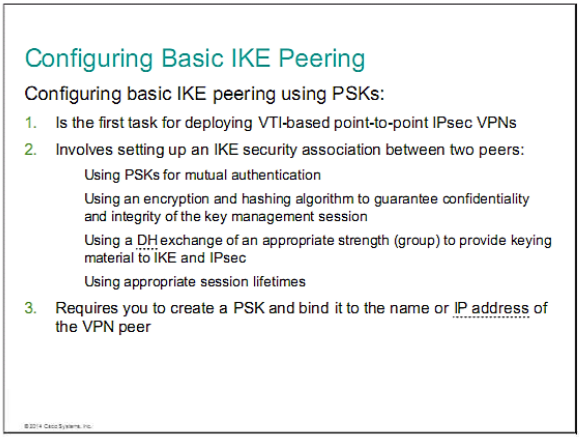
Consider the following guidelines when deploying VTI-based site-to-site IPsec VPNs:

Use VTI-based site-to-site VPNs as the default IPsec technology for individual point-to-point VPN links and for hub-and-spoke VPNs.

Consider deploying Dynamic Multipoint VPN (DMVPN) or Group Encrypted Transport GET VPN for larger environments with partial or fully meshed VPN requirements.

Configuring Basic IKE Peering

This topic describes how to configure basic IKE peering using PSKs.



Configuring IKE peerings between VPN members is the first step when configuring VTI-based IPsec VPNs. You should determine the IKE (ISAKMP) phase 1 policy requirements that you want to use, and then configure those policies on all peers. Having a detailed IKE policy plan lessens the chances of improper configuration.

Here are some planning steps:

1 **Determine the peer authentication method:** Choose the peer authentication method based on the credentials that are provisioned to all peers. Cisco IOS Software supports either PSKs, RSA-encrypted nonces, or RSA signatures to authenticate IPsec peers. This lesson focuses on using PSKs.

2 **Determine the session protection policy for the IKE session:** Specify an encryption and hashing algorithm that will be used to protect IKE packets.

3 **Determine the strength of the session key exchange method:** IPsec uses the DH algorithm to exchange session keys. The length of DH keys is determined by the DH group, which is part of an IKE policy.

4 **Use an appropriate IKE session lifetime:** Typically, the default Cisco IOS Software IKE session lifetime is appropriate for most use scenarios.

The goal of this planning step is to gather the precise data that you will need in later steps to minimize misconfiguration.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configuring Basic IKE Peering (Cont.)

Cisco IOS Software Default IKE PSK-Based Policies

Priority	Authentication Algorithm	Encryption Algorithm	Hash Algorithm	DH Group
65508	PSK	AES	SHA	5
65510	PSK	AES	MD5	5
65512	PSK	3DES	SHA	2
65514	PSK	3DES	MD5	2

Cisco IOS Release 12.4(20)T introduced default IKE policies:

- Avoid policies that use MD5.
- Avoid policies that use DH group 2.
- Use the highest priority PSK-based default policy (65508) for optimal security.

© 2014 Cisco Systems, Inc.

Cisco IOS Software Release 12.4(20)T introduced default IKE policies. There are eight default IKE policies that are supported with protection suites of priorities 65507-65514, where 65507 is the highest priority and 65514 is the lowest priority policy. The highest priority (65508) PSK-based default policy in the figure provides the highest protection of the default policy options. It uses AES as the encryption algorithm, SHA as the hash algorithm and DH group 5.

Avoid policies that use MD5 as the hash algorithm, because SHA-1 is stronger. Avoid policies that use DH group 2, because DH group 5 is stronger and more suitable as the default method strength.

Configuring Basic IKE Peering (Cont.)

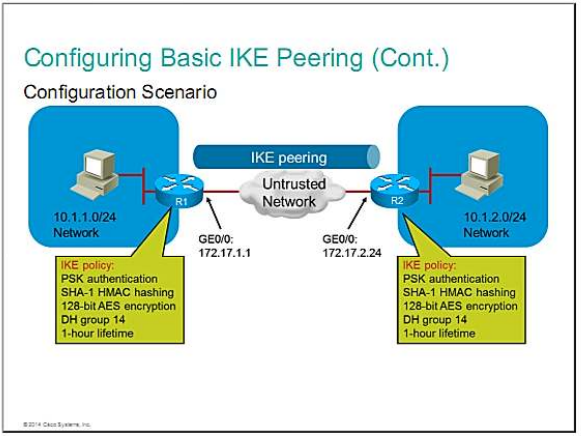
Configuration Choices

Configuration Choice	Criteria
Configure a nondefault IKE (ISAKMP) policy?	The default policies provide adequate security for most environments. You may want to choose a stronger DH group for additional protection, at the expense of tunnel set-up rate performace.
Tune default IKE (ISAKMP) policies?	Default lifetimes are very conservative; for systems that will not agree on them properly during negotiation, change the lifetime to an acceptable value.

© 2014 Cisco Systems, Inc.

When configuring basic IKE peering, you will make configuration choices based on the criteria shown in the figure. The default IKE protection policies provide adequate security for most environments. However, you may want to choose a stronger DH group to match the strength of other algorithms and key lengths for longer-term protection. Note, however, that higher DH groups require more computation effort at the beginning of the IKE exchange, therefore using stronger DH keys will come at the expense of tunnel set-up rate performance.

Default lifetimes are very conservative; for systems that will not agree on them properly during negotiation, change the lifetime to an acceptable value. Cisco IOS Software does not require IKE peers to have a matching IKE lifetime setting. The IKE SA will establish based on the shorter of the two settings.



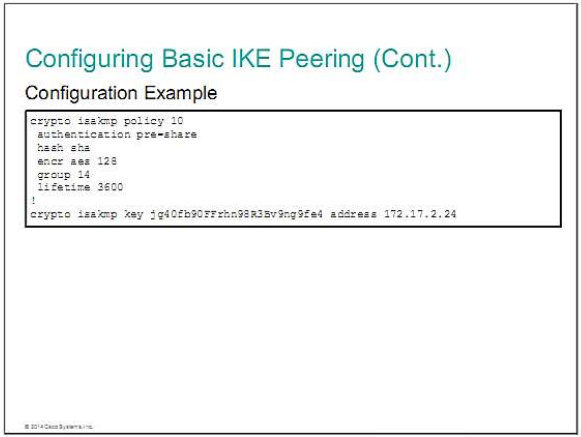
This figure shows the scenario that will be used in the upcoming configuration sequence. You will configure two IKE peers (using the 172.17.1.1 and 172.17.2.24 addresses on their untrusted interfaces respectively) using the following (symmetric) policy on both peers:

- IKE PSK-based peer authentication
- IKE packet authentication and hashing using the SHA-1 HMAC algorithm.
- IKE session encryption using the 128-bit AES encryption algorithm.
- Key exchange using the DH algorithm with its strength defined by DH group 14 (2048-bit).
- A one-hour (3600-second) lifetime of the IKE session.

You will also generate and configure a strong pre-shared authentication key on both peers.

The example assumes that connectivity between the two peers has been established and that IKE sessions are permitted along the path between the two routers.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



When configuring IKE peering, first configure the IKE policy. Then configure the authentication credentials.

Note The R2 router configuration should be a similar mirror configuration as is configured on the R1 router.

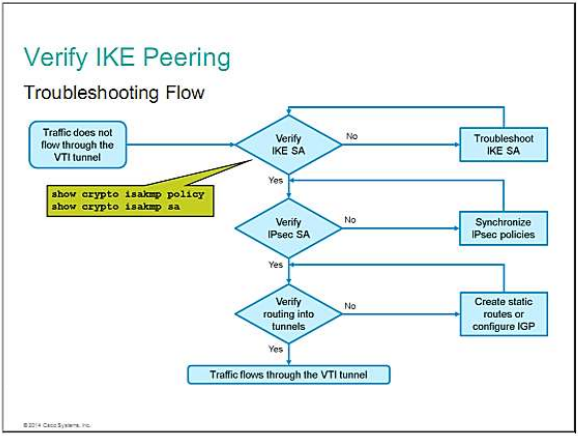
To configure basic IKE peering on the Cisco router, use the commands that are described in this table:

Commands	Description
crypto isakmp policy 10	Creates a new IKE policy.
authentication pre-share	Specifies PSKs for authentication.
hash sha	Sets hash algorithm. Use SHA-1 as the IKE hash (HMAC) algorithm; avoid MD5.
encr aes 128	Sets encryption algorithm. Use 128-bit AES or 3DES as the preferred encryption algorithms; avoid DES.
group 14	Sets DH groups for key exchange. Use DH groups 5, 14, 15, or 16 as the key exchange method; avoid DH groups 1 and 2.
lifetime 3600	Configures a lifetime; the default value is reasonable and does not require tuning.
crypto isakmp key jg40fb90FFrhn98R3Bv9ng9fe4 address 172.17.2.24	Creates a random, long PSK and binds it to the IP address of the peer.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Verify IKE Peering

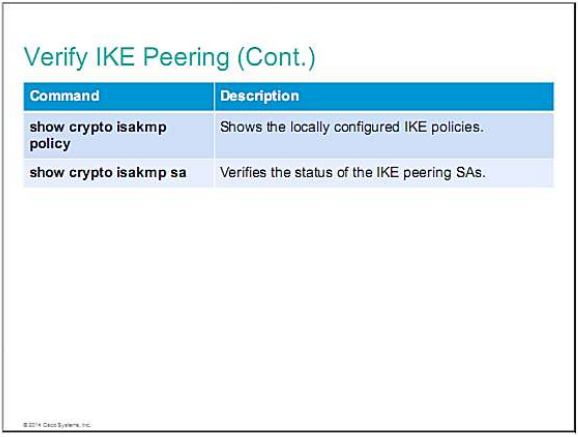
This topic describes how to verify IKE peering.



Before verifying static point-to-point VTI tunnels it is recommended to troubleshoot IKE peering.

Verify locally configured IKE policies with the **show crypto isakmp policy** command.

Verify the establishment of IKE peering with the **show crypto isakmp sa** command. If IKE policies do not match, synchronize them.

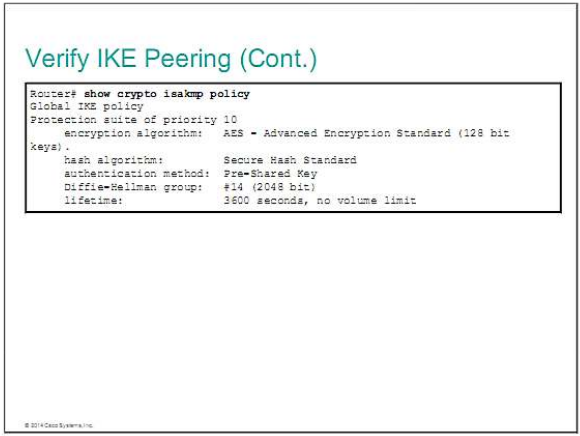


https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco IOS Software includes the following IKE peering validation commands:

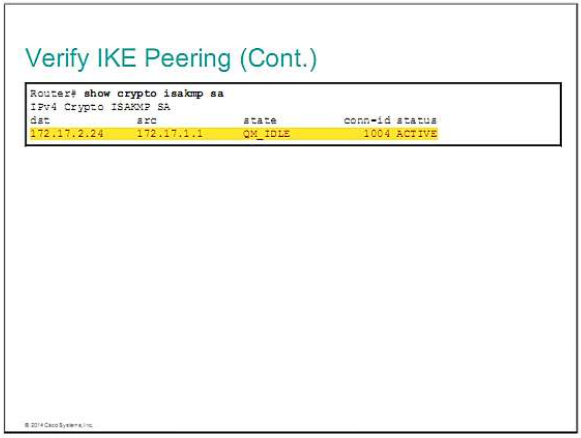
Use the **show crypto isakmp policy** command to display the parameters for each local IKE policy.

Use the **show crypto isakmp sa** command to display all existing IKE peering SAs with a particular peer.



To display the parameters for each local IKE policy, use the **show crypto isakmp policy** command.

If you have neither manually configured IKE policies nor disabled the default IKE policies by issuing the **no crypto isakmp default policy** command, the default IKE policies will be displayed.



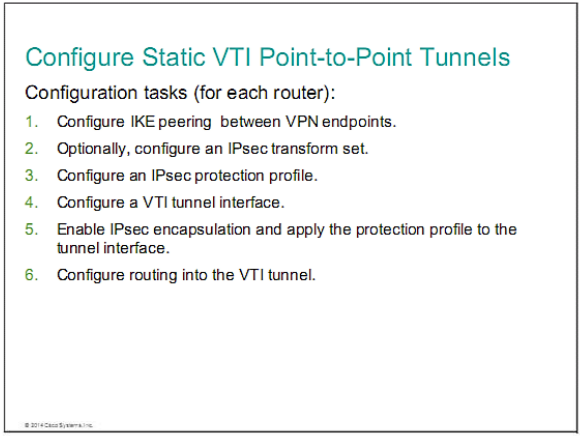
You can directly verify the status of the IKE peering SAs using **show crypto isakmp sa** command, which displays all existing IKE peering SAs.

Note that a session will be established only after IPsec requests it.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

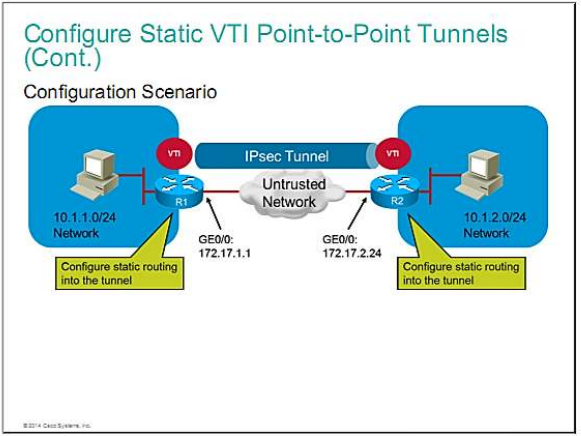
Configure Static VTI Point-to-Point Tunnels

This topic describes how to configure static VTI point-to-point tunnels.



You should follow this sequence to configure a site-to-site VPN peering using a static point-to-point VTI tunnel on each of the two peers that will establish a tunnel:

1. Configure IKE peering between VPN endpoints.
2. (Optional) Configure an IPsec transform set. This step is optional because default transform sets are applied if you do not specify a specific IPsec transform set.
3. Configure an IPsec protection profile to specify a traffic protection policy for the VTI tunnel.
4. Configure a VTI tunnel interface and provide it with IP addressing information.
5. Enable IPsec encapsulation on the VTI interface and apply the IPsec protection profile to the VTI interface.
6. Configure static or dynamic routing across the VTI tunnel.

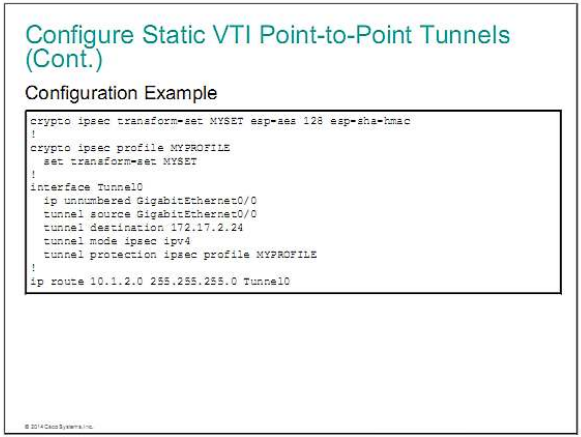


This figure shows the configuration scenario for the remainder of the configuration sequence. You will configure the left peer (R1) in this VTI-based VPN peering to protect all traffic flowing in both directions between the 10.1.1.0/24 and 10.1.2.0/24 networks that are directly connected to the two VPN peers, as shown in the figure. You will configure a static VTI tunnel between the IP addresses of the router untrusted interfaces (GigabitEthernet0/0 on both routers) and configure appropriate static routing to enable traffic exchange over the VTI tunnel.

Note The R2 router configuration should be a similar mirror configuration as is configured on the R1 router.

The example assumes that full IP connectivity is established between the two VPN peers.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



The configuration shows the creation of a nondefault IPsec transform set.

For VTI-based VPN, peering is the mandatory configuration of an IPsec protection profile.

Create a new tunnel interface using the **interface tunnel** command and assign it an IP address and subnet mask. Configure a tunnel source and a tunnel destination addresses.

Configure IPsec encapsulation on the tunnel using the **tunnel mode ipsec ipv4** command in tunnel interface configuration mode. Configure the assignment of the configured traffic protection policy to this tunnel by referencing the configured IPsec profile using the **tunnel protection ipsec profile** command to attach the configured MYPROFILE IPsec protection profile to the Tunnel0 interface.

You should configure routing to all networks reachable over the tunnel. One option is to set the static route or to include the tunnel subnet in the dynamic routing protocol.

To configure a static VTI point-to-point tunnels on the Cisco router, use the commands that are described in this table:

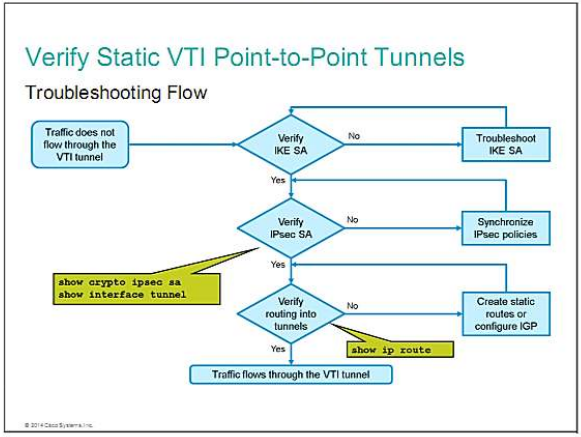
Commands	Description
crypto ipsec transform-set MYSET esp-aes 128 esp-sha-hmac	Creates an IPsec transform set for user traffic protection. Use AES or 3DES as the encryption algorithm. Use SHA-1 as the IKE hash (HMAC) algorithm.
crypto ipsec profile MYPROFILE **set transform-set MYSET**	Creates an IPsec profile and includes nondefault transform set in the profile.
interface Tunnel0 **ip unnumbered GigabitEthernet0/0** **tunnel source GigabitEthernet0/0** **tunnel destination 172.17.2.24**	Creates a new tunnel interface. Assigns an IP subnet and address to the interface. Specifies a tunnel source. Specifies a tunnel destination.

Commands	Description
tunnel mode ipsec ipv4 **tunnel protection ipsec profile MYPROFILE**	Specifies the tunnel encapsulation as **ipsec**. Specifies the traffic protection policy by referencing the configured IPsec profile.
ip route 10.1.2.0 255.255.255.0 Tunnel0	Creates static routes to all networks reachable over this tunnel.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Static VTI Point-to-Point Tunnels

This topic describes how to verify static VTI point-to-point tunnels.



The figure shows the recommended flow for troubleshooting static point-to-point VTI tunnels.

Verify the establishment of IPsec SA with the **show crypto ipsec sa** or **show interface tunnel** command. If IPsec policies do not match, synchronize them.

Verify that user traffic is properly routed into the VTI tunnel by using the **show ip route** command.

Verify Static VTI Point-to-Point Tunnels (Cont.)

Command	Description
show crypto ipsec sa	Verifies the status of the IPsec SAs.
show interface tunnel	Verifies the status of the local tunnel interface.
show ip route	Displays all local routes that point to VTIs.

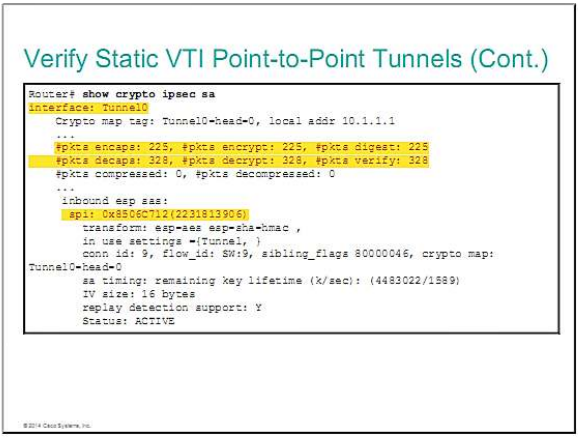
© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Cisco IOS Software includes the following static VTI point-to-point tunnels validation commands:

Use the **show crypto ipsec sa** command to display all existing IPsec SAs with a particular peer.

Use the **show interface tunnel** command to verify the status of the local tunnel interface.

Use the **show ip route** command to display routes pointing to virtual tunnel interfaces.



You can directly verify the status of the IPsec SAs using **show crypto ipsec sa** command, which displays all existing IPsec SAs.

You can see the number of packets that were successfully encapsulated and decapsulated, encrypted, decrypted, and hash-digested and verified by IPsec. Verify that these traffic statistics have nonzero values when traffic is present. Also check that IPsec SAs are populated with SPI numbers, which indicates that specific IPsec SAs have been negotiated.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Static VTI Point-to-Point Tunnels (Cont.)

```
Switch# show interface tunnel 0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered. Using address of FastEthernet0/0 (10.1.1.2)
  MTU 17898 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 5/255, rxload 5/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.1.1.1, destination 10.1.1.2
  Tunnel protocol/transport IPSEC/IP
Tunnel protection via IPSec (profile "MYIPSECPROFILE")
  ...
     67 packets input, 5700 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     67 packets output, 6700 bytes, 0 underruns
  ...
```

© 2014 Cisco Systems, Inc.

Cisco IOS Software tunnel interfaces are always up, unless they are administratively shut down. A more important indicator is the line protocol of the tunnel interface, which goes up on successful IPsec SA negotiation. The line protocol normally indicates whether the software processes that manage the line protocol consider the line usable or not. In the case of IPsec VTI, the line protocol state determines whether the IPsec tunnel has been fully established and is ready to forward traffic.

The output of the **show interface tunnel** command also gives you information about whether any traffic is entering and exiting the tunnel interface.

Verify Static VTI Point-to-Point Tunnels (Cont.)

```
Router# show ip route | include Tunnel
D EX  10.1.2.0/24 [170/26882560] via 10.1.1.2, 00:00:06, Tunnel0
```

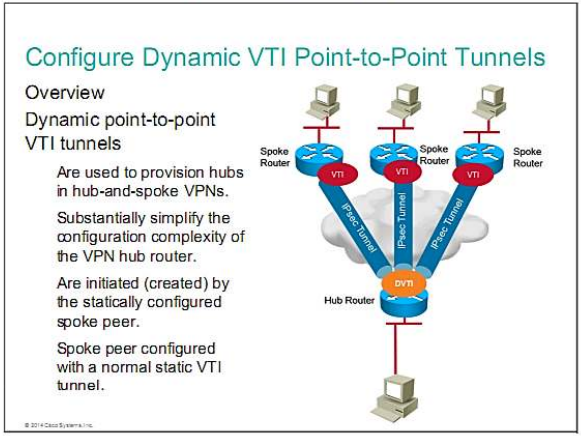
© 2014 Cisco Systems, Inc.

The output shows that the network 10.1.2.0/24 was learned through EIGRP and that traffic to that network is leaving the local router via the virtual Tunnel0 interface.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Dynamic VTI Point-to-Point Tunnels

This topic describes how to configure dynamic VTI point-to-point tunnels.



DVTIs can provide highly scalable hub configuration in hub-and-spoke VPNs for site-to-site and remote-access connectivity. The IPsec DVTI configuration at the hub device does not require a static mapping of IPsec sessions to a physical interface. Instead, VTIs are created dynamically from a preconfigured template just as tunnels to the hub establish. These dynamic e-tunnels provide on-demand separate virtual access interfaces for each VPN session, and their configuration is cloned from a virtual template configuration, which includes the IPsec configuration and any Cisco IOS Software features that are otherwise configured on regular interface, such as QoS, NetFlow, or ACLs.

Dynamic VTI functionality requires minimal configuration on the hub router in hub-and-spoke VPNs to create a large number of VTIs, one for each spoke. The tunnels are created by the spoke peer initiating the VPN connection and triggering the creation of the hub DVTI interface. On the spoke, a normal, static VTI is used to provision its tunnel to the hub.

Configure Dynamic VTI Point-to-Point Tunnels (Cont.)

Virtual Template Interface

```
interface Virtual-template 1
 ip unnumbered FastEthernet0/0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile PR1
 ip flow ingress
 zone-member security VPNZONE
```

Virtual-Access Interface

```
interface Virtual-access 125
 ip unnumbered FastEthernet0/0
 tunnel source 10.1.1.1
 tunnel mode ipsec ipv4
 tunnel destination 10.1.1.2
 tunnel protection ipsec profile PR1
 no tunnel protection ipsec initiate
 ip flow ingress
 zone-member security VPNZONE
```

Dynamic VTI hubs create VTIs automatically after the IKE session establishes.

©2014 Cisco Systems, Inc.

Configure Dynamic VTI Point-to-Point Tunnels (Cont.)

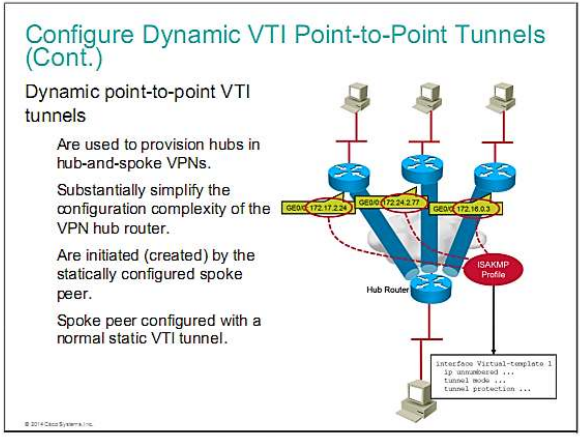
- Dynamic VTIs are represented as virtual-access interfaces.
- These are cloned from manually configured virtual template interfaces.
- A virtual template defines common settings for dynamic VTIs.
- All other dynamic parameters are filled in by the hub as the remote peer connects.

©2014 Cisco Systems, Inc.

On the hub router, dynamic VTIs are created when spoke peers create an IKE session to the hub device and negotiate IPsec policies. In terms of configuration, dynamic VTIs on the hub do not appear as a tunnel interface, but as virtual-access interfaces, which are automatically cloned from virtual template interfaces. A virtual template interface is a set of common settings that are inherited by dynamic VTIs, while all other dynamic parameters (such as the tunnel address of the spoke) are filled in by the hub as the remote peer connects.

https://t.me/learningnets

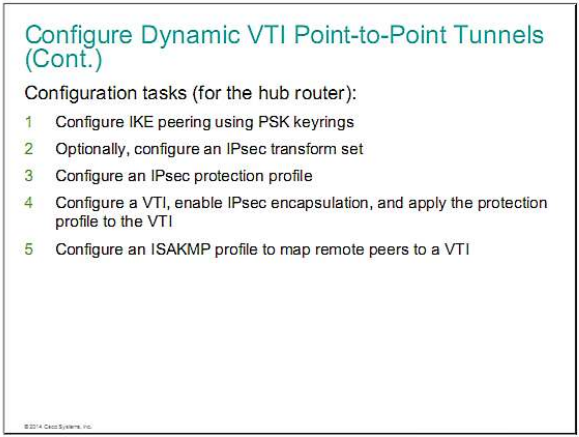
The example configuration in the figure shows a virtual template interface that defines common settings for all dynamic interfaces: the hub IP address inside the tunnel (unnumbered), the tunneling mode (IPsec), the tunnel protection, and, optionally, other Cisco IOS Software services (such as NetFlow accounting or Zone-Based Policy Firewall). From this virtual template, a DVTI is automatically created as a spoke completes IKE or IPsec negotiation, filling in the tunnel source or destination parameters and creating a dedicated interface (virtual-access 125) for a spoke site.



DVTIs require the use of a Cisco IOS Software feature called ISAKMP profiles.

ISAKMP profiles are generally used to bind the identity of a peer to a specific VPN feature. Inside an ISAKMP profile, you define a set of match statements to describe the identity of a remote peer or set of peers. In basic PSK-based VPNs that are the subject of this lesson, peer identity is based on IP addresses; however, the peer identity could also be based on an FQDN, or fields from certificate subject name of a peer. Depending on the granularity of such identity specification, ISAKMP profiles can match no peers, one peer, or an arbitrary large set of peers.

DVTIs use ISAKMP profiles to select remote (spoke) peers that can use a virtual template interface and create DVTIs from it.

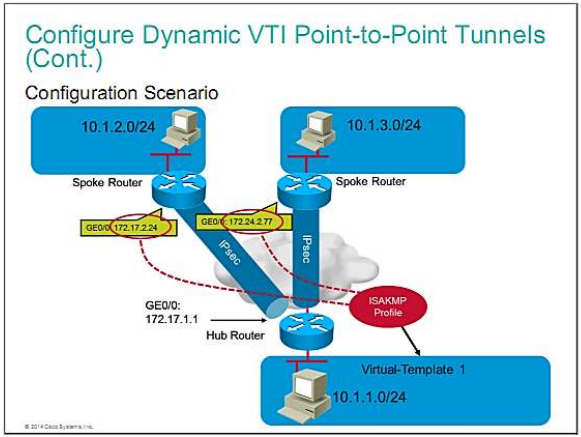


Complete the following tasks to configure a dynamic point-to-point VTI tunnel on a hub router:

1. Configure IKE peering using PSK keyrings instead of per-peer ISAKMP keys.
2. (Optional) Configure an IPsec transform set. This step is optional because Cisco IOS Software will use default transform sets if you do not specify a custom IPsec transform set for the VTI.
3. Configure an IPsec protection profile to be applied to the VTI.
4. Configure a VTI, where you should enable IPsec encapsulation and apply the protection profile to the VTI.
5. Configure an ISAKMP profile to map remote peers to the configured virtual template interface.

Note	When the hub router is configured, you should configure all spoke routers using static VTI tunnels, according to the guidance given.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



This figure shows the configuration scenario for the upcoming configuration sequence. You will configure the hub router to protect all traffic flowing in both directions from the remote 10.1.2.0/24 and 10.1.3.0/24 networks that are directly connected to spoke routers, as shown in the figure, to the central site network (10.1.1.0/24). On the hub, you will configure a dynamic VTI tunnel template that will accept tunnels from the two remote peers and dynamically create point-to-point hub-and-spoke tunnels.

The example assumes that full IP connectivity is established between all VPN peers and that basic IKE peering has been preconfigured between each spoke and the hub.

Configure Dynamic VTI Point-to-Point Tunnels (cont.)

Configuration Example

```
crypto keyring MYPSKS
  pre-shared-key address 172.17.2.24 key kfu7eibf2weHfiwefRbtke
  pre-shared-key address 172.17.2.77 key stvfi03w8eyII782rhMFuR
!
crypto ipsec transform-set MYSET esp-aes 128 esp-sha-hmac
!
crypto ipsec profile MYIPSECPROFILE
 set transform-set MYSET
!
interface Virtual-Template1 type tunnel
 ip unnumbered GigabitEthernet0/0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile MYIPSECPROFILE
!
crypto isakmp profile MYISAKMPPROFILE
 keyring MYPSKS
 match identity address 172.17.2.24 255.255.255.255
 match identity address 172.17.2.77 255.255.255.255
 virtual-template 1
```

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

To configure a dynamic VTI point-to-point tunnel on the Cisco router, use the commands that are described in this table:

Commands	Description
crypto keyring MYPSKS **pre-shared-key address 172.17.2.24 key** **kfu7eibf2weHfiwefRbtke** **pre-shared-key address 172.17.2.77 key** **stvfi03w8eyll782rhMFuR**	Creates authentication PSKs for all remote peers and groups them in a keyring.
crypto ipsec transform-set MYSET esp-aes 128 esp-sha-hmac	Creates an IPsec transform set for user traffic protection.
crypto ipsec profile MYIPSECPROFILE **set transform-set MYSET**	Creates an IPsec profile and includes a nondefault transform set in the profile.
interface Virtual-Template1 type tunnel **ip unnumbered GigabitEthernet0/0** **tunnel mode ipsec ipv4** **tunnel protection ipsec profile MYIPSECPROFILE**	Creates a numbered VTI of the type "tunnel." Use IP unnumbered addressing for the template. Specifies the tunnel encapsulation as IPsec and specifies the traffic protection policy by referencing the configured IPsec profile.
crypto isakmp profile MYISAKMPPROFILE **keyring MYPSKS** **match identity address 172.17.2.24 255.255.255.255** **match identity address 172.17.2.77 255.255.255.255** **virtual-template 1**	Creates a named ISAKMP profile. Specifies the authentication keys associated with this peer group. Specifies the remote peer identities. Specifies the VTI from which to create dynamic tunnels.

The first task of configuring a dynamic point-to-point VTI tunnel on a hub router involves creating authentication PSKs for all remote peers and grouping them into a named keyring.

Then configure a custom IPsec transform set for user traffic protection.

As is done with static VTIs, you need to configure a named IPsec profile that will determine the traffic protection policy for DVTI tunnels.

Then create a virtual template interface, where you will define common settings for dynamic VTIs.

Configure an ISAKMP profile to map remote peers to the created VTI.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Dynamic VTI Point-to-Point Tunnels

This topic describes how to verify dynamic VTI point-to-point tunnels.

Verify Dynamic VTI Point-to-Point Tunnels

Command	Description
show ip interfaces brief	Displays virtual access interfaces on the hub router. The Virtual-Access1 interface is used internally by Cisco IOS Software and is always down.
show interface	Displays the status of the virtual access interface.
show running-config	Displays the parameters of the dynamic VTI instance. Virtual access interfaces are never saved to the startup configuration.

© 2014 Cisco Systems, Inc.

Cisco IOS Software includes the following dynamic VTI point-to-point tunnel validation commands:

Use the **show ip interfaces brief** command to display virtual access interfaces on the hub router.

Use the **show interface** command to verify the status of the virtual access interface.

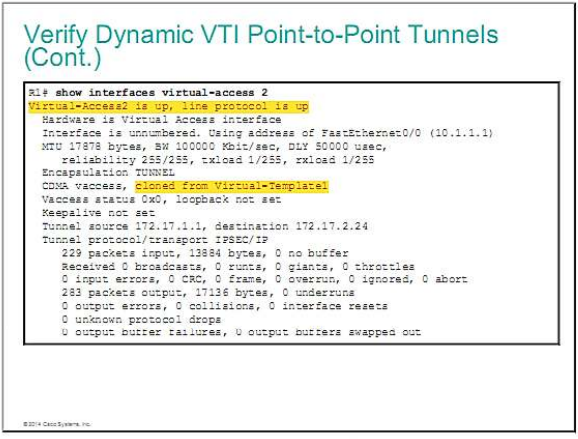
Use the **show running-config** command to display the parameters of the dynamic VTI instance.

Verify Dynamic VTI Point-to-Point Tunnels (Cont.)

```
R1# show ip interface brief | include Virtual-Access
Virtual-Access1            unassigned      YES
unset  down                  down
Virtual-Access2            172.17.1.1      YES unset  up                    up
```

© 2014 Cisco Systems, Inc.

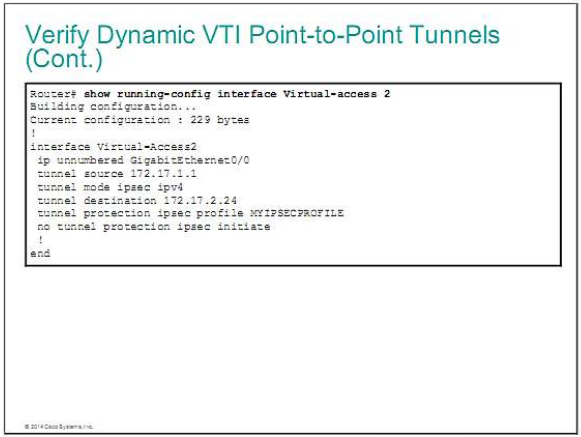
After configuring dynamic VTIs on the hub, you can verify tunnel status on the hub router if spoke peers are already configured and actively attempting to peer with the hub. The **show ip interfaces brief | include Virtual-Access** command will display all instantiated virtual access interfaces on the hub router. You should see a virtual access interface for every dynamically created VTI interface (peering with the spoke) on the hub. The special Virtual-Access1 interface is used internally by Cisco IOS Software and will always be present in the command output and will always be down.



After you have configured dynamic point-to-point VTI, verify that the dynamically created VTI operates properly. Verify the status of the specific virtual access interface using the **show interfaces virtual-access** command.

A virtual access interface is always up. The line protocol of a virtual access interface will be up on successful IPsec SA negotiation. The line protocol indicates whether the software processes that manage the line protocol consider the line usable. The output line "cloned from Virtual-Template1" indicates that the current virtual access interface was cloned from the Virtual-Template1 interface.

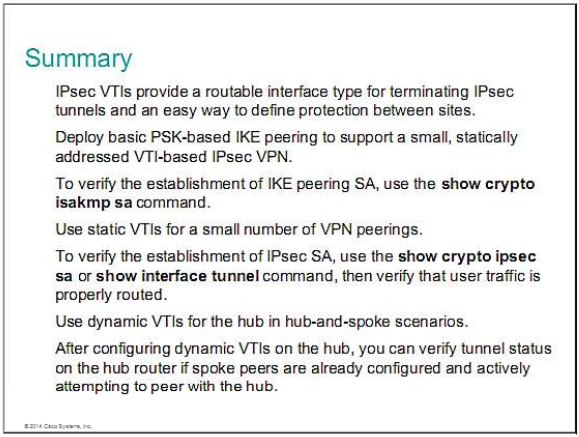
https://t.me/learningnets
© 2014 Cisco Systems, Inc.



The **show running-config interface Virtual-access** command displays the complete configuration for a dynamic VTI virtual access interface that was dynamically created, including the parts that were cloned from the virtual template interface. Note that configuration for virtual access interfaces is never automatically saved to the startup configuration, nor can it be manually saved.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets

Lesson 4

Deploying Cisco IOS DMVPNs

VPN-based security solutions are increasingly popular and have proven to be an effective and secure technology for protecting sensitive data that is traversing insecure channel mediums, such as the Internet.

Traditional IPsec-based site-to-site, hub-to-spoke VPN deployment models do not scale well and are adequate only for small- and medium-sized networks. As demand for IPsec based VPN implementation grows, organizations with large-scale enterprise networks require scalable and dynamic IPsec solutions that interconnect sites across the Internet with reduced latency, while optimizing network performance and bandwidth utilization.

The DMVPN technology is used for scaling IPsec VPN networks by offering a large-scale IPsec VPN deployment model that allows the network to expand and realize its full potential. DMVPN offers scalability that enables zero-touch deployment models.

This lesson provides a complete overview of the DMVPN solution architecture, implementation, and various deployment scenarios.

Upon completing this lesson, you will be able to mee these objectives:

- Describe Cisco IOS DMVPN
- Describe DMVPN solution components
- Describe GRE
- Describe NHRP
- Describe DMVPN operations
- Describe types of authentication
- Describe how to configure DMVPN hub
- Describe how to configure DMVPN spoke
- Describe how to configure routing in DMVPN

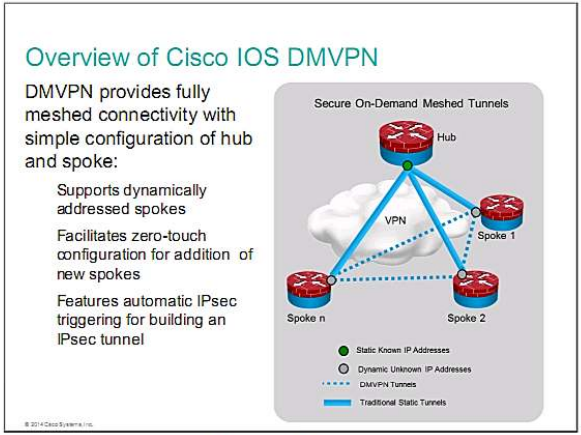
https://t.me/learningnets

Describe how to verify DMVPN

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Overview of Cisco IOS DMVPN

This topic describes Cisco IOS DMVPN.



DMVPN allows IPsec VPN networks to better scale hub-to-spoke and spoke-to-spoke designs, thereby optimizing performance and reducing latency for communications between sites.

The following are the primary benefits of DMVPNs:

Hub router configuration reduction: Traditionally, configuration on the hub router defines the crypto map characteristics, the crypto ACL, and the GRE tunnel interface for each spoke router. The DMPVN feature enables users to configure a single mGRE tunnel interface and a single IPsec profile without requiring crypto ACLs on the hub router to manage all spoke routers. Thus, the size of the configuration on the hub router remains constant even if spoke routers are added to the network.

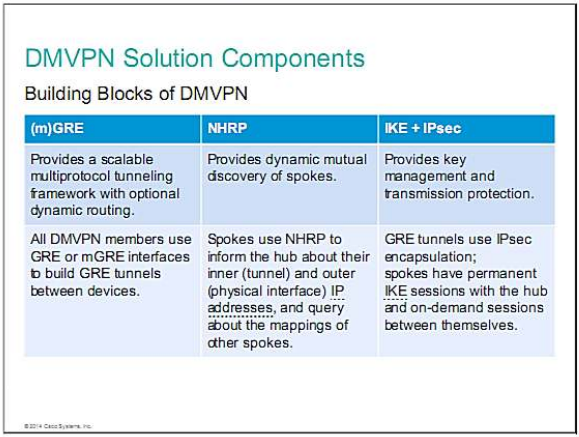
Automatic IPsec initiation: GRE uses NHRP to configure and resolve the peer destination address. Thus, this feature allows IPsec to be immediately triggered to create point-to-point GRE tunnels without any IPsec peering configuration.

Support for dynamically addressed spoke routers: When you use point-to-point GRE and IPsec hub-and-spoke VPN networks, you must know the physical interface IP address of the spoke routers when you configure the hub router because the IP address must be configured as the GRE and IPsec tunnel destination address. DMVPN enables spoke routers to have dynamic physical interface IP addresses and uses NHRP to register the dynamic physical interface IP addresses of the spoke routers with the hub router.

With networks becoming geographically distributed, network availability over a secure channel is becoming a critical factor in designing scalable IPsec VPN solution designs. DMVPN solution architecture is by far the most effective and scalable solution available.

DMVPN Solution Components

This topic describes DMVPN solution components.



DMVPNs use the following mechanisms to combine the best of hub-and-spoke and fully meshed topologies, as well as provide some other features:

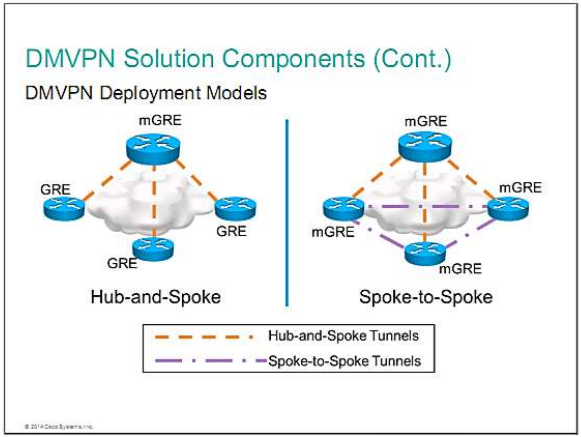
mGRE: mGRE technology enables a single GRE interface to support multiple GRE tunnels and simplifies the complexity of the configuration. GRE tunnels also provide support for IP multicast and non-IP protocols. IP multicast, in turn, enables the designer to use routing protocols to distribute routing information and detect changes in the VPN. All DMVPN members use GRE or mGRE interfaces to build tunnels between devices.

NHRP: NHRP is a client and server protocol in which the hub acts as an NHRP server and the spokes act as NHRP clients. The hub maintains an NHRP database of mappings between the outer (public, physical interface) and tunnel (inside the tunnel interface) IP addresses of each spoke. Each spoke registers its public and internal tunnel addresses when it boots and queries the NHRP database for the addresses of other spokes when building direct spoke-to-spoke tunnels. NHRP is used to reduce the configuration complexity of fully meshed or partially meshed VPNs into something that is comparable to the complexity of hub-and-spoke VPNs.

IPsec: IPsec provides transmission protection for GRE tunnels. DMVPNs form a permanent hub-and-spoke IPsec VPN, which can dynamically reconfigure itself into a partial mesh or full mesh as needed.

Dynamic routing protocol is used to advertise private networks within the DMVPN networks. IP routing updates and IP multicast data packets traverse only hub-and-spoke tunnels. Unicast IP data packets traverse both hub-and-spoke and direct dynamic spoke-to-spoke tunnels. Routing adjacencies are established across the hub-to-spoke links, and the spoke-to-spoke routing logic is performed by NHRP. Routing protocol does not monitor the state of spoke-to-spoke tunnels. Protocols that are supported include RIP, EIGRP, OSPF, ODR, and BGP.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



A DMVPN cloud topology can support either a hub-and-spoke or a spoke-to-spoke deployment model.

Hub and spoke: A strict hub-and-spoke DMVPN deployment model requires each branch to be configured with a point-to-point GRE interface to the hub, and in which all traffic between spoke networks must flow through the hub router. DMVPN provides scalable configuration to the hub router, but does not facilitate direct spoke-to-spoke communication.

Spoke to spoke: A spoke-to-spoke DMVPN deployment model requires each branch to be configured with an mGRE interface in which dynamic spoke-to-spoke tunnels are used for the spoke-to-spoke traffic. In this model, DMVPN provides a scalable configuration model for all involved devices, and also allows spoke devices to dynamically peer and establish optimal routing paths. Note that DMVPN will not immediately produce a partially meshed or fully meshed topology. DMVPN initially establishes a permanent hub-and-spoke topology, from which a partial mesh or full mesh is dynamically generated based on traffic patterns.

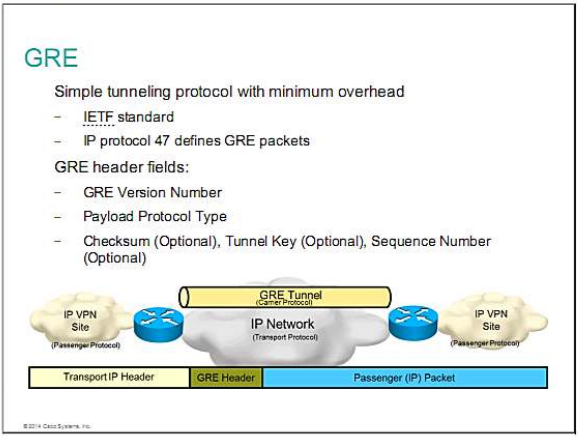
Because VPNs are used for secure enterprise communications across a shared public infrastructure, such as the Internet, you must consider the following two distinct IP address domains:

The enterprise addressing space, sometimes referred to as the private or inside addresses: These are the addresses that are assigned to (m)GRE tunnel interfaces, and are registered to the hub NHRP server by each spoke.

The infrastructure addressing space, also referred to as the service provider, public, or outside addresses: These are the addresses on physical interfaces of the routers, and used as addresses in tunnel envelopes (that is, the source and destination addresses in the GRE/IPsec packet headers).

GRE

This topic describes GRE.



GRE is a simple, general-purpose protocol that is designed to manage the transportation of multiprotocol and IP multicast traffic between two or more sites, which may only have IP connectivity.

GRE uses IP protocol 47 and can be used with IPsec VPNs to pass routing information between connected networks. Tunneling protocols encapsulate packets inside of a transport protocol. A tunnel interface supports a header for each of the following:

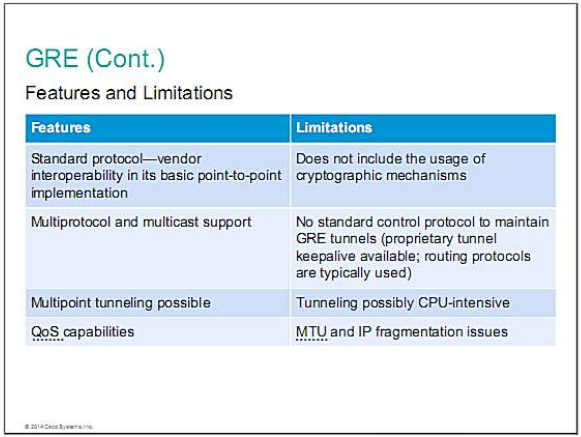
A passenger protocol or encapsulated protocol.

A carrier protocol.

A transport protocol.

GRE may be combined with IPsec, where GRE encapsulates the plaintext packet, and then IPsec cryptographically envelopes the packet.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

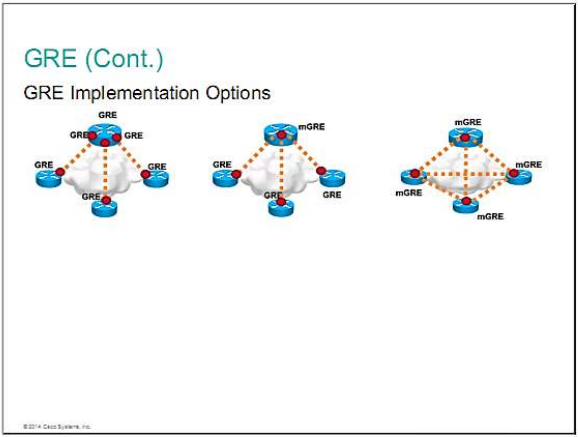


GRE offers many useful features. Many vendors support GRE, so interoperability is usually not a problem when using basic GRE point-to-point connectivity. GRE tunnels support multiple network protocols, including IP multicasting. A significant advantage of GRE tunneling relative to non-virtual tunnel interface (non-VTI) IPsec tunnels is that GRE uses Cisco Internetwork Operating System Software interfaces, which can offer finer-grained QoS features.

Some limitations of GRE tunneling include the following:

- GRE provides no cryptographic traffic protection, and must be combined with IPsec to provide it.
- There is no standard method to determine the end-to-end state of a GRE tunnel. Cisco IOS Software supports proprietary GRE keepalives for this purpose. Alternatively, dynamic routing protocols can verify the end-to-end tunnel state.
- It may be CPU intensive on some platforms.
- Tunnels can exhibit MTU and IP fragmentation-related issues.

The majority of systems installed today leave the default MTU value set to 1500 and generate packets at this MTU size limit. The tunneling system tries to add the 24-bit GRE header onto the packet and this makes the packet too large for the MTU size. The tunneling system then begins to fragment the packets so they will fit under the MTU limit. Where this becomes an issue is when the DF bit is set and the tunneling systems fragment the packet. Once fragmented, some applications cannot process the packets, and this can cause problems. This also becomes an issue when a device wants to send an ICMP packet out to have the sending system retransmit the packets at a smaller size, but ICMP is blocked before reaching the transmitting system (for example, a firewall blocks all ICMP messages).



The three diagrams show three options for implementing GRE networks using GRE or mGRE functionality:

In the leftmost diagram, a hub-and-spoke network uses a set of point-to-point tunnels using only GRE interfaces. On the hub, you would need to create as many GRE interfaces as there are spokes, and on a spoke, you would require one GRE interface. All traffic between spokes flows strictly over the hub.

The middle diagram shows the hub optimized with an mGRE interface. In this setup, only a single interface is required on the hub. However, you must deploy NHRP for the hub to learn spoke addresses and correctly provision the spoke-to-hub GRE tunnels.

In the rightmost diagram, all devices in a hub-and-spoke network use mGRE interfaces. Using NHRP, these devices can establish a partial mesh or full mesh of GRE tunnels by only configuring a single mGRE interface on each device, greatly simplifying configuration and improving manageability.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

GRE (Cont.)

Point-to-Point vs Point-to-Multipoint GRE Tunnels

Point-to-Point GRE	mGRE
Used on point-to-point tunnels or on spokes in hub-and-spoke VPNs	Typically used on the hub in hub-and-spoke VPNs
One tunnel interface for each peer	A single interface on each router for all (m)GRE peers
Does not require NHRP; all other peers statically configured	Requires NHRP to build tunnels to other peers
Supports unicast, multicast, and broadcast	Supports unicast, multicast, and broadcast

© 2014 Cisco Systems, Inc.

This table compares point-to-point and point-to-multipoint GRE tunnels. The following are the main characteristics of a point-to-point GRE configuration:

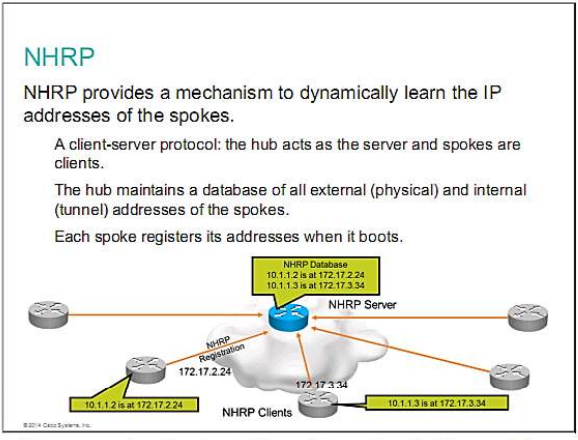
- It is typically used for simple point-to-point tunnels that emulate a point-to-point WAN link, or on spoke routers in hub-and-spoke VPNs.
- On each device, there is a separate GRE tunnel interface configured for each GRE peer.
- It does not require NHRP because other peers (their destination tunnel address) are statically configured.
- It supports unicast, multicast, and broadcast traffic.

The following are the main characteristics of point-to-multipoint GRE configuration:

- Only one tunnel interface needs to be configured on a router to support multiple remote GRE peers; in a hub-and-spoke network, a single mGRE tunnel interface on the hub accommodates many spoke GRE peers. This greatly simplifies the management of the hub device because new spokes can be added without the need to reconfigure the hub.
- In order to learn about the IP addresses of other peer, devices using mGRE require NHRP to build dynamic GRE tunnels. Peers can also use dynamically assigned addresses that will then be used by NHRP when registering with the hub.
- mGRE interfaces also support unicast, multicast, and broadcast traffic.

NHRP

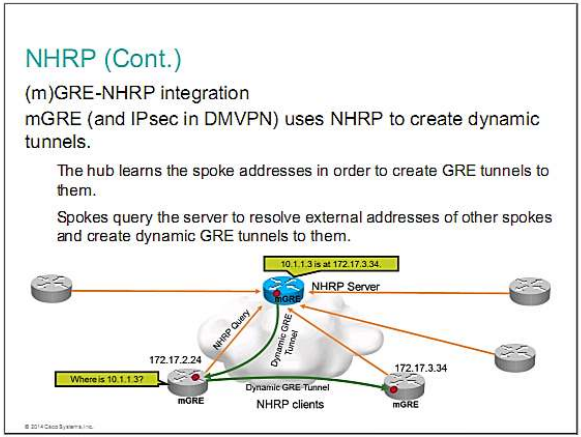
This topic describes NHRP.



NHRP is used by routers to determine the IP address of the next hop (in this case, the inner and outer IP addresses of spoke routers) in IP tunneling networks. When a spoke router initially connects to a DMVPN network, it registers its inner (tunnel) and outer (physical interface) address with the hub router (the NHRP server) whose IP address is statically preconfigured on the spoke router. This registration enables the mGRE interface on the hub router to build a dynamic GRE tunnel back to the registering spoke router without having to know the branch tunnel destination in advance. Therefore, NHRP creates a mapping for a tunnel IP address to the physical interface IP address for each spoke at the hub.

NHRP is a client-server protocol: the hub acts as the server and the spokes are clients. You must configure spoke routers with the physical interface IP address of the hub router as their next-hop server in order for spokes to register with the hub router. The spoke routers send an NHRP registration to the hub router that contains their tunnel IP address and the physical interface IP address. The hub router creates an entry in its NHRP cache and returns a registration reply. The spoke router now views the hub router as a valid next-hop server and uses it as a source to locate any other spokes and networks in the NHRP domain. From the routing protocol perspective, the NHRP domain operates similarly to a NBMA network, such as a multipoint Frame Relay network.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



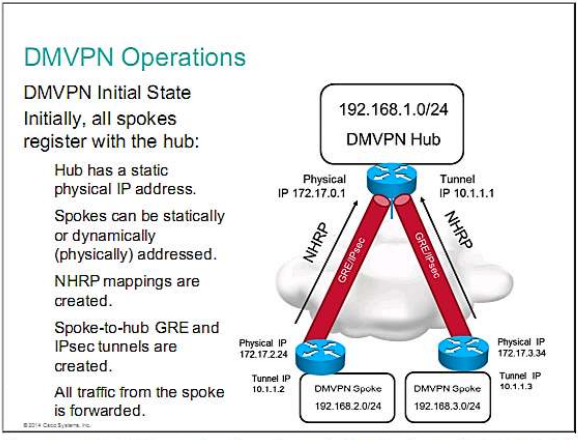
Using NHRP in multipoint GRE networks maps inner tunnel IP addresses to the outer transport IP addresses. In a hub-and-spoke DMVPN deployment, no GRE or IPsec information about a spoke is configured on the hub router. The spoke router for the GRE tunnel is configured (via NHRP commands) with information about the hub router as the next-hop server. When the spoke router starts up, it automatically initiates the IPsec tunnel with the hub route. It then uses NHRP to notify the hub router of its current physical interface IP address. This notification is useful for the following reasons:

- Configuration of the hub router is shortened and simplified because it does not need to have GRE or IPsec information about the peer routers. All of this information is learned dynamically via NHRP.
- When you add a new spoke router to the DMVPN network, you do not need to change the configuration on the hub or on any of the current spoke routers. The new spoke router is configured with the hub information, and when it starts up, it dynamically registers with the hub router. The dynamic routing protocol propagates the routing information for this spoke to the hub. The hub propagates this new routing information to the other spokes. It also propagates the routing information from the other spokes to the new spoke.

In the figure, one spoke wants to send IP traffic to another spoke, which has a tunnel interface that is configured with the IP address of 10.1.1.3. The originating router sends an NHRP query for the 10.1.1.3 IP address to the hub, which is configured as an NHRP server. The hub responds with information that the IP address 10.1.1.3 is mapped to the physical interface (172.17.3.34) of the receiving spoke router.

DMVPN Operations

This topic describes DMVPN operations.

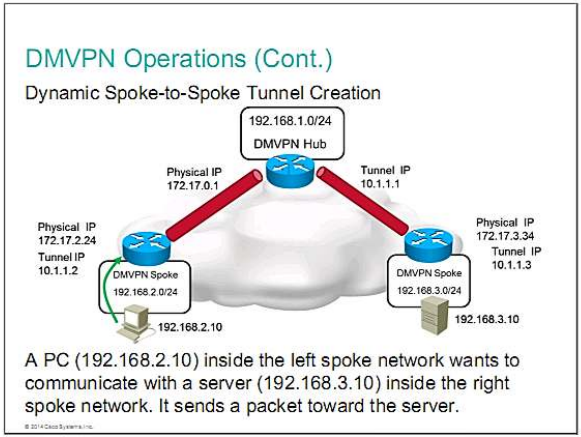


The figure illustrates a DMVPN example, where the top (hub) router is used as DMVPN hub and other three routers are statically configured to connect to the hub using NHRP, GRE, and IPsec. The VPN tunnels between the spokes are not permanent, but are established on demand. Initially, all spokes establish a GRE/IPsec tunnel with the hub and NHRP-register with the hub.

At this point, the network is strictly hub-and-spoke and can remain that way, if so configured. The following are the current network properties:

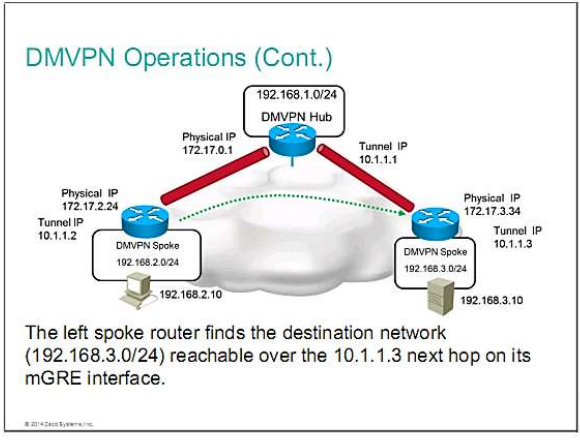
- The hub knows the outer and inner IP addresses of each spoke in its NHRP database.
- Two spoke-to-hub GRE/IPsec tunnels are created.
- All traffic from each spoke (to the hub networks or to the networks of other spokes) is forwarded via the hub.

https://t.me/learningnets

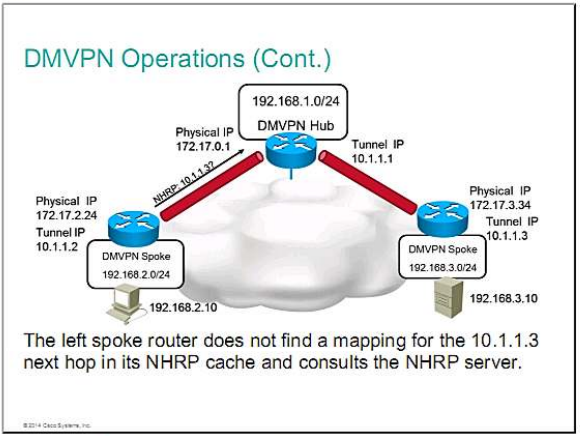


In a spoke-to-spoke-enabled DMVPN, spokes can dynamically locate each other and establish direct tunnels.

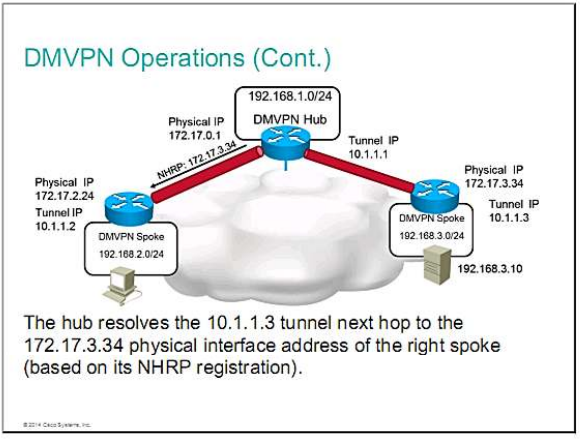
This figure illustrates dynamic spoke-to-spoke tunnel creation when a packet is sent from the left spoke site to the right spoke site. In this figure, static NHRP-enabled GRE tunnels are already configured and pre-established from the hub to all spokes. A PC (192.168.2.10) on the left spoke subnet wants to connect to a web server (192.168.3.10) behind the right spoke router. It sends a packet towards the server.



Dynamic routing (running over the permanent hub-to-spoke links) is used to populate the routing table of the spoke, so each spoke knows about subnets behind other spokes. The left spoke router consults its routing table for a route to the destination network (192.168.3.0) and finds it reachable behind the right spoke router. The routing table provides an IP next-hop of 10.1.1.3 via the mGRE tunnel interface of the left spoke router.

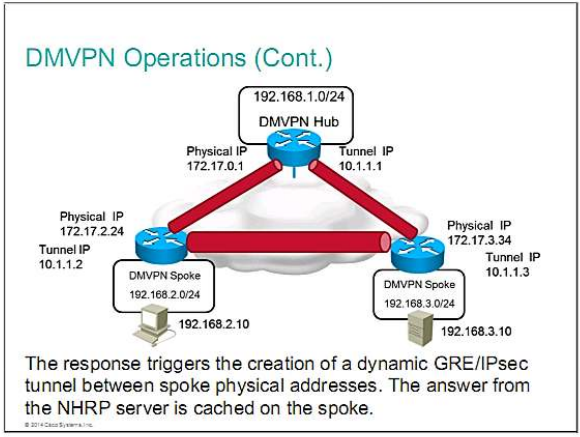


After consulting its local NHRP cache for the 10.1.1.3 next hop and finding no cached entries, the left spoke router sends an NHRP query to the NHRP server (DMVPN hub) to resolve the inner (private) tunnel address 10.1.1.3 to an outer (physical, public) tunnel endpoint address.

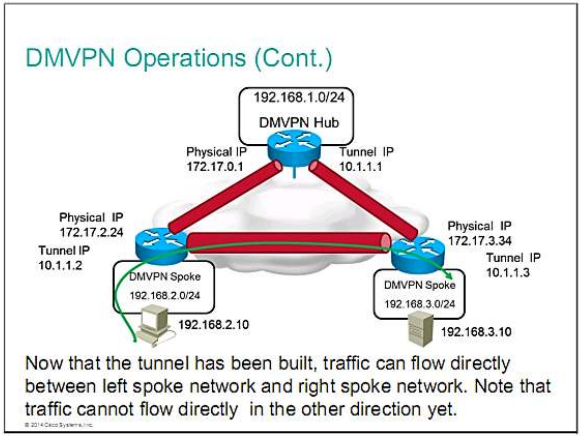


https://t.me/learningnets
© 2014 Cisco Systems, Inc.

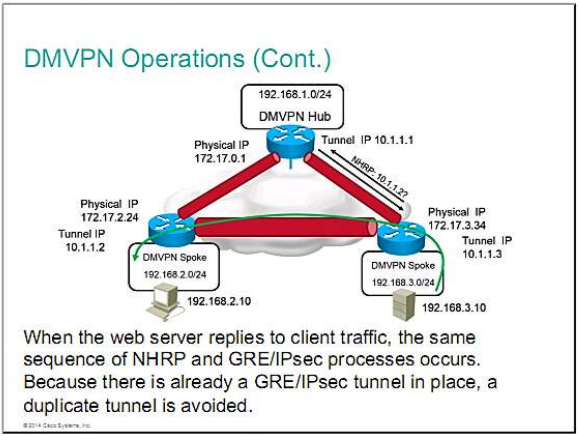
The NHRP server (hub) stores information about the inner (tunnel) and outer (physical) addresses of all spoke routers based on the NHRP registration information that is received from the initial registration of the spokes to the hub. It sends an NHRP reply to the left spoke informing it that the 10.1.1.3 next hop is reachable through the 172.17.3.34 outer (physical) IP address of the right spoke router.



The left spoke receives the NHRP response from the DMVPN hub (the NHRP server) and enters it in local NHRP cache. This triggers IPsec on the left spoke to create an IPsec tunnel directly to the right spoke reachable at 172.17.3.34. The left spoke initiates an IKE session to 172.17.3.34 using its physical IP address (172.17.2.24) as the source, and negotiates IPsec SAs for a spoke-to-spoke GRE tunnel.

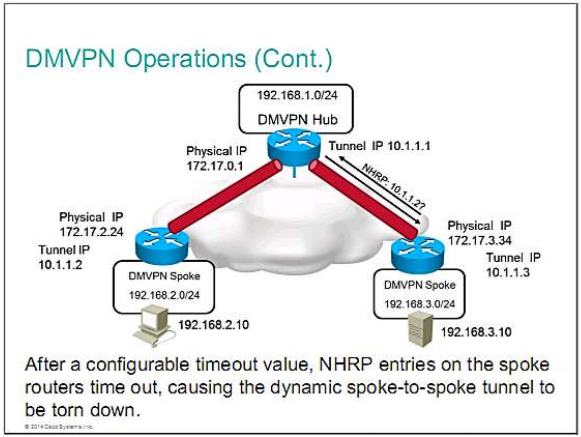


Once the IPsec tunnel is created, the left spoke will create a GRE tunnel to 172.17.3.34, and all subsequent packets from the left spoke network and the right spoke network will bypass the hub. This solution has the performance benefit of relieving the hub from decrypting and encrypting these packets, and also reduces the impact on the hub link. Note that at this point, the spoke-to-spoke tunnel can pass traffic in one direction only: from the left spoke to the right spoke.



To provide bidirectional connectivity, the right spoke router also needs next-hop information for the left spoke router from the NHRP server. The right spoke router sends an NHRP query when the first packet needs to be forwarded from a host in the right spoke network to the network behind the left spoke router. Once the right spoke router has the NHRP mapping for the left spoke router, the corresponding GRE association is built, and a response packet is sent directly to the left spoke router, encapsulated in a GRE/IPsec tunnel. The IPsec tunnel, however, is already in place and is not duplicated.

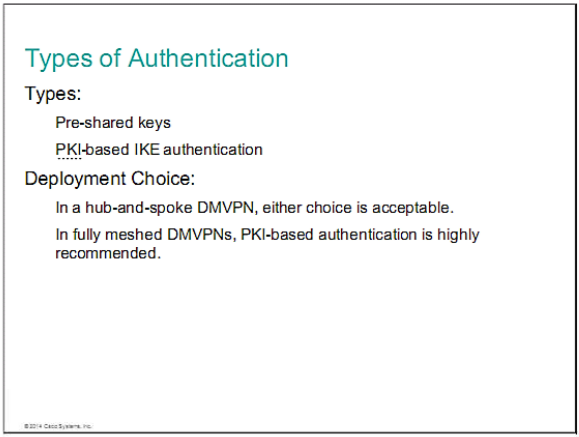
https://t.me/learningnets
© 2014 Cisco Systems, Inc.



The designer may want to tune (lower) the IPsec SA and NHRP timers to prevent too many lingering tunnels after there is no more traffic between spokes. After a timeout period, the NHRP entries will age out, triggering IPsec to break down the dynamic spoke-to-spoke tunnel because it is no longer needed.

Types of Authentication

This topic describes types of authentication.

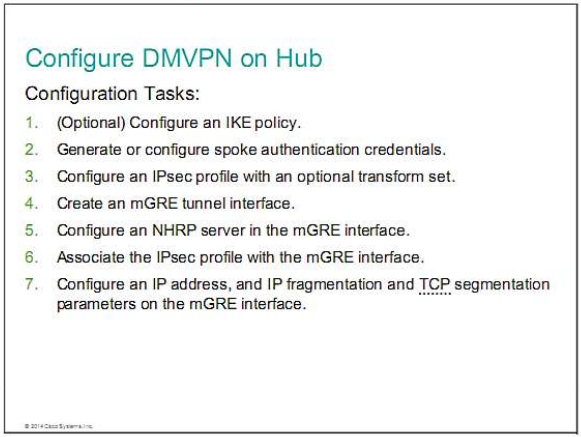


When configuring DMVPN, you will have to make some deployment choices. In a strict hub-and-spoke DMVPN, choose either PSKs or PKI-based IKE authentication, both are acceptable. In fully meshed DMVPNs, PKI-based authentication is highly recommended.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

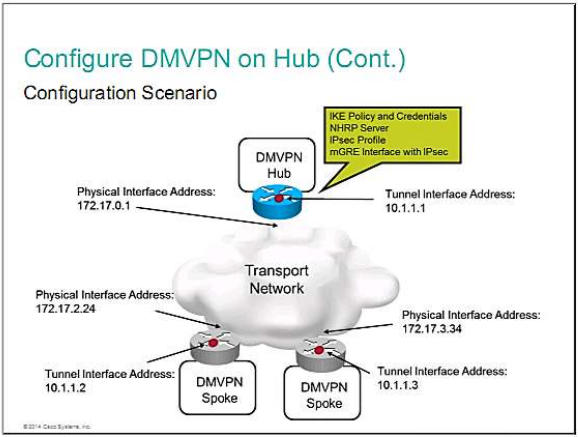
Configure DMVPN on Hub

This topic describes how to configure DMVPN on hub.



These steps represent the DMVPN hub configuration sequence:

1. (Optional) First, you can configure an IKE policy to interface with other IPsec peers. If you do not configure a specific policy, the hub will use default IKE policies.
2. Generate or configure authentication credentials to authenticate spoke devices. Typically, you will enroll the hub into a PKI to obtain a certificate for this purpose.
3. Configure an IPsec profile. and inside the profile, an optional IPsec transform set. If you do not configure a transform set, the router will use default IPsec transform sets. An IPsec profile, however, is mandatory, even if it is empty.
4. Create an mGRE tunnel interface.
5. Configure NHRP server parameters inside the mGRE tunnel interface.
6. Inside the mGRE tunnel interface, configure tunnel protection, referencing the IPsec profile.
7. Inside the mGRE tunnel interface, configure all required and optional IP parameters.

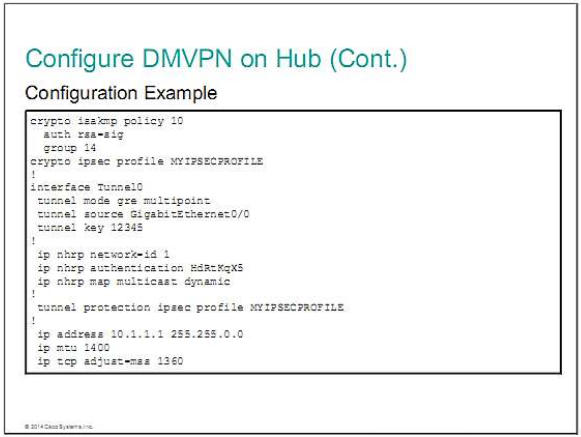


This figure shows the configuration scenario that will be used in the DMVPN hub configuration task sequence. You will configure the following features on the DMVPN hub router:

- An IKE policy and peer authentication credentials
- An NHRP server, using an NHRP network ID and authentication key
- An IPsec profile, specifying the required IPsec SA parameters for transmission protection
- A single Cisco Internetwork Operating System Software mGRE interface that will interface with IPsec to create IPsec/GRE tunnels to spokes

The addressing information that is used for the physical and logical (tunnel) interfaces is given in the figure. The configuration scenario assumes that there is complete IP connectivity between physical interfaces of all routers in the transport network.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.



To configure a DMVPN on Hub on the Cisco router, use the commands that are described in this table:

Commands	Description
crypto isakmp policy 10 **auth rsa-sig** **group 14**	Creates a new IKE policy. This example uses certificates for IKE authentication.
crypto ipsec profile MYIPSECPROFILE	Creates an IPsec profile and includes nondefault transform set in the profile.
interface Tunnel0 **tunnel mode gre multipoint** **tunnel source GigabitEthernet0/0** **tunnel key 12345**	Creates a new point-to-multipoint GRE interface. Specifies the tunnel source; this will be the IKE/IPsec endpoint IP address of the hub. Specifies a tunnel key. You can run multiple DMVPN hubs on the same router interface, each with a different tunnel key.
ip nhrp network-id 1 **ip nhrp authentication HdRtKqX5** **ip nhrp map multicast dynamic**	Creates an instance of an NHRP server and assigns it an NHRP net ID. Creates a strong NHRP authentication string. Specifies that multicast addresses (for routing protocols) will be dynamically learned from the spokes.
tunnel protection ipsec profile MYIPSECPROFILE	Assigns the configured IPsec profile to the mGRE interface.
ip address 10.1.1.1 255.255.0.0 **ip mtu 1400** **ip tcp adjust-mss 1360**	Enables IP processing on this interface by assigning it an IP address. Lowers the MTU and TCP MSS to avoid fragmentation of GRE and IPsec packets.

The first configuration task for IKE peering is the configuration of the IKE policy.

Then create a named IPsec profile, in which you can specify a transform set and encapsulation mode that the hub will use to protect its GRE tunnels.

Create an mGRE tunnel interface, and configure basic GRE settings inside its configuration mode. You have to configure the multipoint GRE mode, tunnel source IP address, and a tunnel key. The tunnel key is required, and will have to match the tunnel key configured on spoke routers, and allows you to run multiple parallel DMVPNs on the same set of devices. A GRE tunnel key therefore uniquely identifies a DMVPN network.

Create an NHRP server inside the mGRE tunnel interface. Specify an NHRP network ID, an NHRP authentication key, and dynamically register multicast capabilities if spokes so desire.

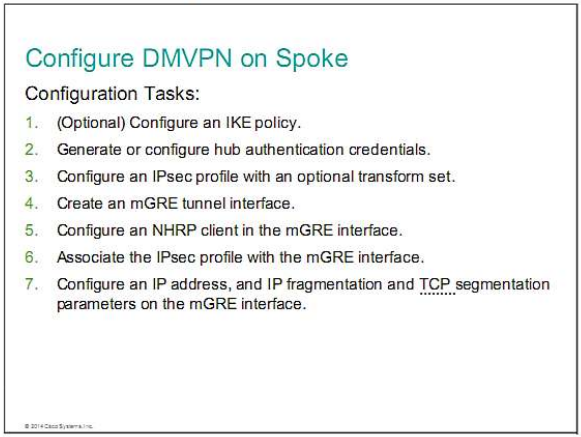
Inside the mGRE interface specify the IPsec traffic protection profile. This will indicate the set of algorithms and the IPsec encapsulation method that IKE will negotiate to protect GRE tunnels with all spokes.

Finally, inside the mGRE interface, configure an IP address, set the IP MTU size and TCP MSS value.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

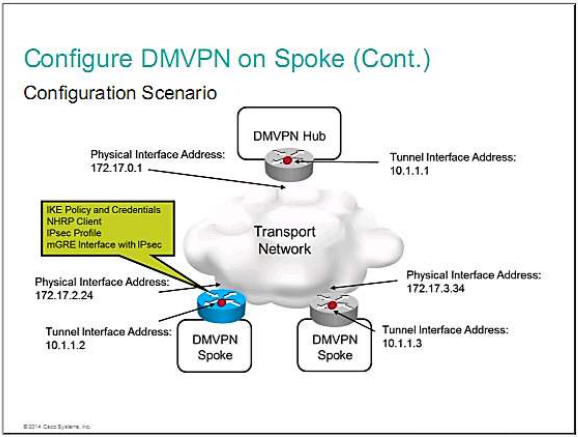
Configure DMVPN on Spoke

This topic describes how to configure DMVPN on spoke.



These steps represent the detailed configuration sequence that is needed to configure a DMVPN spoke:

1. (Optional) Configure an IKE policy to interface with the hub and other spokes. If you do not configure an IKE policy, the spoke will use default IKE policies.
2. Generate or configure authentication credentials to authenticate the hub and other spoke devices. Typically, you will enroll the spoke into a PKI to obtain a certificate for this purpose.
3. Similar to configuring the hub, configure an IPsec profile and inside the profile, an optional IPsec transform set. If you do not configure a transform set, the router will use default IPsec transform sets.
4. Create mGRE tunnel interface.
5. Configure NHRP client parameters inside the mGRE tunnel interface.
6. Inside the mGRE tunnel interface, configure tunnel protection.
7. Inside the mGRE tunnel interface, configure all required and optional IP parameters.



This figure shows the configuration scenario that will be used in the DMVPN spoke configuration task sequence. You will configure the following features on the DMVPN spoke router:

An IKE policy and peer authentication credentials

NHRP client, using an NHRP network ID and authentication key, and the location of the NHRP server.

An IPsec profile, specifying the required IPsec SA parameters for transmission protection.

A single Cisco Internetwork Operating System Software mGRE interface that will interface with IPsec to create IPsec/GRE tunnels to the hub and other spokes; alternatively, you can also configure a point-to-point GRE interface to the hub, if you plan to use this DMVPN in a strict hub-and-spoke mode.

The addressing information that is used for the physical and logical (tunnel) interfaces is given in the figure. The configuration scenario assumes that there is full IP connectivity between physical interfaces of all routers in the transport network.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure DMVPN on Spoke (Cont.)

Configuration Example

```
crypto isakmp policy 10
  auth rsa-sig
  group 14
crypto ipsec profile MYIPSECPROFILE
!
interface Tunnel0
 tunnel mode gre multipoint
 tunnel source GigabitEthernet0/0
 tunnel key 12345
!
 ip nhrp network-id 1
 ip nhrp authentication HdRtKqX5ZdYaZQjSYPqb
 ip nhrp nhs 10.1.1.1
 ip nhrp map multicast 172.17.0.1
 ip nhrp map 10.1.1.1 172.17.0.1
!
 tunnel protection ipsec profile MYIPSECPROFILE
 ip address 10.1.1.2 255.255.0.0
 ip mtu 1400
 ip tcp adjust-mss 1360
```

© 2014 Cisco Systems, Inc.

To configure a DMVPN on Spoke on the Cisco router, use the commands that are described in this table:

Commands	Description
crypto isakmp policy 10 **auth rsa-sig** **group 14**	Creates a new IKE policy. This example uses certificates for IKE authentication.
crypto ipsec profile MYIPSECPROFILE	Creates an IPsec profile and includes nondefault transform set in the profile.
interface Tunnel0 **tunnel mode gre multipoint** **tunnel source GigabitEthernet0/0** **tunnel key 12345**	Creates a new point-to-multipoint GRE interface. Specifies the tunnel source; this will be the IKE/IPsec endpoint IP address of the spoke. Specifies a tunnel key that matches the tunnel key of the hub.
ip nhrp network-id 1 **ip nhrp authentication HdRtKqX5ZdYaZQjSYPqb** **ip nhrp nhs 10.1.1.1** **ip nhrp map multicast 172.17.0.1** **ip nhrp map 10.1.1.1 172.17.0.1**	Creates an instance of an NHRP client.
tunnel protection ipsec profile MYIPSECPROFILE	Assigns the configured IPsec profile to the mGRE interface.
ip address 10.1.1.2 255.255.0.0 **ip mtu 1400** **ip tcp adjust-mss 1360**	Enables IP processing on this interface by assigning it an IP address. Lowers the MTU and TCP MSS to avoid fragmentation of GRE and IPsec packets.

The first configuration task for IKE peering is the configuration of the IKE policy.

Then create a named IPsec profile, in which you can specify a transform set and encapsulation mode that the spoke will use to protect its GRE tunnels.

Create an mGRE tunnel interface, and configure basic GRE settings inside its configuration mode. You have to configure the multipoint GRE mode, tunnel source IP address, and a tunnel key. The tunnel key will have to match the tunnel key configured on the hub.

Create an NHRP client inside the mGRE tunnel interface. Specify an NHRP network ID, an NHRP authentication key, the hub as the NHRP server, and dynamically register the spoke multicast capabilities with the hub to enable dynamic routing protocol peering, if desired.

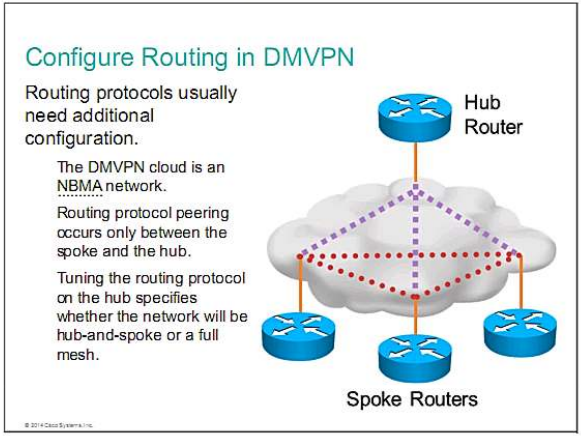
Inside the mGRE interface specify the IPsec traffic protection profile. This will indicate the set of algorithms and the IPsec encapsulation method that IKE will negotiate to protect GRE tunnels with a hub.

Finally, inside the mGRE interface, configure an IP address, set the IP MTU size and TCP MSS value.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Routing in DMVPN

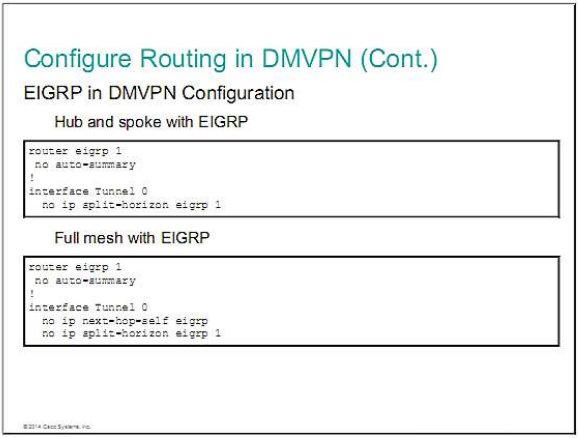
This topic describes how to configure routing in DMVPN.



Some considerations must be made when running dynamic routing protocols across the DMVPN, especially when you are implementing a spoke-to-spoke design, because the DMVPN cloud is like an NBMA network. Many routing protocols have an IP multicast mechanism that is used to discover other participating nodes. For this purpose, you must configure NHRP multicast maps on spoke routers to register this capability with the hub. You configure the hub router with a dynamic multicast map that will replicate multicast packets to all spokes. This configuration allows the hub and spokes to exchange multicast and broadcast information, but does not permit spokes to receive broadcasts from other spokes.

In a DMVPN, routing protocol peering (neighbor relationships) occurs only between each spoke and the hub. There is never spoke-to-spoke peering. Therefore, the hub will pass information about all spoke networks to other spokes.

It is important to note that the tuning of the routing protocol functions on the hub is the main factor that will determine whether your DMVPN will operate in strict hub-and-spoke or partially meshed or fully meshed topology. The only other factor is the choice of the GRE interface for the spokes: a point-to-point GRE interface on the spoke will only allow a strict hub-and-spoke topology, while a mGRE interface on the spokes will allow the use of both topologies.



To configure an EIGRP routing in DMVPN, use the commands that are described in this table:

Commands	Description
router eigrp 1 **no auto-summary**	Disables route summarization.
interface Tunnel 0 **no ip split-horizon eigrp 1**	Enable propagation of spoke routes to other spokes via the hub.
interface Tunnel 0 **no ip next-hop-self eigrp**	Preserve spoke next-hop addresses in advertisements.

You can use EIGRP to deploy either a strict hub-and-spoke DMVPN, or a partially meshed or fully meshed DMVPN.

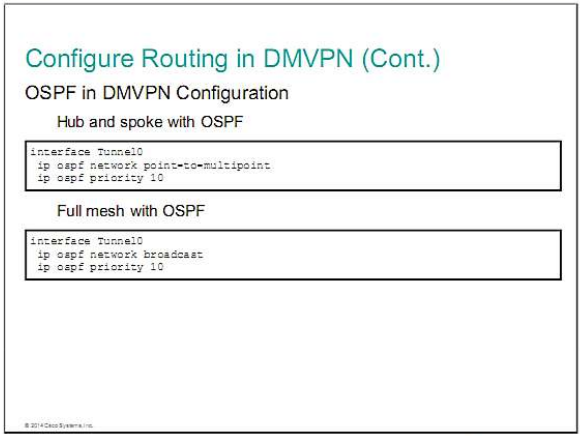
EIGRP deployment is straightforward in a strict hub-and-spoke DMVPN deployment. You should typically disable automatic summarization on the hub, and disable the EIGRP split-horizon feature in order for the hub to propagate spoke network information to other spokes because they are reachable over the same mGRE interface.

Spoke-to-spoke, partially meshed, or fully meshed DMVPN networks present a unique challenge because the spokes cannot directly exchange information with one another, even though they are on the same logical subnet. This limitation first requires the hub router to advertise the subnets from other spokes on the same subnet. In addition, the advertised route must contain the original next hop (which is the spoke router) as learned by the hub router.

Note The outside address space of the tunnel should not be included in any protocol running inside the tunnel.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

As with all EIGRP networks, the number of neighbors should be limited to ensure that the hub router can re-establish communications after a major outage. If the DMVPN cloud is configured with a /24 network prefix, the neighbor count is limited to 254, which is a safe operational limit. Beyond this number, a compromise is required to balance reconvergence with recovery. In large EIGRP networks, it might be necessary to adjust the EIGRP hold time to allow the hub more time to recover without thrashing.



To configure a OSPF routing in DMVPN, use the commands that are described in this table:

Commands	Description
interface Tunnel0 **ip ospf network point-to-multipoint**	Creates the tunnel network as a point-to-multipoint network.
interface Tunnel0 **ip ospf priority 10**	Ensures that the hub will be the designated router.
interface Tunnel0 **ip ospf network broadcast**	Creates the tunnel network as a broadcast network.

Running OSPF over a DMVPN network has some of the same challenges as running OSPF over other types of networks. Because only the hub is in direct communications with all of the branches, it should be configured as the DR on the DMVPN subnet. There is not typically a BDR for this type of configuration. A BDR is possible if a second hub is placed on the same subnet.

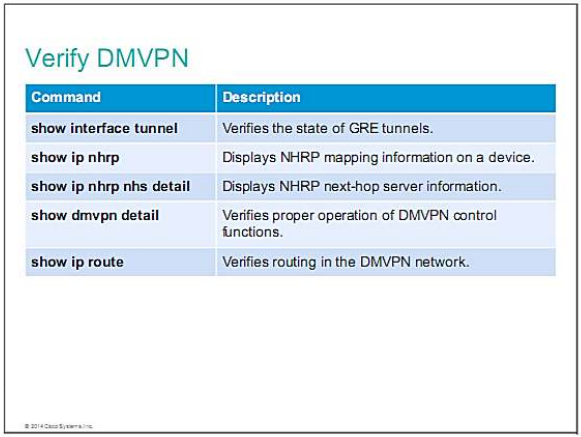
In strict hub-and-spoke DMVPNs, you should include the tunnel interface in the OSPF routing process, and configure the tunnel interface as a point-to-multipoint OSPF network type on the hub router, and as a point-to-point network type on the branch routers. In this case, there is no need to elect a DR on the DMVPN subnet. The hub router serves as the master for the subnet. The branches consider the hub as the only path off the subnet, thus simplifying the Dijkstra's algorithm for the OSPF area. Use a suitable OSPF priority to ensure that the hub is always elected as the DR.

To create a partially meshed or fully meshed DMVPN, configure the mGRE tunnel on the hub router as an OSPF broadcast network. Each spoke router should be configured with an OSPF priority of 0 to prevent a spoke from becoming the DR. In addition, if the spoke is configured with point-to-point GRE and the hub is configured with an mGRE interface, the hello timer on the spoke should be changed from the default of 10 seconds to 30 seconds to match the hello timers on the mGRE interface. The tunnel IP MTU must match on all of the GRE interfaces that are OSPF-adjacent. In addition, OSPF areas running over DMVPN can be not-so-stubby or totally stubby areas to reduce LSA flooding over the WAN.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify DMVPN

This topic describes how to verify DMVPN.



Cisco Internetwork Operating System Software includes DMVPN validation features.

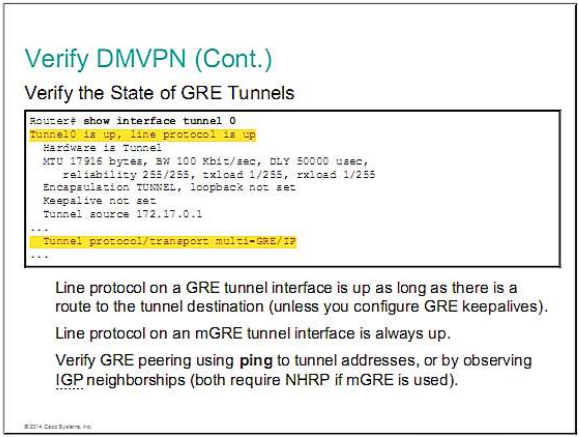
Use the **show interface tunnel** command to verify the state of GRE tunnels.

Use the **show ip nhrp** command to display NHRP mapping information on a device.

Use the **show ip nhrp nhs detail** command to display NHRP next-hop server information.

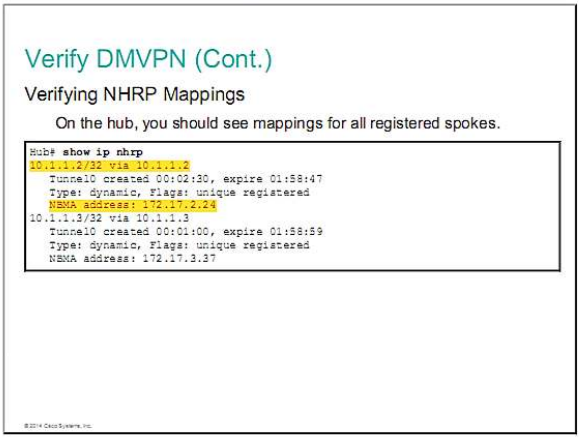
Use the **show dmvpn detail** command to verify proper operation of DMVPN control functions.

Use the **show ip route** command to verify routing in the DMVPN network.

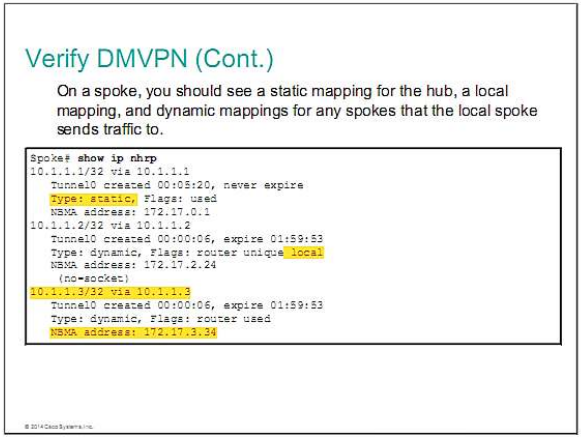


After configuring hub-and-spoke GRE tunnel, verify the state of it by using the **show interface tunnel** command. The line protocol on a GRE tunnel interface is up as long as there is a route to the tunnel destination, unless you configure GRE keepalives (which are not recommended if you will also use a dynamic routing protocol over the tunnel). The line protocol on an mGRE tunnel interface is always up.

You can quickly verify that a GRE tunnel is actually passing traffic by using the ping command and specifying the (inner) tunnel IP addresses of the remote GRE peer as target IP addresses. Another option is to observe if a routing protocol neighbor relationship establishes successfully over the tunnel (if you have configured a routing protocol). Again, both approaches require NHRP if mGRE is used.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.



On the hub (the NHRP server), you should see NHRP mappings for all registered spokes. In this printout, you can see two mappings: 10.254.0.2 is IP address of the tunnel interface of the first spoke, mapped to its physical interface IP address 172.17.2.24, while 10.254.0.3 is the IP address of the tunnel interface of second spoke, mapped to its physical interface IP address 172.17.3.37.

On a spoke (the NHRP client), you should see a static mapping for the hub, a local mapping for the local spoke, and dynamic mappings for any spokes that the local spoke is sending traffic to. In this printout, you can see static server mapping, where 10.254.0.1 is the IP address of the tunnel interface of the hub, mapped to its physical interface IP address 172.17.0.1. You can also see dynamic spoke mapping, where 10.254.0.3 is IP address of the tunnel interface of the second spoke, mapped to its physical interface IP address 172.17.3.37.

Verify DMVPN (Cont.)

Debugging NHRP

On the spoke, shows issues with NHRP registration (most likely a network ID or authentication key mismatch).

```
Spoke# show ip nhrp nhs detail
Legend: E=Expecting replies, R=Responding
Tunnel0:      172.17.0.1  E  req-sent 0  req-failed 120 repl-recv 0
Pending Registration Requests:
Registration Request: Reqid 1201, Ret 64  NHS 172.17.0.1
```

©2014 Cisco Systems, Inc.

If the **req-failed** field displays a number greater than 0, this is a good hint that there is a network ID or authentication key mismatch, or a connectivity problem with the next-hop server.

Verify DMVPN (Cont.)

Verify Registered Spoke Details

Shows the details of all spoke-to-hub tunnels, including traffic statistics

```
Hub# show dmvpn detail
...
Type:Hub, Total NBMA Peers (v4/v6): 1
# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1    172.24.2.24  10.1.1.2        UP    00:01:29 D     10.1.1.2/32
...
Interface: Tunnel0
Session: [0x682CE3A4]
  IKE SA: local 172.17.0.1/500 remote 172.17.2.24/500 Active
          Capabilities:(none) connid:1004 lifetime:23:58:29
  Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 172.17.2.24
  IPSEC FLOW: permit 47 host 172.17.0.1 host 172.17.2.24
        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 11 drop 1 life (KB/Sec) 4570629/3510
        Outbound: #pkts enc'ed 11 drop 0 life (KB/Sec) 4570629/3510
```

©2014 Cisco Systems, Inc.

In the output, you should see entries for all connected spokes. The additional per-spoke information available is as follows:

The spoke NBMA address, on the physical interface or loopback, where tunnel is attached

The spoke tunnel address, as set on the spoke GRE or mGRE tunnel

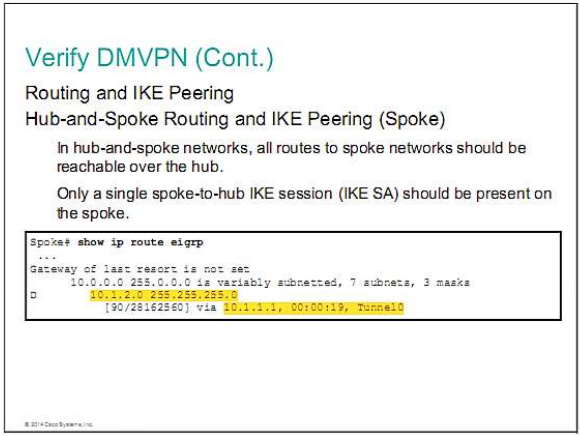
The tunnel state, **UP** for successfully established tunnel

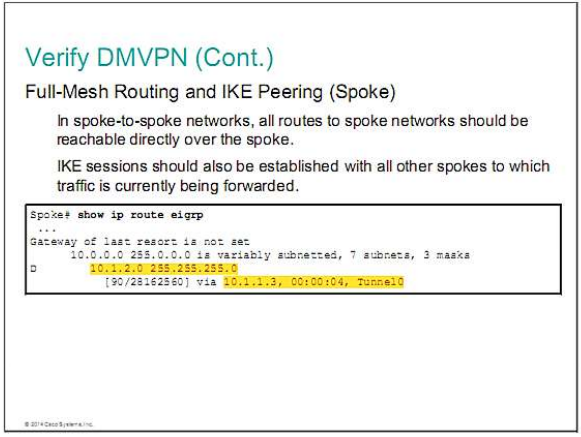
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The tunnel uptime

The target network, which represents the NHRP client address of the spoke

The **show dmvpn detail** command will reveal additional information about DMVPN peers. In addition to information already displayed by **show dmvpn**, detailed output includes IKE SA information: local and remote IP addresses and IKE status. IPsec statistics will show the number of packets that the router has encrypted and decrypted. If these counters are not increasing and traffic is not flowing through the DMVPN, there could be a problem with IPsec SAs or routing.





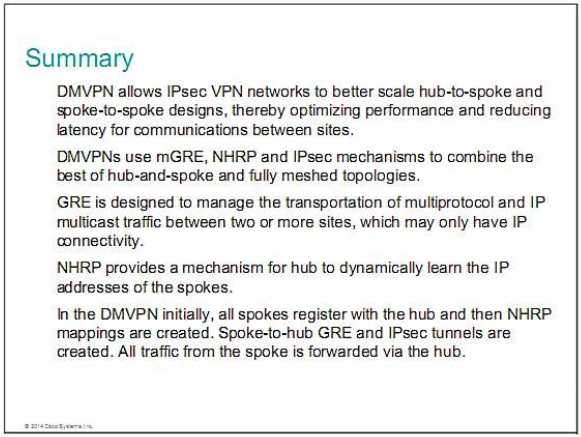
When verifying routing in strict hub-and-spoke DMVPNs, issue the **show ip route** command on a spoke to verify that all other spoke networks are routed through the hub. The next hop for each spoke prefix should be the tunnel (inner) IP address for the hub. You can also check that there is only one IKE SA and that it is established with the hub, by using the **show crypto isakmp sa** command.

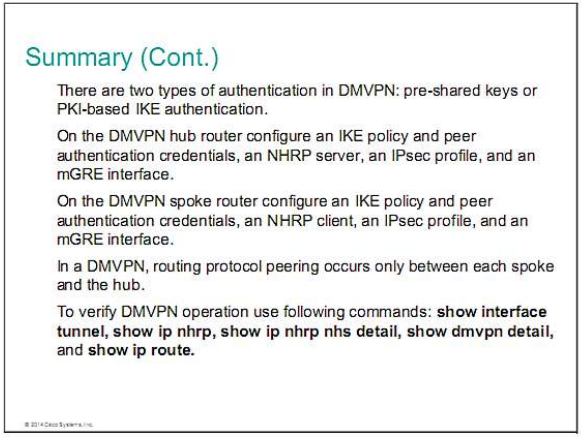
When verifying routing in a partially meshed or fully meshed DMVPN, issue the **show ip route** command on a spoke and check that all spoke networks are routed directly toward the spokes, with the their tunnel (inner) IP addresses shown as next hops for spoke prefixes. There should be an IKE session with the hub, and potentially IKE sessions with some other spokes as well, depending on with which spokes the local spoke is currently exchanging traffic.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.

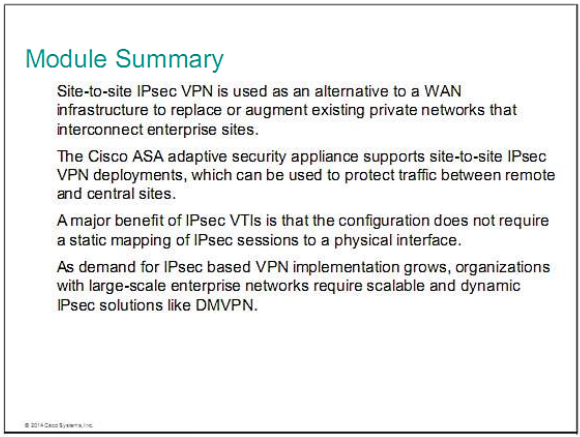




https://t.me/learningnets

Module Summary

This topic summarizes the key points that were discussed in this module.



https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Self-Check

Questions

Use the questions here to review what you learned in this module.

1. Which network topology is in use when every network has a direct VPN connection to every other network? This topology provides any-to-any communication and provides the most optimal direct path for network traffic.

   A. fully meshed network
   B. star topology network
   C. partially meshed network
   D. individual point-to-point VPN connection
   E. hub-and-spoke network

2. Which of the following is the appropriate encapsulation type on a virtual tunnel interface?

   A. Frame Relay
   B. ATM
   C. AH or ESP
   D. ISAKMP
   E. HDLC

3. Where are dynamic point-to-point VTI tunnels deployed?

   A. on the hub router
   B. on each spoke router
   C. on the hub router and on each spoke router
   D. on the VPN concentrator
   E. none of the above

https://t.me/learningnets

4. Which of the following interface commands is used to configure the IP address of a virtual tunnel interface?

   A. **ip address**
   B. **ip address dhcp**
   C. **ip address pppoe**
   D. **ip unnumbered**

5. Which mechanism provides dynamic mutual discovery of spoke devices?

   A. GRE
   B. IKE
   C. NHRP
   D. DHCP
   E. Expired Certificate List

6. When deploying an IPsec site-to-site VPN, what is the recommended method of peer authentication from a security perspective?

   A. pre-shared keys
   B. digital certificates
   C. biometrics
   D. OTP

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Answer Key

1. A
2. C
3. A
4. D
5. C
6. B

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module 3

Deploying Cisco IOS Site-to-Site FlexVPN Solutions

Introduction

The Cisco FlexVPN is a unified VPN solution created to simplify the deployment of VPNs, to address the complexity of multiple solutions, and, as a unified ecosystem, to cover all types of VPNs, such as remote access, teleworker, site-to-site, mobility, and managed security services.

This module describes the FlexVPN implementation for point-to-point, hub-and-spoke, and spoke-to-spoke VPNs on Cisco IOS routers.

Upon completing this module, you will be able to meet these objectives:

- Evaluate site-to-site VPN technologies
- Describe the use of FlexVPN in point-to-point IPsec VPNs
- Describe the hub-and-spoke connectivity scenario that can be implemented using the FlexVPN framework
- Describe the spoke-to-spoke connectivity scenario that can be implemented using the FlexVPN framework

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Introducing Cisco FlexVPN Solution

FlexVPN is a new framework to configure IPsec VPN with IKEv2 on Cisco IOS platforms. FlexVPN combines multiple approaches (crypto maps, Easy VPN, DMVPN) into a single CLI. It offers more flexibility and provides for future functionality. This lesson describes the main components of the FlexVPN solution.

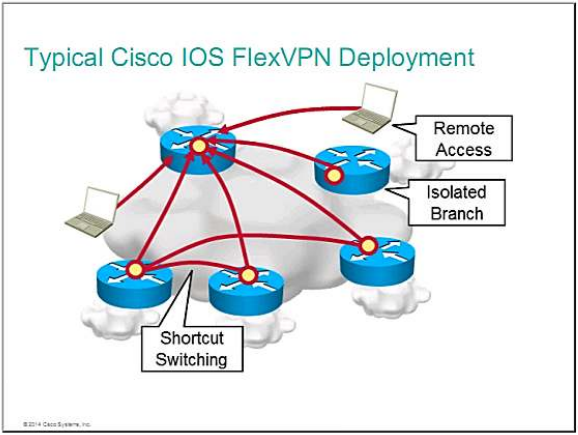
Upon completing this lesson, you will be able to meet these objectives:

- Provide an overview of FlexVPN
- Describe PKI
- Describe FlexVPN architecture
- Describe FlexVPN configuration blocks
- List FlexVPN capabilities
- Provide an overview of IKEv2
- Describe IKEv2 message exchange
- Describe IKEv2 DoS prevention
- Compare IKEv1 and IKEv2 features
- Describe FlexVPN use cases

https://t.me/learningnets

FlexVPN Overview

This topic provides an overview of FlexVPN.

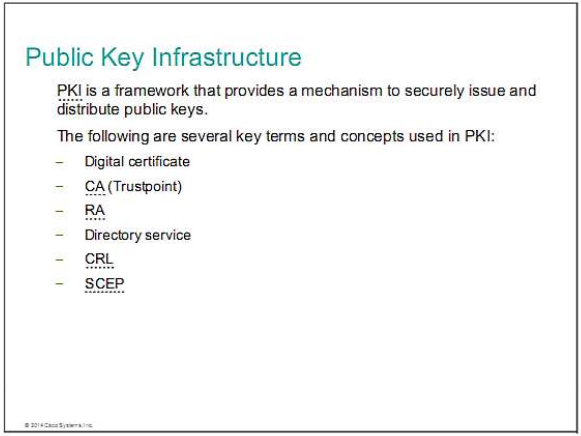


Large organizations deploying IPsec VPN over IP networks are faced with high complexity and high cost of deploying multiple types of VPN to meet different types of connectivity requirements. Network administrators often have to learn different types of VPNs to manage and operate different types of networks. After selecting a technology for deployment, they often avoid migrating or adding functionality to enhance the VPN. FlexVPN was created to simplify the deployment of VPNs, to address the complexity of multiple solutions, and as a unified ecosystem to cover all types of VPN: remote access, teleworker, site to site, mobility, managed security services, and others, as illustrated in the figure.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Public Key Infrastructure

This topics describes Public Key Infrastructure.



PKI is a framework that provides a mechanism to securely issue and distribute public keys. PKI arrangement enables the transmission of secure data over insecure public networks by using trusted public and private cryptographic key pairs, which are obtained and shared through a trusted authority.

PKI provides a digital certificate that can identify an individual or an organization by binding of public keys to users. When used in IPsec VPN, PKI provides authentication, nonrepudiation, and confidentiality services.

PKI consists of various protocols, services, and standards that support applications of the public key infrastructure.

The public key infrastructure assumes the use of the public key cryptography system, also known as asymmetric cryptography, which is based on two-key pairs: one key for encryption and another for decryption.

PKI components include the following:

Digital certificate (also known as identity certificate): A digital certificate uses a digital signature to bind a public key with a user's identity to produce a public key certificate. A certificate contains information, such as the certificate validity period, peer identity information, encryption keys that are used for secure communications, and the signature of the issuing CA. A certificate is a digital equivalent of a passport.

Certificate Authority: A CA is the major component in the PKI system. The CA issues and verifies digital certificates. A certificate includes the public key that is bound to a user identity. A CA is also known as a trustpoint, which manages certificate requests and issues certificates to participating network devices. Many vendors offer CA Server solutions, including Cisco. The Cisco IOS Router can be configured as the certificate server. CA implementation can be either an in-house implementation or be outsourced.

A trustpoint is basically a certificate authority whom you trust. The idea is that by trusting a given self-signed certificate, your PKI system will automatically trust any other certificates signed with that trusted certificate. A trustpoint certificate is a self-signed certificate, hence the name trustpoint, since it does not rely on the trust of anyone else or other party.

Registration Authority (optional): RAs provide an interface between the user and the CA. The RA acts as the verifier for the certificate authority before a digital certificate is issued to a requestor.

Directory service: The certificates with their public keys are held in directory services.

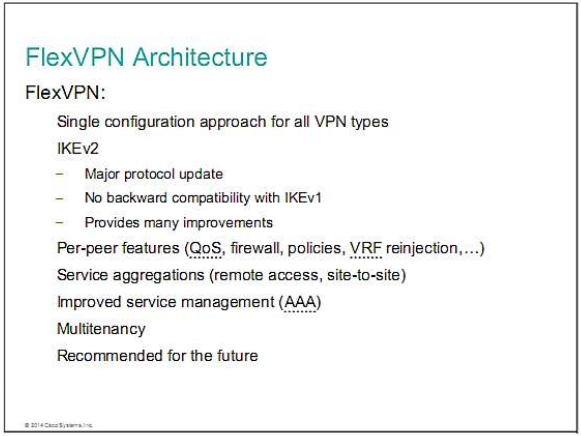
Certificate Revocation List: The CRL is a list of serial numbers of revoked certificates. A CA revokes certificates that are no longer used and publishes the CRL.

Simple Certificate Enrollment Protocol: SCEP is a Cisco proprietary certificate enrollment protocol used on network devices to obtain digital certificates from the CA server. SCEP uses HTTP to communicate with the CA or RA. SCEP is the most commonly used method for sending and receiving requests and certificates.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

FlexVPN Architecture

This topic describes the need for a common VPN configuration approach.



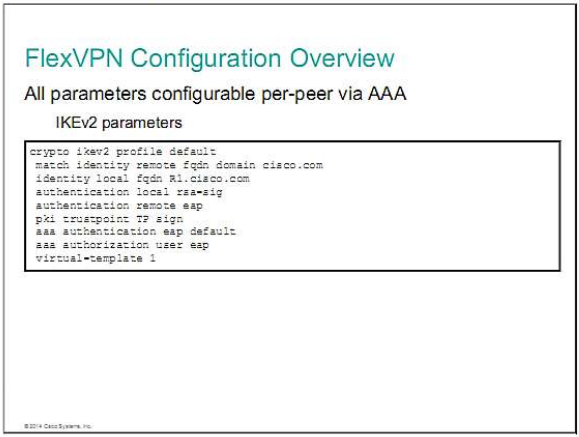
FlexVPN is the Cisco implementation of the IKEv2 standard that features a unified paradigm and CLI that combines site-to-site, remote access, hub-and-spoke topologies, and partial meshes (spoke-to-spoke direct). FlexVPN offers a simple but modular framework that extensively uses the tunnel interface paradigm while remaining compatible with legacy VPN implementations using crypto maps.

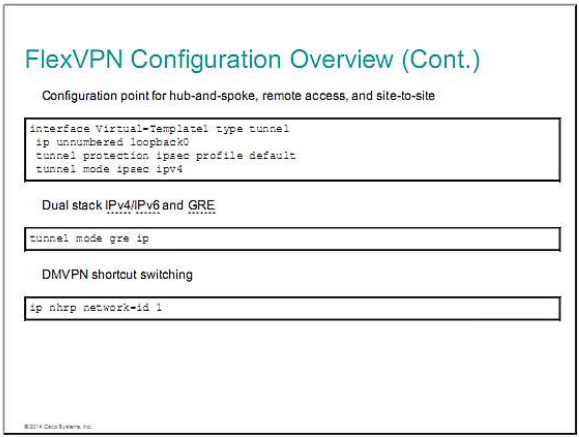
IKEv2, a next-generation key management protocol based on RFC 4306, is an enhancement of the IKE protocol. IKEv2 is used for performing mutual authentication and establishing and maintaining SAs.

FlexVPN supports per-peer configurations of service parameters, such as QoS, firewall mechanisms, policies, and VRF-related settings. It is ideal for service aggregation that encompasses both remote access and site-to-site VPNs. It provides improved service management through integration with external AAA databases, and works well in multitenancy scenarios. FlexVPN is the future technology of choice.

FlexVPN Configuration Overview

This topic provides an overview of the FlexVPN configuration.





The sample configuration in the figure shows the main configuration blocks of the FlexVPN solution.

The first configuration blocks contain the IKEv2 settings that define the authentication, encryption, and integrity proposals in a way that is similar to the methods in IKEv1. In addition, the IKEv2 may refer to an interface or template, in this example to virtual-template 1.

https://t.me/learningnets

The interface or interface template settings are responsible for the configuration of a particular VPN type, such as remote access or site to site.

The tunnel mode configuration enables you to configure various encapsulation methods, including GRE and dual-stack IPv4/IPv6.

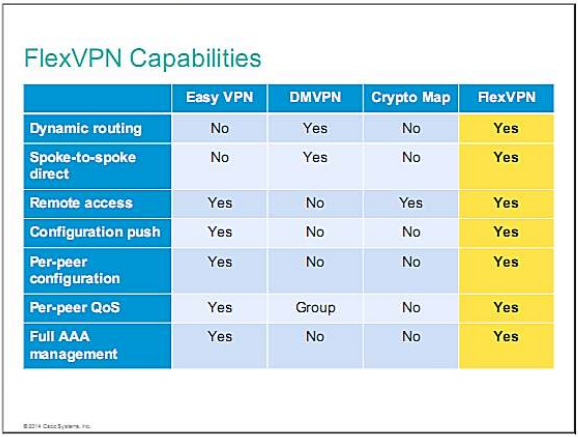
NHRP configuration enables spoke-to-spoke shortcut switching, which is the main characteristic of the DMVPN solution.

All parameters are tunable per-peer via AAA.

https://t.me/learningnets

FlexVPN Capabilities

This topic describes the FlexVPN capabilities.

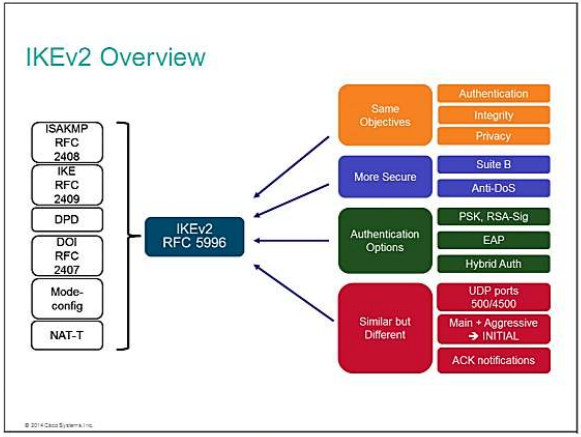


The table summarizes the capabilities of the Cisco IOS FlexVPN solution and contrasts them with traditional configuration approaches.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IKEv2 vs. IKEv1 Overview

This topic provides an overview of IKEv2.



IKE is used to negotiate security associations between a pair of communicating peers. It is defined as part of IPsec. IPsec is currently in its third generation of specification. The first generation (RFCs 1825-1829) did not specify a specific key management protocol. IKEv1 was included in the second generation of specification (RFCs 2401-2412). Specifically, IKEv1 was specified in RFCs 2407, 2408, and 2409. It was considered complex and had some ambiguities that have led to interoperability problems. The third generation (RFCs 4301-4309) introduced IKEv2. IKEv2 attempts to address the specification difficulties, security issues, and protocol inefficiencies of IKEv1.

IKEv1 is more mature and more widely deployed than IKEv2; however, IKEv2 is meant to replace IKEv1. IKEv2 is defined in RFC 5996, which consolidates several RFCs that describe IKEv1 and its add-ons, such as DPD, NAT-T, mode config, and others.

IKEv2 shares the same main objectives with IKEv1: authentication, privacy, and integrity.

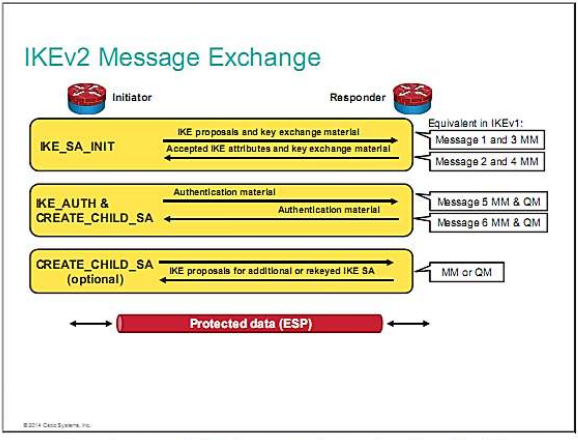
IKEv2 is more secure than IKEv1 because it supports Suite B, a set of cryptographic algorithms promulgated by the U.S. National Security Agency as part of its Cryptographic Modernization Program. It is also more resistant to DoS attacks.

IKEv2 supports more authentication methods than IKEv1. In addition to pre-shared keys, and certificate-based and hybrid authentication types, the responder can use EAP to authenticate the initiator.

Other key differences include the consolidation of the main and aggressive modes into one method, called Initial; introduction of acknowledged notifications; and different packet formats transported over the same two UDP ports: 500 and 4500.

IKEv2 Message Exchange

This topic describes the IKEv2 message exchange.



IKEv2 message exchange is open-ended and may contain more than the minimal four messages shown in the upper part of the figure.

The first two messages represent the IKE_SA_INIT phase and exchange IKEv2 proposals and key material information. The first message corresponds to IKEv1 main mode messages 1 and 3. The second message corresponds to IKEv1 main mode messages 2 and 4. The key exchange material is generated based on the DH group with highest priority.

The third and fourth messages represent the IKE_AUTH and CREATE_CHILD_SA phases. They are used to authenticate the peers to each other and generate the SA parameters. The third message corresponds to IKEv1 main mode message 5 and part of the IKEv1 quick mode (phase 2). The fourth message corresponds to IKEv1 main mode message 6 and the other part of the IKEv1 quick mode.

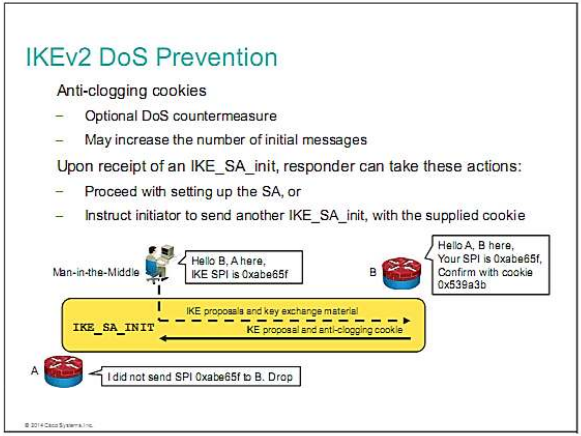
The third exchange shown in the figure (CREATE_CHILD_SA) is optional. This exchange is performed when an additional SA needs to be created, or an existing SA is rekeyed. In IKEv1, the entire IKE exchange or just the quick mode was run to refresh an ISAKMP SA, depending on the PFS configuration.

The aim of the IKEv2 exchange is to create an SA that is used to secure the traffic exchanged between the two sites, using, most commonly, the ESP protocol. A significant difference to IKEv1 is the absence of the quick mode (IKEv1 second phase).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IKEv2 DoS Prevention

This topic describes the IKEv2 defenses against DoS attacks.

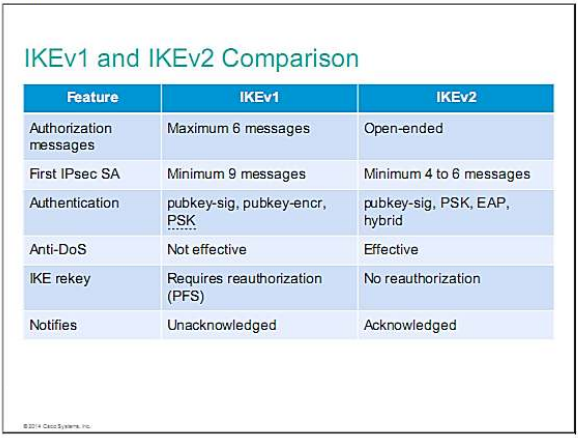


IKEv2 uses stateless cookies. The main purpose of cookies in IKEv2 is for a limited protection from DoS attacks. Receipt of a request to start an SA can consume substantial resources. A likely DoS attack against IKE consists of overwhelming a system with a large number of SA requests from forged IP addresses. Such an attack can consume CPU resources by doing the cryptography, and memory resources by remembering the state of the "half open" connections until they time out. The responder limits the resources that it is willing to devote to new connections by doing a minimal computation and allocating no memory until the Initiator has proven that it can receive messages at the address that it claims to be sending from.

IKEv2 anti-clogging cookies are an optional anti-DoS mechanism. The cookies are exchanged in the first two messages. Upon receipt of an IKE_SA_init, the responder may either proceed with setting up the SA or may tell the Initiator to send another IKE_SA_init, this time providing the cookie supplied in the first packet from the responder to the initiator.

IKEv1 and IKEv2 Comparison

This topic compares IKEv1 and IKEv2.



Some of the major similarities and differences between IKEv1 and IKEv2 are identified below.

Both protocols utilize two phases. The first phase in each is used to create the IKE_SA. The second phase is used to establish child SAs using the IKE_SA. In IKEv2, the first child SA is piggybacked on the IKE_AUTH exchange that is used to complete the mutual peer authentication.

Both protocols provide identify protection, denial-of-service protection mechanism, and perfect forward secrecy.

Negotiation of the IKE_SA required 6 messages in IKEv1 and 4 messages in IKEv2. When EAP is used in IKEv2, an additional 2 messages may be required.

Negotiation of the first CHILD_SA required 3 messages in IKEv1 and no messages in IKEv2 since it is piggybacked onto the negotiation of the IKE_SA.

Subsequent CHILD_SAs require 3 messages in IKEv1 and 2 messages in IKEv2.

IKEv1 and IKEv2 run over UDP port 500.

IKEv2 supports NAT traversal using UDP port 4500.

IKEv2 introduces retransmission and acknowledgement functions.

IKEv1 uses the more generic ISAKMP protocol specification so support for protocols other than IPsec is more evident.

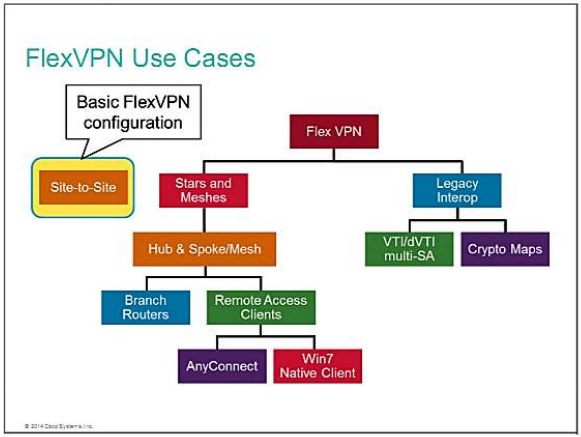
IKEv2 is written specifically for IPsec. However, it maintains support for other protocols by allowing for the private specification of SA payloads between peers.

Security Association lifetimes were explicitly negotiated in IKEv1 but are not in IKEv2. Each peer in IKEv2 maintains its own local policy for Security Association lifetime. When the lifetime is about to expire, a rekeying operation is initiated.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

FlexVPN Use Cases

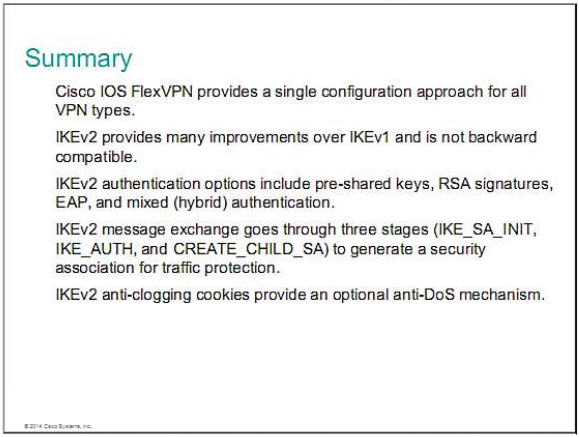
This topic describes the applications of FlexVPN in various VPN environments.



The figure shows the applications of FlexVPN in different VPN topologies and scenarios. The diagram shows that Cisco IOS FlexVPN provides a universal configuration approach that can be used to replace the traditional implementation methods for crypto maps, VTIs, Easy VPN, DMVPN and GET VPN.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 2

Deploying Point-to-Point IPsec VPNs Using Cisco IOS FlexVPN

This lesson describes the IKEv2 CLI that is required for FlexVPN site-to-site functionality. The lesson explains IKEv2 commands, describes IKEv2 smart defaults, and the mandatory IKEv2 commands that are required for FlexVPN site-to-site functionality.

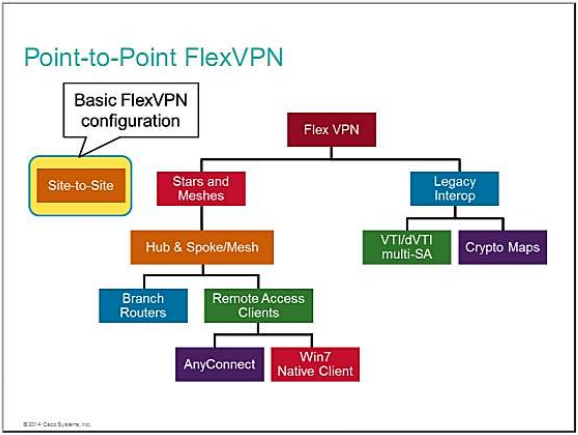
Upon completing this lesson, you will be able to meet these objectives:

- Describe the various applications of Cisco IOS FlexVPN
- Describe Cisco IOS FlexVPN configuration blocks
- Describe and modify IKEv2 smart defaults
- Describe a basic FlexVPN scenario: Point-to-Point VPN with IPv4 Static Routes
- Configure a basic FlexVPN: Point-to-Point VPN with IPv4 Static Routes
- Describe a basic FlexVPN scenario: Point-to-Point FlexVPN with OSPFv3
- Configure a basic FlexVPN: Point-to-Point FlexVPN with OSPFv3
- Describe an advanced FlexVPN scenario between a Cisco IOS router and a Cisco ASA adaptive security appliance
- Configure an advanced FlexVPN between a Cisco IOS router and a Cisco ASA adaptive security appliance
- Verify FlexVPN configurations

https://t.me/learningnets

Point-to-Point FlexVPN

This topic describes the various applications of FlexVPN.

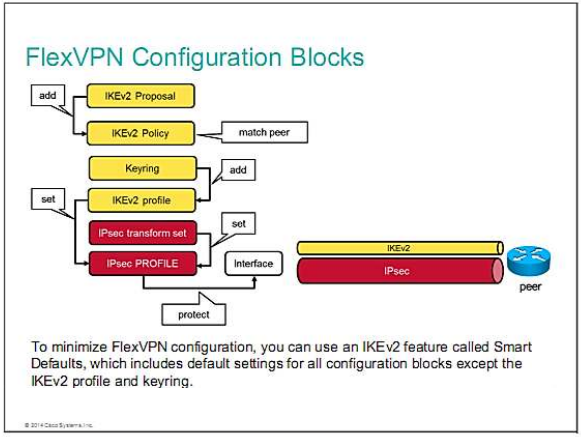


The figure illustrates the various applications of the Cisco IOS FlexVPN framework. A point-to-point VPN is the most basic scenario of the cases that are shown.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

FlexVPN Configuration Blocks

This topic describes the FlexVPN configuration blocks.



The figure shows the main FlexVPN configuration blocks:

IKEv2 proposal: Defines the protection attributes to be used in the negotiation of the IKEv2 SA. After you create an IKEv2 proposal, attach it to an IKEv2 policy so that the proposal is picked for negotiation.

IKEv2 policy: Binds the proposal to a VPN peer. The IKEv2 policy references the IKEv2 proposal.

IKEv2 keyring: Used to define preshared keys, which can be asymmetric

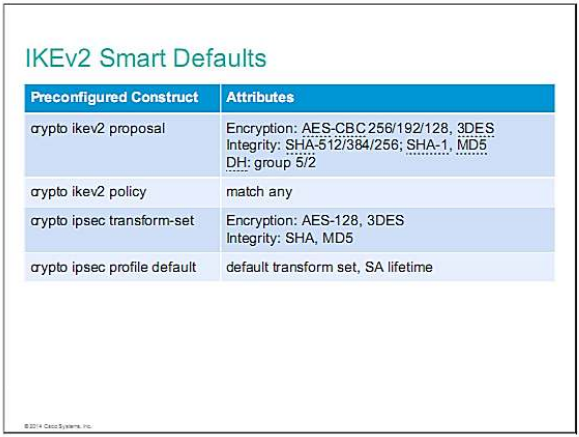
IKEv2 profile: A repository of nonnegotiable parameters of the IKE SA, such as the VPN peer address and authentication methods to be used. There is no default IKEv2 profile, so you must configure one and attach it to an IPsec profile on the initiator. If pre-shared key authentication is used, the IKEv2 profile references the IKEv2 keyring.

IPsec transform set: Specifies an acceptable combination of security protocols and algorithms for the IPsec SA

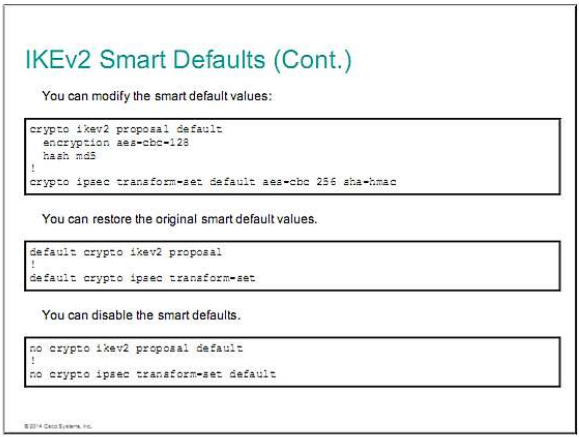
IPsec profile: Summarizes FlexVPN settings into a single profile that can be applied to an interface. The IPsec profile references the IPsec transform set and the IKEv2 profile.

IKEv2 Smart Defaults

This topic describes the IKEv2 smart default values.



The figure shows the IKEv2 smart defaults and their attributes. The IKEv2 smart defaults feature minimizes the FlexVPN configuration by covering most of the use cases. IKEv2 smart defaults can be customized for specific use cases, though this practice is not recommended.



You can use the default values without any configuration. Although not recommended, you can also change the default parameters by overwriting the values of the default constructs, as shown in the first example.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

You can restore the parameters to their default values using the **default crypto ikev2 proposal** and **default crypto ipsec transform-set** commands. This is shown in the second example.

You can also disable the use of smart default by negating the commands, as shown in the last example.

The following rules apply to the IKEv2 Smart Defaults feature:

A default configuration is displayed in the corresponding **show** command with **default** as a keyword and with no argument. For example, the **show crypto ikev2 proposal default** command displays the default IKEv2 proposal and the **show crypto ikev2 proposal** command displays the default IKEv2 proposal, along with any user-configured proposals.

A default configuration is displayed in the **show running-config all** command; it is not displayed in the **show running-config** command.

A default configuration can be disabled using the no form of the command; for example, **no crypto ikev2 proposal default**. A disabled default configuration is not used in negotiation but the configuration is displayed in the **show running-config** command. A disabled default configuration loses any user modification and restores system-configured values.

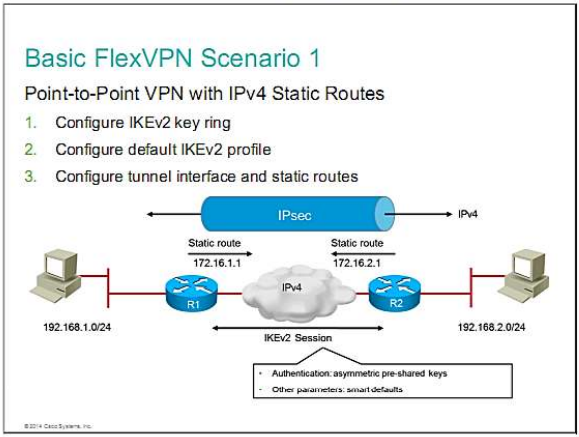
A default configuration can be re-enabled using the default form of the command, which restores system-configured values; for example, **default crypto ikev2 proposal**.

The default mode for the default transform set is **transport**; the default mode for all other transform sets is **tunnel**.

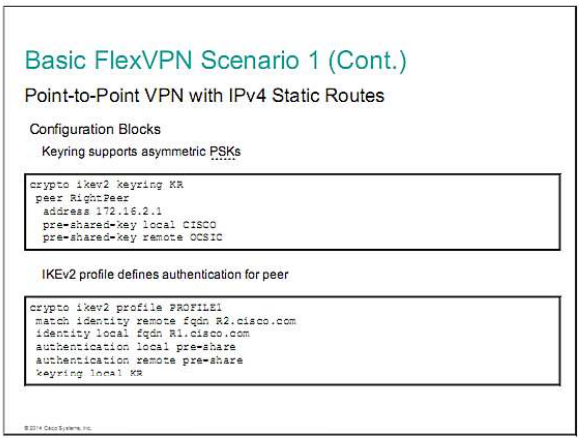
https://t.me/learningnets

Basic FlexVPN Scenario 1

This topic describes a basic FlexVPN scenario, a point-to-point VPN scenario with IPv4 static routes.



This figure illustrates the configuration scenario consisting of two sites connected over an IPsec tunnel. IPv4 is both the transport and the passenger protocol. Static routes should be set up to provide IPv4 connectivity. In this scenario, you configure an IKEv2 keyring, an IKEv2 profile, a tunnel interface, and the required static routes.

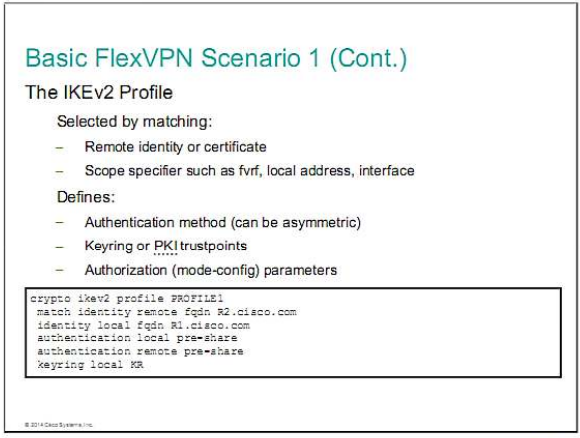


https://t.me/learningnets
© 2014 Cisco Systems, Inc.

A basic implementation of FlexVPN may use only one or two of the FlexVPN configuration blocks. An IKEv2 profile is always needed to define the authentication for the peer, but this Point-to-Point VPN with IPv4 Static Routes example uses two, the IKEv2 profile and the IKEv2 keyring.

IKEv2 keyring: Used to define preshared keys, which can be asymmetric. An IKEv2 keyring is simply a repository of symmetric and asymmetric preshared keys. It is independent of the IKEv1 key ring. The IKEv2 key ring is associated with an IKEv2 profile and hence supports a set of peers that match the IKEv2 profile. The IKEv2 key ring gets its VRF context from the associated IKEv2 profile.

IKEv2 profile: Defines the authentication for a peer. The IKEv2 profile is a repository of nonnegotiable parameters of the IKE SA.

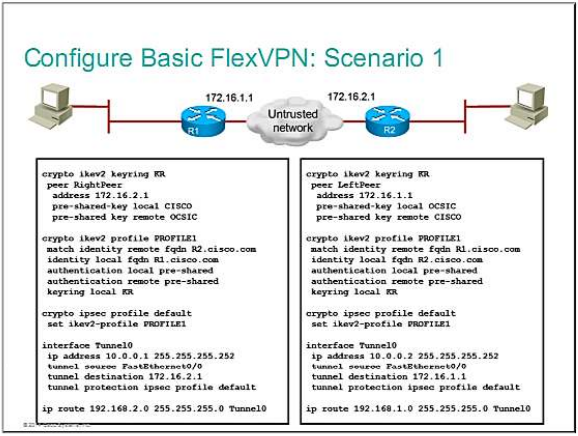


The IKEv2 profile is the repository for the nonnegotiable parameters of the IKE SA, such as local or remote identities and authentication methods and services that are available to authenticated peers that match the profile. An IKEv2 profile must be attached to either a crypto map or an IPsec profile on the initiator. An IKEv2 profile is not mandatory on the responder.

An IKEv2 profile is selected by matching a remote identity or certificate, or any one scope specifier such as FVRF, local address, or interface. It specifies the authentication methods, which can be asymmetric, the keyring or PKI trustpoints, and authorization parameters, also known as mode-config.

Configure Basic FlexVPN: Scenario 1

This topic explains how to configure a basic FlexVPN scenario, a point-to-point VPN scenario with IPv4 static routes.



The figure shows the configuration detail for both routers in the Point-to-Point VPN with IPv4 Static Routes scenario. The first configuration block creates an IKEv2 key ring named KR that defines asymmetric pre-shared keys. The key configuration on each side is provided for the peer identified by its IPv4 address.

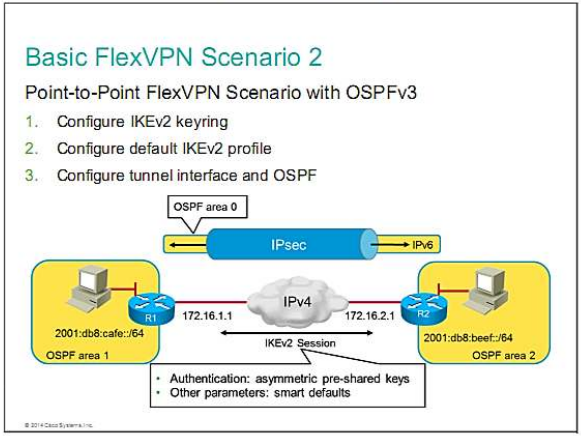
The second configuration block in this example creates an IKEv2 profile named PROFILE1, which specifies that each peer is identified by its FQDN. The profile also specifies the use of asymmetric pre-shared keys and the local identity name, and it references the IKEv2 keyring.

Although the IPsec profile smart default is used in this scenario, it is modified to reference the IKEv2 profile PROFILE 1. The **tunnel protection** command is then used to apply the ipsec profile and a default IPsec transform set to the tunnel interface, which has an IPv4 address. The tunnel interface is will be protected by the IPsec profile named "default." A static route is configured on each router to send traffic (traffic that is destined for the network behind the peer router) to the tunnel interface.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Basic FlexVPN Scenario 2

This topic describes another basic FlexVPN scenario, a point-to-point VPN scenario with OSPFv3.

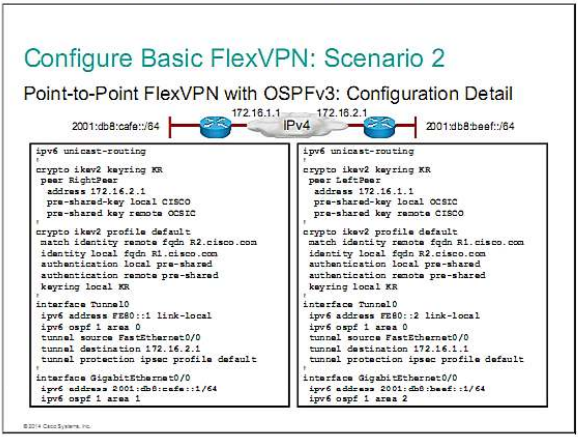


This figure illustrates the configuration scenario consisting of two sites connected over an IPsec tunnel. The transport protocol is IPv4. The passenger protocol is IPv6. OSPFv3 should be used to ensure IPv6 connectivity between the sites.

In this configuration, you can configure an IKEv2 keyring, the default IKEv2 profile, the tunnel interface, and the OSPFv3 routing.

Configure Basic FlexVPN: Scenario 2

This topic describes the configuration of a FlexVPN point-to-point IPsec tunnel with OSPFv3.

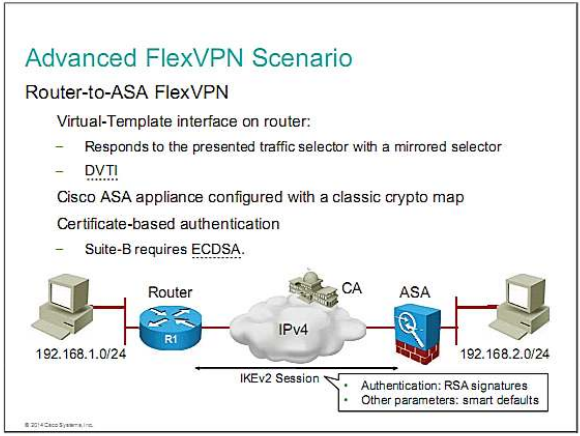


This configuration differs from the previous scenario only in the configuration of the passenger protocol and the routing process.The tunnel is configured with IPv6 addresses and connects via IPv4 only network. The OSPFv3 routing protocol is used to exchange IPv6 routes between tunnel endpoints.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Advanced FlexVPN Scenario

This topic describes an advanced FlexVPN scenario, a VPN between a Cisco IOS router and a Cisco Adaptive Security Appliance.



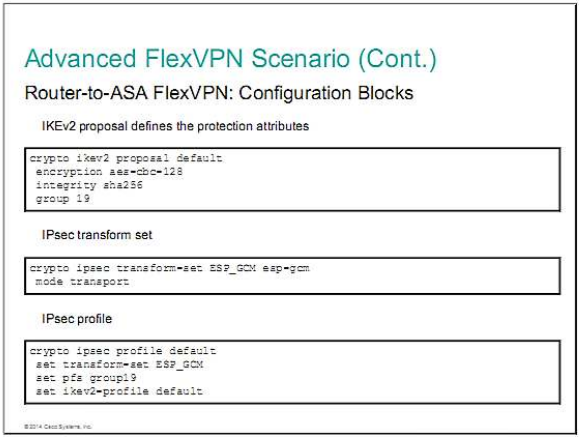
This figure illustrates another configuration scenario consisting of two sites connected over an IPsec tunnel. The IPsec gateways are a Cisco IOS router and the Cisco ASA.

The Cisco ASA creates IPsec SAs based on traffic defined in an ACL. Therefore, the router must respond to the IKEv2 initiation with a mirror of the proposed traffic selectors. The use of DVTI on the FlexVPN router allows this device to respond to the presented traffic selector with a mirror of the traffic selector that was presented.

This example encrypts traffic between both internal networks. When the Cisco ASA presents the traffic selectors of the Cisco ASA internal network to the Cisco IOS internal network, 192.168.2.0/24 to 192.168.1.0/24, the DVTI interface responds with a mirror of the traffic selectors, which is 192.168.1.0/24 to 192.168.2.0/24.

Currently, Cisco IOS Software and Cisco ASA do not support a local CA server with ECDSA certificates, which is required for Suite-B. Therefore, a third-party CA server must be implemented. For example, you can use OpenSSL to act as a CA.

This scenario requires a Cisco Integrated Services Router G2 router with the security license running Cisco IOS Software Release Version 15.2-3.T2 or later. Any release of M or T for releases later than Cisco IOS Software Release Version 15.1.2T can be used because this is included with the introduction of GCM. The ASA must run Software Release 9.0 or later to support NGE. Only multi-core platforms support AES GCM.



The advanced FlexVPN scenario (router-to-ASA FlexVPN) illustrates the use of FlexVPN configuration blocks that were not used in the basic FlexVPN scenarios. The advanced scenario uses the following additional FlexVPN configuration blocks:

Internet Key Exchange Version 2 proposal: A collection of transforms used in the negotiation of IKE SAs as part of the IKE_SA_INIT exchange. The transform types used in the negotiation are as follows:

- Encryption algorithm
- Integrity algorithm
- PRF algorithm
- DH group: 1 (768-bit DH group), 2 (1024-bit DH group), 5 (1536-bit DH group), 14 (2048-bit DH group), 15 (3072-bit DH group), 16 (4096-bit DH group), 19 (256-bit Elliptic Curve DH group), 20 (384-bit ECDH group), 24 (2048-bit DH/DSA group)

Note Diffie-Hellman algorithm allows two nodes to securely agree on a shared secret key over an insecure network. The group chosen must be strong enough to protect the IPsec keys during negotiation. A generally accepted guideline recommends the use of a 2048-bit group after 2013 (until 2030). Either group 14 or group 24 can be selected to meet this guideline. Even if a longer-lived security method is needed, the use of Elliptic Curve Cryptography is recommended, but group 15 and group 16 can also be considered.

You must configure at least one encryption algorithm, one integrity algorithm, and one DH group for the proposal to be considered complete. The PRF algorithm is the same as the integrity algorithm, and hence, it is not configured separately.

IPsec transform set: Specifies an acceptable combination of security protocols and algorithms for the IPsec SA

IPsec profile: Protects the interface to which it is applied by pointing to the IKEv2 profile and the IPsec transform

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

You can use an **IKEv2 policy** to bind an IKEv2 proposal to a specific entity. IKEv2 policies contains proposals (which are used to negotiate the encryption, integrity, PRF algorithms, and DH group in the IKE_SA_INIT exchange) and may also contain match statements (which are used as selection criteria to select a policy during negotiation).

The following rules apply to the match statements:

An IKEv2 policy without any match statements will match all peers in the global FVRF.

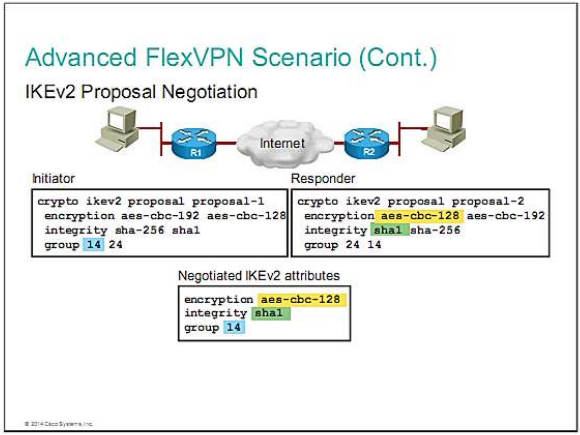
An IKEv2 policy can have only one match FVRF statement.

An IKEv2 policy can have one or more match address local statements.

When a policy is selected, multiple match statements of the same type are logically ORed and match statements of different types are logically ANDed.

There is no precedence between match statements of different types.

Configuration of overlapping policies is considered a misconfiguration. In the case of multiple, possible policy matches, the first policy is selected.



The configuration of the advanced FlexVPN scenario involves the configuration of an IKEv2 proposal. An IKEv2 proposal is a set of transforms used in the negotiation of IKEv2 SA as part of the IKE_SA_INIT exchange. An IKEv2 proposal is regarded as complete only when it has at least an encryption algorithm, an integrity algorithm, and a DH group configured. If no proposal is configured and attached to an IKEv2 policy, the default proposal in the default IKEv2 policy is used in negotiation.

You may configure an IKEv2 proposal with multiple transforms for each transform type, such as the following:

```
crypto ikev2 proposal proposal-2
 encryption aes-cbc-128 aes-cbc-192
 integrity sha1
 group 14
```

In this case, the IKEv2 proposal, proposal-2, translates to the following prioritized list of transform combinations:

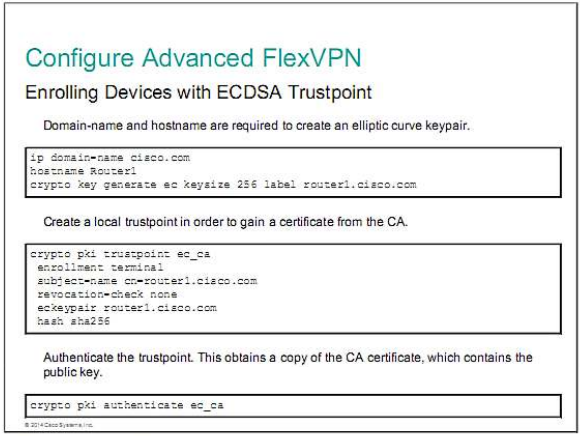
1. aes-cbc-128, sha1, 14
2. aes-cbc-192, sha1, 14

The figure shows proposals for the initiator and responder. The initiator and responder have conflicting preferences. For most attributes, the responder is preferred over the initiator. The DH group preference of the initiator is chosen because the initiator sends the key exchange material generated with the prime and generator values from this preferred group.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Advanced FlexVPN

This topic shows how to configure an advanced FlexVPN scenario, a VPN between a Cisco IOS router and a Cisco Adaptive Security Appliance.



This figure illustrates the enrollment procedure of the Cisco IOS router with the PKI. In this example, the EC is used to generate the key pair. To generate RSA key pair use the **crypto key generate rsa** command.

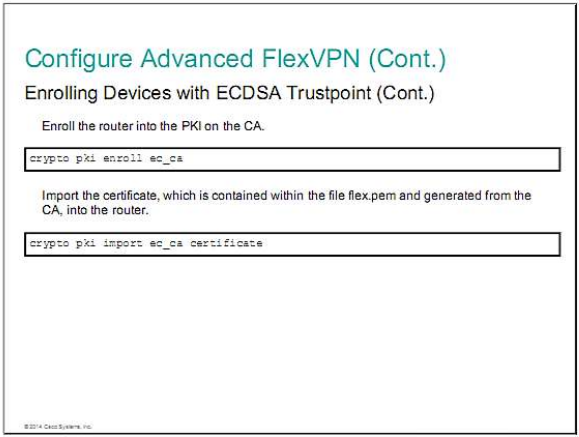
To declare the trustpoint that your router should use, use the **crypto pki trustpoint** command in global configuration mode. Additionally, you can include fully-qualified domain name (**fqdn** subcommand), serial number (**serial-number** sub-command), or IP address (**ip-address** sub-command) into trustpoint configuration. To specify EC key pair for trustpoint identity, use the **eckeypair** sub-command.

To authenticate the trustpoint, use the **crypto pki authenticate** command. This obtains a copy of the CA certificate, which contains the public key.

If you want to enable a Cisco IOS certificate server, use the **crypto pki server** command. This command takes you to certificate server configuration mode where you can define the behavior of the certificate server as shown in the following example:

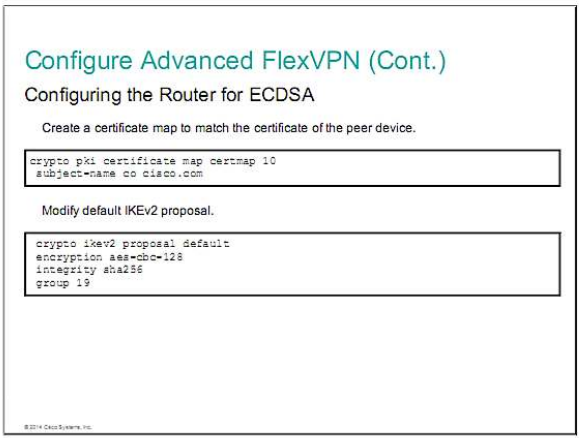
```
crypto pki server Partner-CA
 issuer-name CN=Partner-CA, O=partner.public
 grant auto
 hash sha256
```

The **issuer-name** sub-command specifies the DN as the CA issuer name for the certificate server, the **grant** sub-command specifies the granting method, and the **hash** command specifies the cryptographic hash function the certificate server uses to sign certificates.

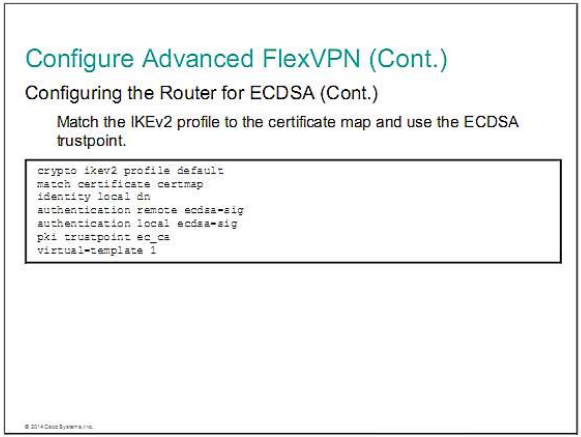


This figure illustrates the remaining enrollment procedure of the Cisco IOS router with the PKI. To complete the enrollment procedure, enroll the router into the PKI on the CA. Then import the certificate, which is contained within the file flex.pem and generated from the CA, into the router. After you enter this command, enter **quit**.

Follow the same procedure on the Cisco ASA. There are only minor differences in the command syntax, which are not shown here.

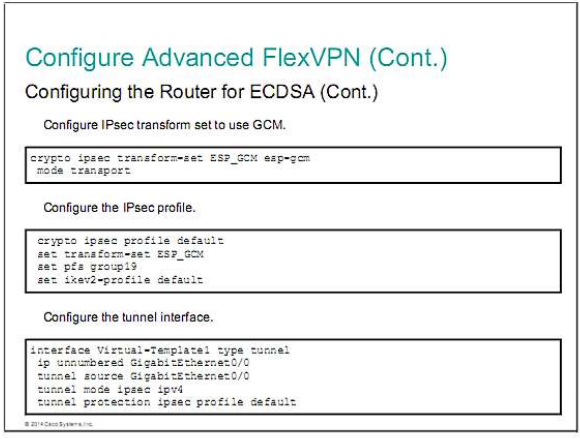


https://t.me/learningnets
© 2014 Cisco Systems, Inc.

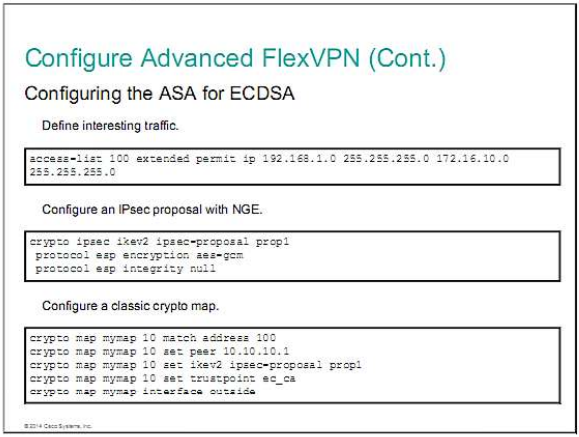


The figure illustrates the FlexVPN configuration on the Cisco IOS router.

First, you create a certificate map to match the certificate of the peer device using the **crypto pki certificate map** command. Then, you define the IKEv2 proposal for Suite-B configuration using the **crypto ikev2 proposal default** command. For maximum security, you use the **aes-cbc-256** and the **sha512** parameters. You match the IKEv2 profile to the certificate map and use ECDSA with the defined trustpoint in the **crypto ikev2 profile default** configuration.

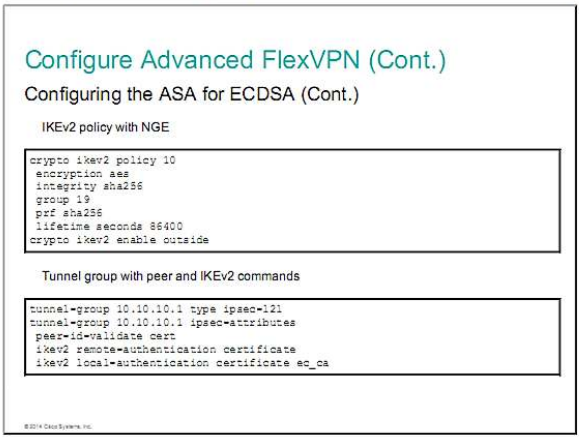


This figure shows the remaining tasks for configuring the router for ECDSA. The next task is configuring the IPsec transform set to use GCM, using the **crypto ipsec transform-set ESP_GCM esp-gcm** command. Finally, you define the IPsec profile and the tunnel interface.



This figure illustrates the ECDSA configuration on the Cisco ASA.

First, you create an access list to define the traffic to be encrypted. Then, you configure an IPsec proposal for NGE. You configure the crypto map using the classic approach, based on the **crypto map** commands.



This figure illustrates the remaining ECDSA configuration on the Cisco ASA.

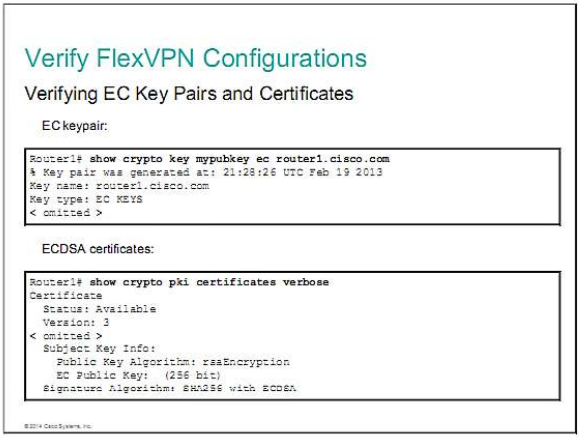
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

To complete the configuration, you use the **crypto ikev2 policy** command to configure the IKEv2 policy for NGE. The **crypto ikev2 enable outside** command enables IKEv2 on the outside interface. Finally, you create a tunnel group with the required **peer** commands.

https://t.me/learningnets

Verify FlexVPN Configurations

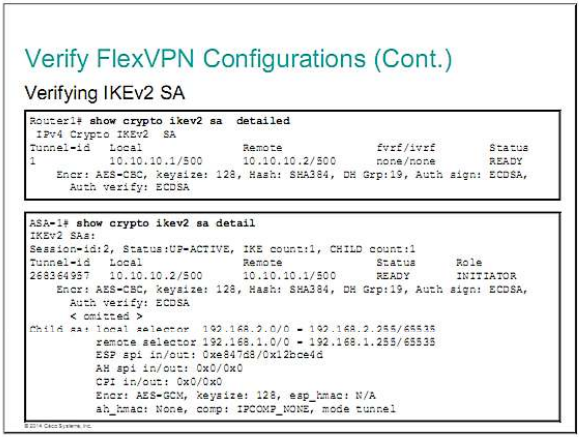
This topic explains how to verifyFlexVPN configurations.



The command outputs show how to verify the generated key pairs and the installed certificates and shows that the ECDSA is used in this scenario.

To verify that the ECDSA keys have been successfully generated use the **show crypto key** command.

To verify that the certificate has successfully been imported and that ECDSA is used use the **show crypto pki certificates verbose** command.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The **show crypto ikev2 sa detailed** command output shows how to verify the IKEv2 SAs on both IPsec gateway devices.

Verify FlexVPN Configurations (Cont.)

Verifying IPsec SA

```
Router1# show crypto ipsec sa
interface: Virtual-Access1
    Crypto map tag: Virtual-Access1-head-0, local addr 10.10.10.1
   protected vrf: (none)
   local  ident (addr/mask/prot/port): (192.168.1.0/255.255.255.0/0/0)
   remote ident (addr/mask/prot/port): (192.168.2.0/255.255.255.0/0/0)
   current_peer 10.10.10.2 port 500
     PERMIT, flags={origin_is_acl,}
< omitted >
     inbound esp sas:
      spi: 0x12BCE4D(19648077)
        transform: esp-gcm ,
        in use settings ={Tunnel, }
```

© 2014 Cisco Systems, Inc.

Verify FlexVPN Configurations (Cont.)

Verifying IPsec SA (Cont.)

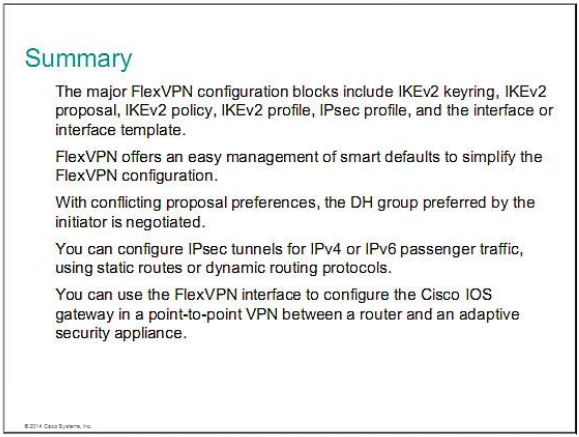
```
ASA-1# show crypto ipsec sa detail
interface: outside
    Crypto map tag: mymap, seq num: 10, local addr: 10.10.10.2
      access-list 100 extended permit ip 192.168.2.0 255.255.255.0
192.168.1.0 255.255.255.0
 < omitted >
    inbound esp sas:
      spi: 0x0E847D8 (15222744)
         transform: esp-aes-gcm esp-null-hmac no compression
         in use settings ={L2L, Tunnel, IKEv2, }
```

© 2014 Cisco Systems, Inc.

The **show crypto ipsec sa** and **show crypto ipsec sa detail** command outputs show how to verify the IPsec SAs on both IPsec gateway devices.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets

Lesson 3

Deploying Hub-and-Spoke IPsec VPNs Using Cisco IOS FlexVPN

This lesson describes the hub-and-spoke connectivity scenario that can be implemented using the FlexVPN framework.

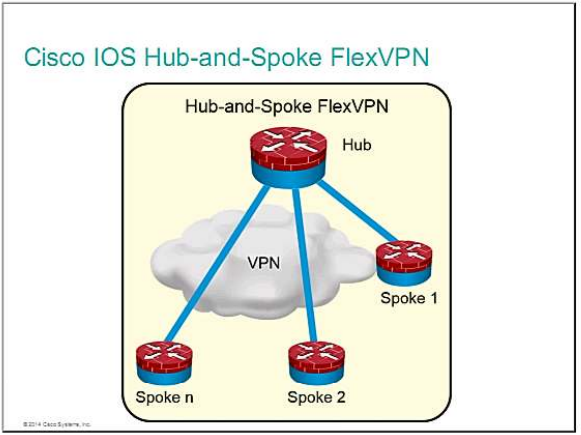
Upon completing this lesson, you will be able to meet these objectives:

- Describe the FlexVPN hub-and-spoke scenario
- Describe IKEv2 configuration payload
- Describe a locally managed hub-and-spoke scenario
- Describe how to configure FlexVPN on the spoke
- Describe how to configure FlexVPN on the hub
- Describe the configuration exchange in the hub-and-spoke scenario
- Describe the verification and troubleshooting of hub and spoke FlexVPN

https://t.me/learningnets

Cisco IOS Hub-and-Spoke FlexVPN

This topic introduces the FlexVPN hub-and-spoke scenario.

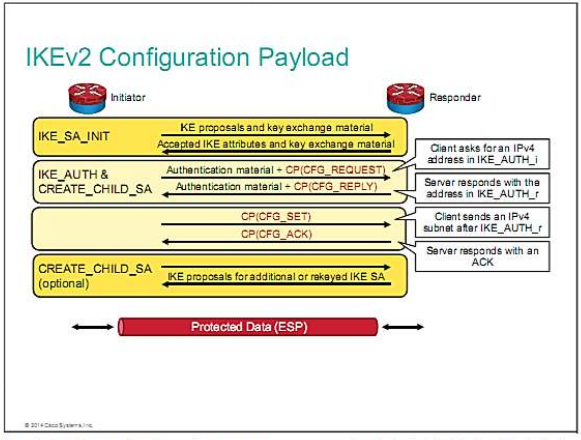


In a hub-and-spoke FlexVPN deployment model, each spoke router must be configured with a point-to-point interface to the hub. All traffic between spoke networks must flow through the hub router.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

IKEv2 Configuration Payload

This topic describes the IKEv2 configuration payload feature.

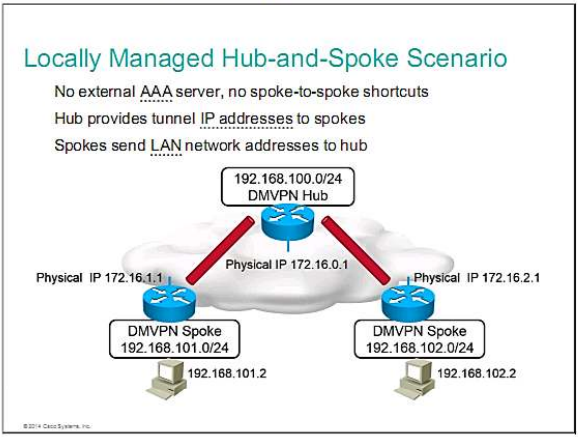


In hub-and-spoke VPN topologies and remote access scenarios, it is desirable for the hub device to provide configuration data, such as an internal IP address, to the spokes before Child-SAs are created. IKEv2 uses additional configuration payload in message 3 and 4 (IKE_SA_AUTH) of IKE-SA creation in order to provide configuration data necessary for the creation of Child-SAs. The use of the configuration payload feature is optional.

The predefined configuration attributes include IP address and mask, IP subnet, DHCP, DNS, and NBNS information. These attribute sets are available for both IPv4 and IPv6.

Locally Managed Hub-and-Spoke Scenario

This topic describes a locally managed hub-and-spoke scenario.

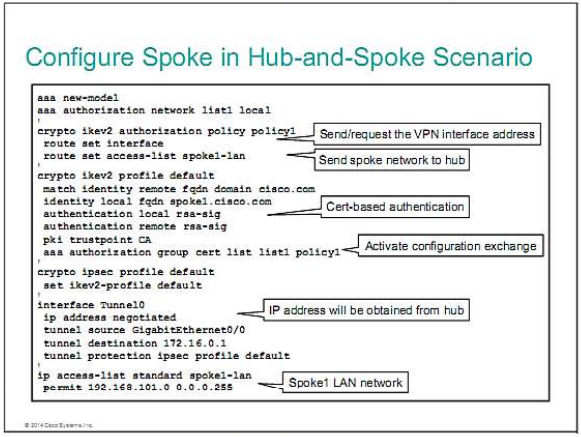


The scenario in the figure shows spokes connecting to the hub router. The hub router will provide IP addresses to the spoke VPN tunnel interfaces. The spokes will advertise their LAN networks to the hub. No external AAA server is used in the hub site.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure a Spoke in a Hub-and-Spoke Scenario

This topic describes the configuration of the spoke router in the hub-and-spoke scenario.



The cryptographic IKEv2 authorization policy, policy1, allows the spoke router to obtain the IP address of the VPN tunnel interface from the hub router. It also sends the IP address of the local LAN network, configured later using the ACL named spoke1-lan.

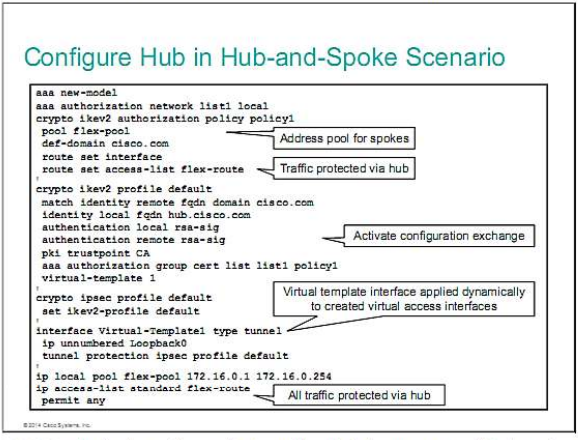
The cryptographic IKEv2 profile default is configured for certificate-based authentication and enables the configuration exchange using the IKEv2 authorization policy, policy1. The keyword **cert** in the **aaa authorization group cert list** command indicates that certificates must be used for the authentication. "list1" references the network authorization method that uses the local router database.

The cryptographic IPsec profile default is configured to reference the crypto IKEv2 profile default only for presentation purposes. That is a FlexVPN default setting.

The tunnel interface is configured to obtain the IP address from the hub router. It is protected using the default IPsec profile.

Configure a Hub in a Hub-and-Spoke Scenario

This topic describes the configuration of the hub router in the hub-and-spoke scenario.



The crypto IKEv2 authorization policy, policy1, enables the hub router to provide the spokes with IP addresses taken from the local pool, flex-pool. It also enables protection for the LAN subnets. The respective ACL, named flex-route, is configured at the bottom.

The crypto IKEv2 profile default is configured for certificate-based authentication and enables the configuration exchange using the policy, policy1, configured in the first section. The keyword **cert** in the **aaa authorization group cert list** command indicates that certificates must be used for the authentication. The profile points to the virtual template interface that will be applied dynamically to the virtual access interfaces. The virtual access interfaces will be created when spokes connect to the hub.

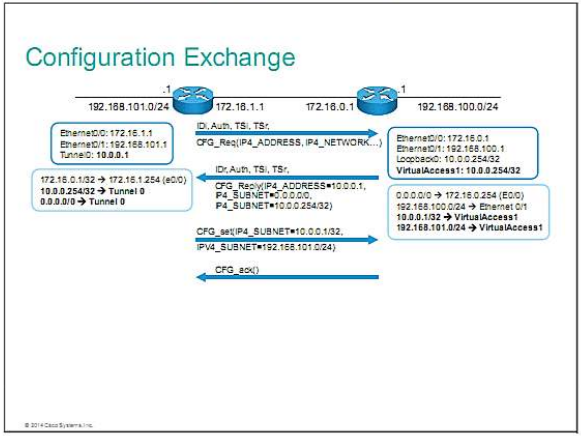
The crypto ipsec profile default is configured to reference the crypto IKEv2 profile default only for presentation purposes.

The virtual template interface borrows its address from the interface Loopback 0. It is protected using the default IPsec profile.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configuration Exchange

This topic describes the configuration exchange in the hub-and-spoke scenario.



The figure shows the exchange of the tunnel IP address and LAN subnet information in the IKEv2 exchange.

The following table describes the conditions when the initiator and the responder send different configuration payload types:

Configuration Payload Types

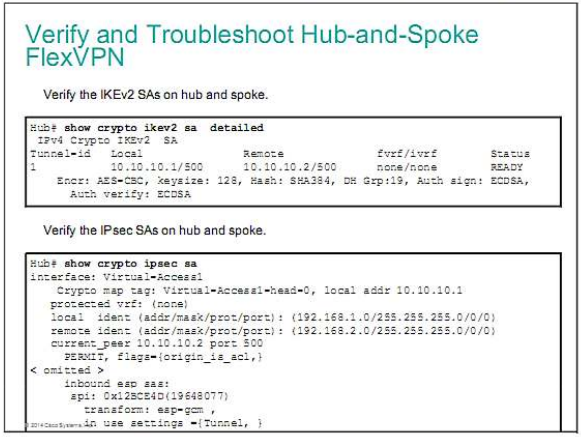
Configuration Payload Type	Sent By...	When...
CFG_REQUEST	Initiator	The initiator is the FlexVPN client or if the **config-exchange request** command is enabled in the IKEv2 profile.
CFG_REPLY	Responder	The responder receives the CFG_REQUEST.

Configuration Payload Type	Sent By...	When...
CFG_SET	Initiator and responder	Initiator—The **config-exchange set send** command is enabled in the IKEv2 profile. Responder—The CFG_REQUEST is not received, the configuration data is available, and the **config-exchange set send** command is enabled in the IKEv2 profile.
CFG_ACK	Initiator and responder	Initiator—The **config-exchange set accept** command is enabled in the IKEv2 profile. Responder—The **config-exchange set accept** command is enabled in the IKEv2 profile.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify and Troubleshoot Hub-and-Spoke FlexVPN

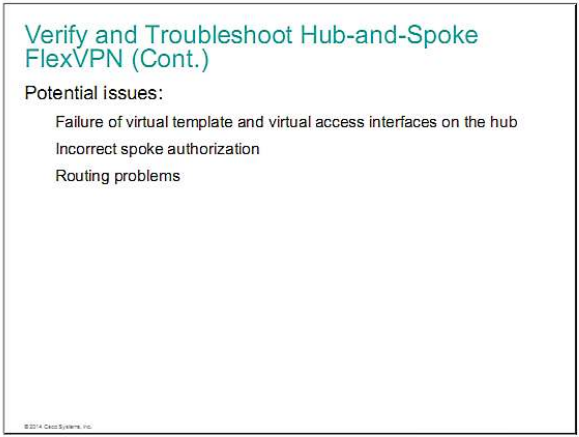
This topic describes the verification and troubleshooting of hub and spoke FlexVPN.



The command outputs above show how to verify the IKEv2 SAs and IPsec SAs on hub and spokes.

To display the information about the current IKEv2 security associations, use the **show crypto ikev2 sa** command in privileged EXEC mode. Optionally, use the **detailed** keyword to display detailed information about the current security associations.

To display the settings used by current IPsec security associations, use the **show crypto ipsec sa** command in privileged EXEC mode.

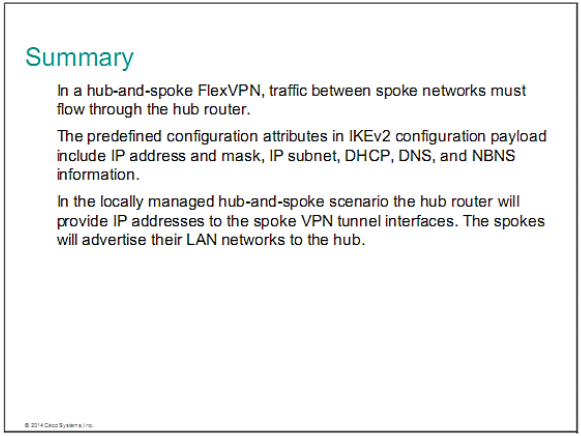


In a hub-and-spoke FlexVPN topology you can experience all the problems of a point-to-point scenario. You should also consider other potential points of failure, such as virtual template and virtual access interfaces, the authorization of the spoke tunnel interfaces, and network reachability.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets

Lesson 4

Deploying Spoke-to-Spoke IPsec VPNs Using Cisco IOS FlexVPN

This lesson describes the spoke-to-spoke connectivity scenarios that can be implemented using the FlexVPN framework.

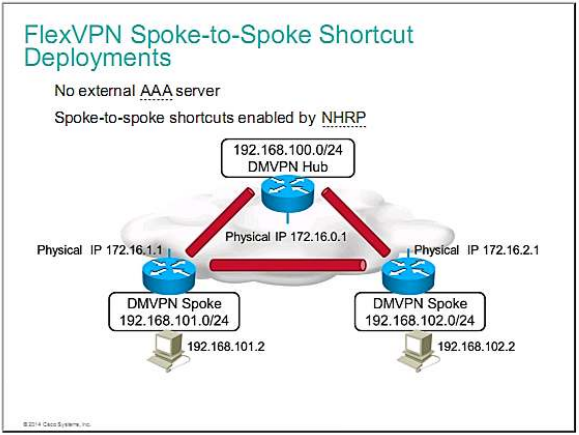
Upon completing this lesson, you will be able to:

- Describe a locally managed spoke-to-spoke shortcut scenario
- Describe the NHRP operation in the FlexVPN environment
- Describe the configuration of the spoke router in the spoke-to-spoke shortcut scenario
- Describe the configuration of the hub router in the spoke-to-spoke shortcut scenario
- Describe the verification and troubleshooting of spoke-to-spoke shortcut switching

https://t.me/learningnets

FlexVPN Spoke-to-Spoke Shortcut Deployments

This topic describes FlexVPN spoke-to-spoke shortcut deployments.



The Cisco Dynamic Multipoint Virtual Private Network feature enables users to better scale large and small IPsec VPNs. The Cisco DMVPN feature combines IPsec tunneling and NHRP to provide simple provisioning of many VPN peers. DMVPN also easily supports dynamically addressed spoke routers by its design, if you use an appropriate peer authentication method such as PKI-enabled peer authentication.

There are two connectivity models: hub and spoke (which was discussed in a previous lesson) and shortcut switching, which is also known as spoke-to-spoke deployment. A spoke-to-spoke DMVPN deployment model requires each branch to be configured with a tunnel interface in which dynamic spoke-to-spoke tunnels are used for the spoke-to-spoke traffic. In this model, DMVPN provides a scalable configuration model for all involved devices, and also allows spoke devices to dynamically peer and establish optimal routing paths. Note that DMVPN will not immediately produce a partially meshed or fully meshed topology. DMVPN initially establishes a permanent hub-and-spoke topology, from which a partial mesh or full mesh is dynamically generated based on traffic patterns. This scenario can also be augmented with RADIUS management.

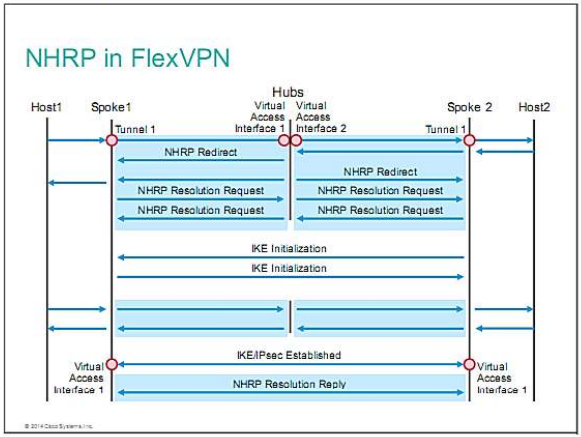
Independent of the connectivity patterns, you can deploy locally-managed or RADIUS-managed DMVPNs. In the RADIUS-managed approach, the topology is enhanced by RADIUS-provided management options that the hub obtains from the AAA database and provisions to the spokes. This approach is used in large DMVPN deployments. In the locally managed environment, by contrast, the hub router must locally know all the parameters provisioned to the spoke routers.

In the scenario shown in the figure, the spokes will establish direct tunnels triggered by the traffic. The hub router will provide IP addresses to the spoke VPN tunnel interfaces. The spokes will advertise their LAN networks to the hub. No external AAA server is used in the hub site.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

NHRP in FlexVPN

This topic describes the NHRP operation in the FlexVPN environment.



The figure shows the NHRP resolution request and reply in FlexVPN.

Due to bidirectional traffic, similar events occur in both directions at Spoke1, Spoke2, and hub. For clarity, events from Host1 to Host2 are discussed. Assume that there is a network N1 (192.168.1.0/24) behind Spoke1 and another network N2 (192.168.2.0/24) behind Spoke2. The network between the two spokes is matched through an ACL. This is because ACLs are applied on the IKEv2 policies on both spokes.

The network along with its prefix information from both the spokes is conveyed to the hub via IKEv2 information payload exchanges. This causes a route addition in the routing table by IKEv2 at the hub as follows:

192.168.1.0/24—Connected via virtual access interface1

192.168.2.0/24—Connected via virtual access interface2

The hub will push a summarized route via IKEv2 to both spokes, and the spokes will install the route in their routing table as follows:

192.168.0.0/16—next hop <tunnel address of the hub> - interface Tunnel 1

Note The routing protocol can also add the route to the routing table.

Assuming that traffic moves from N1 to N2, the traffic flow is as follows:

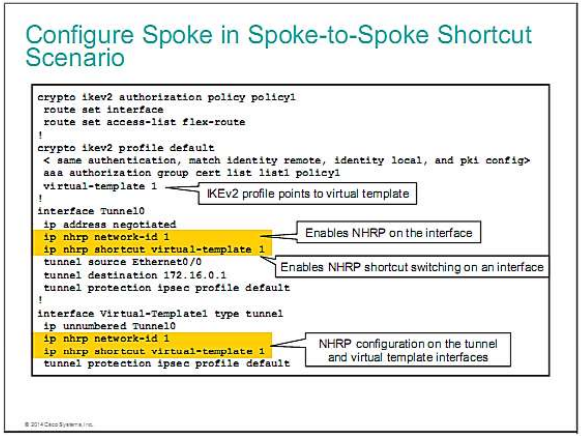
1. Host1 sends traffic destined to Host2. The traffic reaches the LAN interface of spoke1, looks up the route, hits the summarized route, and routes the packet to interface tunnel 1.
2. When the traffic reaches the hub's virtual access interface1, the traffic looks up the route table for a route entry for N2, either directly connected over virtual access interface 2 or via a point-to-point tunnel interface.

3. The traffic from Host1 to Host2 traverses the hub through virtual access interface1 and virtual access interface2. The hub determines that ingress and the egress interfaces (virtual access interface1 and virtual access interface2) belong to same NHRP network (network D configured on both the interfaces). The hub sends out an NHRP redirect message to spoke1 on virtual access interface1.
4. On receiving the redirect, Spoke1 initiates a resolution request for Host2 over the point-to-point tunnel interface (the same interface over which it received the redirect). The resolution request traverses the routed path (Spoke1-hub-spoke2). On receiving the resolution request, Spoke2 determines that it is the exit point and needs to respond to the resolution request.
5. Spoke2 receives the resolution request on the tunnel interface and retrieves the virtual template number from the tunnel interface. The virtual template number is used to create the virtual access interface to start a crypto channel and establishes IKEv2 and IPsec SAs. Once the crypto SAs between the two spokes are up, Spoke2 installs the necessary NHRP cache entries for Spoke1 and its network under the newly created virtual access interface and sends out the resolution reply over the virtual access interface.
6. After receiving the resolution request over the virtual access interface, Spoke1 installs the necessary cache entries for Spoke2 and its network. Spoke1 also deletes the temporary cache entry pointing to the hub to resolve the network under tunnel interface1.
7. NHRP adds shortcut routes as NHO or H route.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure a Spoke in a Spoke-to-Spoke Shortcut Scenario

This topic describes the configuration of the spoke router in the spoke-to-spoke shortcut scenario.

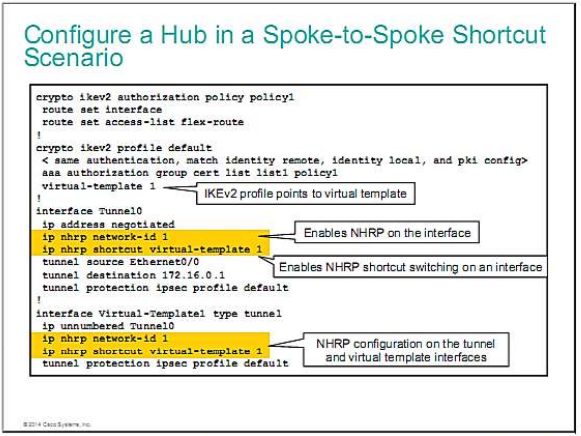


The spoke router configuration resembles the settings from the previous scenario. In contrast to the previous configuration, you see two interfaces: a tunnel interface and a virtual template. The Tunnel 0 interface is used for the connection to the hub. The virtual template settings are applied dynamically to the virtual access interfaces that are created for spoke-to-spoke connections.

Both interfaces are configured for NHRP. You can enable NHRP with the **ip nhrp network-id** command. NHRP shortcut switching is enabled with the **ip nhrp shortcut virtual-template** command. The **ip nhrp redirect** command is used to enable NHRP redirects on an interface. This is useful when networks move from one spoke to another.

Configure a Hub in a Spoke-to-Spoke Shortcut Scenario

This topic describes the configuration of the hub router in the spoke-to-spoke shortcut scenario.



The configuration of the hub router differs from the previous scenario by the additional NHRP settings on the virtual template interface.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify and Troubleshoot Spoke-to-Spoke Shortcut Switching

This topic describes the verification and troubleshooting of spoke-to-spoke shortcut switching.

Verify and Troubleshoot Spoke-to-Spoke Shortcut Switching

Make sure that spoke-to-spoke traffic triggers shortcuts.

Use trace to verify the shortcut.

On spokes, examine IP routes pointing via NHRP interfaces.

On spokes, verify IKEv2 and IPsec SAs with other spokes.

```
Partner-ISR# show ip route nhrp
...
10.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
H 10.11.11.0/24 [250/1] via 192.0.10.1, 00:07:11, Virtual-Access1
...
Partner-ISR# show crypto ikev2 sa
 IPv4 Crypto IKEv2  SA

Tunnel-id Local                 Remote                fvrf/ivrf       Status
6         198.51.100.1/500      192.0.2.3/500         none/none       READY
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:20, Auth sign: RSA,
Auth verify: RSA

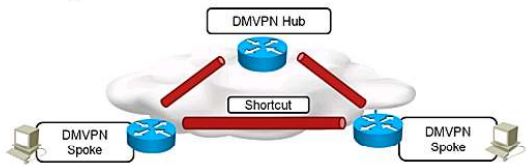
Tunnel-id Local                 Remote                fvrf/ivrf       Status
2         198.51.100.1/500      203.0.113.1/500       none/none       READY
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:20, Auth sign: RSA,
Auth verify: RSA
```

When the traffic between the spokes triggers the spoke-to-spoke shortcuts, the packets will take the direct path between the spokes. You can verify it using various tools, such as trace. You should see IP routes to other spoke subnets pointing via NHRP interfaces. You can also verify that direct spoke-to-spoke IKEv2 and IPsec SAs have been negotiated.

Verify and Troubleshoot Spoke-to-Spoke Shortcut Switching (cont.)

Potential issues:

1. Inconsistent or missing NHRP configuration
2. Virtual Templates unavailable for spawning Virtual Access interfaces
3. Unsupported tunnel encapsulation



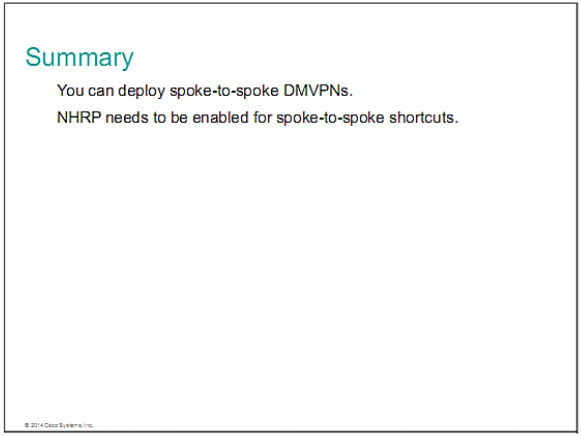
© 2014 Cisco Systems, Inc.

In a hub-and-spoke FlexVPN topology with the spoke-to-spoke shortcut switching you can experience all the problems of a hub-and-spoke scenario, and you should consider several additional potential problems: inconsistent or missing NHRP configuration, missing Virtual Template interfaces that cannot be used for spawning Virtual Access interfaces, as well as unsupported tunnel encapsulation modes.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

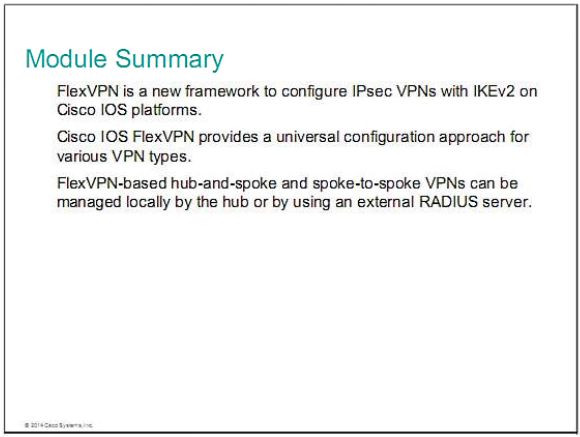
This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Summary

This topic summarizes the key points that were discussed in this module.



References

For additional information, refer to this resource:

Cisco Systems, Inc. *FlexVPN and Internet Key Exchange Version 2 Configuration Guide*. http://www.cisco.com/en/US/docs/ios-xml/ios/sec_conn_ike2vpn/configuration/15-mt/sec-flex-vpn-15-mt-book.html

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Self-Check

Questions

Use the questions here to review what you learned in this module.

1. What are the benefits of using IKEv2 instead of IKEv1 when deploying remote-access IPsec VPNs? (Choose two.)

   A. IKEv2 supports EAP authentication methods as part of the protocol.
   B. IKEv2 inherently supports NAT traversal.
   C. IKEv2 messages use random message IDs.
   D. The IKEv2 SA plus the IPsec SA can be established in six messages instead of nine messages.
   E. All IKEv2 messages are encryption-protected.

2. Which of the following configuration elements reference the virtual template interface? (Choose two.)

   A. IKE profile
   B. IPsec profile
   C. tunnel interface on hub
   D. tunnel interface on spoke
   E. transform set on hub
   F. transform set on spoke

3. Where do you configure virtual template interfaces in a spoke-to-spoke shortcut scenario?

   A. only on spokes
   B. only on a hub
   C. on spokes and a hub
   D. on a hub, if you use config mode
   E. on spokes and a hub, if you use config mode

https://t.me/learningnets

4. Which of the following options describe how an IKEv2 profile is selected? (Choose three.)

   A. by matching a remote certificate
   B. by matching a local certificate
   C. by matching remote identity
   D. by matching local identity
   E. by matching local address
   F. by matching cryptographic protection
   G. by matching a proposal

5. Which of teh following statements describe the interaction of IKEv2 constructs correctly? (Choose two.)

   A. The IKEv2 profile references an IPsec profile.
   B. The IPsec profile references an IKEv2 profile.
   C. The IKEv2 proposal references an IKEv2 policy.
   D. The IKEv2 policy references an IKEv2 proposal.
   E. The IKEv2 proposal references an IKEv2 profile.
   F. The IKEv2 profile references an IKEv2 proposal.

6. Which option correctly describes IKEv2 smart defaults?

   A. noneditable default values
   B. always-active default values
   C. default IKE values
   D. default IKE and IPsec values

7. Which of the following options describe the DoS mitigation capabilities of IKEv2 correctly? (Choose two.)

   A. are based on next-generation cryptographic algorithms
   B. can overwhelm the IKEv2 node
   C. use anti-clogging cookies
   D. are optional
   E. are based on EAP

8. Which of the following methods are available in IKEv2 but are not supported in IKEv1? (Choose two.)

   A. PSK
   B. pubkey-sig
   C. pubkey-encr
   D. EAP
   E. hybrid

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Answer Key

1. A, B
2. A, D
3. C
4. A, C, E
5. B, D
6. D
7. C, D
8. D, E

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module 4

Deploying Clientless SSL VPNs

Clientless Secure Sockets Layer (SSL) VPN solutions provide browser-based access to resources behind the Cisco ASA adaptive security appliance. In clientless SSL VPNs, users can access resources without any special client software. Users can also access web-based applications, Common Internet File System (CIFS) file shares, and FTP servers when using clientless SSL VPNs. Using application plug-ins, port forwarding, and smart tunnels, users can access almost any application that uses static TCP ports.

This module describes deployment of the basic clientless SSL VPN, as well as advanced application access and advanced authentication. This module also describes how to customize the clientless SSL VPN portal to the needs of the organization.

Upon completing this module, you will be able to meet these objectives:

- Describe clientless SSL VPN and provide a general description of the SSL/TLS protocol
- Configure and verify baseline clientless SSL VPN remote access features of the Cisco ASA security appliance
- Deploy and manage advanced application-access features of a clientless Cisco SSL VPN
- Deploy and manage advanced authentication and authorization features of a clientless Cisco SSL VPN

https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 1

Clientless SSL VPN Overview

Cisco clientless SSL VPN enables users to access web and some nonweb application over a web browser. This lesson describes clientless SSL VPN and provides a general description of the SSL/TLS protocol.

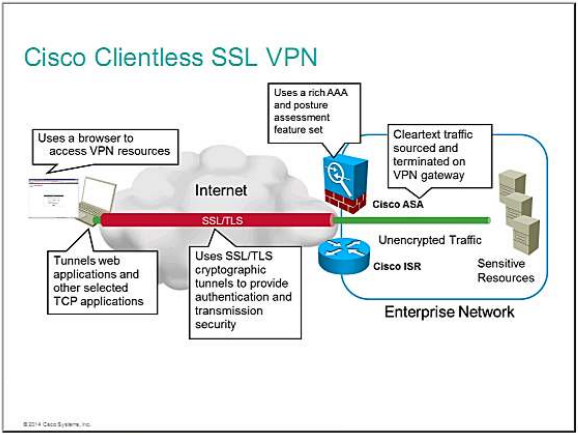
Upon completing this lesson, you will be able to meet these objectives:

- Describe Cisco Clientless SSL VPN
- Describe Cisco Clientless SSL VPN Use Cases
- Describe the application access methods in Cisco Clientless SSL VPNs
- Describe SSL and TLS protocols
- Describe SSL session setup and key management
- Describe SSL server authentication
- Describe SSL client authentication
- Describe SSL transmission protection
- Describe Cisco ASA remote access configuration concept
- Describe Cisco ASA connection profiles
- Describe Cisco ASA group policies

https://t.me/learningnets

Cisco Clientless SSL VPN

This topic describes Cisco Clientless SSL VPN.



Remote access VPNs generally connect individual remote users to a set of protected resources in the enterprise internal network or a perimeter, extranet, or demilitarized zone buffer network. Remote access VPNs use strong client authentication and cryptographic data transmission protection across an untrusted transport network, usually the Internet. Most remote access VPNs require various security controls in addition to authentication and transmission protection. Often, VPN gateways will strictly limit access to sensitive resources; may request that remote users download additional security controls from the enterprise network; and may assess the security posture of remote clients before allowing access. Additionally, remote access VPNs must support any kind of client connectivity such as wireless, wired, DSL, or cable modem access, as long as it has routed IP access to the VPN gateway. Remote access VPN users often operate in environments that do not natively support or allow VPN connectivity, such as behind the firewalls of other enterprises.

The Cisco Adaptive Security Appliance and Cisco IOS Integrated Services Router support the following types of remote access VPNs:

- Remote access clientless SSL VPNs
- Remote access AnyConnect full-tunnel client-based SSL VPNs
- Remote access AnyConnect full-tunnel client-based IPsec VPNs

These VPNs differ in the technologies they use to encapsulate user traffic and provide transmission protection.

There are two major modes of remote access user encapsulation:

- **Full-tunnel client-based VPNs:** Full-tunnel client-based VPNs require VPN client software on the remote computer. The VPN client software supports IPsec and SSL VPNs.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Clientless VPNs: Clientless SSL VPNs give organizations the additional flexibility of providing remote access to protected network resources even when the remote device is not managed and has no VPN client software. In clientless SSL VPN deployments, remote users use a web browser to establish a Secure Sockets Layer or Transport Layer Security (SSL/TLS) VPN connection with a VPN gateway, which can be a Cisco ASA or a Cisco IOS router. The VPN gateway acts as a proxy device to protected network resources. After bidirectional authentication between the VPN gateway and the user, the gateway presents the user with a web portal interface through which remote users can navigate the network and applies a set of authorization and accounting rules for the user session. Although clientless SSL VPNs are easier to deploy than full-tunnel remote access VPNs, they provide only limited application access and may pose additional security risks to the enterprise.

The full-tunnel site-to-site IPsec VPN connections behave similarly to dedicated point-to-point links over the untrusted transport network. The dedicated VPN devices that are used for site-to-site VPNs are able to complete routed IP access to protected resources.

There are two primary methods for providing transmission protection over VPNs: the first is either SSL, TLS, or DTLS, and the second is IPsec. Both methods provide comparable transmission security features.

Clientless SSL VPNs do not offer the complete network access that is provided by full-tunnel VPNs. To enable remote users to access enterprise applications, the VPN gateway acts as a proxy. It transforms web applications and some nonweb applications so that they can be SSL-protected. The Cisco ASA and Cisco IOS routers offers several techniques to facilitate resource and application access:

URL and CIFS file access is provided. When the client browser establishes the SSL session and the user authenticates, the gateway can present a page with resource bookmarks. These bookmarks allow the user to access preconfigured web pages or file shares. The user can also enter the address of the resource and be redirected to it.

Smart tunnels are connections between an application and a remote site, using a clientless SSL VPN session as the pathway. A successor technology to port forwarding, smart tunnels do not require administrator privileges on the remote computer. Clientless VPN users can run Winsock2 TCP applications just like full-tunnel SSL VPN client access.

Application plug-ins (Cisco ASA only), which are also referred to as thin-client connections, exist for well-known applications, such as RDP, Telnet, Citrix, VNC, and SSH. They allow the user to launch a Java application plug-in inside the browser and connect to an internal server that is running that application.

The clientless remote access SSL VPN architecture has the following benefits:

It can traverse most firewalls and NAT devices, because the SSL VPN encapsulation uses the HTTPS port (TCP port 443) and is indistinguishable from an HTTPS session to the transport network operator.

Does not require special-purpose desktop VPN client software to be preinstalled on the system.

Does not require administrative rights on clients

Offers flexibility to establish VPN connections from any Internet-enabled system, be it a company-managed or non-company-managed system.

Uses the native SSL encryption of a web browser to provide data confidentiality.

However, the clientless remote access SSL VPN architecture has some limitations:

It does not support all IP applications, although the majority of web-based client-server enterprise applications are supported.

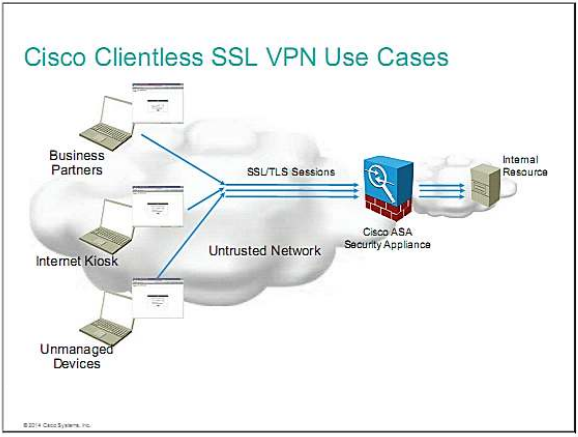
It may require user training, because the manner in which users access some resources may change.

It does not support low-latency forwarding and the use of real-time applications due to its proxying nature.

https://t.me/learningnets

Cisco Clientless SSL VPN Use Cases

This topic describes the cases of Cisco Clientless SSL VPN use.



Whether you use a Cisco Adaptive Security Appliance clientless SSL VPN solution depends on the type of clients that need to access the network and the level of connectivity that they require. When using the clientless solution, the Cisco ASA acts as a proxy for external devices, providing access to internal components. This process protects internal devices because the external devices do not have complete connectivity to them.

Clientless SSL VPN connections on the Cisco ASA differ from remote access IPsec connections in two ways. The first difference is the way in which clientless SSL VPN connections interact with SSL-enabled servers. The second difference involves the precautions that you must take to reduce security risks.

When a clientless SSL VPN user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate. The browser never receives the presented certificate, so it cannot examine and validate the certificate.

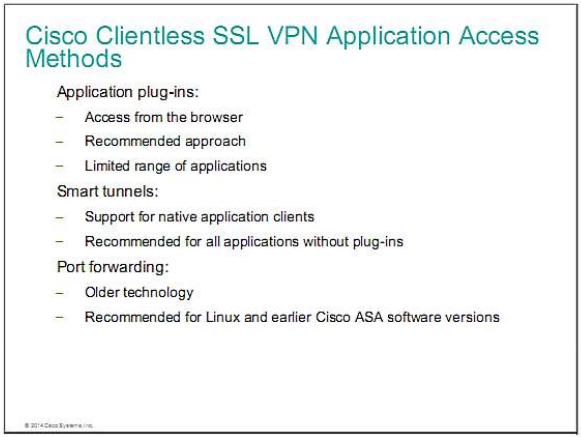
The current implementation of clientless SSL VPNs on the Cisco ASA does not permit communication with sites that present expired certificates. Even if a user has added a certificate to the trusted certificates list on the operating system, the security appliance does not perform trusted CA certificate validation. Because the Cisco ASA acts as a proxy, users cannot benefit from certificate validation of pages that are delivered from an SSL-enabled web server until the server has a valid certificate.

By default, the Cisco ASA permits all portal traffic to all web resources, such as HTTPS, CIFS, RDP, and plug-ins. The Cisco ASA clientless service rewrites each URL to a URL that is meaningful only to the security appliance. Users cannot confirm that they are on the site that they requested by looking at the rewritten URL that displays in the browser.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Clientless SSL VPN Application Access Methods

This topic describes the application access methods in Cisco clientless SSL VPNs.



You can choose several options when deploying advanced application access:

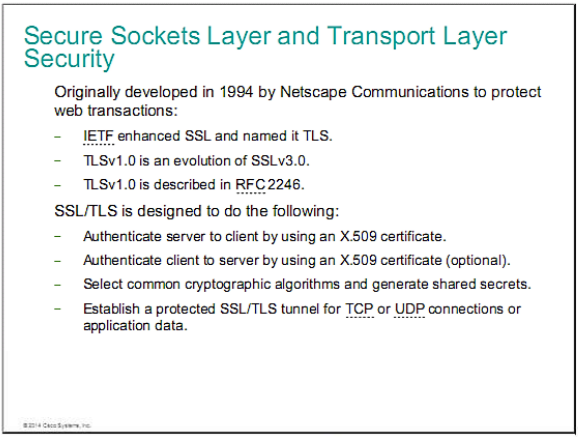
Application plug-ins: This option provides users with thin-application client access to enterprise resources. This approach is the recommended approach but supports a limited set of applications.

Smart tunnels: This option provides users with native-application client access to enterprise resources. This approach is the recommended approach for all applications that do not have a plug-in.

Port forwarding: This option provides users with native-application client access to enterprise resources. This method should be used only if smart tunnels cannot be deployed, such as on Linux workstations or when running older Cisco Adaptive Security Appliance software versions.

Secure Sockets Layer and Transport Layer Security

This topic describes SSL and TLS protocols.



TLS and its predecessor, SSL, are cryptographic protocols that provide transmission security for communications over networks such as the Internet. Applications such as web browsing, electronic mail, and VoIP commonly use these protocols.

The SSL protocol was originally developed by Netscape in 1994 to protect web transactions. Netscape continued to develop SSL until version 3.0, which was released as an Internet draft in 1996. In January 1999, the IETF adopted the SSL protocol and called it Transport Layer Security, or TLS. TLS is also known as SSL version 3.1.

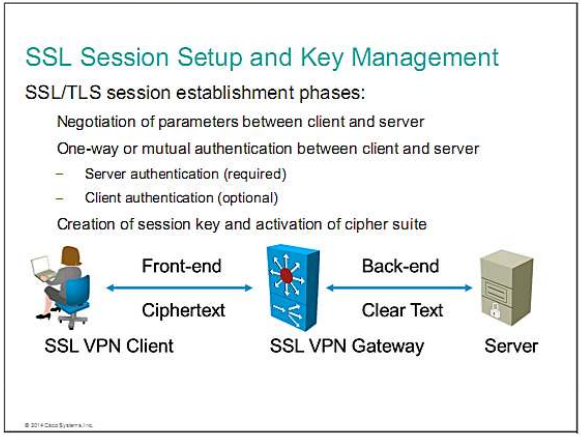
Currently, almost all web browsers have implemented SSLv3 or TLSv1. Many other applications, including Cisco Secure Socket Layer VPN Clients, also use SSL/TLS for transmission protection.

SSL/TLS provides endpoint authentication both for the client and the server; data encryption to ensure that it is only readable by the intended recipient; and data integrity and data authentication to ensure that the data has not been modified in transit. These services allow traffic to be protected as it traverses public network segments such as the Internet.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

SSL Session Setup and Key Management

This topic describes SSL session setup and key management.



The SSL and TLS protocols work in two distinct planes: the session establishment phase, in which the negotiation of parameters and peer authentication takes place, and a data transfer phase that provides a protected path between the client and the server. Both of these phases occur inside the SSL/TLS Record Protocol.

The following describes the SSL/TLS session establishment and key management:

Phase 1: Establish client-server security capabilities.

In Phase 1 (initiated by the client), hello messages are exchanged between the client and the server to negotiate parameters, including authentication and encryption algorithms. These messages include whether RSA or DH will be used as the key exchange method.

Phase 2: Server certificate and key exchange

In Phase 2 (initiated by the server), one- or two-way authentication between the client and the server is performed. The server begins this phase by sending its certificate to the client. Unlike most authenticated solutions that consider server authentication to be optional, HTTP and SSL consider client authentication to be optional. A premaster key is also sent by the client by using the public key of the server to start protecting the session if RSA is used.

Phase 3: Client certificate (if required) and key exchange

In Phase 3 (initiated by the client), if client authentication is required, then the client sends its certificate to the server, the session key is calculated, and the cipher suite is activated. HMAC is used to protect the integrity of the data transfer by using either keyed Algorithm SHA-1 or keyed MD5. Confidentiality is provided using DES-40, DES-CBC, 3DES-EDE, 3DES-CBC, RC4-40, or RC4-128.

Phase 4: Finish

In Phase 4, the data transfer starts, beginning with the exchange of session keys.

https://t.me/learningnets

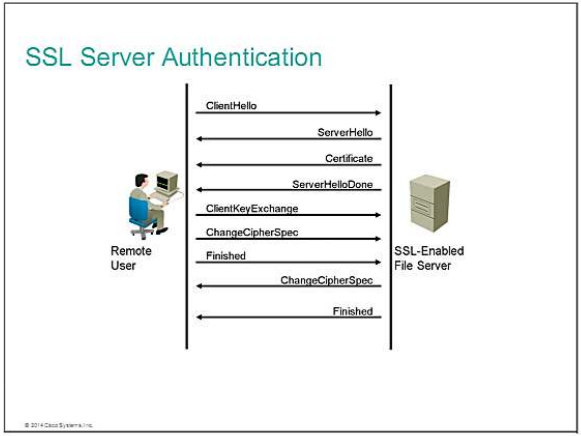
Each SSL session is uniquely identified by a session ID that is exchanged during the authentication process. The SSL session ID is used to differentiate between new and old sessions. Old sessions might exist because session IDs can be cached, but usually not for longer than 24 hours.

SSL also has the ability to resume a previously negotiated session if TCP communication between a client and a server is interrupted. During the initial exchange between the SSL client and the server, the server gives the client a session ID that helps the server to keep track of that SSL session. The client can ask that the server resume the SSL session without having to renegotiate parameters.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

SSL Server Authentication

This topic describes SSL server authentication.



SSL server-side authentication is used when a client needs to verify the identity of a server. This type of authentication is commonly used for servers that require secured transactions to protect user data or account information for online purchases.

The client starts the exchange by sending a client hello message to the server to let the server know that the client wants to establish secure communications. It also proposes a list of security parameters (cipher suites) that the client would like to use.

The server sends a server hello message that informs the client of the preferred cipher suite. In addition, it tells the client that the server is willing to proceed with the SSL negotiation. The server replies immediately with a certificate message (which carries the public key certificate of the server).

A trusted CA, typically issues this server-side certificate. The client must verify the certificate to ensure that the certificate has been issued by a trusted CA, has not expired, or has not been revoked.

The server then sends a server hello done message to the client. This message tells the client that the server has finished its part of the initial negotiations.

The client generates random numbers to use as a shared session key. In a client key exchange message, the client encrypts the session key with the public key of the server and sends the result to the server. The client depends on the ability of the server to decipher the client key exchange message to verify the identity of the server.

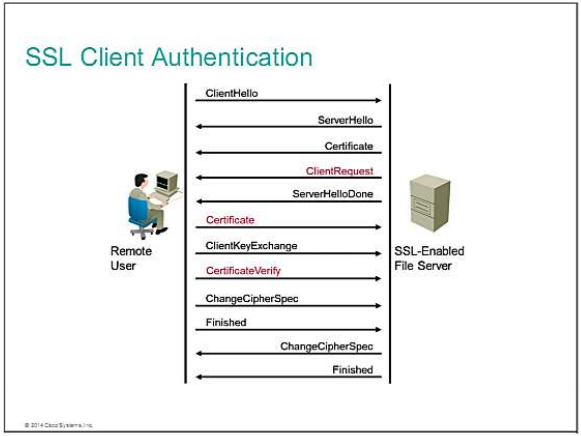
The client now sends a change cipher specification message that tells the server to activate the negotiated cipher suite. After this exchange, all messages from the client are encrypted by using the algorithm from the server hello and the session key in the client key exchange. The client sends an encrypted finished message, which ensures that both parties use the same cryptographic algorithms and parameters. The negotiation is successful when the message that follows the change cipher spec of the client is successfully decrypted as a finished message.

The server sends a change cipher spec message that tells the client that all subsequent messages from the server will use the negotiated security. The server then sends its own finished message, which enables the client to confirm that the negotiated security is in place.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

SSL Client Authentication

This topic describes SSL client authentication.



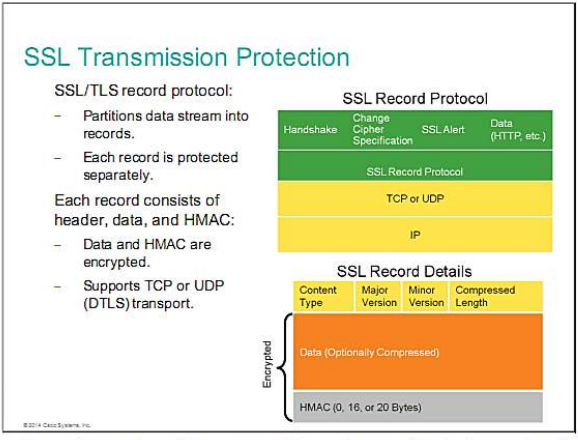
The server can authenticate the client, in addition to the server being authenticated. Authenticating the client ensures that the user is allowed to connect to the requested network resource. Client-side authentication is common when connecting end users to a remote access VPN that is using SSL.

When an SSL client must be authenticated in addition to the server, the protocol flow is almost the same. There are only a few additional required exchanges.

After the server sends its certificate to the client, the server will request that the client send its certificate in return. After the client sends the client key exchange, the client sends a certificate verify message. The client encrypts a known piece of plaintext by using its private key. The server uses the client certificate to decrypt the message. Successful decryption of the message indicates to the server that the client has the correct private key.

SSL Transmission Protection

This topic describes SSL transmission protection.



The SSL/TLS Record Protocol provides cryptographic envelopes to data that is transferred between the SSL/TLS client and the SSL/TLS server. Application data is broken into discrete chunks that are cryptographically protected, and a header and trailer are added. A series of SSL records is then forwarded to the TCP or UDP layer that provides a transmission for the session. The UDP encapsulation is defined in the DTLS standard that provides a connectionless transport service that can efficiently transport real-time data.

Four record types are defined to implement signaling and secure data exchange. One record contains up to 16,384 data bytes. Special record types are used for handshake, signaling cipher change, and alert messages. Each record consists of a header, a data portion, and an HMAC. Both the data and the HMAC are encrypted. Optionally, the data can also be compressed.

The SSL connections can be described as made of two main phases:

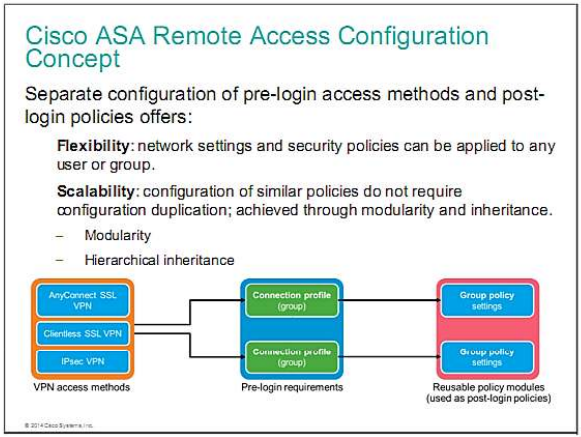
SSL session negotiation (or **handshake**)—This phase ensures authentication of the server and, optionally, of the client. It is based on public key cryptography, and it allows server and client to agree on a secret master key used for the data encryption. This phase typically uses RSA as the public key algorithm.

SSL application data transfer—In this phase, SSL encrypts the data traffic and ensures data confidentiality and integrity. Confidentiality is achieved by using symmetric algorithms such as DES, 3DES, RC4. Data integrity is achieved by using message digest algorithms such as MD5 and SHA-1.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ASA Remote Access Configuration Concept

This topic describes Cisco Adaptive Security Appliance remote access configuration concept.



VPNs involve various classes of remote users or sites that connect to a protected network over an untrusted transport network. The security policies for these users or sites may be very complex, and there may be many different user groups, each with different requirements. For this reason, VPNs require a flexible and scalable configuration method to make the provisioning of such a configuration manageable.

The configuration method should allow you to apply network settings and security rules to any specific user, user group, or site. In addition, the configuration method should enable you to configure similar policies for different groups or to create small exceptions to group security policies, with little effort and no duplication of workload. Mechanisms that provide appropriate modularity of configuration and inheritance in a configuration hierarchy should provide these abilities.

To provide you with flexibility and scalability when you configure the VPN policy, the Cisco ASA uses a separation of access methods and policies. You should be aware of two major configuration constructs:

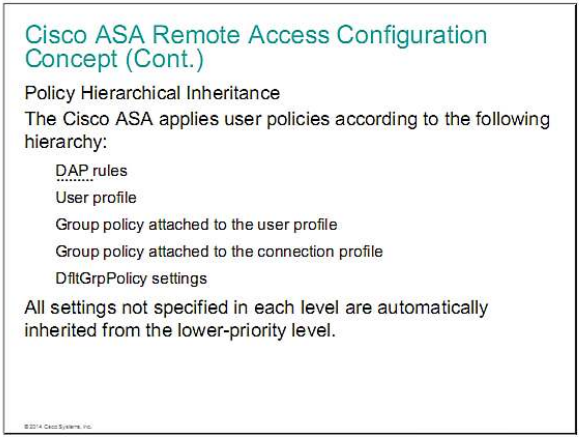
Connection profiles: Connection profiles define the prelogin requirements for a particular access method.

Group policies: Reusable group policies define the postlogin settings and security policies that must be applied to a remote access session after user authentication.

The Cisco ASA uses the same configuration philosophy for all three major access methods: full-tunnel SSL VPNs, clientless SSL VPNs, and full-tunnel IPsec VPNs. A remote user that connects to a remote access VPN is classified into a connection profile, which determines how the user should authenticate to the VPN gateway. After successful authentication, group policy settings dictate the special VPN features or policy restrictions that apply to that user session.

Note	Some documentation also refers to connection profiles as tunnel groups.

https://t.me/learningnets



A major feature that the Cisco ASA provides is policy inheritance. Policy inheritance connects connection profiles and group policies to enable configuration reuse. To make the configuration of VPN policy easy and consistent, you should use this feature as much as possible.

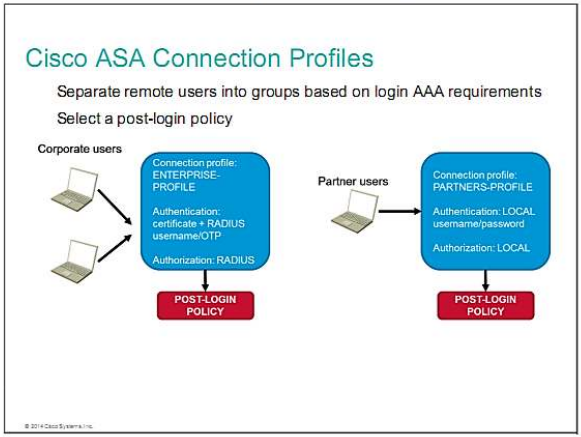
The Cisco ASA uses a hierarchical policy inheritance model, which follows this priority philosophy:

1. The most important postlogin settings and policies that apply to the remote access user VPN connection are the rules that the DAP creates. To determine access rights, DAP can take into account dynamic events, such as the posture of the client.
2. If DAP does not configure a particular parameter, then settings that are configured in the local user profile (or pushed from the AAA server as part of user authentication and authorization) determine the postlogin policy.
3. If the user profile does not specify a particular parameter, then a group policy that is attached to the user profile can specify that parameter.
4. If no group policy is attached to the user profile, then the Cisco ASA uses the settings from the group policy that is attached to the user connection profile.
5. If the group policy that is attached to the user connection profile does not specify a particular setting, then the security appliance uses the setting from the default group policy (DfltGrpPolicy).

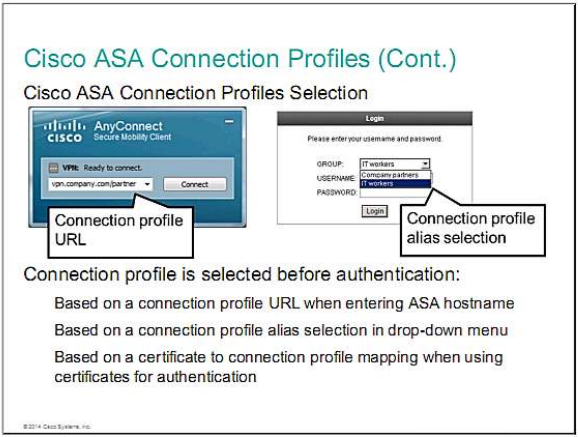
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ASA Connection Profiles

This topic describes Cisco Adaptive Security Appliance connection profiles.



Connection profiles are configuration objects that the Cisco ASA uses to separate users or sites into groups, based on your requirements for possible AAA login and postlogin policies. Therefore, you should create different connection profiles for users who require different AAA methods and for users who require different postlogin policies. You should also create different connection profiles for sites with different policy attributes. Often, these requirements are interdependent. You typically require stronger AAA methods for users who require access to more resources (or more-sensitive resources). You can allow weaker AAA controls for users who require only limited access. You should not use any AAA method if you are configuring site-to-site VPNs.

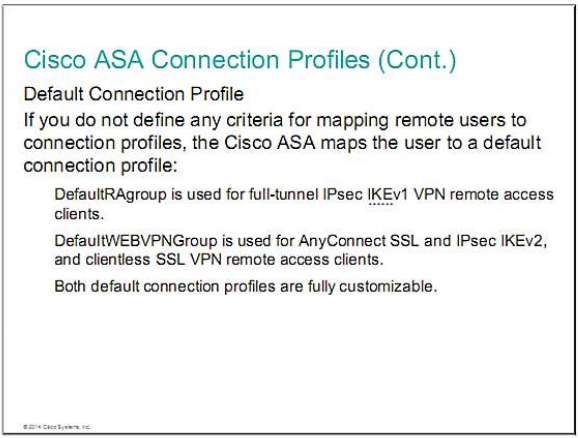


The Cisco ASA enables you to choose connection profiles according to one of these methods:

You can allow remote users to choose a connection profile. A user can choose a connection profile in two ways:

- By connecting to a VPN group URL that is associated with the connection profile
- By choosing a connection profile from a drop-down list and then successfully authenticating against the profile to obtain network access

Remote system identity enables the ASA to use a property of the remote system, such as a field in a digital (identity) certificate, to choose the connection profile for the remote client.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The Cisco ASA includes two default connection profiles, which the security appliance uses if there is no connection profile-mapping configuration:

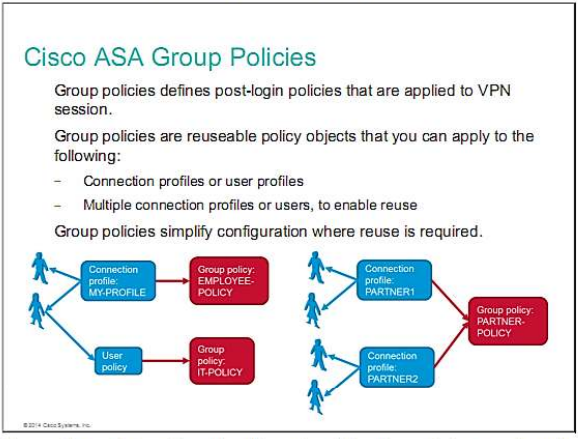
The default remote access group (DefaultRAGroup) connection profile:
The Cisco ASA assigns this profile to remote users of full-tunnel IPsec IKEv1 remote access VPNs who do not have rules that assign them to any other connection profile.

The default Cisco WebVPN group (DefaultWebVPNGroup) connection profile: The Cisco ASA assigns this profile to remote users of AnyConnect SSL and IPsec IKEv2, and clientless SSL VPNs who do not have rules that assign them to any other connection profile.

https://t.me/learningnets

Cisco ASA Group Policies

This topic describes Cisco Adaptive Security Appliance group policies.



Group policies are Cisco ASA configuration objects that define the postlogin network settings and security policies that you can apply to a remote access VPN user. These settings and policies include IP addresses (or pools), DNS servers, ACLs, and so on.

You can apply group policies to connection profiles or to individual locally configured user profiles if you are using local AAA. Group policies are reusable; therefore, you can apply the same group policy to multiple connection profiles or user profiles. Group policies can greatly simplify configuration when you need to deploy the same or similar policies to different users.

In the example in the figure, three connection profiles and two group policies are configured on the Cisco ASA. The first two users belong to the MY-PROFILE connection profile and will be assigned to the EMPLOYEE-POLICY group policy by default. However, the second user is assigned to a per-user policy after login and therefore is assigned a different group policy (IT-POLICY). The rest of the users use different connection profiles but are assigned to the same group policy (PARTNER-POLICY).

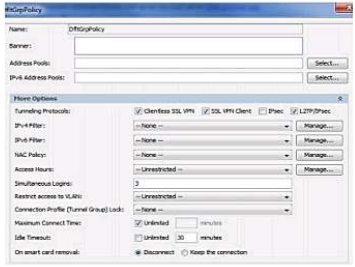
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco ASA Group Policies (Cont.)

Default Group Policy

The Cisco ASA by default includes a policy named DfltGrpPolicy:

- DfltGrpPolicy is applied to the default DefaultRAgroup and DefaultWEBVPNGroup connection profiles.
- The DfltGrpPolicy is fully customizable.



Newly created group policies inherit settings from the DfltGrpPolicy.

© 2014 Cisco Systems, Inc.

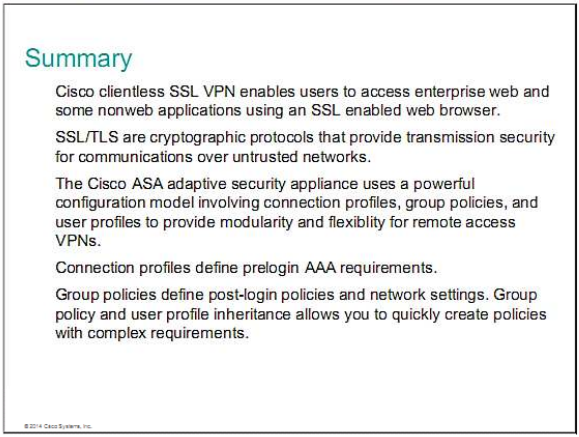
The Cisco ASA includes a default group policy named DfltGrpPolicy. By default, this policy is assigned to the DefaultRAGroup and DefaultWebVPNGroup connection profiles. You cannot delete the DfltGrpPolicy group policy, but you can customize it. You can also disassociate this policy from the two default connection profiles and attach a group policy with a different name to default connection profiles.

The DfltGrpPolicy group policy plays an important role in policy inheritance. You can use policy inheritance to configure scalable complex policies without duplicating your work. With policy inheritance, your custom policies can inherit certain common parameters from the DfltGrpPolicy group policy. Therefore, correctly configuring and tuning the DfltGrpPolicy group policy is important.

If you do not specify a value for a parameter in your custom policies, then that parameter can inherit its value from the DfltGrpPolicy group policy.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets

Lesson 2

Deploying Basic Cisco Clientless SSL VPN on Cisco ASA

A basic clientless Cisco Secure Sockets Layer VPN solution enables users to use a web browser to access sensitive resources over a remote access SSL VPN gateway that is implemented on the Cisco Adaptive Security Appliances. A basic Cisco ASA clientless SSL VPN uses basic user authentication with usernames and passwords and basic SSL VPN portal features, such as URL entry and bookmarks. This lesson enables you to configure and verify a basic clientless SSL VPN solution on the Cisco ASA.

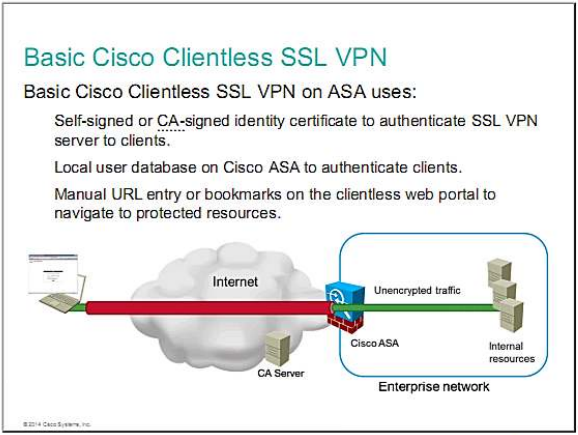
Upon completing this lesson, you will be able to meet these objectives:

- Describe Cisco clientless SSL VPN on Cisco ASA
- Describe different options to authenticate Cisco ASA when implementing clientless SSL VPN
- Describe client authentication options in basic clientless SSL VPN
- Describe how to access the protected resources using the SSL VPN clientless web portal
- Describe a configuration scenario for deploying basic Cisco clientless SSL VPN on Cisco ASA
- Describe required configuration tasks to configure basic clientless SSL VPN
- Configure basic clientless SSL VPN on Cisco ASA
- Verify Cisco clientless SSL VPN on Cisco ASA

https://t.me/learningnets

Basic Cisco Clientless SSL VPN

This topic provides an overview of Cisco clientless SSL VPN on Cisco Adaptive Security Appliance.



In a basic Cisco ASA clientless SSL VPN solution, remote users use a standard web browser to establish an SSL or TLS (SSL/TLS) session with the Cisco ASA. The basic solution uses bidirectional authentication. With bidirectional authentication, the client uses certificate-based authentication to authenticate the Cisco ASA, and the appliance authenticates the user using a username and password from its local user database. After authentication, the Cisco ASA applies a set of authorization rules to the user session and presents the user with a web portal. This portal allows the user to access internal resources using only a browser, using a URL input field or using pre-configured bookmarks.

In the basic clientless solution, the client supports a subset of services that are available to an SSL VPN client-based session, such as web application access and browser-based file-share access to internal resources. A clientless SSL VPN provides secure and easy access to a range of web resources and web-enabled applications from almost any computer on the Internet:

- Internal websites
- Web-enabled applications
- Microsoft Windows NT and Active Directory file shares
- Microsoft Outlook Web Access

The clientless SSL VPN configuration is divided into following blocks:

- **Pre-SSL VPN configuration -** enrolling digital certificates, setting up Cisco Adaptive Security Device Manager, setting up tunnel and group policies, setting up user authentication, clientless SSL VPN end user setup.
- **Basic SSL VPN configuration -** enable SSL VPN on an interface, configuring local user database on Cisco ASA to authenticate clients, configure bookmarks, creating and applying clientless SSL VPN policies and connection profile attributes for accessing resources, configure application access.

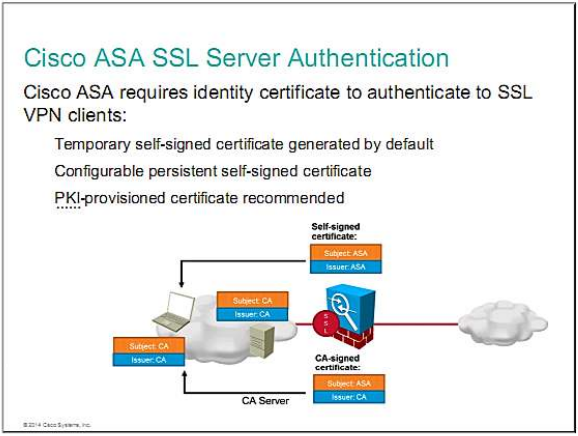
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Advanced SSL VPN configuration – advanced authentication, configuring browser access to plug-ins, configuring application access, configure SSL VPN portal customization, configure web-type ACLs.

https://t.me/learningnets

Cisco ASA SSL Server Authentication

This topic describes different options to authenticate Cisco Adaptive Security Appliance when implementing clientless SSL VPN.



The Cisco ASA requires a server identity certificate, which the appliance sends to remote SSL VPN clients in order for remote clients to authenticate the Cisco ASA.

By default, the security appliance will create a self-signed X.509 certificate on each reboot, resulting in many client warnings when attempting SSL VPN access, as the certificate cannot be verified by any means. You can address this issue using one of the two approaches:

- By creating a permanent self-signed certificate which is persistent across reboots, and which you can save on the client, if the remote users have the option of initially accessing the Cisco ASA over a trusted network and saving the identity certificate of the Cisco ASA to their local storage. This is usually not the case with most clients and is generally not recommended in production environments.
- By enrolling your Cisco ASA with an external CA. Your clients can obtain the root certificate of the CA from this external CA and then authenticate the identity of the certificate that is installed on the Cisco ASA. This method implies that the Cisco ASA must also obtain a root certificate and identity certificate from a CA server.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

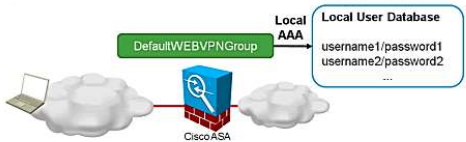
SSL VPN Clients Authentication

This topic describes client authentication options in basic clientless SSL VPN.

SSL VPN Clients Authentication

The simplest client authentication uses local passwords:

- Local user database
- Locally configured static passwords



© 2014 Cisco Systems, Inc.

SSL VPN Clients Authentication (Cont.)

Clientless password-based users:

- May be permitted to select connection profile from the selection menu or group URL
- DefaultWEBVPNGroup used by default which uses local AAA authentication

Client authentication type is configured in a connection profile.

© 2014 Cisco Systems, Inc.

When a clientless SSL VPN users connect to the Cisco Adaptive Security Appliance, the users may be permitted to select their connection profile by either choosing the desired profile from a drop-down list or connecting to the group URL. If no specific connection profile has been chosen, Cisco ASA will assign them to the DefaultWEBVPNGroup connection profile. This profile is by default configured to use user authentication using the local user database on the Cisco ASA.

When the user selects a connection profile (or uses the DefaultWEBVPNGroup connection profile), the user is required to authenticate using a method defined in the connection profile. The simplest user authentication method is using usernames and static passwords, which are stored in the local user database on the Cisco ASA. When a client enters a username and password into the login prompt, both are sent to the Cisco ASA. On the Cisco ASA, the provided username and password are compared to the locally stored username and password. If the provided and stored username and password match, the ASA establishes an SSL VPN session and applies a set of policies to the session.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

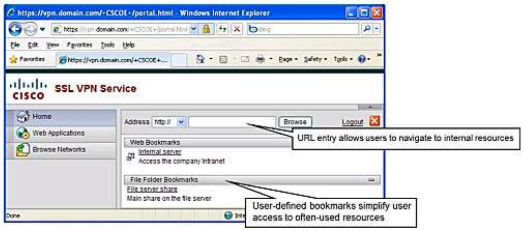
Clientless SSL VPN URL Entry and Bookmarks

This topic describes how to access the protected resources using the SSL VPN clientless web portal.

Clientless SSL VPN URL Entry and Bookmarks

In clientless SSL VPN portal the following applications are supported by default:

- Web browsing using HTTP/HTTPS
- File share using CIFS and FTP



Clientless SSL VPN URL Entry and Bookmarks (Cont.)

These applications can be accessed by:

- Selecting a protocol from the drop-down menu and entering URL into input field.
- Clicking the pre-configured bookmark

Bookmarks are configured as a part of a group policy.

When a clientless SSL VPN user successfully authenticates to the Cisco Adaptive Security Appliance, the user is presented with a web portal where he can access protected resources behind the Cisco ASA.

By default, the user can access the following types of applications:

- Internal web pages and web-based applications using HTTP and HTTPS.

File shares using CIFS and FTP.

By default, the Cisco ASA security appliance SSL VPN web portal does not restrict users and allows authenticated users access to all the internal resources. The SSL VPN portal also provides several ease-of-use interface features. You can use basic portal tuning to control basic portal appearance by enabling or disabling some of its user-interface functions.

This figure shows two main basic features of the Cisco ASA security appliance SSL VPN portal:

URL entry: This feature allows users to specify the URL of the resource that they want to access by selecting a protocol from the drop-down menu and entering a resource URL.

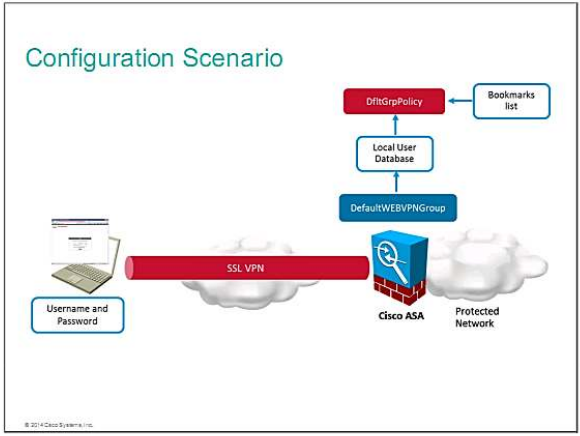
User-defined bookmarks: This feature enables an administrator to specify a list of often-used resources to which users can quickly navigate.

Bookmarks are configured as a part of a group policy. When a user logs into the SSL VPN portal, all the bookmarks that are enabled in the group policy of the user are displayed on the portal home page. In addition to the bookmark list, the URL entry field is displayed at the top of the interface. By tuning portal functions, you can enable or disable these two features.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configuration Scenario

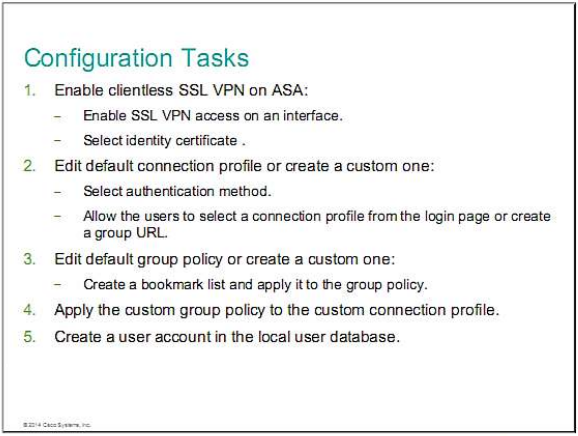
This topic provides a configuration scenario for deploying basic Cisco clientless SSL VPN on Cisco Adaptive Security Appliance.



This figure presents a configuration scenario, where you configure clientless SSL VPN on a Cisco ASA. For this scenario, you will provision an identity server SSL/TLS certificate for the Cisco ASA. You will use the default connection profile, DefaultWEBVPNGroup. Then you will create a bookmark list and assign it to the default group policy, DfltGrpPolicy. Finally, you will create one user in the local user database.

Configuration Tasks

This topic describes required configuration tasks to configure basic clientless SSL VPN on Cisco Adaptive Security Appliance.



Five general deployment tasks are necessary to configure basic clientless SSL VPN:

1. Enable clientless SSL VPN on ASA:
   - Enable SSL VPN traffic termination on a Cisco ASA security appliance interface, which enables the security appliance SSL VPN server function.
   - You must assign the installed identity certificate of the Cisco ASA to the chosen VPN traffic termination interface.
2. Edit default connection profile or create custom one:
   - Select the authentication method. Default options is local AAA authentication method.
   - Allow the user to choose a connection profile at login. This setting is required if you want to assign users to a specific connection profile rather than having them use the default DefaultWebVPNGroup profile.
3. Edit default group policy or create custom one:
   - Make sure that support for clientless SSL VPNs is enabled.
   - Create a bookmark list and apply it to the group policy.
4. If you use custom connection profile and group policy, you have to apply the custom group policy to the custom connection profile.
5. Create a user account in the Cisco ASA local database. You should ensure that these user accounts can only log into the VPN, and not to the security appliance management user interfaces (the Cisco Adaptive Security Device Manager and the CLI).

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Before implementing a basic clientless SSL VPN on Cisco ASA, you must obtain and analyze several pieces of information that relate to the network and system environment:

The IP addressing plan for the VPN gateway: This information enables you to assign an appropriate IP address to the VPN-terminating interface of the Cisco ASA.

The enterprise naming plan for the VPN gateway: This information enables you to assign a correct name in the VPN gateway SSL/TLS identity certificate.

The enterprise certificate policy and certificate settings: With this information, you can enroll the Cisco ASA into a PKI (if desired) and you can include all of the relevant fields inside a PKI-provisioned certificate.

The username and password policy of the enterprise: This information enables you to correctly create the local user database on the Cisco ASA.

Consider the following guidelines when deploying a basic Cisco clientless SSL VPN solution on Cisco ASA:

Use PKI-provisioned certificates for SSL VPN identity certificate.

Use static passwords only in small, single-device, low-risk environments.

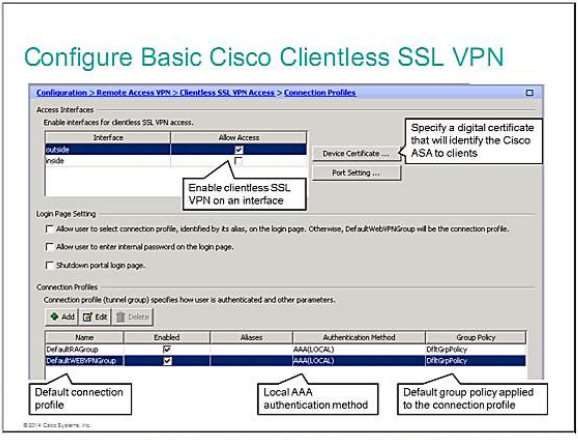
Strictly set the service type of local VPN user accounts to prevent these accounts from using management access.

You can use the DefaultWebVPNGroup instead of a specific group; however, this choice makes it more difficult to differentiate users later.

To increase authentication strength provided by static passwords, consider using client certificates or one-time passwords, at the expense of management complexity and cost.

Configure Basic Cisco Clientless SSL VPN

This topic describes how to configure basic clientless SSL VPN on Cisco Adaptive Security Appliance.



The figure shows the Cisco Adaptive Security Device Manager Connection Profiles panel, from which you can begin several of the clientless SSL VPN configuration tasks, including selecting or creating an identity certificate that will identify the ASA to the client when it attempts to create a connection. After provisioning an identity server SSL/TLS certificate to the Cisco ASA, you can begin the other clientless SSL VPN configuration tasks.

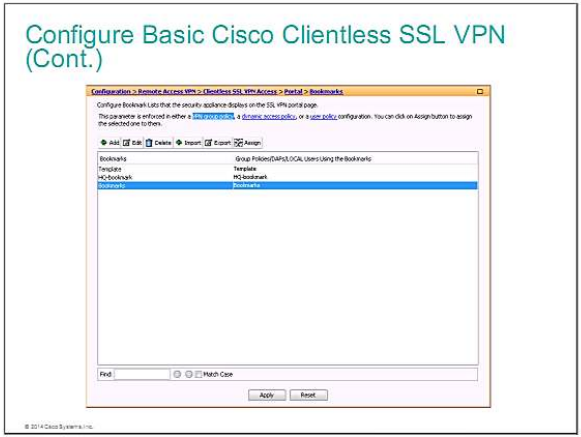
The figure presents a summary of these configuration tasks:

Clientless SSL VPN function on the Cisco ASA is enabled on the outside interface.

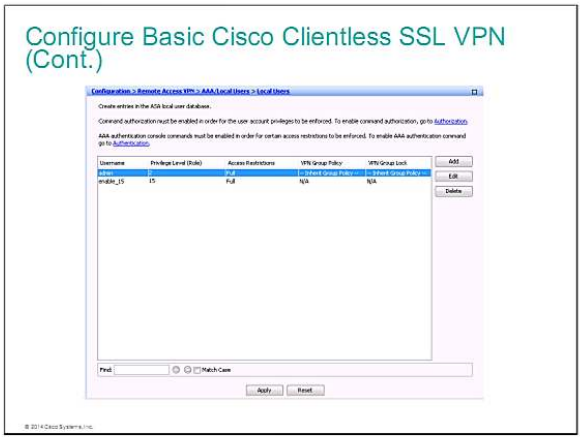
The default connection profile DefaultWEBVPNGroup is used. Authentication method is set to local AAA, as local database on Cisco ASA is used to authenticate SSL VPN users.

The deafult group policy DfltGrpPolicy is applied to the connection profile. This group policy is configured to support clientless SSL VPN.

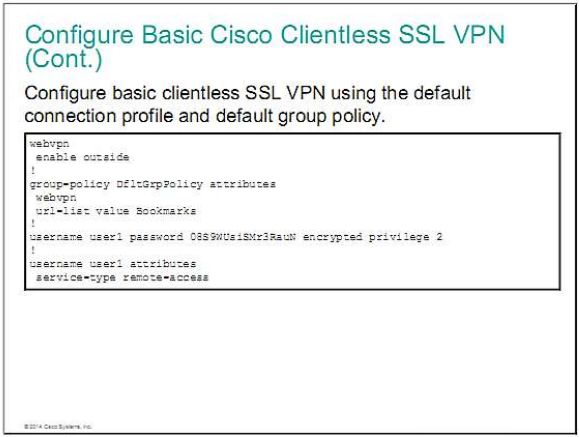
https://t.me/learningnets
© 2014 Cisco Systems, Inc.



This figure shows the Cisco ASDM Bookmarks panel, which enables you to create bookmark lists and add bookmarks to them. From this panel, you can also assign the bookmark list to a group policy.



This figure shows the Cisco ASDM Local Users panel, which enables you to create user accounts in the local user database of the Cisco ASA.



To configure basic clientless SSL VPN using the CLI, use the commands that are described in the following table:

Commands	Description
webvpn **enable** *ifname*	The **webvpn** command enters SSL configuration mode. The **enable** *ifname* command enables SSL VPN on an interface.
group-policy *group_policy_name* **attributes** **webvpn** **url-list value** *list_name*	The **group-policy** *group_policy_name* **attributes** command enters group policy configuration mode. The **webvpn** command enters SSL VPN configuration submode. The **url-list value** *list_name* command applies the bookmark list to the group policy.
username *username* **password** *password*	The **username** *username* **password** *password* command creates a user in the local user database.
username *username* **attributes** **service-type remote-access**	The **username** *username* **attributes** command enters configuration mode for the user. The **service-type remote-access** command specifies that the username can be used for VPN only and not for management access to the ASA.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

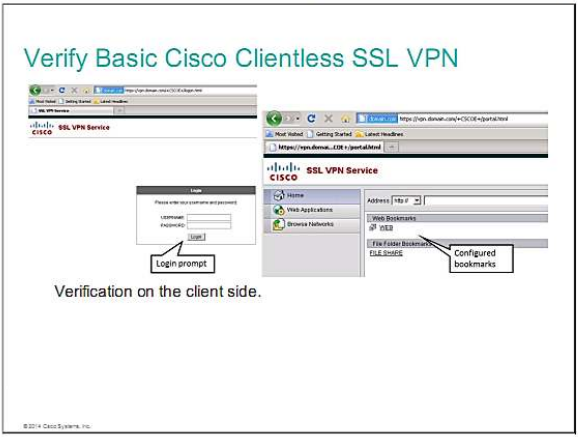
Commands	Description
crypto ca trustpoint **revocation-check** **id-usage ssl-ipsec** **fqdn** **subject-name** **enrollment self**	The **crypto ca trustpoint** command enters trustpoint configuration mode. The **revocation-check** command specifies whether revocation checking is needed for a trustpool policy. The **id-usage** command specifies how the enrolled identity of a certificate can be used. The **fqdn** command includes the indicated FQDN in the Subject Alternative Name extension of the certificate during enrollment. The **subject-name** command includes the indicated subject DN in the certificate during enrollment. The **enrollment self** command generates a self-signed certificate.
crypto ca enroll	The **crypto ca enroll** command starts the enrollment process with the CA.

Note Bookmark list is described using the XML language and cannot be created using the CLI. You can only import already created bookmark list using the CLI. It is highly recommended to use ASDM to create a bookmark list.

https://t.me/learningnets

Verify Basic Cisco Clientless SSL VPN

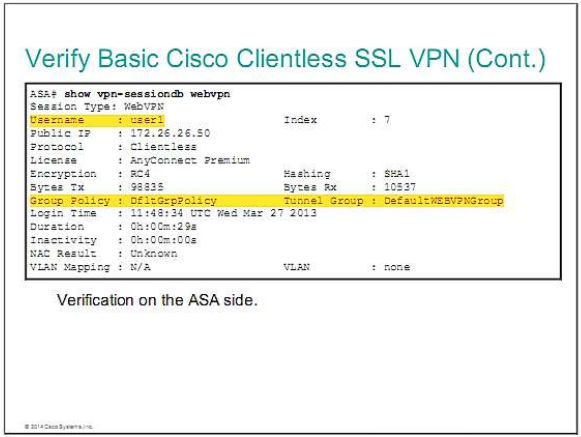
This topic describes how to verify Cisco clientless SSL VPN on Cisco Adaptive Security Appliance.



To verify clientless SSL VPN from the client side, open a web browser and navigate to the hostname of the ASA using HTTPS. Enter a username and password when prompted to log in. After the successful login you should be presented with the web portal. You should see the URL entry field, as well the configured bookmarks. In the example, two bookmakrs have been configured. The first one provides a link to a web server, and the second provides a link to a file server.

This topic describes how to verify Cisco clientless SSL VPN on Cisco ASA.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

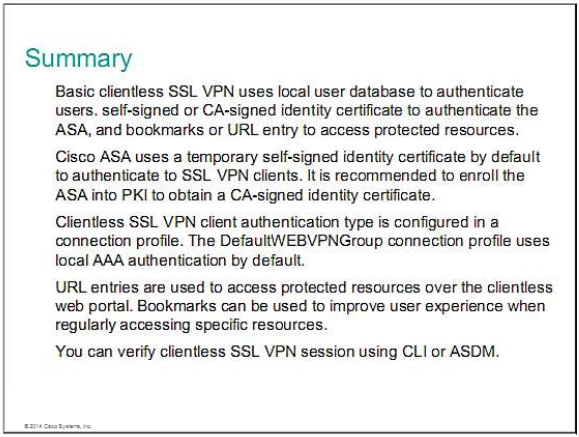


To verify clientless SSL VPN on the ASA using the CLI, use the **show vpn-sessiondb webvpn** command. The output of this command will show you all established clientless SSL VPN sessions. The output displays a username of the connected user, and the connection profile and group policy a user is assigned to.

Verification of SSL VPN session establishment on Cisco ASA could also be done using Cisco Adaptive Security Device Manager. To verify the connection of the client on the Cisco ASA, choose **Monitoring > VPN > VPN Statistics > Session** from the Cisco ASDM.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 3

Deploying Application Access in Cisco ASA Clientless SSL VPN

Many enterprise applications are not web-based and use other standard or proprietary protocols to communicate over IP networks. Therefore, clientless SSL VPN gateways must provide some alternative possibilities for users to access these application resources.

Further, the enterprises need scalable authentication schemes in which the network devices offload the authentication process to back-end user databases such as LDAP, TACACS+, or RADIUS. In clientless SSL VPNs, PKI offers a scalable and secure authentication method.

This lesson discusses the application plug-ins and smart-tunnel features of the Cisco Adaptive Security Appliance SSL VPN gateway. These features provide clientless access to a wide range of thin- and thick-client applications.

Upon completing this lesson, you will be able to meet these objectives:

- Describe the application access methods in Cisco clientless SSL VPNs
- Describe the SSL VPN application access solution components
- Describe the application plug-ins
- Describe the application plug-ins that are supported by the Cisco ASA adaptive security appliance
- Describe a configuration scenario for Cisco SSL VPN plug-ins
- Describe Cisco SSL VPN plug-in configuration
- Describe how to verify plug-in operation
- Describe smart tunnels
- Describe the smart tunnel configuration options and scenario

https://t.me/learningnets

Describe smart tunnel configuration

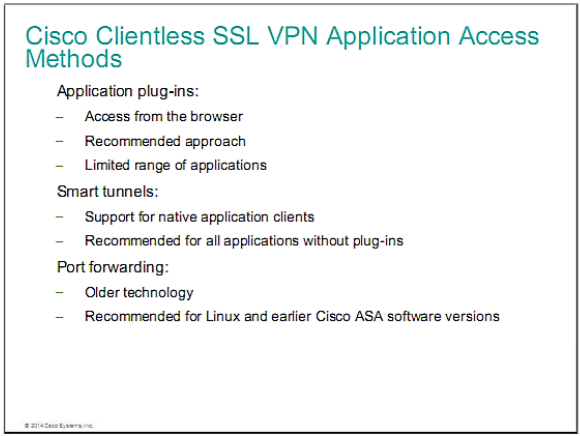
Describe how to verify smart tunnel operation

Describe the common problems with clientless SSL VPN configuration

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Cisco Clientless SSL VPN Application Access Methods

This topic describes the application access methods in Cisco clientless SSL VPNs.



You can choose several options when deploying advanced application access:

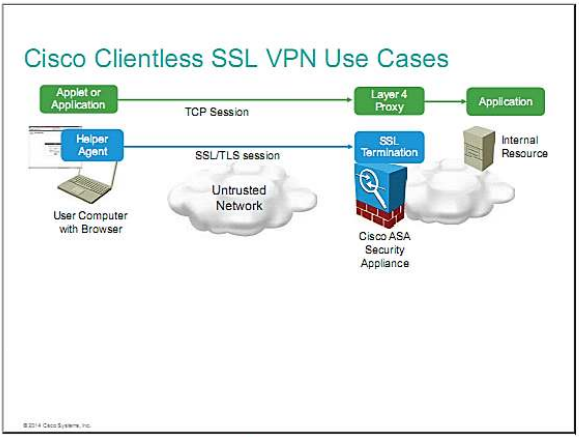
Application plug-ins: This option provides users with thin-application client access to enterprise resources. This approach is the recommended approach but supports a limited set of applications.

Smart tunnels: This option provides users with native-application client access to enterprise resources. This approach is the recommended approach for all applications that do not have a plug-in.

Port forwarding: This option provides users with native-application client access to enterprise resources. This method should be used only if smart tunnels cannot be deployed, such as on Linux workstations or when running older Cisco Adaptive Security Appliance software versions.

Cisco Clientless SSL VPN Application Access Solution Components

This topic describes the SSL VPN application access solution components.



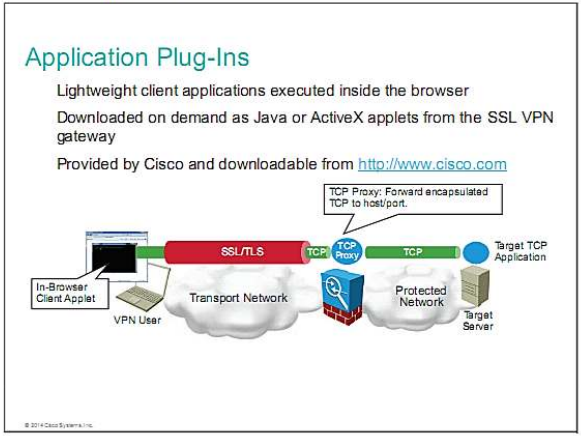
The SSL VPN rewriting proxy in the Cisco Adaptive Security Appliance provides clientless, transparent access to web and CIFS resources behind the Cisco ASA to clients, using only a web browser. To provide access to other enterprise applications, including terminal access, database clients, or IM applications, the Cisco ASA provides several methods of relaying data over the SSL VPN gateway, without requiring an SSL VPN client installation.

To enable such advanced application-access methods, the Cisco ASA loads some additional software in the web-browsing session of the client. This software acts as a lightweight client or a helper relay agent that enables additional communications over an SSL VPN session. The Cisco ASA does not require any other components to enable advanced application access.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Application Plug-Ins

This topic describes the application plug-ins.



An application plug-in is a separate program that a web browser invokes to perform a dedicated function, such as connecting a client to a server within the browser window. Application plug-ins are lightweight client applications that provide basic application functionality inside the browser of a user. These plug-ins are downloaded on demand from the SSL VPN gateway, and all their communications with internal protected resources are transparently encapsulated within the SSL VPN session. The user does not need to use any local applications, but instead uses only the thin-application plug-ins.

The figure illustrates plug-in-based application access:

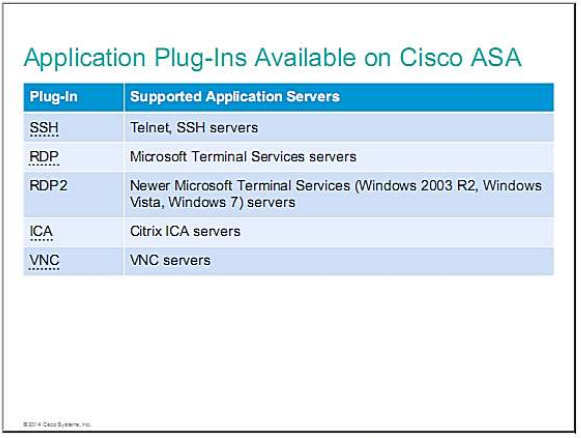
1. A user connects and authenticates to the SSL VPN portal and then runs an application plug-in that allows access to a server that is running on the internal network.
2. The application plug-in runs inside the browser and reuses the SSL session between the browser and the SSL VPN gateway to forward a TCP connection.
3. The SSL VPN gateway extracts the TCP session from the SSL VPN session, establishes a TCP connection with the destination server, and acts as a data relay between the two TCP sessions.

Cisco distributes the application plug-ins through the Cisco Adaptive Security Appliance (remote access plug-ins) software download page, which you must import into the appliance. Each plug-in contains multiple files that are packed into a JAR file.

https://t.me/learningnets

Application Plug-Ins Available on Cisco ASA Security Appliances

This topic describes the application plug-ins that are supported by the Cisco Adaptive Security Appliance.



Cisco provides several application plug-ins for download from http://www.cisco.com:

SSH: This plug-in provides a standard terminal emulator with support for Telnet and SSH versions 1 and 2 (SSHv1 and SSHv2) protocols. You can control the SSH version by providing command-line parameters to the application plug-in in its URL.

RDP: This plug-in provides a Microsoft RDP client that you can use to connect to older versions of Microsoft Terminal Services servers.

RDP2: This plug-in provides an RDP client that you can use to connect to newer versions of Microsoft Terminal Services (Microsoft Windows 2003 Server R2, Windows Vista, and Windows 7) servers. The Microsoft ActiveX version of the plug-in also supports advanced functionality, such as sharing remote drives or printers on the terminal server. You can control this functionality by providing command-line parameters to the application plug-in in its URL.

ICA: This plug-in provides an ICA client that you can use to connect to Citrix WinFrame, XenApp, and XenDesktop terminal services.

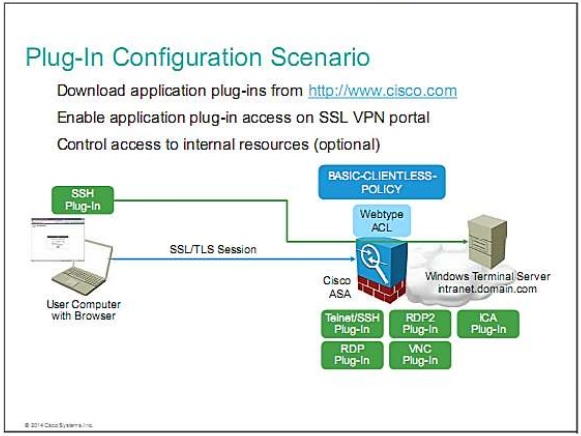
VNC: This plug-in provides a VNC client that you can use to connect to VNC display servers. The VNC plug-in does not support encryption over the internal network.

Note You can import both RDP and RDP2 plug-ins to make them both available to clientless users.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Plug-In Configuration Scenario

This topic describes a configuration scenario for Cisco SSL VPN plug-ins.



To configure application plug-in-based access through a Cisco Adaptive Security Appliance SSL VPN gateway, follow these configuration tasks:

1. Download and import application plug-in files from http://www.cisco.com to the Cisco ASA security appliance FFS.
2. Enable application plug-in access on the SSL VPN portal, using optional application plug-in bookmarks.
3. Optionally, deploy access-control features on the Cisco ASA security appliance SSL VPN gateway to control access from application plug-ins to internal resources.

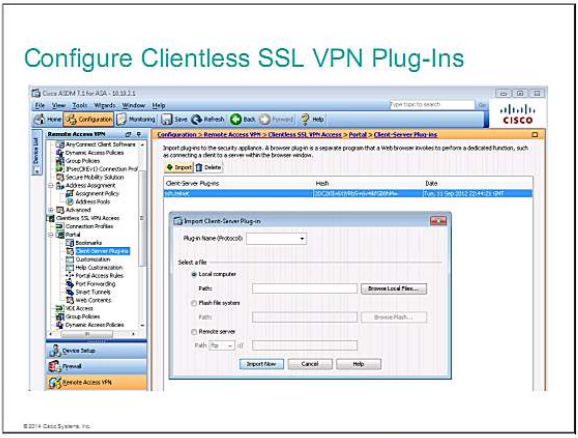
The figure shows the configuration scenario for the plug-ins. You can import all the required application plug-ins on the Cisco ASA and configure an SSH bookmark, which the client will choose after login. The client will then load the SSH plug-in. You can also configure the Cisco ASA to limit access to the protected network by using a webtype ACL, which will allow only the specific SSH connection to the internal network.

The configuration scenario assumes that the Cisco ASA is already configured with a basic clientless SSL VPN gateway functionality.

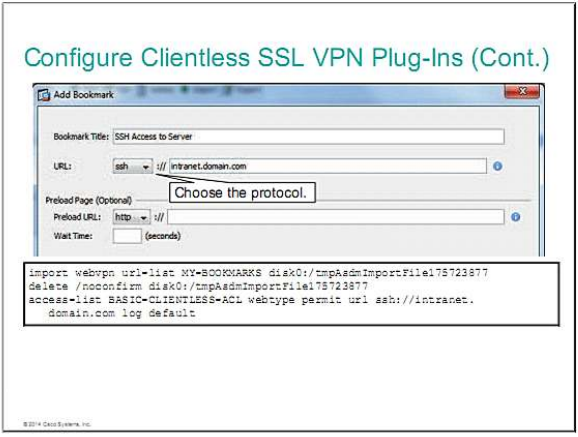
https://t.me/learningnets

Configure Clientless SSL VPN Plug-Ins

This topic describes the Cisco Secure Socket Layer VPN plug-in configuration.



To configure Clientless SSL VPN Plug-ins go to the **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Client-Server Plug-ins** in the Cisco Adaptive Security Device Manager and click **Import** button, to import client server plug-ins.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

When you have imported the plug-ins to the Cisco Adaptive Security Appliance, you do not need to authorize their use in group policies. If you are using SSL VPN portal bookmarks, you have the option of choosing additional protocols—as supported by plug-ins—in the bookmark URL specification. If you allow free URL entry on the SSL VPN portal, users can enter URLs with protocols that the imported plug-ins support. The figure shows the second task in the configuration procedure, in which you create bookmarks that use application plug-ins and you enable users to start these plug-ins from the SSL VPN portal.

The configuration output shows the CLI commands that Cisco ASDM applies to configure application plug-in access.

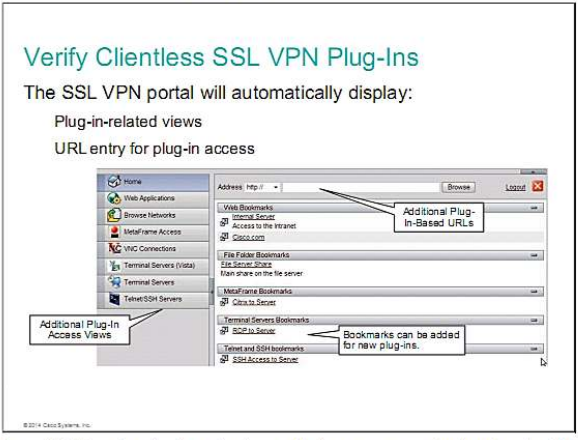
The **import webvpn url-list** command uses the XML file to apply the bookmark list.

The **access-list BASIC-CLIENTLESS-ACL webtype permit url** command configures the ACE for access to the SSH server.

https://t.me/learningnets

Verify Clientless SSL VPN Plug-Ins

This topic describes how to verify the plug-in operation.

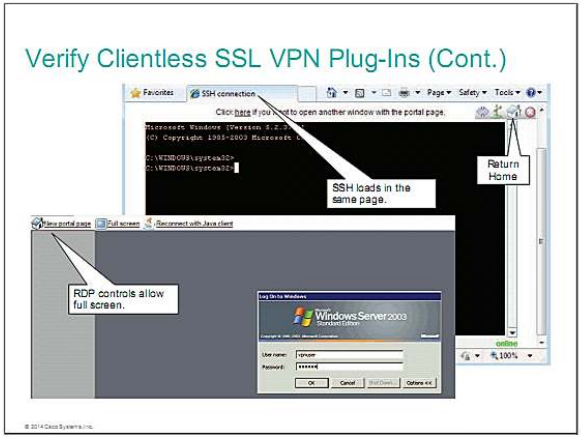


To verify the availability of application plug-ins and their correct operation, log into the SSL VPN portal. You should see additional plug-in-related views on the left side of the SSL VPN portal home page. These views, if selected, display all plug-in-related bookmarks on the right side of the portal page. By default, all bookmarks are listed. In the example, the newly created SSH bookmark is available in the right pane of the portal window.

Users can also use the URL entry control (if allowed) to select URLs for protocols that imported plug-ins enable.

The example also shows other imported plug-ins, such as Telnet and SSH, Citrix MetaFrame, and RDP.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

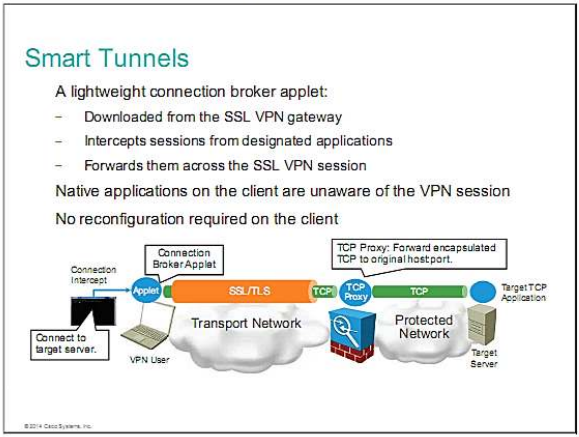


From the VPN portal home page, click the plug-in bookmark or enter a plug-in-related URL in the URL entry control (if allowed) to start an application plug-in. The portal starts a plug-in in a new window.

In the example, when the user clicks the **RDP2 plug-in**, the RDP2 application plug-in starts and opens a Java RDP2 client in the same window. This window shows the login page of an internal server. The example also shows that when the SSH bookmark is clicked, a window opens and gives access to the intranet server through an embedded CLI.

Smart Tunnels

This topic describes the smart tunnel technology.



Smart tunnels enable users to use native client applications, without the need for administrative rights or application reconfiguration. Smart tunnels work by downloading a smart-tunnel agent (connection broker) applet to the client system. This applet intercepts all local socket calls (that is, connection requests of the application to the operating system kernel) from Winsock2 TCP applications and automatically redirects them into the SSL VPN session.

The figure illustrates smart tunnel-based application access:

1. A user connects and authenticates to the SSL VPN portal. A smart-tunnel configuration is in place for that user and automatically downloads and executes the smart-tunnel agent on the client system. For example, the user starts the IBM Lotus Sametime IM client.
2. The smart-tunnel applet intercepts the TCP connection of the Lotus Sametime client to the real server and forwards it over the existing SSL VPN session.
3. The SSL VPN gateway extracts the TCP session from the SSL VPN session, establishes a TCP connection with the real target server, and acts as a data relay between the two TCP sessions.

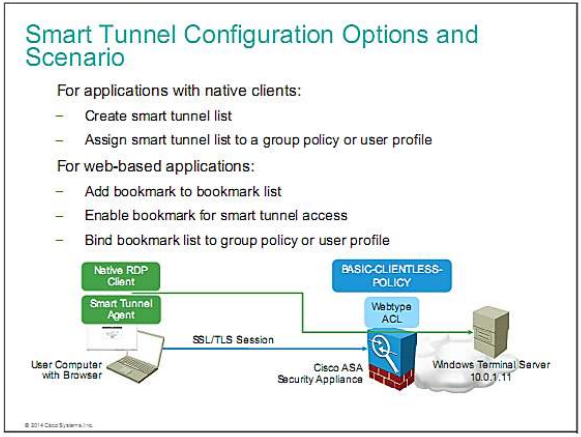
As long as the clientless SSL VPN session of the user is established, specific local applications can access protected resources, with the agent relaying their communications through the browser session.

Smart tunnel offers better performance than plug-ins. Unlike port forwarding, smart tunnel simplifies the user experience by not requiring the user connection of the local application to the local port. Unlike port forwarding, smart tunnel does not require users to have administrator privileges. The advantage of a plug-in is that it does not require the client application to be installed on the remote computer.

https://t.me/learningnets

Smart Tunnel Configuration Options and Scenario

This topic describes the smart tunnel configuration options and scenario.



You have several options when deploying smart-tunnel access through a Cisco Adaptive Security Appliance SSL VPN gateway:

Applications with native clients:

- Create a smart-tunnel list.
- Assign the smart-tunnel list to a group policy or user profile.

Web-based applications:

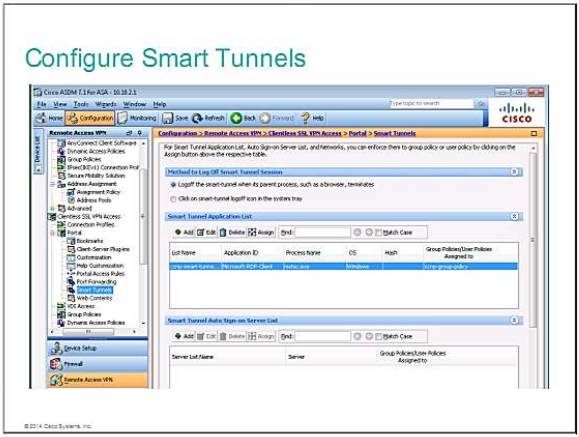
- Add a bookmark to the bookmark list.
- Enable a bookmark for smart-tunnel access.
- Bind a bookmark list to a group policy or user profile.

Control access to internal resources.

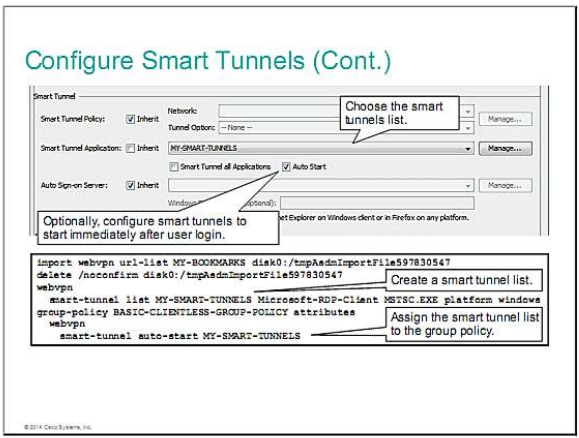
© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Configure Smart Tunnels

This topic describes the configuration of smart tunnels.



To enable smart tunnels in the Clientless SSL VPN access go to the **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels** in the Cisco Adaptive Security Device Manager and click **Add** button, to add smart tunnel configuration.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

To enable smart-tunnel access for applications with a native client, create a list of applications that are subject to smart tunnels relaying on the remote client. Then, enable smart-tunnel access for applications with a native client by applying a smart-tunnel policy to a group policy or user profile, as shown in the figure. The list of applications is not shown in the Cisco ASDM configuration.

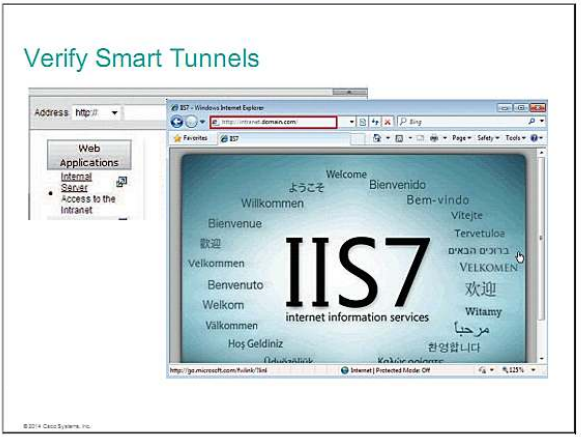
In the CLI configuration, you see that the **webvpn** command has been used to enter SSL VPN server configuration submode. The **smart-tunnel list** command is used to configure smart-tunnel application lists. In the example, a **smart-tunnel list** that is named MY-SMART-TUNNELS, which specifies the Windows RDP application, is defined.

Next, assign the configured smart-tunnel application list to a group policy by entering group-policy attributes mode and then switching to its webvpn submode. Use the **smart-tunnel** command to assign the smart-tunnels list. Use the optional *auto-start* argument to automatically download and start the smart-tunnel agent. In the example, the MY-SMART-TUNNELS smart-tunnel list is applied to the BASIC-CLIENTLESS-GROUP-POLICY group policy.

https://t.me/learningnets

Verify Smart Tunnels

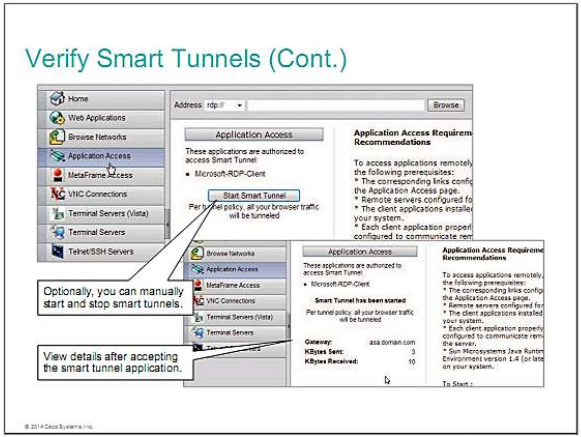
This topic describes how to verify smart tunnel operation.



When you access the SSL VPN portal and the **Auto Start** option is configured for the smart tunnel, you are required to allow the browser to run the ActiveX control. Windows also prompts you to allow the application to take action on your system. After you accept the Windows configuration changes, you can click a link in the portal to launch access to the internal resource via smart-tunnel access. Notice that the URL in the example is not mangled as before. You can see the URL of the server as configured in the portal bookmark.

When you choose the bookmark, the application displays in a separate browser window. You must accept the smart-tunnel applet that is signed by Cisco before using the application.

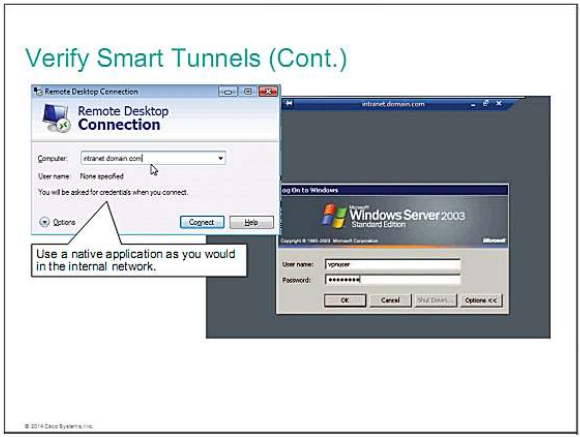
https://t.me/learningnets
© 2014 Cisco Systems, Inc.



To verify the smart tunnels and their correct operation when they are configured to start manually, log into the SSL VPN portal. You should see an **Application Access** view on the left side of the SSL VPN portal home page.

Click the **Application Access** view and click the **Start Smart Tunnel** button. The browser downloads and executes the smart-tunnel relay agent, on demand.

In the **Application Access** view, you can also observe basic smart-tunnel statistics and settings, such as the list of applications that are enabled for smart tunneling and the amount of data that is exchanged through the smart-tunnel session since its start.



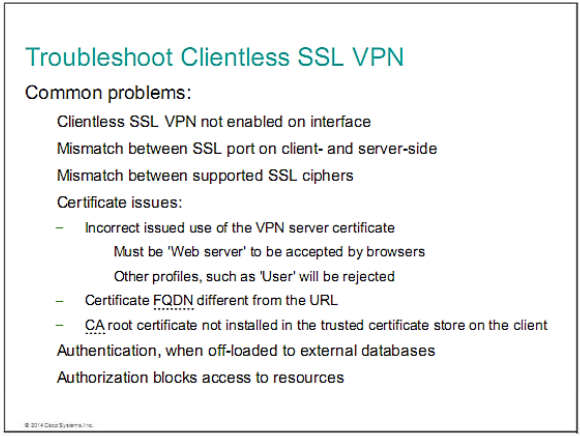
https://t.me/learningnets

After the smart-tunnel relay agent is running, you can start the native applications that are present in the smart-tunnel list of your group policy and directly access internal resources (from an end-user perspective). In the example, the native Windows RDP client is used as if it were started in a non-VPN environment (or a full-tunneling Cisco AnyConnect session), displaying the login page of an internal server.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Troubleshoot Clientless SSL VPN

This topic describes the common problems with clientless SSL VPNs.



This list describes the common issues with clientless SSL VPNs:

The most basic requirements for a functional clientless SSL VPN are that the SSL VPN is enabled on an appropriate VPN server interface, that it uses the same TCP port that the clients access, and that IP/TCP connectivity is ensured along the path. If any of these conditions is not met, the browser will return an error page.

Another potential reason for a clientless VPN failure is the SSL/TLS protocol or cipher mismatch. The client browser will typically not give you any hint as to the nature of the problem. If you investigate the logging information on the security appliance, you will see the SSL/TLS cipher negotiation failure:

```
SSL lib error. Function: SSL3_GET_CLIENT_HELLO Reason: no shared cipher
```

There are many different special certificate types, such as user, web server, root signing, wildcard, Extended Validation, code signing, and email certificates. The browsers will accept and validate SSL/TLS sessions to holders of the web server certificates. If the Cisco Adaptive Security Appliance identity certificate was issued using a user certificate template, the clients will reject SSL/TLS connections to the Cisco ASA.

The clients may not be able to validate the server certificate for a number of reasons. The browser will return a warning if the subject FQDN in the server certificate is different from the FQDN in the VPN URL. The validation will fail if neither the root CA certificate nor the VPN server certificate are in the respective certificate stores. Other certificate problems may be related to incorrect system time. More specifically, the client time must fall within the validity range of the server certificate. Otherwise the client will decline the server certificate as invalid.

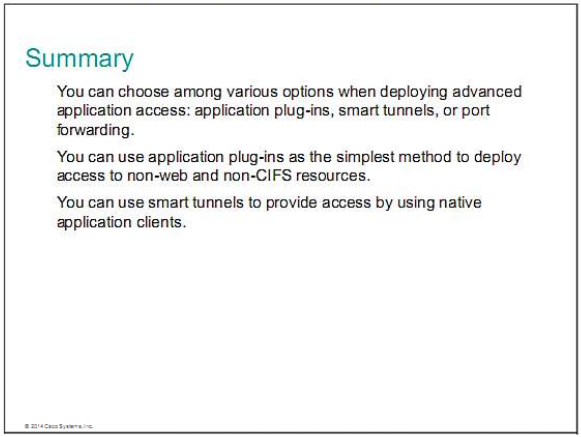
You may encounter user authentication problems in clientless SSL VPNs. User authentication can be based on user certificates. It can be performed against the local or an external database. Troubleshooting of the authentication problems often involves the analysis of the integration with the external authentication server, such as the Cisco Identity Services Engine.

Authorization in clientless SSL VPNs in performed using the WebVPN ACLs, which can be configured either locally on the Cisco ASA, or downloaded from the RADIUS server. An additional method is provided by a hybrid approach where the RADIUS Class attribute references a local group policy by name. Upon obtaining the Class attribute, the VPN server applies the referenced group policy to the session.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Summary

This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Lesson 4

Deploying Advanced Authentication and Authorization in Clientless SSL VPN

In clientless SSL VPNs, PKI offers a scalable and secure authentication method. The lesson discusses the various authentication and authorization approaches that you should evaluate when designing a clientless SSL VPN solution, including the option of combining multiple authentications in a single process.

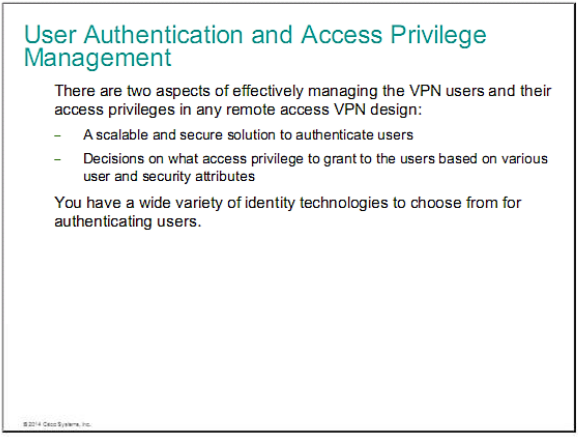
Upon completing this lesson, you will be able to meet these objectives:

- Describe the user authentication and access privilege management in any remote access VPN design
- Describe the local CA user authentication scenario
- Describe how to configure local CA user authentication
- Describe how to verify client authentication using the local CA
- Describe a client authentication scenario using an external CA
- Describe a client authentication scenario using an external CA
- Describe client-side authentication and authorization using a AAA server
- Configure client-side authentication and authorization using a AAA server
- Describe how to verify client-side authentication and authorization using a AAA server
- Describe a scenario with multiple client authentications
- Describe how to configure the multiple client authentication
- Describe how to verify multiple client authentication

https://t.me/learningnets

User Authentication and Access Privilege Management

This topic provides an overview of user authentication and access privilege management in any remote access VPN design.



Effectively managing the VPN users and their access privileges is the core consideration in any remote access VPN design. There are mainly two aspects:

A scalable and secure solution to authenticate users

Decisions on what access privilege to grant to the users based on various user and security attributes

AAA stands for authentication (which defines who you are), authorization (which defines what you are allowed to do), and accounting (which provides a record of what you did). User authentication is a key step in an SSL VPN solution. Aside from validating users' credentials, user authentication allows an SSL VPN gateway to assign the user to a policy group. The assignment is made by using a user's organization group information, which is derived during the authentication phase, along with other attributes, such as endpoint security posture and time of day. The policy group defines the authorization privileges of the users.

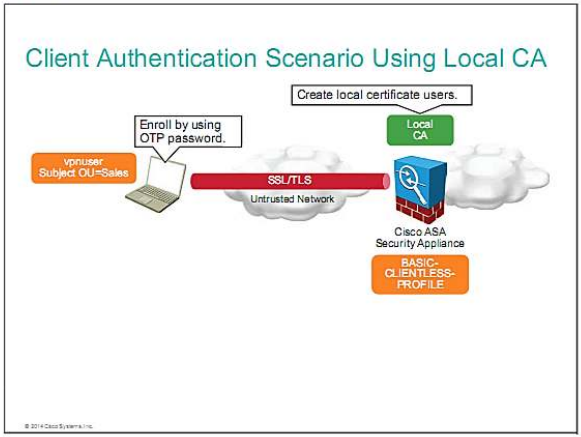
You have a wide variety of identity technologies to choose from for authenticating users. The common choices are passwords, RADIUS, TACACS+, OTP systems, PKI, smart cards, and so on. For remote access VPN authentication, a two-factor OTP system provides the strongest security and manageability combination. It is also common for small to medium-sized companies to leverage existing user directory infrastructure such as LDAP, Windows NTLM, or Windows XP/2000 Active Directory for VPN user authentication. To use this, you need to apply and enforce strong password policies because the strength of the security relies on those policies.

The design of the AAA system can vary depending on the size of your network and the disparity of access methods.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Client Authentication Scenario Using Local CA

This topic presents the local CA user authentication scenario.



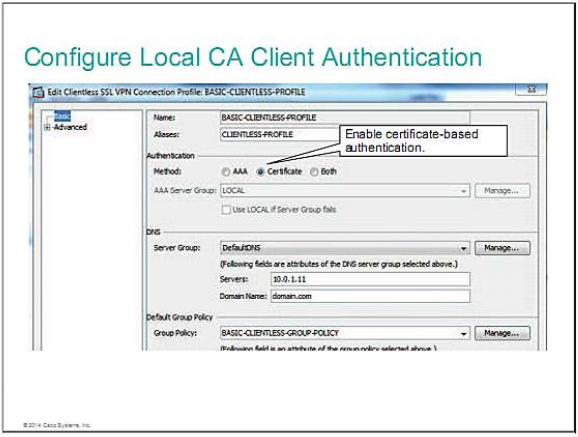
The figure presents the configuration scenario that uses the local CA function on the Cisco Adaptive Security Appliance and CA user accounts.

The FQDNs of the users include the organizational unit to which the users belong. When the users connect to the SSL VPN, the Cisco ASA checks the organizational unit information to apply the correct connection profile (BASIC-CLIENTLESS-PROFILE in this configuration scenario) to the user session.

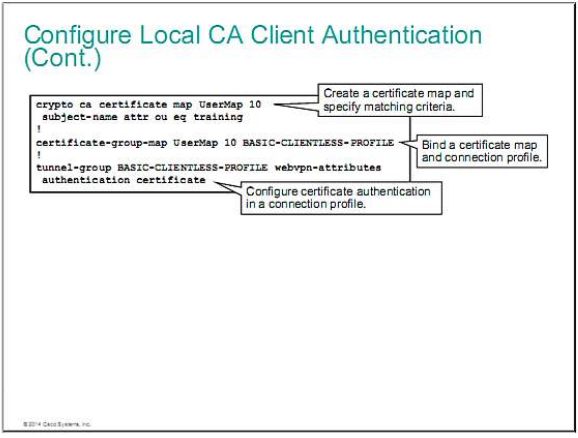
https://t.me/learningnets

Configure Local CA Client Authentication

This topic describes how to configure client authentication against the local CA.



When you configure the Cisco Adaptive Security Appliance local CA and CA user accounts, enable certificate-based authentication in the required connection profiles. The figure illustrates how to enable certificate-based authentication.

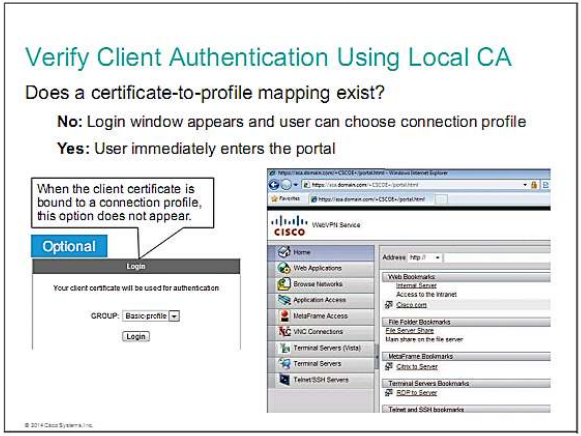


This figure shows the CLI commands that you use to enable client authentication that uses the local CA. The figure does not show the configuration for the local CA nor the creation of the local CA user.

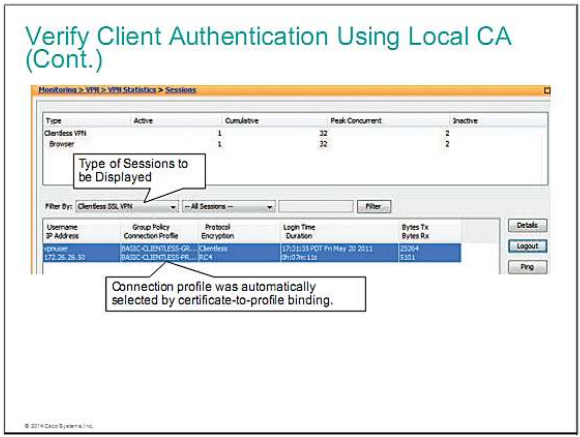
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Client Authentication Using Local CA

This topic describes how to verify client authentication using the local CA.



To verify the certificate-based client authentication, connect to the SSL VPN. If the security appliance has a certificate-to-connection profile map that binds the user certificate to a connection profile, then the specified connection profile is automatically chosen. If the map does not match, then the user may choose a connection profile, as shown in the figure. The **WebVPN Service** window in the figure shows a successful certificate-based authentication in which the user is not prompted for any credentials.

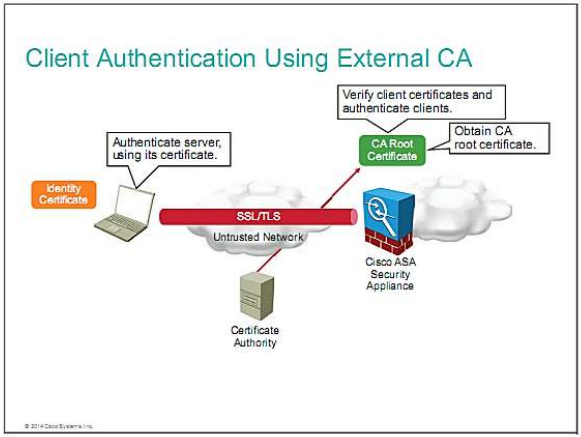


To view the user session parameters, including the selected connection profile, choose **Monitoring > VPN > VPN Statistics > Sessions**, and then choose **Clientless SSL VPN** from the **Filter By** drop-down list. This verification shows whether the Certificate to Connection Profile Map has activated the correct profile.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Client Authentication Using External CA

This topic presents a client authentication scenario using an external CA.



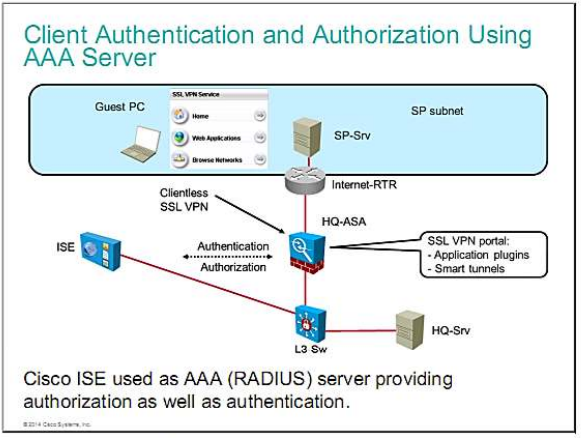
The figure illustrates a configuration scenario for using an external CA to deploy browser (client) certificates. The security appliance receives the root certificate of the external CA and enrolls with the CA to obtain its identity certificate.

The users may have the root certificate preinstalled in their browsers, depending on which CA provider is selected in a particular enterprise environment. Alternatively, the enterprise can manage its own CA. In either case, users must enroll with the CA to obtain their identity certificates.

Both server authentication and client authentication are certificate-based. The entities use the self-signed root certificate to verify the authenticity of the peer certificate. After the peer certificate is validated, the entities extract the embedded public key to verify the peer signature.

Client Authentication and Authorization Using AAA Server

This topic describes client authentication and authorization using a AAA server.



The figure shows a scenario in which client-side authentication is performed by a external AAA (RADIUS) server. In this scenario, the Cisco Identity Services Engine is used as the RADIUS server.

The Cisco ISE is a RADIUS-based product that offers a centralized control point for comprehensive policy management and enforcement. It includes rigorous identity enforcement with an automatic device feed service to keep the profiling engine up-to-date with the latest smartphones, tablets, laptops, and even specialized network-enabled devices used in retail, healthcare, and manufacturing industries. Cisco ISE automates and simplifies access control and security compliance for wired, wireless, and VPN connectivity. It is primarily used to provide secure access and guest access, support the Cisco Bring Your Own Device initiatives, and enforce usage policies in conjunction with Cisco TrustSec.

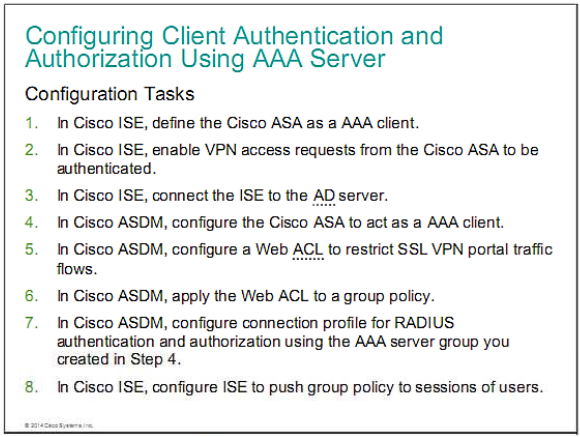
Cisco ISE can use most Cisco network devices as device profiling sensors and access enforcement points. It is also capable of extending authentication services on most 802.1X-compliant devices, although profiling may require a specialized architectural design. Additional deployment components include Cisco Network Access Control Agent, Cisco AnyConnect™, or the native 802.1X supplicant on the endpoint; Cisco partner MDM solutions, Cisco Catalyst® switches and Cisco wireless LAN controllers acting as policy enforcement points for the LAN; and Cisco Adaptive Security Appliances for secure remote access. Cisco ISE also integrates with directory services such as Microsoft Active Directory and Sun ONE Directory Server as policy information points.

The Cisco ISE is available as either a physical or virtual appliance.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Client Authentication and Authorization Using AAA Server

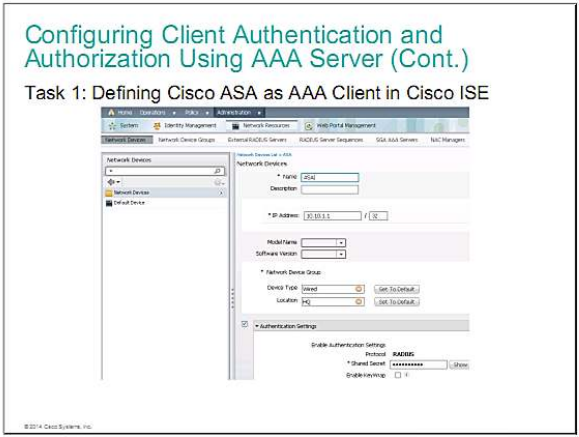
This topic explains how to configure client authentication and authorization using a AAA server.



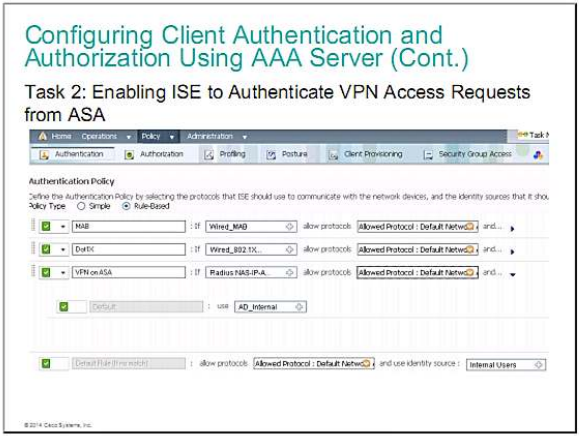
The figure lists the tasks that are required to configure client authentication for clientless SSL VPN using Cisco Identity Services Engine as a RADIUS AAA server. In this example, authorization is configured as well as authentication. Both services are provided by the Cisco ISE acting as RADIUS server. Authorization configuration involves configuring a Web ACL to allow only FTP to HQ-SRV.

Note The screen shots in this topic show the Cisco ISE and Cisco Adaptive Security Device Manager windows and panels that must be used to complete the configuration tasks for client authentication and authorization using the Cisco ISE as a AAA server. Screen shots are not provided for configuration tasks that have already been covered in this course.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets



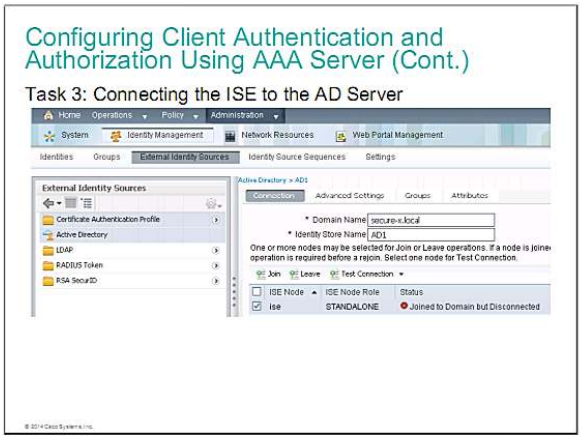
The figure shows the screen in the Cisco ISE GUI that is used to define the Cisco Adaptive Security Appliance as a AAA client in Cisco ISE. To access this screen, go to Administration > Network Resources > Network Devices. Then specify the attributes (such as device name, IP address, and authentication settings) for the device (Cisco ASA in this case) in the Network Devices panel.



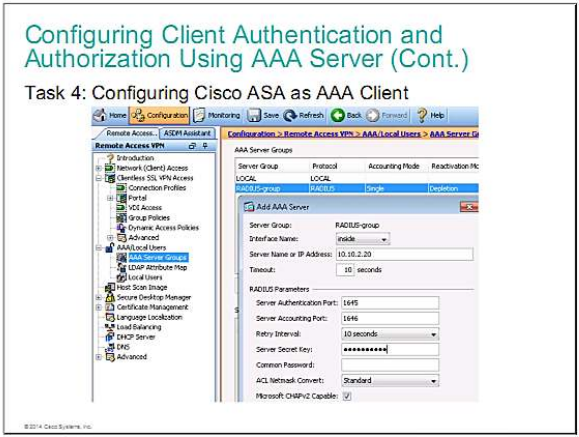
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

The figure shows the Cisco ISE Authentication Policy panel, which enables you to create an authentication policy. Authentication policies define the protocols that Cisco ISE uses to communicate with the network devices, and the identity sources that it uses for authentication. A policy is a set of conditions and a result. A policy condition consists of an operand (attribute), an operator (equal to, not equal to, greater than, and so on), and a value. Compound conditions are made up of one or more simple conditions that are connected by the AND or OR operator. At runtime, Cisco ISE evaluates a policy condition and then applies the result that you have defined based on whether the policy evaluation returns a true or a false value.

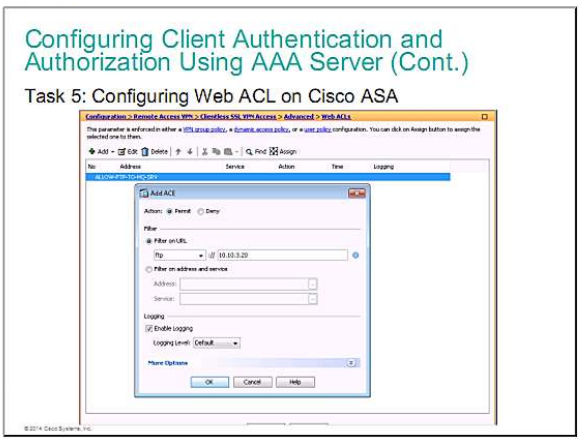
You can create a simple authentication policy or a rule-based policy. In rule-based policies, you can define multiple rules as illustrated in the figure. The identity database is selected based on the first rule that matches the criteria. In the example in the figure, an authentication rule named VPN on ASA has been created. This rule configures the Cisco ISE to process authentication requests coming from a specific IP address (not shown in figure) on the Cisco ASA.



The figure shows the Cisco ISE Active Directory panel, where you can ensure that the ISE is joined and connected to the AD server.



This figure shows the Cisco ASDM AAA Server Groups panel, which enables you to create a AAA server groups and add servers to the server groups. In the example shown in the figure, a AAA server group of type RADIUS has been created on the Cisco ASA. The Cisco ISE (10.10.2.20), which is located on the inside interface of the security appliance, has been added as a member of the server group, and a server secret key has been specified to authenticate the Cisco ISE to the security appliance. The server secret key that you configure here on the AAA client should match the key configured on the AAA server, which in this case, is the Cisco ISE.

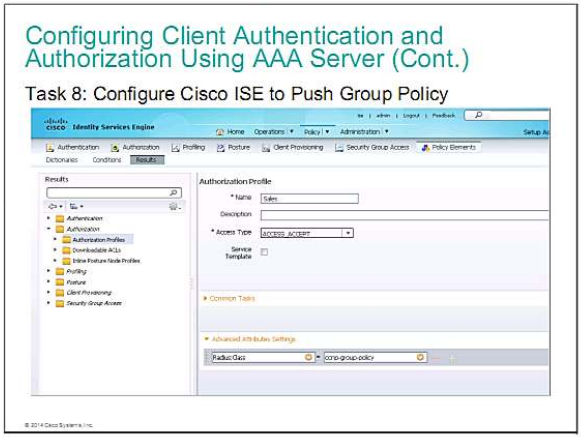


https://t.me/learningnets
© 2014 Cisco Systems, Inc.

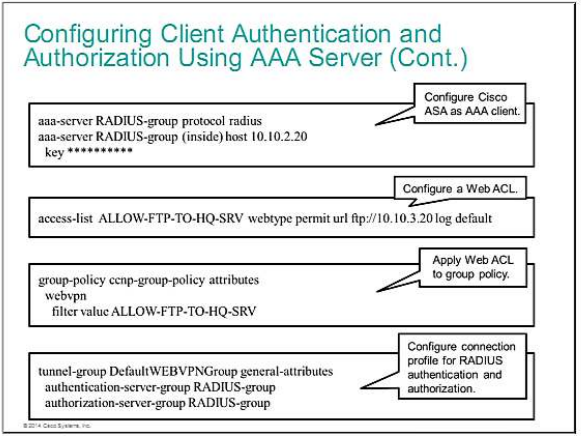
The figure shows the Cisco ASDM Web ACLs panel, where you can create Web ACLs to restrict SSL VPN portal traffic flows. In the example in the figure, a Web ACL called ALLOW-FTP-TO-HQ-SRV has been created and an ACE is being added to it. After you configure the Web ACL, complete Task 6 (applying the Web ACL to a group policy) by using the More Options menu in the Edit Internal Group Policy window. Then complete Task 7 (configuring the connection profile for RADIUS authentication and authorization) by completing the following steps:

In the Basic page of the Edit Clientless SSL VPN Connection Profile window, ensure that the AAA radio button is selected in the Authentication area. Then select the AAA server group you created from the AAA Server Group drop-down list.

Go to **Advanced > Authorization** and select the AAA server group from the Server Group drop-down list within the Authorization Server Group area.



This figure shows the Cisco ISE Authorization Profile panel, which enables you to configure the Cisco ISE to push group policy to sessions of users. In the example shown in the figure, the Cisco ISE is being configured to push the ccnp-group-policy for Sales domain users. The attribute value that you enter in the Advanced Attribute Settings section of this panel must match exactly the policy configured on the ASA.

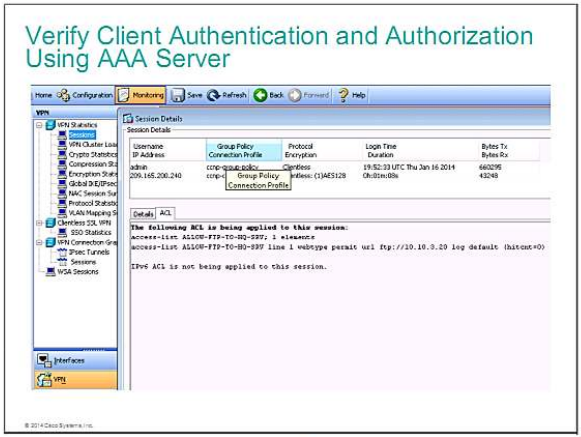


This figure shows the CLI commands that are used to complete the Cisco ASA configurations in this topic. It does not show the CLI commands for configuring the Cisco ISE.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Client Authentication and Authorization Using AAA Server

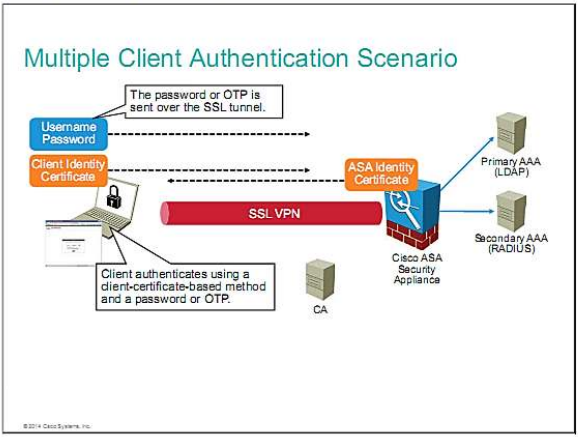
This topic explains how to verify client authentication and authorization by an AAA server.



To verify that a user session has been authenticated and authorized for specific network access by the AAA server (in this case, Cisco Identity Services Engine), you can use the **show access-list** command at the CLI of the Cisco Adaptive Security Appliance to display the Web ACL downloaded by the ISE to the security appliance. You can also view the ACL in the ACL tab or the Cisco Adaptive Security Device Manager Session Details window. To access this window, go to **Monitoring > VPN > VPN Statistics > Sessions** and view the details for the session.

Multiple Client Authentication Scenario

This topic presents a scenario with multiple client authentications.



You can deploy one of several multiple-authentication combinations. The multiple client-side authentication options in clientless SSL VPN are identical to the authentication options in Cisco AnyConnect VPN clients:

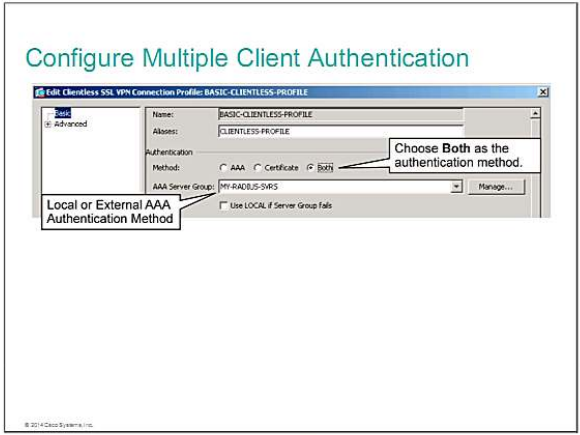
- Certificate-based and one AAA authentication
- Certificate-based and one AAA authentication, where the username for AAA authentication can be extracted from the certificate subject field and hidden from users
- Double AAA authentication (no certificate) with optional username reuse

The figure shows the configuration scenario with the certificate-based authentication and an AAA password-based authentication.

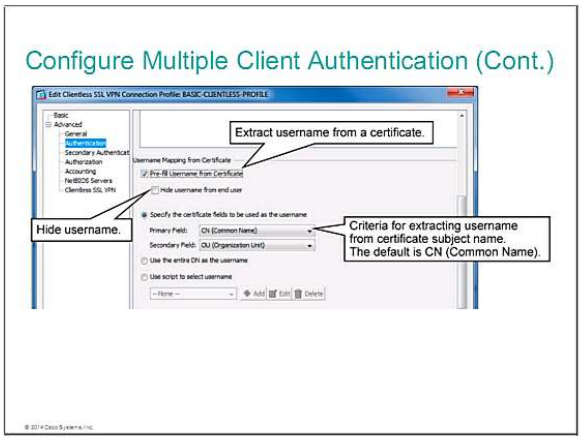
https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Configure Multiple Client Authentication

This topic describes how to configure multiple client authentication, the first using certificates, and the second based on an external AAA server.



The figure shows how to configure a connection profile for certificate-based authentication and a single AAA password-based authentication.

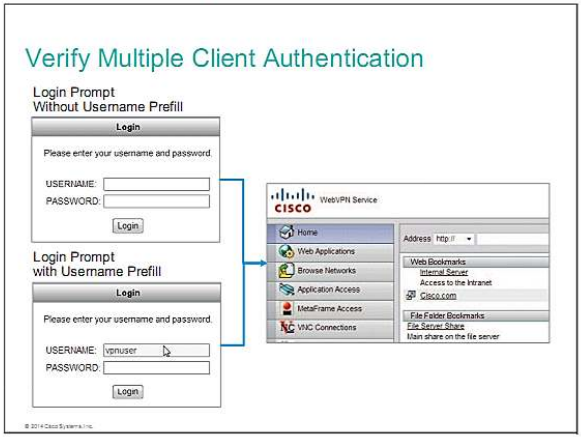


Next, you can optionally configure the username prefill feature, which allows the username for AAA authentication to be extracted from a certificate subject field. The figure shows how to configure this feature.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Verify Multiple Client Authentication

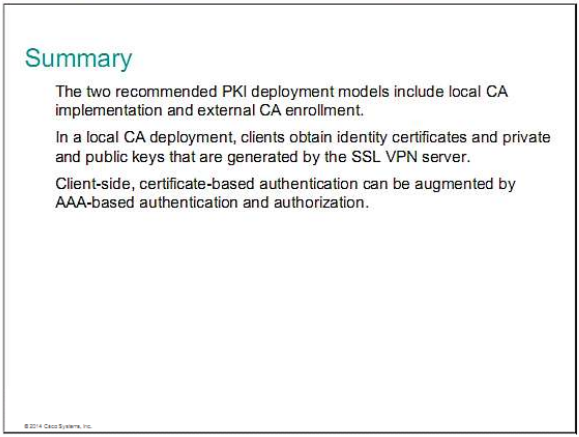
This topic shows how to verify multiple client authentication.



To verify double authentication that uses certificates and a primary AAA method, use a browser to connect to the SSL VPN. If a correct certificate is installed in the certificate store, then the browser submits the certificate for certificate-based authentication. In the second step, the SSL VPN server requests the user credentials. If the username prefill feature is enabled, then the username field is already filled in (the name is pulled from the certificate subject field) and made unavailable. If the field is unavailable, then the client cannot alter the username. After the user submits the correct static password, then access to the SSL VPN portal is granted.

Summary

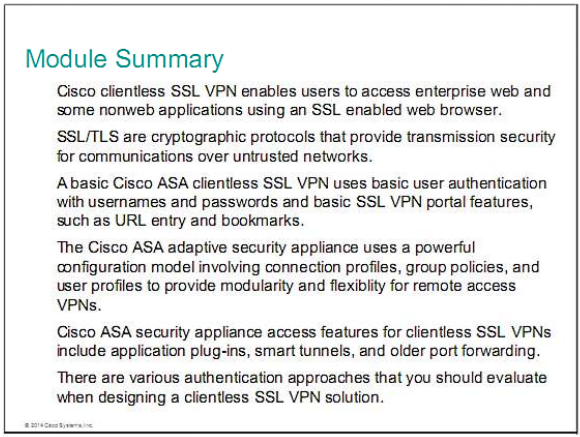
This topic summarizes the key points that were discussed in this lesson.



https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Summary

This topic summarizes the key points that were discussed in this module.



https://t.me/learningnets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Module Self-Check

Questions

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

1. From the following, choose the required and the optional authentication method in the SSL session establishment phases. (Choose two.) (Source: Clientless SSL VPN Overview)

   A. Server authentication is required.
   B. Client authentication is optional.
   C. Client authentication is required.
   D. Server authentication is optional.

2. Which of the following Cisco ASA configuration constructs define the postlogin settings? (Source: Clientless SSL VPN Overview)

   A. group policies
   B. connection profiles
   C. group profiles
   D. user profiles

3. What does Cisco Clientless SSL VPN on ASA use to authenticate an SSL VPN server to clients? (Choose two.) (Source: Deploying Basic Cisco Clientless SSL VPN on Cisco ASA)

   A. self-signed identity certificate.
   B. CA-signed identity certificate.
   C. local user database on Cisco ASA.
   D. local user database on Cisco IOS router.

https://t.me/learningnets

4. Which of the following identity certificate methods is recommended to authenticate a SSL VPN client? (Source: Deploying Basic Cisco Clientless SSL VPN on Cisco ASA)

   A. temporary self-signed certificate.
   B. configurable persistent self-signed certificate.
   C. PKI-provisioned certificate.
   D. default generated certificate.

5. What are the simplest SSL VPN client authentication local passwords? (Choose two.) (Source: Deploying Basic Cisco Clientless SSL VPN on Cisco ASA)

   A. local user database passwords.
   B. locally configured static passwords.
   C. configurable persistent self-signed passwords.
   D. temporary self-signed passwords.

6. Which of the following clientless SSL VPN portal applications are supported by default? (Choose two.) (Source: Deploying Basic Cisco Clientless SSL VPN on Cisco ASA)

   A. web browsing using HTTP
   B. web browsing using HTTPS
   C. web browsing using HTTP and HTTPS
   D. file share using CIFS and FTP
   E. file share using CIFS
   F. file share using FTP

7. Which of the following access methods would you use in clientless SSL VPNs for a custom Windows application that is transported over TCP? (Source: Deploying Application Access in Cisco ASA Clientless SSL VPN)

   A. plug-ins
   B. smart tunnels
   C. port forwarding
   D. applets

8. Which of the following access methods would you use in clientless SSL VPNs for a Linux application that is transported over TCP? (Source: Deploying Application Access in Cisco ASA Clientless SSL VPN)

   A. plug-ins
   B. smart tunnels
   C. port forwarding
   D. applets

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

9. Which of the following statements are true for the Cisco clientless remote access SSL VPN architecture? (Choose two.) (Source: Clientless SSL VPN Overview)

   A. It offers the complete network access that is provided by full-tunnel VPNs.
   B. It allows remote users to access enterprise applications only from enterprise-managed endpoints.
   C. It can traverse most firewalls and NAT devices.
   D. It enables a remote user to use a web browser to access internal resources.
   E. It provides secure remote access to internal resources, including web-enabled applications and other native client-server applications, by default.
   F. The VPN gateway for clientless SSL VPN must be a Cisco ASA security appliance.

10. Which of the following options are available for authenticating clientless SSL VPN user sessions in the Cisco implementation of clientless SSL VPN? (Choose four.) (Source: Deploying Advanced Authentication and Authorization in Clientless SSL VPN)

   A. client authentication using a local CA on the Cisco ASA
   B. client authentication using an external CA
   C. client authentication by an external AAA server, such as a RADIUS server
   D. client authentication by a local Cisco ISE module on the Cisco ASA
   E. multiple client authentication using combinations of certificate-based and AAA server-based authentication

https://t.me/learningnets

Answer Key

1. A, B
2. A
3. A, B
4. C
5. A, B
6. C, D
7. B
8. C
9. C, D
10. A, B, C, E

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Glossary

Term	Definition
AAA	authentication, authorization, and accounting. Pronounced "triple a."
ACE	access control entry.
ACL	access control list. A list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router).
AD	Active Directory
AES	Advanced Encryption Standard.
AH	Authentication Header. A security protocol that provides data authentication and optional antireplay services. AH is embedded in the data to be protected (a full IP datagram).
ATM	Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3.
BDR	backup designated router.
BGP	Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.
CA	certification authority. Entity that issues digital certificates (especially X.509 certificates) and vouches for the binding between the data items in a certificate.
CA	certification authority. Entity that issues digital certificates (especially X.509 certificates) and vouches for the binding between the data items in a certificate.
CBC	Cipher Block Chaining
CIFS	Common Internet File System.
CLI	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.
CRL	certificate revocation list. Data structure that enumerates digital certificates that have been invalidated by their issuer prior to when they were scheduled to expire.
CRL	certificate revocation list. Data structure that enumerates digital certificates that have been invalidated by their issuer prior to when they were scheduled to expire.
DAP	dynamic access policy.

https://t.me/learningnets

Term	Definition
DES	Data Encryption Standard. Standard cryptographic algorithm developed by the U.S. National Bureau of Standards.
DF	default forwarding. A differentiated services per-hop behavior that provides best-effort forwarding and is used for all traffic that does not match any other defined traffic class.
DH	Diffie-Hellman. The Diffie-Hellman algorithm, introduced by Whitfield Diffie and Martin Hellman in 1976, was the first system to utilize public key or asymmetric cryptographic keys. Today, Diffie-Hellman is part of the IPsec standard. A protocol known as OAKLEY uses Diffie-Hellman, as described in RFC 2412. OAKLEY is used by the Internet Key Exchange (IKE) protocol (see RFC 2401), which is part of the overall framework called Internet Security Association and Key Management Protocol (ISAKMP; see RFC 2408).
DHCP	Dynamic Host Configuration Protocol. Provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them.
DMVPN	Dynamic Multipoint VPN.
DMVPN	Dynamic Multipoint VPN.
DMZ	demilitarized zone.
DN	distinguished name. Global, authoritative name of an entry in the OSI Directory (X.500).
DN	Domain name
DNS	Domain Name System. System used on the Internet for translating names of network nodes into addresses.
DoS	denial of service. An intentional or unintentional attack on a device that makes the resource unavailable to perform its normal function.
DPD	dead peer detection.
DR	designated router.
DSA	Digital Signature Algorithm
DSL	digital subscriber line. Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. There are four types of DSL: ADSL, HDSL, SDSL, and VDSL. All are provisioned via modem pairs, with one modem located at a central office and the other at the customer site. Because most DSL technologies do not use the whole bandwidth of the twisted pair, there is room remaining for a voice channel.
DTLS	Datagram Transport Layer Security.
DVTI	dynamic virtual tunnel interface.
EAP	Extensible Authentication Protocol. Framework that supports multiple, optional authentication mechanisms for PPP, including cleartext passwords, challenge-response, and arbitrary dialog sequences.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
EC	elliptic curve.
EC	elliptic curve.
ECC	elliptic curve cryptography.
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm. Standard [A9062] that is the elliptic curve cryptography analog of the Digital Signature Algorithm.
ECDSA	Elliptic Curve Digital Signature Algorithm. Standard [A9062] that is the elliptic curve cryptography analog of the Digital Signature Algorithm.
EIGRP	Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency, and combines the advantages of link-state protocols with those of distance vector protocols.
ESP	Encapsulating Security Payload. Security protocol that provides data privacy services, optional data authentication, and antiplay services. ESP encapsulates the data to be protected.
FFS	flash file system
FIPS 140	Federal Information Processing Standards 140 Publication Series. US government computer-security standards that specify requirements for cryptography modules.
FQDN	fully qualified domain name. FQDN is the full name of a system, rather than just its host name. For example, aldebaran is a host name, and aldebaran.interop.com is an FQDN.
FTP	File Transfer Protocol. Application protocol, part of the TCP/IP protocol stack, used for transferring files between network nodes. FTP is defined in RFC 959.
FVRF	front virtual routing and forwarding.
G2	Generation 2.
GCM	Galois Counter Mode.
GCM	Galois Counter Mode.
GDOI	Group Domain of Interpretation
GET	Group Encrypted Transport.
GRE	Generic Routing Encapsulation. Tunneling protocol that was developed by Cisco and that can encapsulate a variety of protocol packet types inside IP tunnels. This process creates a virtual point-to-point link to Cisco routers at remote points over an IP network.
GRE	Generic Routing Encapsulation. Tunneling protocol that was developed by Cisco and that can encapsulate a variety of protocol packet types inside IP tunnels. This process creates a virtual point-to-point link to Cisco routers at remote points over an IP network.

https://t.me/learningnets

Term	Definition
GUI	graphical user interface. A user environment that uses pictorial as well as textual representations of the input and the output of applications and the hierarchical or other data structure in which information is stored. Such conventions as buttons, icons, and windows are typical, and many actions are performed using a pointing device (such as a mouse). Microsoft Windows and the Apple Macintosh are prominent examples of platforms using a GUI.
HDLC	High-Level Data Link Control. Bit-oriented synchronous data link layer protocol developed by ISO. Derived from SDLC, HDLC specifies a data encapsulation method on synchronous serial links using frame characters and checksums.
HMAC	Hash-based Message Authentication Code. HMAC is a mechanism for message authentication using cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function (for example, MD5 and SHA-1) in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function.
HMAC-MD5	Hashed Message Authentication Code-Message Digest 5 (RFC 2104). A keyed version of MD5 that enables two parties to validate transmitted information using a shared secret.
HMAC-SHA1	Hashed Message Authentication Code-Secure Hash Algorithm 1.
HTTPS	Hypertext Transfer Protocol Secure
IANA	Internet Assigned Numbers Authority. Organization operated under the auspices of the ISOC as a part of the IAB. IANA delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains a database of assigned protocol identifiers used in the TCP/IP stack, including autonomous system numbers.
ICA	Independent Computing Architecture. A proprietary protocol developed by Citrix Systems, used for client-server communication.
ICMP	Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792.
IETF	Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.
IETF	Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.
IGP	interior gateway protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
IKE	Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPsec) that require keys. Before any IPsec traffic can be passed, each router, firewall, or host must verify the identity of its peer. Verification can be done by manually entering pre-shared keys into both hosts or by a CA service.
IKE	Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPsec) that require keys. Before any IPsec traffic can be passed, each router, firewall, or host must verify the identity of its peer. Verification can be done by manually entering pre-shared keys into both hosts or by a CA service.
IKE	Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPsec) that require keys. Before any IPsec traffic can be passed, each router, firewall, or host must verify the identity of its peer. Verification can be done by manually entering pre-shared keys into both hosts or by a CA service.
IM	instant messaging.
IP address	32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address.
IPsec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.
IPsec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

© 2014 Cisco Systems, Inc.
https://t.me/learningnets

Term	Definition
IPv4	IP version 4. Internet Protocol version 4 is the fourth version in the development of the IP and the first version of the protocol to be widely deployed. Along with IPv6, IPv4 is at the core of standards-based internetworking methods of the Internet. IPv4 is still used to route most traffic across the Internet. IPv4 is a connectionless protocol for use on packet-switched link layer networks (for example, Ethernet). It operates on a best-effort delivery model in that it does not guarantee delivery and does not assure proper sequencing or avoidance of duplicate delivery.
IPv4	IP version 4. Internet Protocol version 4 is the fourth version in the development of the IP and the first version of the protocol to be widely deployed. Along with IPv6, IPv4 is at the core of standards-based internetworking methods of the Internet. IPv4 is still used to route most traffic across the Internet. IPv4 is a connectionless protocol for use on packet-switched link layer networks (for example, Ethernet). It operates on a best-effort delivery model in that it does not guarantee delivery and does not assure proper sequencing or avoidance of duplicate delivery.
IPv6	IP version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation).
IPv6	IP version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation).
ISAKMP	Internet Security Association and Key Management Protocol. Internet IPSec protocol [RFC 2408] that negotiates, establishes, modifies, and deletes security associations. It also exchanges key generation and authentication data (independent of the details of any specific key generation technique), key establishment protocol, encryption algorithm, or authentication mechanism.
ISDN	Integrated Services Digital Network. Communication protocol offered by telephone companies that permits telephone networks to carry data, voice, and other source traffic.
ISP	Internet service provider. Company that provides Internet access to other companies and individuals.
JAR	Java archive
L2TP	Layer 2 Tunnel Protocol. An Internet Engineering Task Force (IETF) standards track protocol defined in RFC 2661 that provides tunneling of PPP. Based upon the best features of L2F and PPTP, L2TP provides an industry-wide interoperable method of implementing VPDN.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
LAN	local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
LDAP	Lightweight Directory Access Protocol. Protocol that provides access for management and browser applications that provide read/write interactive access to the X.500 Directory.
LSA	link-state advertisement. Broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables. Sometimes called an LSP.
MAC	message authentication code
MD5	Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and Secure Hash Algorithm (SHA) are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. Cisco uses hashes for authentication within the IPsec framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness.
MDM	Mobile Device Management
mGRE	multipoint Generic Routing Encapsulation.
MPLS	Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information.
MSS	maximum segment size. A TCP protocol parameter that specifies the maximum amount of data that a device can receive in a single TCP segment.
MTU	maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.
NAT	Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. Also known as Network Address Translator.
NAT-T	NAT Traversal.
NBMA	nonbroadcast multiaccess. Term describing a multiaccess network that either does not support broadcasting (such as X.25) or in which broadcasting is not feasible (for example, an SMDS broadcast group or an extended Ethernet that is too large).
NBNS	NetBIOS Name Service.

https://t.me/learningnets

Term	Definition
NetFlow	A feature of some routers that allows them to categorize incoming packets into flows. Because packets in a flow often can be treated in the same way, this classification can be used to bypass some of the work of the router and accelerate its switching operation.
NGE	next-generation encryption.
NGE	next-generation encryption.
NHO	next-hop override
NHRP	Next Hop Resolution Protocol. Protocol used by routers to dynamically discover the MAC address of other routers and hosts connected to an NBMA network. These systems then can communicate directly without requiring traffic to use an intermediate hop, increasing performance in ATM, Frame Relay, SMDS, and X.25 environments.
NIST	National Institute of Standards and Technology. U.S. government organization that supports and catalogs a variety of standards. Formerly the NBS.
NSA	U.S. National Security Agency.
NTLM	NT LAN Manager. Also known as Windows Challenge/Response, NT LAN Manager is the authentication protocol that is used on Windows systems and networks.
OCSP	Online Certificate Status Protocol.
ODR	On-Demand Routing
OSI	Open Systems Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.
OSPF	Open Shortest Path First. Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing. OSPF was derived from an early version of the IS-IS protocol.
OTP	one-time password
PCMCIA	Personal Computer Memory Card International Association. Standard used for credit-card-sized computer peripherals. Type I devices are very thin memory cards, Type 2 devices include most modems and interfaces, and Type 3 devices are used for disk drives and thicker components.
PFS	Perfect Forward Secrecy.
PIN	places in the network

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
PKI	public-key infrastructure. System of CAs (and, optionally, RAs and other supporting servers and agents) that perform some set of certificate management, archive management, key management, and token management functions for a community of users in an application of asymmetric cryptography.
POP	Post Office Protocol. Protocol that client email applications use to retrieve mail from a mail server.
PPP	Point-to-Point Protocol. Successor to SLIP that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: LCP and NCP.
PRF	Pseudo-Random Function
PSK	pre-shared key. Shared secret key that is used during IKE authentication.
PSK	pre-shared key. Shared secret key that is used during IKE authentication.
QoS	quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.
RA	Registration Authority
RADIUS	Remote Authentication Dial-In User Service. Database for authenticating modem and ISDN connections and for tracking connection time.
RDP	Remote Desktop Protocol.
RFC	Request for Comments. Document series that is used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some RFCs are humorous or historical. RFCs are available online from numerous sources.
RIP	Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.
ROM	read-only memory. Nonvolatile memory that can be read, but not written, by the microprocessor.
RSA	Acronym stands for Rivest, Shamir, and Adleman, the inventors of the technique. Public-key cryptographic system that can be used for encryption and authentication.
S/MIME	Secure/Multipurpose Internet Mail Extensions.
SA	security association. Instance of security policy and keying material applied to a data flow.

https://t.me/learningnets

Term	Definition
SA	security association. Instance of security policy and keying material applied to a data flow.
SCEP	Simple Certificate Enrollment Protocol.
SEAL	simple and efficient AAL. Scheme used by AAL5 in which the SAR sublayer segments CS PDUs without adding additional fields.
SHA	Secure Hash Algorithm.
SHA-1	Secure Hash Algorithm 1. Algorithm that takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest provides security against brute-force collision and inversion attacks. SHA-1 [NIS94c] is a revision to SHA that was published in 1994.
SHA-256	Secure Hash Algorithm. SHA-256 is part of the SHA-2 set of cryptographic hash functions.
SMTP	Simple Mail Transfer Protocol. Internet protocol providing email services.
SNMP	Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.
SPI	security parameter index. This is a number that, together with a destination IP address and security protocol, uniquely identifies a particular security association. When using IKE to establish the security associations, the SPI for each security association is a pseudo-randomly derived number. Without IKE, the SPI is manually specified for each security association.
SSH	Secure Shell Protocol. Protocol that provides a secure remote connection to a router through a TCP application.
SSL	Secure Socket Layer. Encryption technology for the Web used to provide secure transactions, such as the transmission of credit card numbers for e-commerce.
TACACS+	Terminal Access Controller Access Control System Plus. Proprietary Cisco enhancement to TACACS. Provides additional support for authentication, authorization, and accounting.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.
TCP/IP	Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite.

https://t.me/learningnets
© 2014 Cisco Systems, Inc.

Term	Definition
Telnet	Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log into remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854.
TFTP	Trivial File Transfer Protocol. Simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, username and password).
TLS	Transport Layer Security. A future IETF protocol to replace SSL.
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
URL	uniform resource locator. Type of formatted identifier that describes the access method and the location of an information resource object on the Internet. [RFC 1738]
VNC	Virtual Network Computing.
VoIP	Voice over Internet Protocol. A way to carry phone calls over an IP data network, whether on the Internet or an internal network. A primary attraction of VoIP is its ability to reduce expenses, because phone calls travel over the data network rather than over the phone company network.
VoIP	Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323.
VPN	virtual private network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses tunneling to encrypt all information at the IP level.
VRF	A VPN routing/forwarding instance. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.
VTI	virtual tunnel interface.
WAN	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.

https://t.me/learningnets

Term	Definition
Winsock2	Window Sockets 2
XML	extensible markup language. A standard maintained by the World Wide Web Consortium (W3C). It defines a syntax that lets you create markup languages to specify information structures. Information structures define the type of information, for example, subscriber name or address, not how the information looks (bold, italic, and so on). External processes can manipulate these information structures and publish them in a variety of formats. Text markup language designed to enable the use of SGML on the World Wide Web. XML allows you to define your own customized markup language.
3DES	Triple Data Encryption Standard.
3DES-EDE	Triple Data Encryption Standard-Encrypt-Decrypt-Encrypt

https://t.me/learningnets
© 2014 Cisco Systems, Inc.