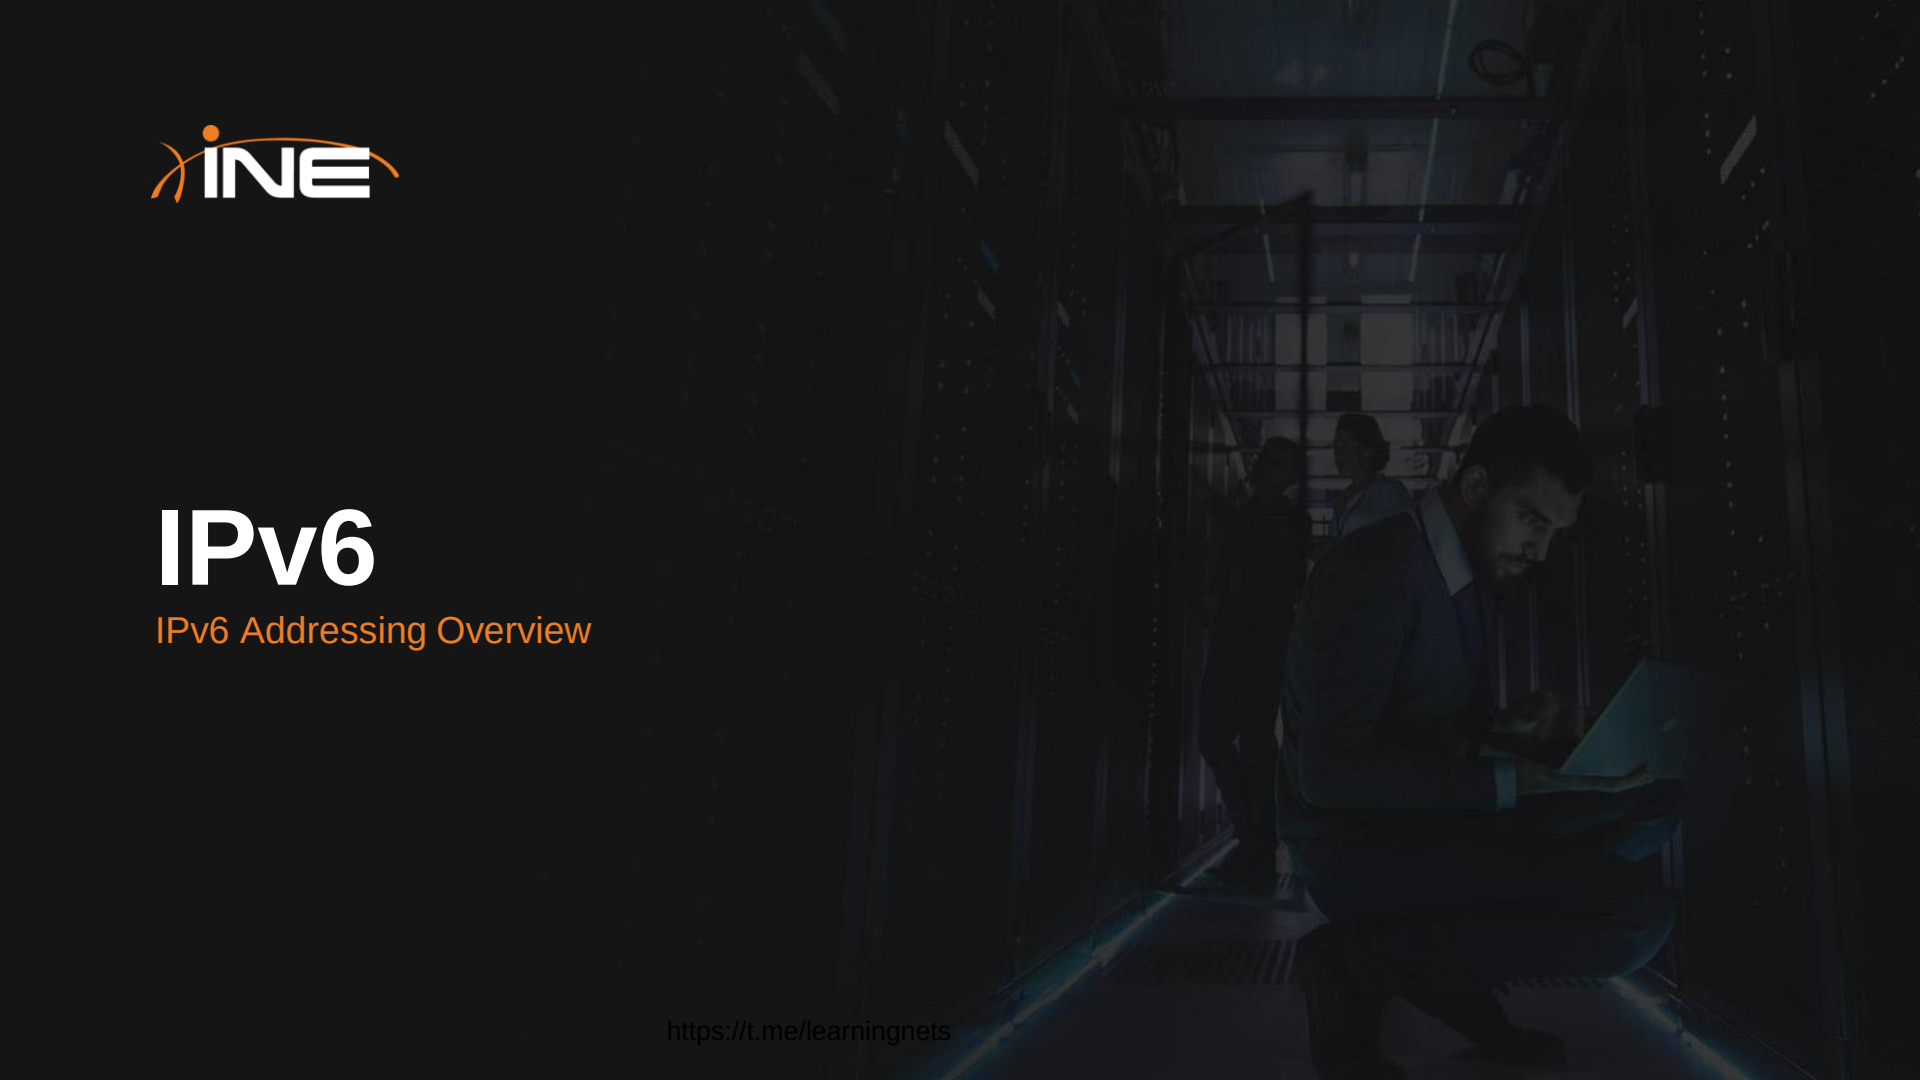

IPv6

IPv6 Addressing Overview

In This Section

- Recommended IPv6 Resources
- IPv6 Addressing Format
  - Global, link-local, site-local, unique local, EUI-64
- IPv6 Neighbor Discovery
- IPv6 SLAAC
- IPv6 General Prefix

Recommended IPv6 Resources

- Books
  - IPv6 Theory Protocol & Practice
  - Microsoft Press - Understanding IPv6
- Online Resources
  - Microsoft MSDN Knowledgebase
    - Internet Protocol (IP) for IPv6
  - Cisco IPv6 Knowledge Base Portal

IPv6 Overview

- RFC 2460 - Internet Protocol, Version 6 (IPv6) Specification
- Much larger address space than IPv4
  - IPv4 uses 4 Byte (32 Bit) Addresses
    - 2^{32} = 4,294,967,296 addresses
  - IPv6 uses 16 Byte (128 Bit) Addresses
    - 2^{128} ~ 340 undecillion (3.4×1038)

IPv4 vs. IPv6 Addressing Format

- IPv4 Dotted Decimal
  - d.d.d.d
  - d = one byte
- IPv6 Hexadecimal
  - hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh
  - hh = one byte

IPv6 Address Types

- RFC 4291 - IP Version 6 Addressing Architecture
  - Section 2.4. Address Type Identification
    - *The type of an IPv6 address is identified by the high-order bits of the address, as follows*

IPv6 Link Local Addresses

- Addresses locally significant to a link
  - FE80::/10 (1111 1110 10)
  - Equivalent of 169.254.0.0/16 in IPv4
- Never routable between interfaces
- Used for…
  - Stateless Address Auto-configuration (SLAAC)
  - Neighbor Discovery
  - Router Discovery

IPv6 Site Local Addresses

- Addresses locally significant to a “site”
  - FEC0::/10 (1111 1110 11)
- No one could agree on what a “site” is
  - RFC 3879 - Deprecating Site Local Addresses

IPv6 Unique Local Addresses

- RFC 4193 - Unique Local IPv6 Unicast Addresses
- ULA is Private Use IPv6 addressing
  - FC00::/7 (1111 110)
- Equivalent to RFC 1918
  - 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
- Likely unique but not routable via global BGP

IPv6 Global Unicast Addresses

- Technically everything else
  - IANA currently allocating 2000::/3
- Per RFC end hosts must…
  - Have 64-bit interface ID
    - nnnn:nnnn:nnnn:nnnn:hhhh:hhhh:hhhh:hhhh/64
  - Use EUI-64 format for interface ID

Modified EUI-64 Addressing

- RFC 4291 - IP Version 6 Addressing Architecture
  - Appendix A: Creating Modified EUI-64 Format Interface Identifiers
- Ethernet MAC to EUI-64 conversion
  - Invert Universal/Local (U/L) bit
    - 7th most significant bit
  - Insert padding 0xFF 0xFE in the middle

IPv6 Temporary Addresses

- RFC 4941 - Privacy Extensions for Stateless Address Autoconfiguration in IPv6
- MSDN – IPv6 Interface Identifiers
  - Randomly generated interface identifiers
- Prevents tracking users based on MAC address
  - I.e. EUI-64 result is always the same for a given MAC address

IPv6 Address Resolution

- RFC 4861 - Neighbor Discovery for IP version 6 (IPv6)
- ICMPv6 ND used for layer 3 to layer 2 resolution
  - Equivalent of IPv4 ARP

IPv6 Neighbor Discovery Messages

- Neighbor Solicitation (NS)
  - Equivalent of ARP request
- Neighbor Advertisement (NA)
  - Equivalent of ARP reply
- Router Solicitation (RS)
- Router Advertisement (RA)

IPv6 SLAAC

- RFC 4862 – IPv6 Stateless Address Autoconfiguration
- Automatically assigns IPv6 address for every on-link prefix
  - Works only with /64s on the link
  - Host uses Duplicate Address Detection (DAD) to verify uniqueness of generated address
- SLAAC is not DHCPv6
  - Does not include options, e.g. DNS server

SLAAC and DHCPv6

- SLAAC commonly works along with DHCPv6
  - Options are set in Router Advertisement (RA) messages
  - Tells the host that there is a DHCPv6 server available for addressing options
- Other-Config-Flag
  - Use DHCPv6 to receive just addressing options (DNS, TFTP, etc.)
  - **ipv6 nd other-config-flag**
- Managed-Config-Flag
  - Use DHCPv6 for both addressing and options
  - **ipv6 nd managed-config-flag**

ND For Host Default Routing

+ Hosts don’t receive default gateway from DHCPv6
+ First Hop IPv6 router deduced from ND RA Messages

IPV6 General Prefix

- General prefix can be defined to act as shortcut
  - E.g. If organization is assigned a /32, then all prefixes should be derived from this /32
  - Helps in renumbering scenarios
- Defined globally as
  - **ipv6 general-prefix [NAME_OF_PREFIX] 2001:123::/32**
- Applied to link as
  - **ipv6 address [NAME_OF_PREFIX] ::1/64**
- General prefix and subnet are merged to derive full address

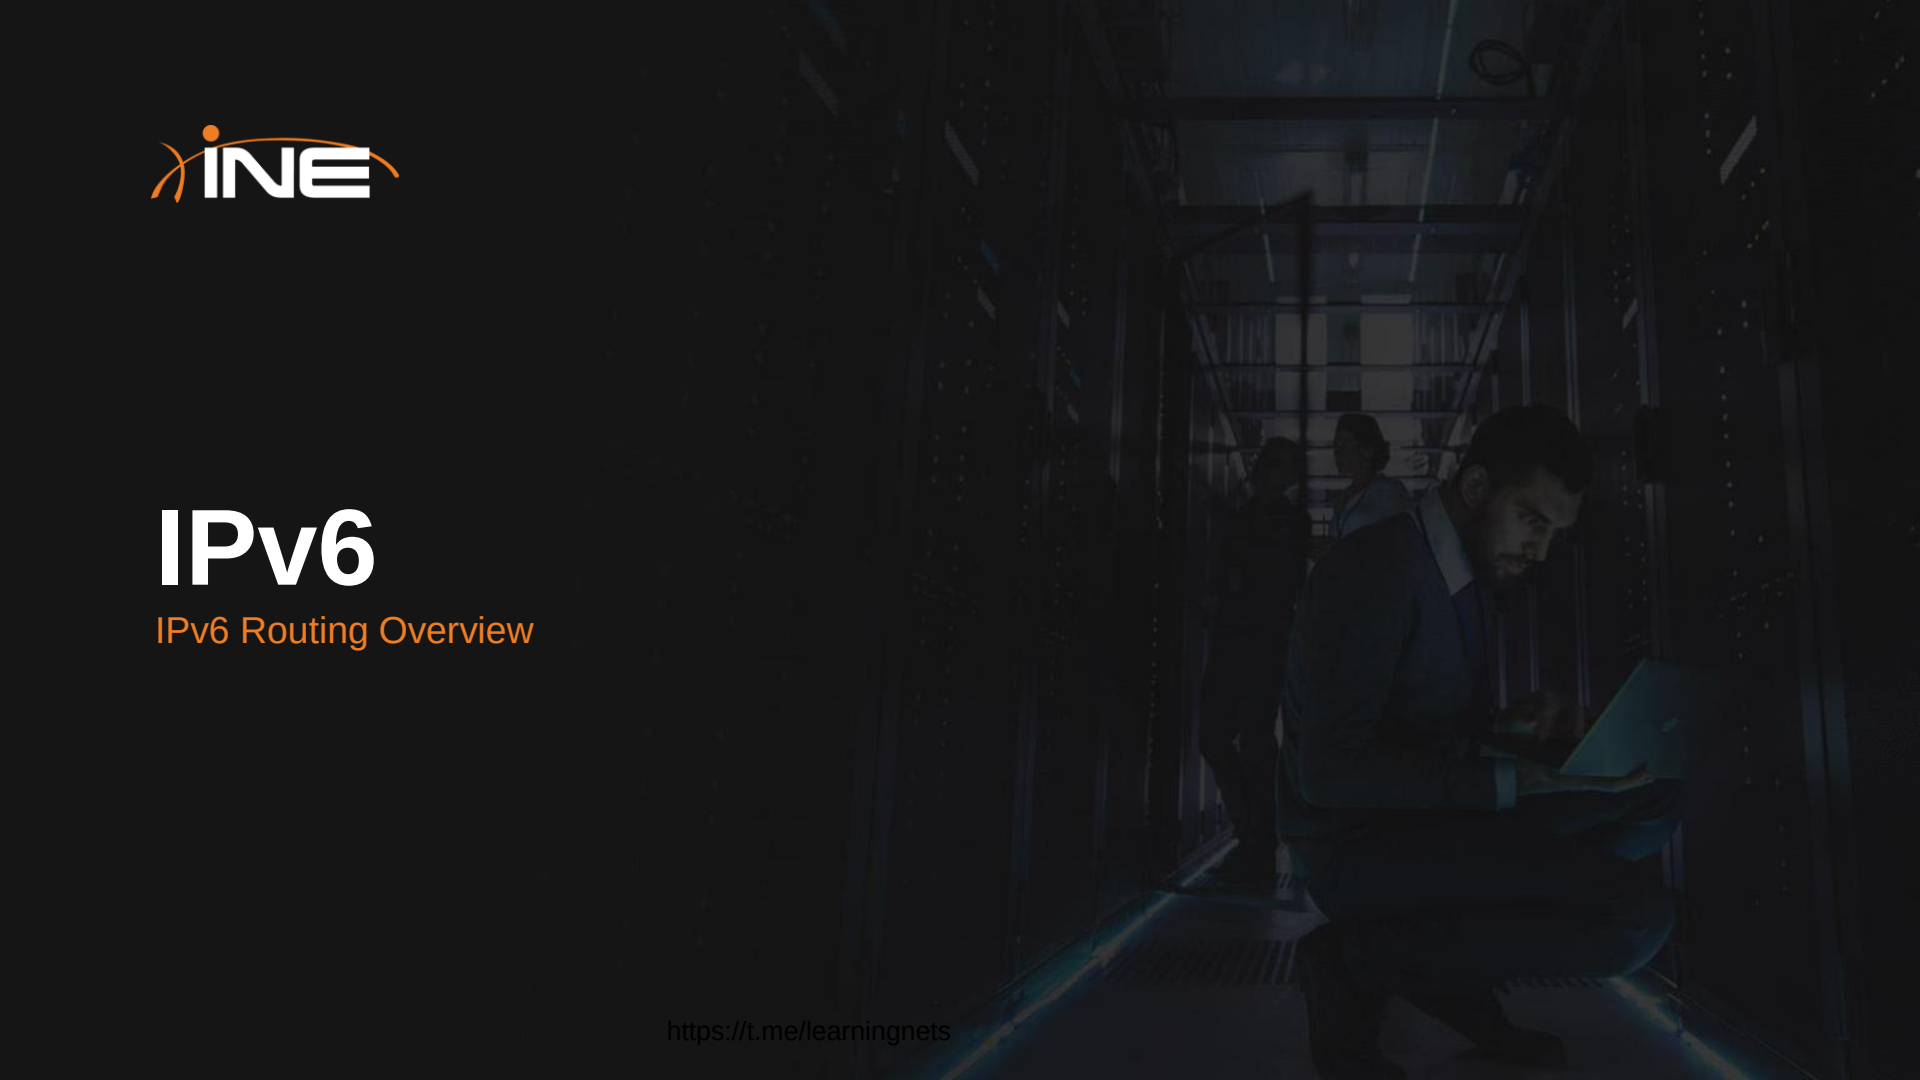

IPv6

IPv6 Routing Overview

In This Section

+ IPv6 Routing Overview
+ IPv6 Static Routing
+ IPv6 VRF Support

IPv6 Routing Overview

- IPv6 routing off by default
  - Enabled with **ipv6 unicast-routing**
  - Enables routing and ICMPv6 ND RAs
- IPv6 routing is supported by all protocols
  - Static, RIPng, EIGRPv6, OSPFv3, IS-IS, MP-BGP, & Policy Routing
- Dynamically learned routes recurse to remote link-local
  - Implies Global Unicast isn’t required
  - Implies static layer 3 to layer 2 resolution for Link Local is needed on multipoint NBMA

IPv6 Static Routing

- Same static routing implications as IPv4
  - Routing to a next-hop
    - Resolve layer 2 address of next-hop
  - Routing to a multipoint interface
    - Resolve layer 2 address of final destination
    - Proxy ND and Proxy IND design issues
  - Routing to a point-to-point interface
    - No layer 2 resolution required

VRF Support

- VRFs can include both IPv4 and IPv6 interfaces
  - Not all dynamic protocols are IPv6 VRF aware in all versions
- Must use “vrf definition [name]” syntax, and enable the IPv6 address family
  - **vrf upgrade-cli** also an option
- Interface Command
  - **vrf forwarding [VRF Name]**
  - Places the interface and both IPv4/IPv6 in specified VRF

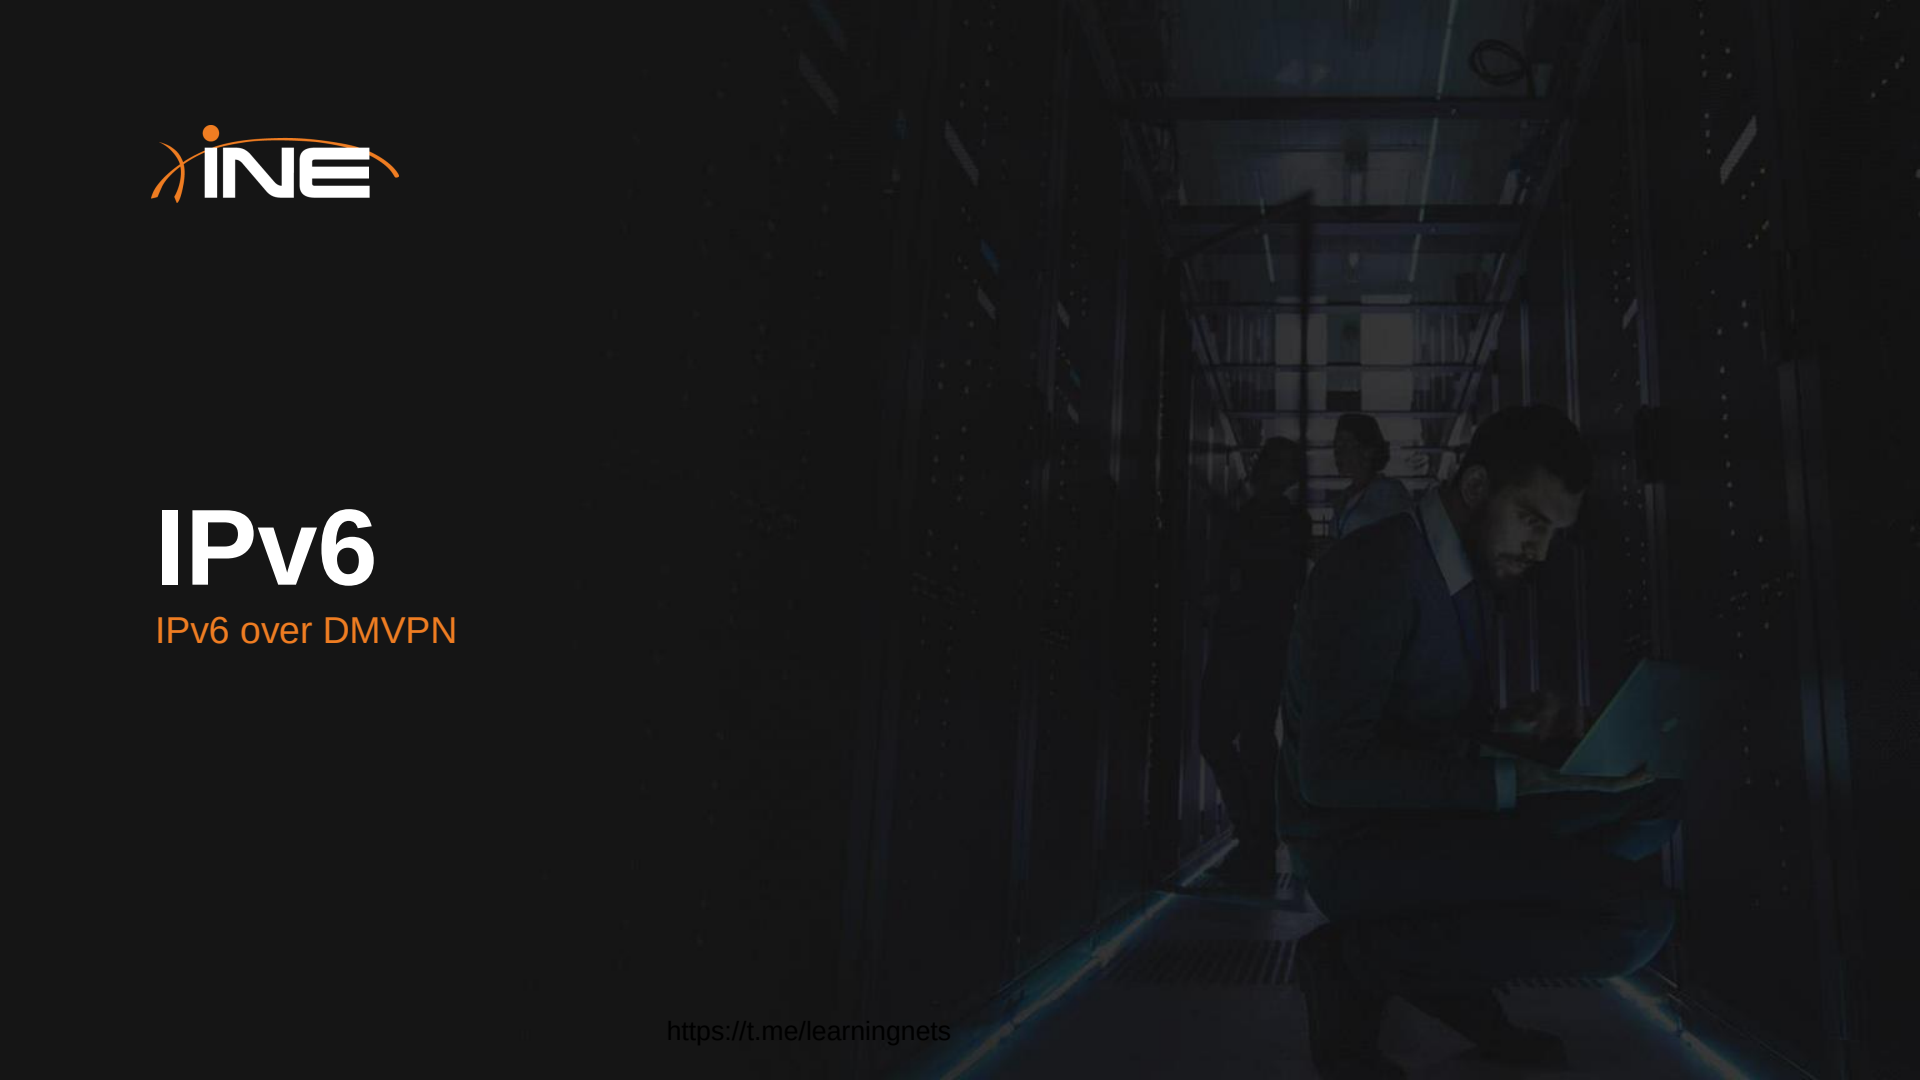

IPv6

IPv6 over DMVPN

In This Section

+ IPv6 over IPv4 DMVPN
+ DMVPN Over IPv6 Transport

IPv6 over IPv4 DMVPN

- DMVPN uses IPv4 mGRE for transport
  - GRE is a multi-protocol encap
  - Implies IPv6 can be supported
- How IPv6 over IPv4 DMVPN works
  - Build the IPv4 underlay and overlay
    - E.g. IPsec spoke to hub
  - Use IPv6 NHRP to resolve IPv6 IGP or IPv6 BGP

DMVPN Over IPv6 Transport

- DMVPN can also use IPv6 mGRE for transport
  - Eventually IPv4 will be the legacy protocol
  - Transport will be native IPv6
- Configured as…
  - **tunnel mode gre multipoint ipv6**
  - Uses IKEv2 for optional IPsec

IPv6

EIGRPv6

In This Section

- EIGRPv6 Overview
- EIGRPv6 Configuration

EIGRPv6 Overview

- Similar in operation to IPv4 EIGRP
  - Transport via protocol 88 to unicast and multicast FF02::A
- “shutdown” by default in EIGRPv6 Classic Mode
  - “no shutdown” under process
- Enabled on all IPv6 links in EIGRPv6 Named Mode
  - **address-family ipv6 unicast <VRF> autonomous-system <ASN>**
  - Process is NOT shutdown by default
  - To exclude a link:
    - “af-interface X”
    - “shutdown”
- Like OSPFv3 & BGP, IPv4 formatted Router-ID needed

IPv6

OSPFv3

In This Section

+ OSPFv3 Overview
+ OSPFv3 Configuration

OSPFv3 Overview

- RFC 5340 - OSPF for IPv6
- Similar in many operations to OSPFv2
  - Transport via protocol 89 to unicast and multicasts FF02::5 & FF02::6
- Normal OSPF rules still apply
  - Adjacency parameters
  - OSPFv3 network types

OSPFv3 LSAs 8 & 9

- New LSAs introduced in OSPFv3
  - Type 8 (Link) - Used for Link Local addresses
  - Type 9 (Intra Area Prefix) - Used for prefixes on the links
- New LSAs separate topology graph from NLRI
  - OSPFv2 has subnet info in LSAs 1 & 2 for Intra-Area
    - If prefix add or remove, run full SPF
  - OSPFv3 uses LSAs 8 & 9 to reference LSA 1
    - If a stub network is add or remove, full SPF not required

OSPFv3 Configuration

- Enabled at link level
  - **ipv6 ospf [process-id] area [area-id]**
  - Automatically enables global process
- Like EIGRPv6 & BGP, IPv4 formatted Router-ID needed

OSPFv3 IPsec Authentication

- OSPFv3 uses IPsec for…
  - Authentication
    - IPsec AH or ESP
  - Encryption
    - IPsec ESP
- No ISAKMP support, so keys must be manually configured
  - Copy and paste is your friend

Multiprotocol OSPFv3

- OSPFv3 can advertise both IPv4 & IPv6 NLRI
  - Two separate trees, v4 and v6 run independently but use same data structures
  - Similar to ISIS Multi-Topology
- To advertise IPv4 NLRI, must have both v4 and v6 on the link
  - “ipv6 enable” on transit links is the minimum required for v6
  - For IPv4 – can have all links unnumbered to Loopback 0 (arbitrary)

Multiprotocol OSPFv3 Configuration

- Enabled on Link-Level
  - ospfv3 [process-id] [ipv4 | ipv6] area [area num]
- Options under OSPFv3 routing process
  - router ospfv3 [process-id]
- Show commands
  - show ospfv3…

IPv6

IPv6 BGP

In This Section

+ IPv6 BGP Overview
+ IPv6 BGP Configuration

IPv6 BGP Overview

- RFC 2545 - Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
  - Uses MP_REACH_NLRI and MP_UNREACH_NLRI for IPv6
  - AFI = 2 (IPv6) SAFI = 1 (unicast)
- Normal BGP rules apply
  - Requires underlying IGP for TCP transport
  - EBGP loop prevention via AS-Path
  - iBGP loop prevention via full mesh, RRs, Confederation
  - Same attributes and best-path selection process

IPv6 BGP Transport

- Transport and NLRI are independent
  - Transport can be IPv4 or IPv6
  - NLRI advertised via AFI 2 SAFI 1
    - address-family ipv6 unicast
- Advertising IPv6 NLRI via IPv4 BGP session results in next-hop issues
  - Can be manually fixed with a route-map
  - Also **no bgp default ipv6-nexthop**
    - If peer is directly connected, use IPv6 address of the peer (not the encoded ::FFFF: address)

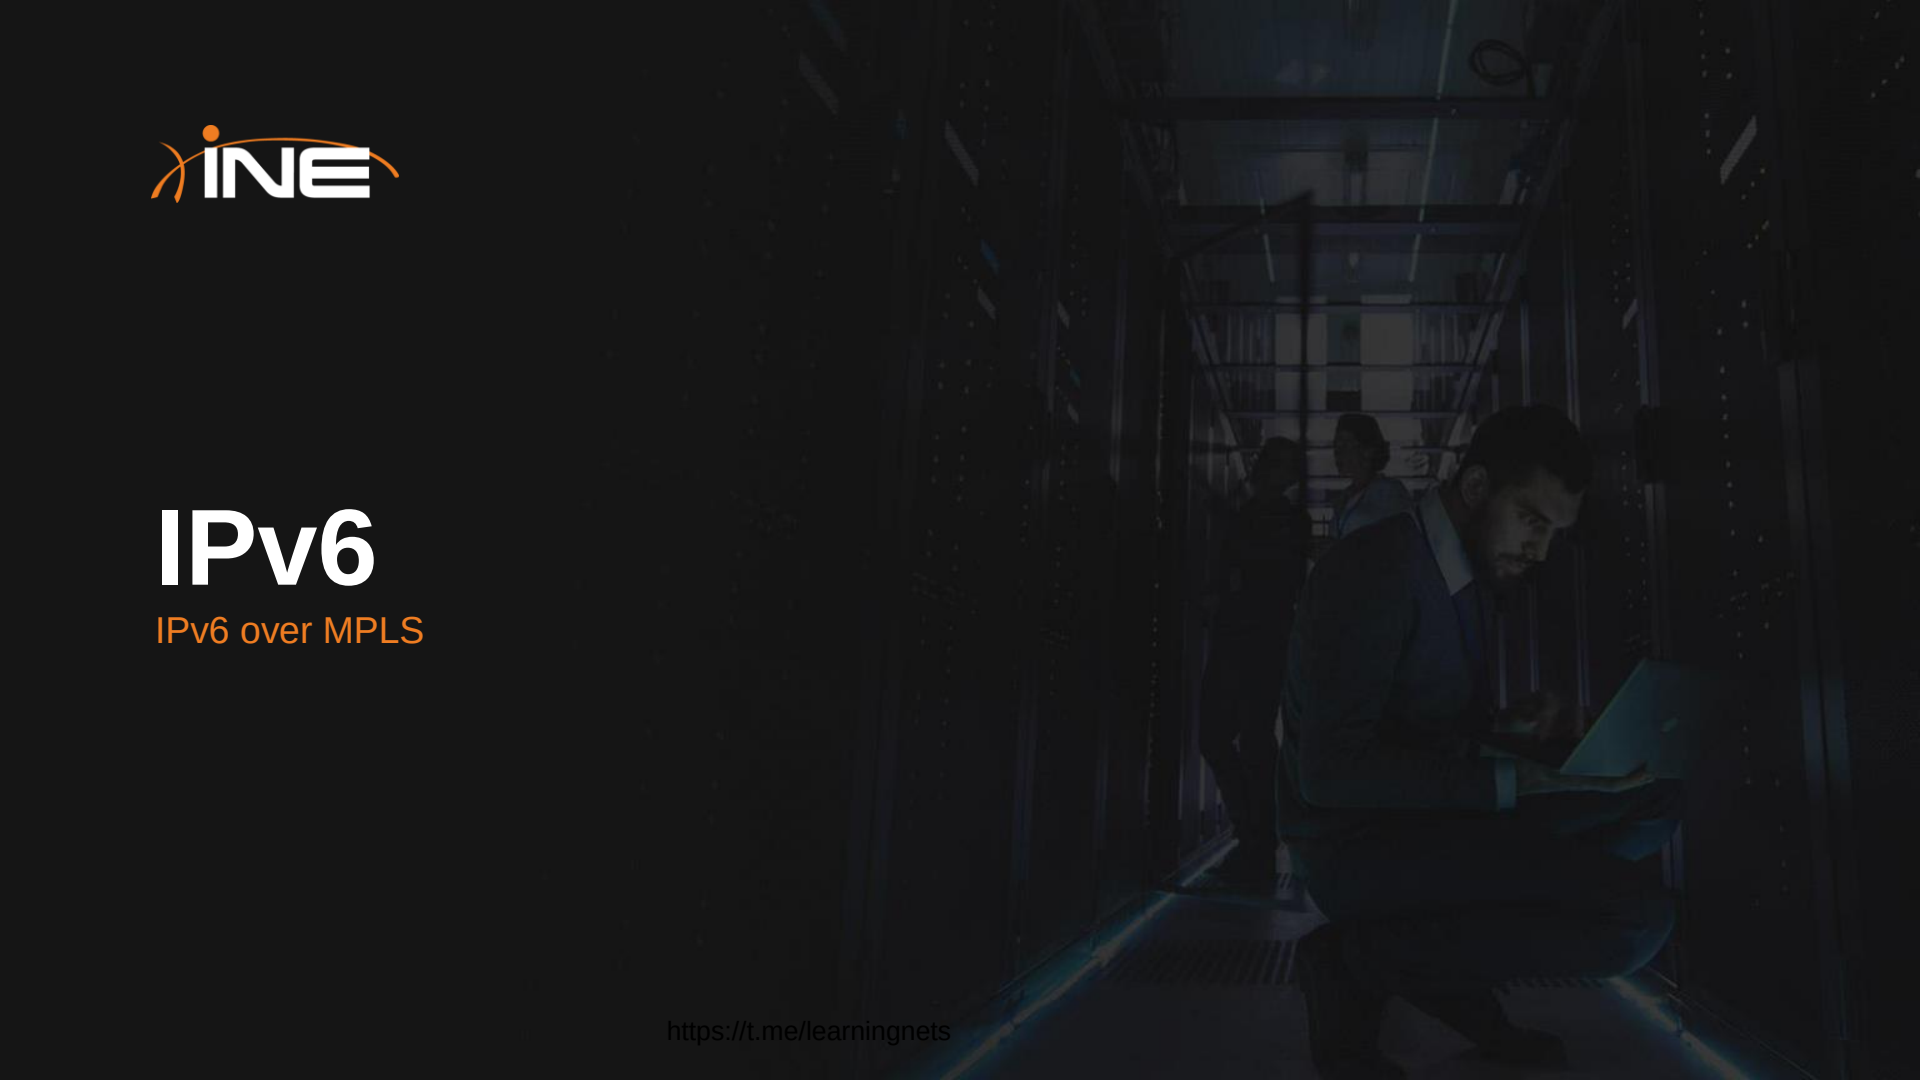

IPv6

IPv6 over MPLS

In This Section

- IPv6 over MPLS
  - 6PE
  - 6VPE

6PE

- Provider Edge Router (6PE) over MPLS
  - Allows tunneling of traffic in global IPv6 table over IPv4 MPLS core
- How 6PE Works
  - Form IPv6 BGP peering from PE to CE in global
  - Form IPv6 BGP + Label peering from PE to PE in global
  - Traffic in core follows IPv4 label path between PEs

6VPE

- IPv6 VPN Provider Edge Router (6VPE) over MPLS
  - Tunneling of traffic in IPv6 VRF table over IPv4 MPLS core
  - Similar to regular MPLS L3VPN
- How 6VPE Works
  - Form IPv6 peering from PE to CE in VRF
  - Form VPNv6 BGP peering from PE to PE in global
  - Traffic in core follows IPv4 label path between PEs